David Rosenboom
Center for Contemporary Music
Mills College
Oakland, California 94613 USA

The Performing Brain

Introduction

The primary purposes of this article are twofold: (1) to describe work done in the 1970s on the spontaneous generation of musical structures with a computer music system by using a detailed analysis of a performer's brainwaves, and (2) to speculate on extensions of these ideas. Aside from its historical interest, this work has resulted in ideas for high-level musical input structures—new ways of playing intelligent, programmable musical instruments. Furthermore, my own work in biofeedback and the arts, begun over twenty years ago, is experiencing a revival due to the fact that advances in technology now permit realization of musical concepts in performance that depend on complex, real-time analysis of electroencephalogram (EEG) signals, previously achievable only with cumbersome, non-real-time, laboratory-bound methods. Consequently, ideas that were impractical when they were proposed many years ago are now practical.

It is beyond the scope of this article to include an explanation of the principles and techniques of EEG analysis on which much of the work depends. I refer the reader to a recent monograph entitled *Extended Musical Interface with the Human Nervous System: Assessment and Prospectus* (Rosenboom 1990). In this monograph, a great deal more information is presented concerning (1) biofeedback modeling and its history both inside and outside of the arts; (2) the varieties of bioelectromagnetic phenomena that have been explored in feedback paradigms; (3) a detailed model for classification of EEG phenomena with particular emphasis on event-related potentials (ERPs) and their significance to the study of the mechanisms of attention in musical experience; and (4) applications of new technology for sensing biomagnetism, such as superconducting quantum interference devices (SQUIDs), multichannel brain imaging, and other developments in hardware and software arenas. In this paper I will concentrate on algorithms developed to facilitate the evolution in performance of a musical structure in response to shifts in selective attention, as evidenced by phenomena detected in the electroencephalogram (EEG).

Computer Music Journal, Vol. 14, No. 1, Spring 1990,

Background

In a now famous paper published in 1934 the pioneering physiologists Adrian and Mathews reported on experiencing a translation of the human electroencephalogram (EEG) into audio signals. While listening to his own alpha rhythm presented through a loudspeaker, Adrian tried to correlate the subjective impression of hearing the alpha come and go with the activity of looking or not looking with his eyes (Adrian and Mathews 1934). Inevitably, artists with an experimental bent would come to apply this—and subsequent developments in brain science—to both artistic production and research in artistic perception. During the past 25 years, composers and artists like Alvin Lucier, Richard Teitelbaum, myself, and numerous others have produced major works of music, as well as visual and kinetic art using EEG and other bioelectronic signals. Lucier's 1965 work *Music for Solo Performer* achieved a direct mapping of a soloist's alpha rhythms onto the orchestrational palette of a percussion ensemble (Lucier 1976, 1982; Lucier and Simon 1980). Teitelbaum's *Organ Music* and *In Tune*, both realized in 1968, added heartbeat and breath sounds—sensed with contact microphones—to EEG signals in the creation of an electronic music texture (Teitelbaum 1976). My own work with brainwaves began with experiments in musical production using alpha rhythms and explorations of the relation of alpha wave production to music perception and the various states of awareness and consciousness associated with music performance. Initially, this took place in 1969 in the laboratory of Les Fehmi, an early biofeedback researcher at the

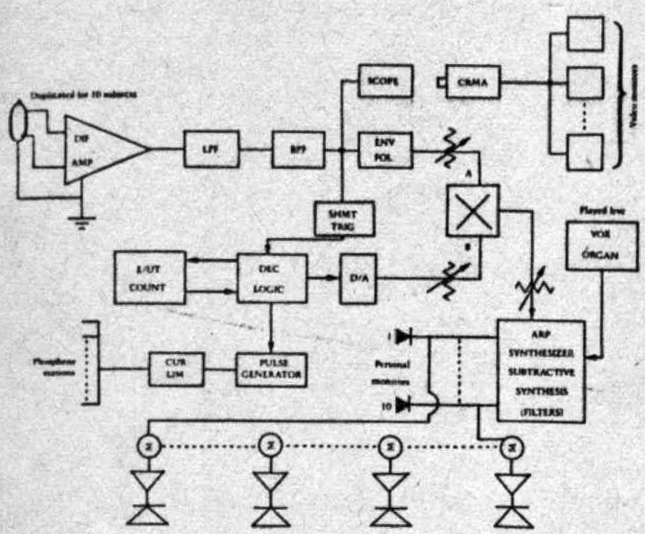

Fig. 1. System configuration diagram for the author's 1970 installation/performance/participation work, Ecology of the Skin. *DIF AMP = differential brainwave amplifier, LPF = lowpass filter, BPF = bandpass filter, ENV FOL = envelope follower, SHMT TRIG = Schmit trigger circuit, E/UT = events per unit time, DEC LOGIC = modular logic circuit system, D/A = digital-to-analog convertors, CUR LIM = current limiting circuit, CMRA = video camera.*

State University of New York at Stony Brook, after a suggestion by E. E. Coons of New York University. I developed an environmental demonstration/participation/performance event entitled *Ecology of the Skin*, in 1970–71. It involved biofeedback monitoring of brainwaves and heart signals from performers and audience members, their translation into a musical texture, along with synchronous, electronic stimulation of visual phosphenes at cerebral light show viewing stations for the audience. The electronic setup for this work included the capability of adjusting the degree of brainwave control over sound for each of ten participants according to a simple statistical measure—the amount of time spent per minute producing alpha waves. It is shown in Fig. 1.

My work led in 1972 to the creation of the Laboratory of Experimental Aesthetics at York University in Toronto, where the intent was to study information processing modalities of the nervous system in relation to aesthetic experience and states of consciousness within an environment of artistic production. Work was carried out there for a number of years under sponsorship by York University, the Canada Council Explorations Programme, and the Aesthetic Research Centre of Canada. Many individuals carried out experiments or produced works of art there over a seven-year period. Some of these are documented in my book *Biofeedback and the Arts, Results of Early Experiments* (Rosenboom 1976a). They included works involving music, visual arts, dance, and kinetic arts.

Another early experimenter was Manfred Eaton, who carried out experiments in music and bioelectric phenomena at the ORCUS Research Center in Kansas City during the 1960s and early 1970s. Eaton described extensive explorations in applying various, bioelectrically derived signals to artistic projects and the study of aesthetic responses to stimuli. These signals resulted from measuring the EEG, pulse rate, respiration, galvanic skin resistance (GSR), blood flow volume, and the electrocardiogram (EKG). A variety of multisensory display systems were devised to follow changes in these measurements. Eaton also speculated on the possibility of employing sensory-evoked responses, requiring more sophisticated analysis capability than what was readily available at that time in order to generate complex patterns for music, kinetic arts, and television.

Many of the earliest experiments involved a kind of simple parametric following in which some aspect of sound or a visual display is made to track changes detected in a bioelectronic signal. It soon became apparent, however, that data from deeper statistical analyses, particularly of EEG trends, would provide more meaningful signals with which to control musical forms. In *Portable Gold and Philosophers' Stones* (1972) (Rosenboom 1976a, 1976c), a battery of such techniques was employed by a technician, who performed with the analysis equipment, along with an ensemble of four biofeedback musicians. Most significant among these was a measure of the *coherence time* of EEG waveforms in various spectral bands, extracted by means of the autocorrelation function. This determined the range of direct control over elements of the sound texture given to each performer. As coherence times for a particular performer increased, the degree of influence over the sound texture allotted to that performer also increased. Cross-correlations on signals from pairs of performers were sometimes used as well. Fourier analysis was used to extract EEG power spectra. These were mapped onto a series of weights applied to a group of resonant bandpass filters, in a device known as a *holophone*. This helped deter-

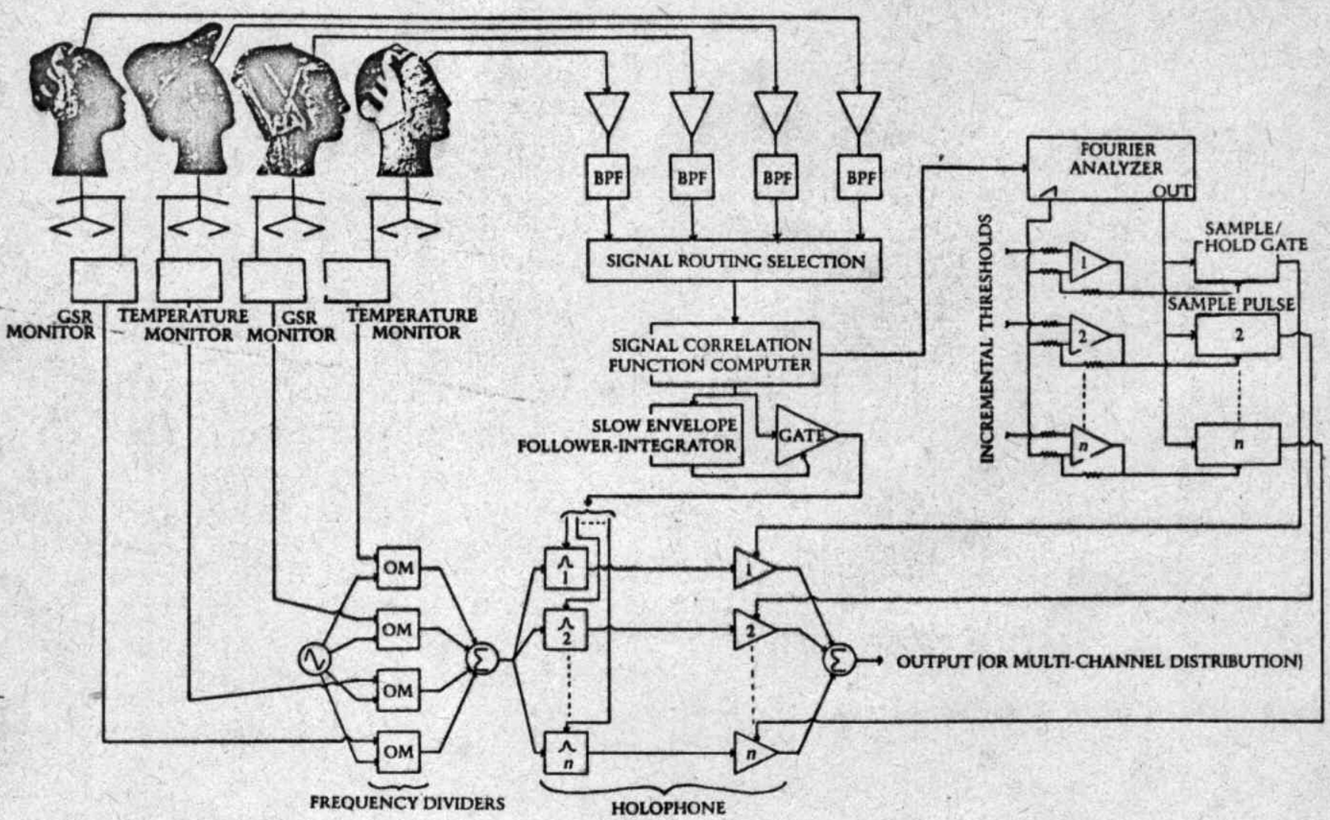

Fig. 2. System configuration diagram for the authors Portable Gold and Philosopher's Stones (Music from Brains in Fours) *(1972).*

mine the spectral composition of the music, as shown in Fig. 2.

Measures of body temperature and GSR were also used to direct the tonality of the musical texture. Subsequent to this, extensive work was carried out in our laboratory to explore how analysis of event-related potentials (ERP) might both elucidate processes of musical perception and cognition—particularly with respect to contemporary styles—and be applied to musical production in a feedback paradigm. Unlike the coherent EEG waves—known as alpha, beta, theta, and delta—ERPs are transient, nonrepetitive waveforms associated with the presentation of clearly defined stimulus events. Complex statistical procedures, such as signal averaging, template matching, and adaptive filtering are required to extract them from the ongoing EEG. ERPs contain a number of peaks, the size and latency of which can provide evidence for the occurrence of brain processes associated with hierarchical information processing. These include various attentional gating effects, shifts in selective attention, degrees of recognition or surprise contained in a stimulus, and effects leading to the formation of a mental image and memory engram associated with the event. Some of the peaks are exogenous in origin, varying with the physical aspects of the stimulus; some are endogenous in origin, varying with psychological or cognitive processes.

Many other art works were produced and research programs carried out during the 1970s (e.g., Grayson 1973; Malina 1974; Rosenboom 1975, 1976a, 1976b, 1976c, 1977a, 1977b, 1984; Paul 1986). For me, the culmination of the musical applications was the production of *On Being Invisible* in 1976–77. In this work (described in detail below), complete musical forms are constructed as a result of the self-organizing dynamics of a system in which both ongoing EEG parameters and ERPs—indicative of shifts in selective attention on the part of a solo performer—are analyzed by computer and used to direct the stochastic evolution of an adaptive interactive electronic music system.

On Being Invisible: Using ERPs to Build Formal Musical Holarchies in Real-Time

On Being Invisible is a self-organizing, dynamical system, not a fixed musical composition. The title refers to the role of the individual within an evolving, dynamical environment—one who makes decisions concerning when and how to be a conscious initiator of action and when simply to allow his/her individual, internal dynamics to co-evolve within the macroscopic dynamics of the system as a whole. Consequently, the work is ongoing. Within the corpus of my music, the title labels a period of work with these ideas from about 1976 to 1979. Recently I have begun new work with this system, titling it *On Being Invisible II.* Note that I've invoked the word *holarchy* as an alternative to *hierarchy* in discussing musical form. It is intended to encourage a way of thinking that gives equal weight to both the top down and bottom up views of how forms arise and evolve. Personal discussions with colleagues, notably J. Tenney and the writings of E. Jantsch have influenced this concept.

The Basic Paradigm: Attention-dependent Sonic Environments

One of the primary objectives in this research was to achieve the technical capability necessary to create an attention-dependent sonic environment. I wanted to create a situation in which the syntax of a sonic language orders itself according to the manner in which sound is perceived. In a sense, *On Being Invisible* has at the core of its structure a model for a way in which language can be acquired. It produces the direct result that aspects of attention—as reflected in electroencephalographic signals—have the immediate physical consequence of changing some aspect of the sound and, more importantly, affecting the way in which the sonic stream orders itself in time on several hierarchical levels. As a biofeedback model, it involves with what might be called "the cybernetics of language and cognition."

Fig. 3. Statistical distribution of behaviors: goal state (a) and behavior distribution with lateral inhibition (b).

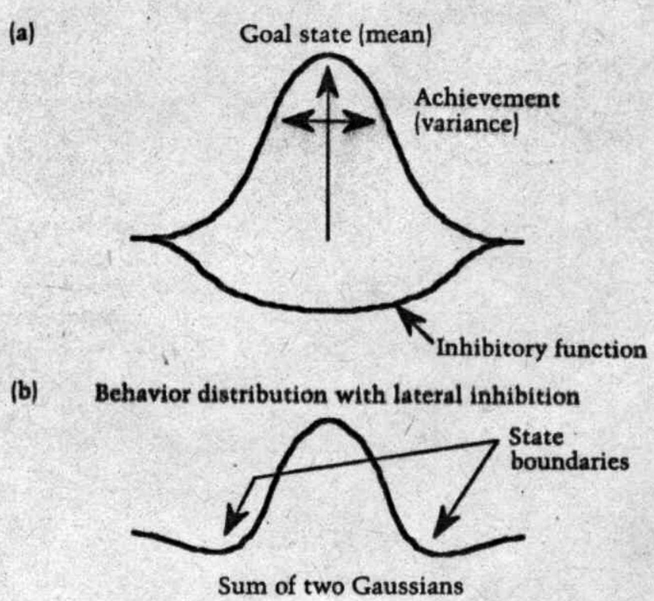

In this feedback model, the desired goal state for a self-organizing process is represented as a statistical mean of system behaviors, while the variance of actual behaviors around the mean represents the level of achievement of the goal state. A lateral inhibition function, separating neighboring goal states that might otherwise be generalized, is produced by summing a narrowly tuned, positive Gaussian function with a broadly tuned, negative one, as can be seen in Fig. 3. This idea was originally suggested by Heinz von Foerster (1981).

The exact shape of these Gaussian functions is associated with the degree of uncertainty involved in representing the goal states. Thus, a stochastic shaping of the unfolding language structure results.

The Early Versions

In each version of *On Being Invisible,* a system was employed that contained the following major components:

1. A *musical-structure-generating mechanism* coupled to a sound synthesis system
2. A *model of musical perception* that detects and makes predictions about the perceptual

Fig. 4. On Being Invisible *(1976). Schematic diagram of the early, nonbrainwave version.*

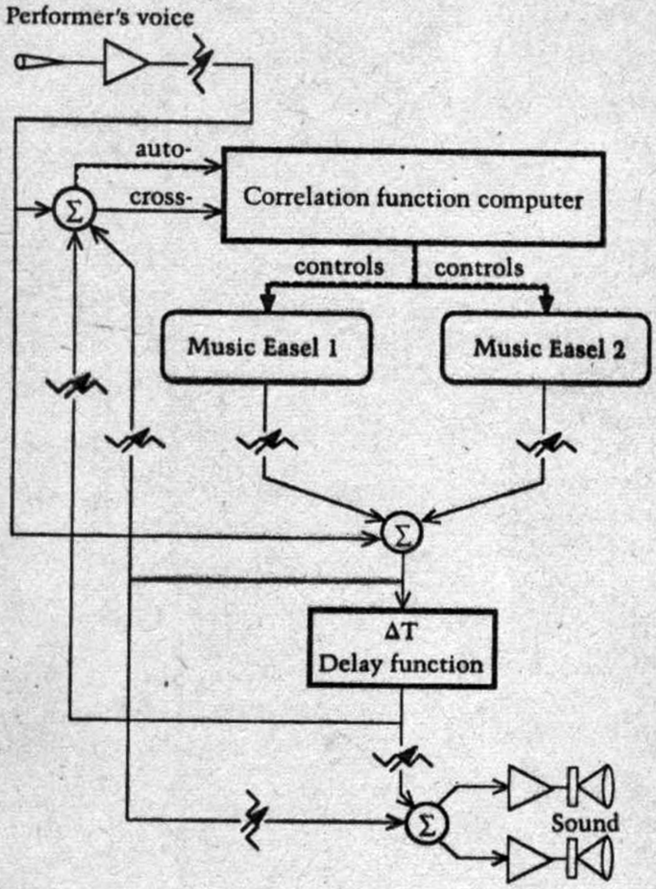

effect of various phenomena in an unfolding musical structure

3. A *perceiving, interacting entity* (i.e., human performer)
4. An *input analysis system* for detecting and analyzing bioelectromagnetic and other input signals
5. A *structure-controlling mechanism* that directs (1) and updates (2) in response to correspondences in information from (4) and (2)

An early, prebrainwave version of *On Being Invisible* is worth noting for its relevance to evolving systems concepts. Its functional diagram is shown in Fig. 4.

A structure-controlling feedback loop is established here that accepts occasional influence from outside the loop in the form of vocal sounds. To begin with, two Buchla Music Easel synthesizers constitute the structure-generating mechanism and sound synthesis system. Structure generation is contained in the qualities of the patch setup in each synthesizer, with the most important structural information residing in short-segment, preset sequencers and pseudo-random pattern generators, which were implemented with feedback shift registers. The results are combined and sent to a time-delay mechanism. Both the delayed and nondelayed signals are added to vocal sounds in a connection and summing matrix (mixer), and sent to the input analysis system. The analysis consists of performing an ongoing autocorrelation or cross-correlation on the audio signals from this mixer, chosen in real-time by the performer. This analysis system also served as the structure-controlling mechanism. The correlation function could be scanned repetitively along its delay axis and read out as a series of points converted into voltage values. These voltages were used to control relatively global parameters in the patch programs set up on the two synthesizers.

The human performer acted both as the perceiving, interacting entity and as the model of musical perception that made predictions. The entire system was activated and directed by vocal sounds. The performer watched a continuously updated display of the chosen correlation function, very often the autocorrelation function of the voice itself. He or she was required to learn through practice the ability to predict the form of the correlation function that would result from a particular type of vocal sound—such as smooth or raspy—containing particular harmonic or noise contents. This would stimulate and shape the behavior of the structure-generating and synthesis setup. Particular relationships among the forms of correlation functions, combined with certain behaviors of the synthesis patch and timings in the delay system, could produce life-like sound forms. These would often persist for some time, then seemingly spontaneously evolve their morphologies in a highly organic manner. In a sense, the results of analyzing vocal waveshapes would determine the content of a long sequencer, which in turn directed global parameters in a synthesis patch.

Several important systems principles can be seen upon closer analysis of the schematic for this piece in Fig. 4. First, there is a primary feedback loop capable of a certain degree of self-organization. The system's long-term memory resides in the interconnections of elements constituting the patch and, to a lesser extent, in patterns resident in sequencers. Short-term memory, which is most important for the maintenance of musical patterns, resides in two places—the delay system and the correlation function. Note that both of these involve time delay and show striking similarities. The delay system involves recombination of original and time-delayed versions of the musical sound. The correlation function involves integrating the results of comparisons of successively delayed time-slices of a signal with the current, incoming signal. In both systems, repetitive features of a signal—which fit neatly into multiples of the delay interval—will periodically reinforce each other. Self-maintaining patterns can result, the viability of which depends upon both the coherence of the signals being correlated and the stability of musical patterns being delayed. With proper tuning of delay parameters and integration time constants, most resulting musical patterns will last for a while and then decay. This decay is due to inherent instabilities, irregularities, or lack of long-term coherence. To produce a presentation for an audience, activity from nearly any set of points around the loop could theoretically be tapped and projected by means of amplifiers, loudspeakers, and possibly visual displays.

The entire configuration acts like a nonequilibrium system capable of organizing itself into patterns with relatively short-term stability and subject to pattern evolution by means of energy exchanges with its environment. These exchanges take the form of perturbations introduced by the vocal performer from outside the loop, whose signals—once analyzed by the system's short-term correlation memory inside the loop—push the internal pattern evolution in new directions.

At the time of its creation, this system was conceived as an experiment in alternative performance input structures—a new way of playing a synthesis system. This way of developing an improvisationally articulated relationship with a complex accompanying system was extremely rewarding. The piece was first performed at The Music Gallery in Toronto on 13 March, 1976.

Performance systems of this kind exhibit very different behaviors from what we normally expect from musical instruments. Most instruments are nonevolving equilibrium systems, or at least the result of attempts to create them. These instruments are constructed so as to tend towards an equilibrium state represented by an even distribution of potential energy in an elastic medium, like the tension in a string or drum head. A performer's actions are to disturb this state, moving the system far from an even distribution of tension, and then to observe, listen to, and sometimes try to influence the way in which the system returns to its equilibrium state of even tension distribution. Along the path of this return to equilibrium, some of the potential energy gained by the disturbance from outside the system is dissipated in the production of sound waves. Wind instruments are somewhat different, in that they contain no potential energy until it is applied from outside in the form of air pressure. They then channel the dissipation of this energy in resonant, usually harmonic, modes producing the compressions and rarefactions of sound waves.

Instrument structures of the type used in *On Being Invisible*, on the other hand, are self-organizing, evolving, nonequilibrium entities. Performance techniques for them tend more along the lines of developing creative influences on their behavior and evolution, rather than traditional technical, physiological, and proprioceptive mastery. Mastery of musical thinking, on the other hand, becomes all the more essential in this new kind of performance. There is no score to guide the performer's actions. Instead, there is a co-evolution of the performer with his or her performance system, the structure of which is an extension of his or her holarchic musical mind and body. In subsequent versions of *On Being Invisible*, bioelectromagnetic inputs were added and the technology changed. Many of the general principles of this performance paradigm continued to be applied, however.

Figure 5 shows a schematic for the next version of *On Being Invisible*, created in 1976–77 and first performed at The Music Gallery in Toronto on

Fig. 5. *On Being Invisible (1976–77). Schematic diagram.*

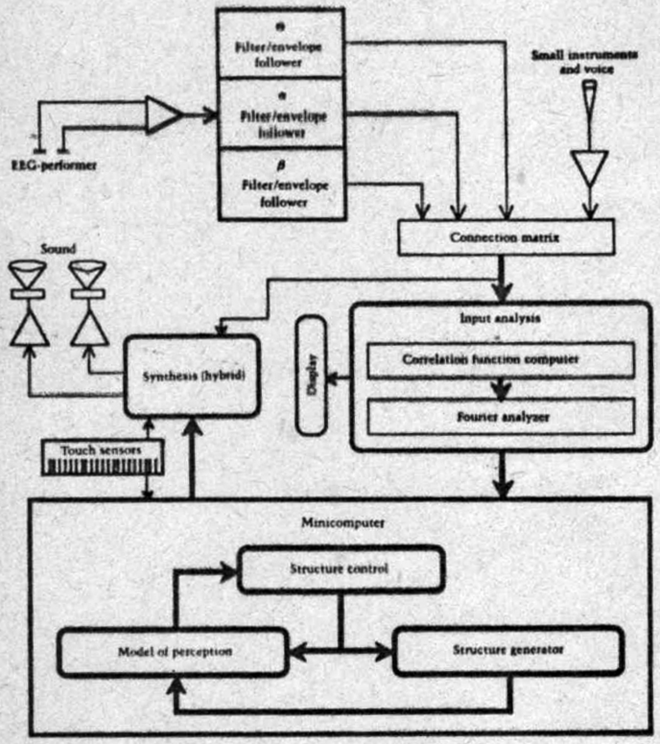

12 February, 1977. In this version, the feedback loop encloses the perceiving, interacting entity—the human performer. Energy and information exchanges with the environment take place through the performer, and the sound output is processed by the performer inside the primary loop. The model of musical perception, the structure-generating mechanism, and the structure-controlling mechanism all reside inside the software of a minicomputer. The input analysis system is comprised of a correlation function computer and Fourier analyzer that produce and continuously update displays of correlation functions along with power density spectra and phase plots. Changes in the values of individual points along the frequency axes of the power density and phase plots can also be generated at varying voltages. In this way the highly selective amplitude envelopes of any desired frequency component in the signal can be made available to the computer or synthesis system. One or two channels of brainwave inputs were derived from some combination of electrodes located at the vertex and occipital or temporal lobes of the performer's brain. Usually, the raw EEG was input to the analysis system. Sometimes, however, theta-, alpha-, and beta-band filters and envelope followers were used independently. The outputs of these were patched into the synthesis system to derive rhythmically synchronous triggers for sound events. Capability of feeding sounds into the analysis system from the voice or small acoustic instruments—monkey drums, Tibetan cymbals, snake charmer's pipes, etc.—was maintained. A keyboard array of pressure sensitive touch sensors was also available. One novel use of these touch sensors involved autocorrelating pressure contours from successive touch epochs. In this way, regular features from touch shapes were extracted and used as another form of input to the structure-controlling mechanism.

The synthesis system was a hybrid one—i.e., digitally controlled analog synthesis hardware—consisting of Buchla 200 Series modules and a Music Easel. The control voltages for these were generated by the minicomputer. In addition, a system was devised wherein Music Easel patches could be stored and set up by the minicomputer so that the system could be reconfigured very rapidly. The structure-generating mechanism was stochastic in nature, the global variables of which were set by the structure-controlling mechanism. Gaussian distributed random values were generated and assigned to several parameters for each of about five voices in the synthesis system. The variances in these values were relaxed or tightened according to directives from the structure-controlling mechanism. The mean values of these parameters were initially allowed to move according to a random walk algorithm at a rate somewhat slower than that at which discrete values were generated. This constituted a second level of control in the structural hierarchy. These mean values, however, would eventually be constrained by the structure-controlling mechanism.

The model of musical perception represents a major addition. Its purpose in this version was to make predictions about the arousal value of types of changes in various acoustic parameters of the musical voices. This was assumed to be strongly related to the likelihood that shifts of attention on the part of the perceiving, interacting entity—the human performer—would accompany such parametric changes. Control signals applied to pitch, amplitude, envelope duration, and timbral complexity—as measured by modulation index, the bandwidth of a filter being imposed on complex waveforms, and a nonlinear waveshaping parameter in the Music Easel oscillators—were tracked and analyzed according to a unidirectional, rate-sensitive (URS) difference detector. This URS model is based on assumptions about the behavior of sensory input channels in the nervous system (Clynes 1972). It can be expressed as follows:

$$D(t) \leftarrow |P'(t)|'(t) \text{ if } |P'(t)|'(t) \geqslant 0, \text{ else } D(t) \leftarrow 0.$$

A simple, discrete-time version of this can be implemented by performing:

$$D(t) \leftarrow |P(t) - P(t-1)| - |P(t-1) - P(t-2)|$$

and then:

$$\text{if } D(t) < 0, \text{ then } D(t) \leftarrow 0,$$

where $D(t)$ is the difference function applied to successive values of an acoustic parameter P. If $D(t)$ is greater than zero, then the element at $P(t)$ is considered to be a potential *initiator* of a shift in attention. The value of $D(t)$ indicates the relative strength of the element at $P(t)$ as an initiator. A threshold value T was set with which to make a determination about a particular element as an initiator. If a threshold crossing of $D(t)$ occurred, its value could then be compared with corresponding values from other parametric contours to determine if this element will be predicted to initiate the forming of a group on the next hierarchical level of perception, made up of lower-level musical elements. The most important quality of this function is that it is sensitive to changes of rate-of-change in the positive direction. This follows the observation about the nervous system's reaction to incoming sense data—that it is most sensitive to increases in the rate-of-change of some aspect of the environment. For example, the most sensitive situation would be a departure from a state of relative rest, that is, something starts to move or starts to move at a faster rate.

In this, still early version of *On Being Invisible*, time steps were simply considered to correspond to the occurrence of envelope triggers. Changes in envelope duration—roughly analogous to element or note duration—were considered simply to comprise another parametric contour. Later versions refined this considerably. Values of $D(t)$ from the several parameters of each voice were summed and tested against an overall threshold T. The current value of $D(t)$ above the current threshold level was used as a measure of the strength S, as follows:

$$S = D(t) - T.$$

Here S is the strength of the prediction being made by the model of musical perception that the event in question would be an initiator, and that it would be attention securing. This prediction would then be tested by interrogating the inupt analysis system to see if evidence of attention shift was present in the EEG. Significant EEG desynchronization, interruption of ongoing coherent waves, and various EEG state changes all contributed to this determination. See Rosenboom (1989) for more information about this. The structure-controlling mechanism was responsible for making the determination. If the prediction was confirmed, then the probability was increased that the kinds of change in musical parameters associated with the prediction would occur again. If the prediction was denied, then the probabilities associated with these kinds of change was decreased.

The structure-controlling mechanism was also responsible for updating the model of musical perception. At the beginning of a session the threshold T was initialized to zero, guaranteeing that the first event would be an initiator. On the occurrence of an initiator—i.e., a prediction concomitant with a successful measure of attention shift as seen in the EEG analysis—T would be set equal to $D(init)$, the difference function value associated with the initiator. This T would be applied at $D(init + 1)$. Thereafter, T was allowed to float according to an accumulating time average of $D(t)$ values with limited history, i.e., number of samples contributing to the average. T was always reinitialized on the occurrence of a successful initiator to equal $D(init)$, and $D(t)$ values prior to a given initiator were not

Fig. 6. Application of the difference function to a simple parametric contour.

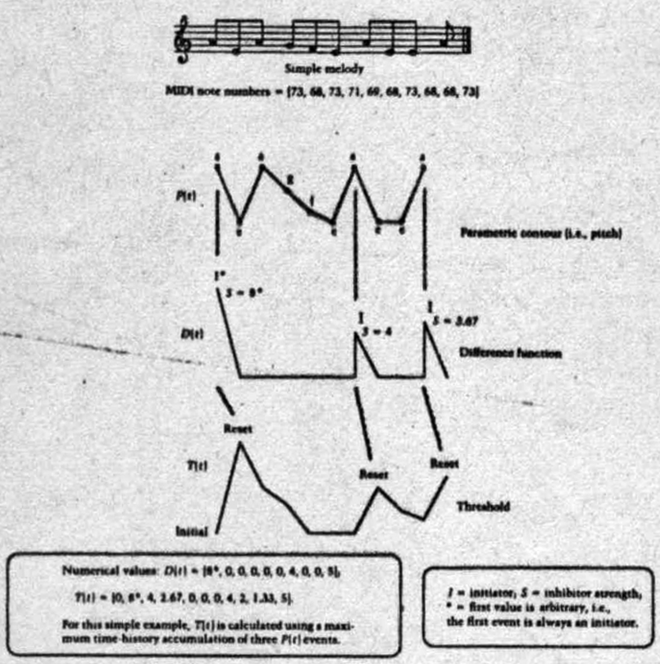

included in the new average calculation. This produced a behavior for T consisting of an upward step immediately subsequent to the occurrence of each successful initiator, followed by a decay, the rate of which depended on the activity of P after the initiator. The assumption was made that events following an initiator would be grouped with that initiator into a perceptual unit and that a succession of initiators would mark off a sequence of these units. Such perceptual units have been described by Tenney (1988) as *temporal gestalts* (TGs) (Tenney 1988).

The application of this difference detector and threshold-updating mechanism to simple melodic sequences is illustrated in Fig. 6. Here the initiators parse the note series into groups or chunks. In Tenney's terminology, the notes would be termed *elements*, the note groups or chunks *clangs*, and a string of clangs a *sequence*. From there, any number of higher ordered sequences may combine, until the highest TG level of interest—known as the *threshold of formal concern*—is reached.

Further Refinements: Real-Time, Convergence and Divergence of Patterns

In subsequent versions of *On Being Invisible*, many refinements were made. First, the ability to account for the flow of real-time was added. This replaced the arbitrary use of element (e.g., note) envelope triggers to mark off a sequence of time steps. Threshold levels T were no longer time-averaged after the occurrence of an initiator. Rather, they were subject to an exponential decay, the rate of which was adjusted to maximize the effectiveness of the system. This decay rate is analogous to the persistence of perceptual phenomena to which attention is shifted in the nervous system and the tendency for these to mask the effects of smaller events in the immediately succeeding time vicinity, as shown in Fig. 7. Remember that an event here is defined as a unidirectional (increasing) change of rate-of-change $D(t)$ in some parameter $P(t)$ that is of sufficient size to cross a threshold T and which successfully initiates a shift of attention, the evidence for which is extracted from the ongoing EEG. This process—combined with the exponentially decaying threshold—takes care of an important effect of duration. The longer the duration of an element for which a given parameter is unchanging (for example, a long sustained pitch), the smaller the subsequent change in that parameter that will be required to produce a threshold crossing event and, thus, an initiator is. See Fig. 7 for a sketch of the behavior of T with exponential decay.

The initiator strength factor S (as defined above) was brought into play to serve as a measure of confidence level for predictions. This could be used to determine the degree of change brought about in the stochastic structure-generating mechanism in response to feedback. The structure-control mechanism has two objectives. The first is to increase the probability that the kinds of musical change associated with successful predictions will recur. Thus, if certain changes continue to evoke attention shifts, they will converge into patterns. The second objective is that if predictions are unsuccessful, the musical structure is made more open to stochastic influences. Consequently, if successful predictions

Fig. 7. Effect of exponential decay of the threshold function T(t) *on initiators in successive time regions.*

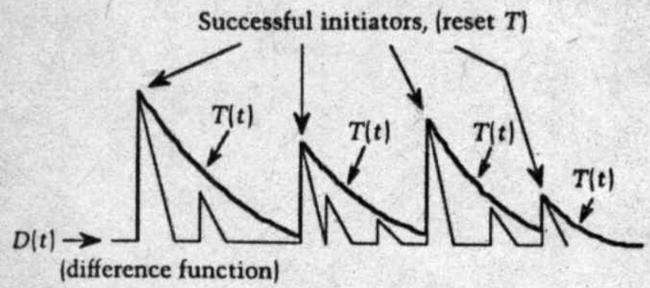

associated with ongoing patterns begin to fail, the patterns are allowed to diverge—to evolve by means of random mutations into new patterns or possibly to be dissipated entirely. Previously successful initiators can fail for a variety of reasons. Repeating patterns may fail to elicit attention shifts because of boredom, volitional shifting of attention focus to other patterns or aspects of the environment, volitional redistribution of attention, distractions from the external or internal environment, shifts in states of consciousness, and many other factors. The structure-control mechanism used S to direct the rate of convergence or divergence of patterns. If a high-S prediction was successful, convergence to repeating patterns was more rapid than if a low-S prediction was successful. Correspondingly, a high-S unsuccessful prediction would cause relatively rapid divergence and a low-S unsuccessful prediction would cause less rapid divergence.

Convergence and divergence is achieved by adjusting variables in various stochastic canons. For example, to create divergence with a Gaussian distributed canon, the range could be widened and the mean allowed to wander according to a random walk with increasing variance. Convergence could be created by restricting the range and variance or by narrowing the window size of a filter applied to the output of some random generator. Convergence was also dependent on processes used to build hierarchical structures.

Hierarchical Structure Building

At a certain point in a session or performance of *On Being Invisible*, the performer could activate a hierarchical structure-building part of the structure-control mechanism. Usually, the performer would use his or her discretion in judging when this was appropriate. At the outset of a performance, it made sense to keep things simple, with converging and diverging processes focused on just one structural level. This way, the biofeedback processes involved could be clearer and more evident. Further on, however, it would usually become desirable to interact with an evolving musical environment of increasing richness. To accomplish this, the structure-control mechanism could be directed to store sequences of parametric values that were delineated by successful initiators. These correspond to what Tenney describes as clangs, with the exception that the *On Being Invisible* system stored all parametric sequences separately—a pitch sequence was stored separately from its associated amplitude, timbre, and duration sequences. These were not kept bound to each other a priori. Consequently, parametric sequences could be recombined with those of other clangs to create transformations up to the limit of the available combinatorial possibilities. This recombination potential was made available to the performer to select at will. The default behavior was that parametric sequences from a given clang remained bound, unless otherwise indicated. Each parametric clang was labeled and assigned a probability value determining its likelihood of being replayed exactly as stored. The performer triggered the system as to when to begin filling memory with clangs and when to stop.

Musical Inference

The hierarchical structure builder contained the beginnings of a simple musical inference engine, but one with a difference. It had to make predictions "on the fly" as to how a growing, evolving structure was being perceived. Again, on a trigger from the performer, the model of musical perception shifted its prediction process to the second hierarchical level of the growing musical structure. At this point, the nature of the prediction process changed. An analysis of the sequence of clangs was carried out, inspired originally by concepts from information theory.

Edgar Coons and David Kraehenbuehl published a stimulating paper in 1958 in which a method was presented for quantifying the information value of events in a sequence—with particular reference to musical structure (Coons and Kraehenbuehl 1958). Rather than being based on probability values assigned a priori to events from a repertoire of possible events—as would normally be the case in traditional information theory analyses—this method involved tabulating all possible predictions that could be made at a particular point in an event stream and then calculating the degree to which each prediction is nonconfirmed by the events that actually take place. A notion of structural hierarchy was also contained in these calculations. Not only were predictions for specific events examined, but predictions involving the occurrence of dissimilarity (maximally informed events) versus similarity (minimally informed events), along with sequences of these, were also taken into account. Initially, I tried to incorporate a variant of this method into the *On Being Invisible* programs. A practical problem prevented a full realization, however. The analysis method of Coons and Kraehenbuehl produces considerable insight into the nature of patterns and how they might be perceived. During real-time algorithmic musical performance situations, however, the computations required by this method can soon grow out of hand. A full-scale analysis of this type requires tabulation of all possible predictions that can be made at a given point in a piece on the basis of past events. This is necessary in order to be able to arrive at a relative-informedness value for the event that eventually does take place. In even moderately complex music, this can involve an enormous number of possible predictions. Although the actual calculations are quite simple, their number becomes unwieldy for small computers and the memory requirements become quite large. Furthermore, the results must be obtained very quickly in order to keep up with a spontaneously emerging and evolving musical fabric.

Fortunately, another set of stimulating experiments was carried out in the Cognitive Psychophysiology Laboratory at the University of Illinois at just the right time (Squires et al. 1976). This study showed enhancements in the amplitudes of two important peaks—known as *N200* and *P300*—and late, slow-wave components of ERPs for particular stimuli as a function of their position in a sequence of events. Furthermore, a relationship was shown between the contents of the sequence preceding the event in question and the degree to which that event was a discriminant one associated with large ERP amplitude peaks. This depended on developing an expectancy function associated with each event, bearing an important conceptual relationship to the Coons and Kraehenbuehl method of tabulating predictions. This expectancy function was derived from the linear combination of a memory function for past events, a probability value associated with each particular stimulus type, and an alternation factor. The alternation factor reflected an attempt to take into account the effects of simple hierarchical groupings perceived by the subject. This study dealt only with sequences of two event types, and was therefore limited in its applicability to complex music. Nevertheless, a modification of ideas from these two studies led to a practical implementation of an information/expectancy function algorithm for hierarchical musical patterns used in *On Being Invisible.*

As mentioned above, the system begins to store clangs when triggered by the performer. At this time, each clang is labeled and assigned a probability value P initially at random. Clangs are then played back stochastically, their probability values being affected by the degree of attention shift they seem to elicit in the EEG signals. High attention shift when a clang is played will result in increasing that clang's probability value, enhancing the likelihood that it will be played again. At the same time, the system begins to group clangs into sequences. A record is kept of the order in which clangs are played. A memory/expectancy function is evaluated on each clang playing. The form of the memory/expectancy function is adapted from Squires, et al. as follows:

$$M_{EN} = \sum_{i=N-1}^{N-m} \alpha^{N-i} S_i ,$$

where M_{EN} is the memory/expectancy function for event E at position N as a function of the sequence

Fig. 8. Illustration of the hierarchical structure builder operating on a simple string of events.

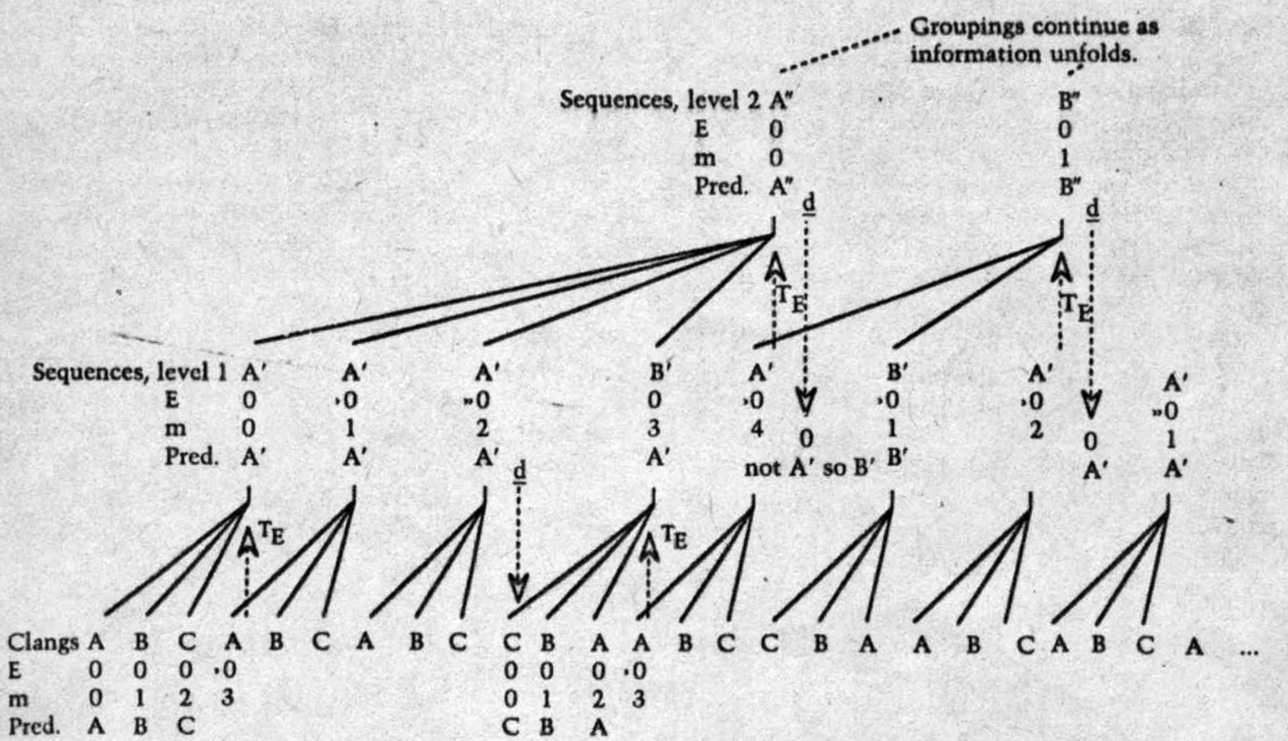

Note: E = expectancy values, shown as 0, ›0, or »0 for simplicity; m = memory order factor; Pred. = predictions made at particular points; d = disconfirmations of predictions. Dashed lines indicate shifts in the hierarchical level on which information processing proceeds. Note that for an event to *initiate* a higher level grouping, its expectancy value must exceed a threshold, T_E; though, not all such events will produce *initiators* unless other conditions are also met.

of past events S_i. As the algorithm runs, S_i is set equal to zero if the event at the current index i is not equal to E and set to one if it is equal to E. An exponential decay of memory for past events results from the evaluation of α where $0 \leq \alpha \leq 1$. A memory order factor m is used to determine how far back in the sequence to go when calculating the expectancy function for a given position in the sequence. Various values of m were chosen purely experimentally for different *On Being Invisible* sessions. It is assumed that the probability values reflect relatively global aspects of the sequence, while the memory/expectancy function corresponds to stimulus processing in more short-term memory.

Perhaps a good way to describe this algorithm is to examine its operation on the simple example illustrated in Fig. 8. Consider the following primitive sequence of three events, or musical objects, simply labeled, *A*, *B*, and *C*, which we will call clangs. Higher order groupings of clangs will be called sequences.

A B C A B C A B C C B A A B C C B A A B C A B C A . . .

Initially, this ordering would be the result of making selections from the three stored clangs simply by applying assigned probabilities for their occurrence. For the purposes of this example, the sequence has been made more regular than would probably be the case at first. We will assume that the likelihood of detecting ERP or other EEG parameters indicative of shifts in selective attention will be associated with predictions (made on the basis of past experience) being disconfirmed, i.e., their expectancy function is low and their global probability of occurrence is low. The detection of such attention shifts will be used to cause objects to be labeled as initiators of higher level groupings. Note that simple, regular sequences—like that shown in this example—tend to produce well-ordered, hierarchical structure trees. In actuality, such well-ordered trees tend to be less interesting than not so well-ordered trees.

Event *A* is of course an initiator by default. The only prediction that can be made on the basis of experience accumulated in the system so far is that *A* will recur. Consequently, the occurrence of *B* is disconfirming, as is the next event, *C*. Expectancy values for all these events are zero at this point. The memory order parameter (m) grows with the length of the pattern being analyzed. On the second occurrence of *A*, a nonzero expectancy value is obtained. It is compared with an expectancy threshold set experimentally in order to tune the behavior of the algorithm. A threshold crossing triggers the algorithm to attempt the formulation of a higher level grouping—a sequence in this case—and to move the level of its analysis up. All clang objects encompassed by the memory order parameter, up to but not including the current event, are gathered into a tentative proposal for a higher-level sequence grouping. It is labeled A' and assigned a probability value for its recurrence. A memory order parameter is kept and updated for each hierarchical level in the unfolding structure.

Now the algorithm operates on the next hierarchical level above the clang. A tentative prediction is made that A' will occur again. Actually, this is the only prediction that can be made at this point on sequence level 1. The order of subsequent clangs is then compared with the contents of A'. After the second *ABC* group occurs, the tentative prediction about the recurrence of A' is confirmed. The expectancy value for this second A' is nonzero. However, only one type of event has yet occurred on this first sequence level, so no higher-level grouping is possible. Note that a higher-level grouping must contain at least two different kinds of objects. In other words, the repetition of the same object over and over is not considered to produce candidates for higher-level pattern groupings. A second kind of object must occur to act as a pattern delimiter.

A prediction for the occurrence of another A' is made. This one is also confirmed, so yet another A' is predicted. The expectancy values for A' are growing. However, on occurrence of the next clang (C), this prediction is immediately disconfirmed. Now the algorithm must drop its level of analysis back down to the clang level. The memory order parameter (m) for clangs, which has been growing from the beginning of the first clang, is reset to zero, referencing the beginning of what may eventually become a new sequence grouping. As was the case at the beginning of the whole sequence, the only predictions that can be made are for the recurrence of each clang, the expectancy values for which are all zero. On the fifth occurrence of *A*, however, a nonzero expectancy value is obtained, and again the preceding clangs, up to the limit of m, are gathered into a tentative sequence grouping and labeled B'. The algorithm returns to the first sequence level again and predicts another A', the strongest prediction it can make at this point. This prediction is confirmed by the occurrence of the next three clangs, *ABC*. Now, a second-level sequence grouping can be made because we have had two types of events on the sequence level 1, A' and B'. Events on this level are then collected up to the limit of m, forming the group $A'A'A'B'$. We can now move up to the second sequence level, label the group A'', and assign it a probability. The tentative prediction of a second A'' is made. The occurrence of the very next clang (C), however, immediately disconfirms this. The algorithm must drop back down one level and make another prediction. We already know at this point that the ensuing sequence cannot be A'. So we predict a sequence with the next highest expectancy, B'. Subsequent clangs confirm it. The next A' will trigger a second-level sequence grouping ($A'B'$) with just two objects if the expectancy threshold is set to facilitate this. It is labeled B'' and given a probability. The last A' does disconfirm the prediction of another B'', but it cannot trigger a new grouping because only one type of object has occurred since the last second-level sequence was formed. No further groupings can be made until more information is provided by the continued unfolding of the main sequence.

Transition probabilities can also be used in building the structural hierarchy. As the system moves up to higher and higher levels of grouping, nthorder transition probability tables (Markov chains) can be built to reflect the likelihood that particular clangs or sequences will tend to follow each other or remain bound in high-level groupings. The probabilities in this table can be skewed by successful detections of attention shifts to increase the likelihood that particular transitions will recur, or by unsuccessful tests to decrease the corresponding transition probabilities. This method alone, however, is not sensitive to certain kinds of attention-securing events that do not reflect grouping boundaries. For instance, attention shifts can be stimulated by events that represent the occurrence of incongruous or surprising endings of groups as well as the beginnings of new groups. It is sometimes difficult to tell—on the basis of EEG concomitants of attention shift alone—on which side of a grouping boundary a particular event lies. Consequently, further inference rules are required.

Inference Rules and Musical Knowledge

A retrospective analysis of the primary sequence may suggest alternative groupings. It is important to recognize that many such alternatives result from an out-of-time analysis. The kind of groupings produced by this algorithm result from what can be known at each point in the sequence as it unfolds in time. One significance of the hierarchical level on which the algorithm operates at a given time is that this level represents what we know at that time about the structure of the sequence. Experienced listeners apply many strategies of analytical listening based on a large knowledge base containing information about musical structure and musical transformations. This algorithm doesn't know about things like retrogrades, inversions, transformations on parametric contours, etc. If it did, the range of predictions about musical objects in an unfolding structure would be considerably widened. In addition, one could design a system that used knowledge about sequences of global features. For example, if the *ABC* labels used in this example referred to actual note names (elements), then clangs on the next hierarchical level could be labeled as to their sequence of ascending or descending pitch content, i.e.,

up, up, up, down, up, down, up, up, . . .

In the preceding example, these labels were meant to represent musical objects or events, the contents of which are undetermined. Consequently, they were labeled clangs to suggest that each may contain lower-level formal features.

We can now list several important principles or rules on which this structure builder operates.

Principle: The system always attempts to operate on the highest level of hierarchical grouping possible in order to obtain a description of the most global features of the unfolding patterns. These are assumed to have the highest predictive value.

Rule: Any event for which the expectancy function is above a threshold is considered to be a potential initiator of a new sequence grouping.

Rule: A sequence grouping must contain at least two event types.

Principle: A search for attention shifts via concomitant EEG phenomena is triggered by disconfirmation of predictions at the current hierarchical level and by events for which expectancy values are low.

Rule: Successful detection of attention shifts results in increasing the global probability that the currently referenced sequence grouping or—if possible—the newly formed sequence will recur. If a triggered search for attention shift is unsuccessful, the corresponding probability is reduced.

A result of this last principle is that as attention shifts are followed, musical patterns will continuously converge and diverge from ordered relationships; musical contexts will appear and dissolve. In a way this is hardly any different from the way music naturally evolves. In this case, however, the potential for such evolution is imbedded in the structure of an artificially intelligent musical instrument.

Global Parametric States

The analysis of parametric values via the difference function—referred to as $D(t)$ above—continues on the second and subsequent hierarchic levels as well, but with a difference. Here we look at changes in the global qualities for each parameter of a clang—what Tenney calls state variables—instead of the individual element values. These are the averages of the parametric values for the elements of a clang, adjusted to reflect element durations, following Tenney's suggestion:

$$\sum_{0}^{n} \left(\frac{\text{parametric values} \cdot \text{durations}}{\sum_{0}^{n} \text{durations}} \right)$$

The difference detector now signals increases in the rates-of-change of these global variables. This is particularly useful in making predictions about shifts in attention concomitant with offsets in one or more global variables, such as pitch transposition or changes in the loudness or timbre of an entire clang. Under this kind of transposition, clangs retain their labels—a transposed clang-*A* is still considered clang-*A* for purposes of calculating the expectancy function described earlier. However, the difference detector will always catch significant changes in a state variable and make predictions.

Parametric Weighting

A significant problem in temporal gestalt (TG) analysis involves the question of parametric weighting. How important is a change of a given size in one parameter in relation to another (e.g., pitch vs. loudness) in determining where to predict TG boundaries? In traditional Western music, pitch tends to be the parameter assumed to carry most of the information articulating form. This assumption cannot be made for twentieth-century Western music or for many other kinds of music. Tenney points out, for example, that in the music of Edgar Varèse, TGs are delimited more by amplitude events than pitch (Tenney 1988). The solution in general is empirical—one must adjust the weighting values until the model behaves as one thinks it should. Parametric weights vary substantially in different contexts, particularly with respect to relative degrees of variance among parameters. The *On Being Invisible* system offers potential as an interesting tool with which to explore how parametric weighting seems to work. A time-record of the self-adjusting, difference detector thresholds ($T(t)$) for a set of parameters provides an indicator of how such weights shift through a musical experience. High thresholds in a given parameter (e.g., loudness) indicate that relatively large changes in the rate-of-change of loudness are required for it to be effective as a formative parameter for perception in the particular musical context being examined.

Psychophysiological Parallels of TG Analysis

Another potential of this system is offered in the possibility of carrying out research into the psychophysiological parallels of TG analysis in musical form perception. To date, at least some preliminary work has been carried out, and rich possibilities for investigation remain. One may apply the system to the investigation of form perception in precomposed, fixed musical works by using only some parts of the *On Being Invisible* feedback loop—the model of musical perception and the input analysis system. Instead of a spontaneously generated musical structure, a precomposed, fixed one is produced. A time history of the results of the analysis could then provide the parallel data. Furthermore, it is very interesting to begin with a fixed composition rather than with a random starting point, and to allow the fixed structure to evolve according to the self-organizing behavior of the complete system. Surprising aspects of the way musical attention behaves often result in fascinating transformations of the initially fixed original. The focus of musical attention traverses a structural landscape in complex and possibly even highly individual ways. A wealth of experience is required to make even partially accurate predictions about how a particular formal architecture will be perceived. Observing the behavior of a system such as the one described above often reveals inspiring surprises.

Algorithmic Improvisation with ERPs and Other Inputs

Although inspired by investigations into models of perception and musical cognition, the above system is not the result of an attempt to create an artificial listener that behaves exactly like a human listener. Instead it is an interesting generative musical tool with which to produce creative results. As such, it has been subjected to continuous refinement, modification, and expansion to serve new goals and reflect evolving knowledge and insight.

Since the use of this system in *On Being Invisible* involves a real-time evolution of both performer and musical system, it is representative of a form of musical improvisation. Placing the system inside a group context involving improvisation with an attention-dependent sonic environment can be quite exciting—though results can often become very complex. In addition, the methods of input signal analysis—originally focused on the EEG—can, with minimal modification, be applied to other inputs as well. This has been tried with other physiological signals, such as touch contours and EMG signals, and even with acoustic signals. In these contexts, this system functions as an intelligent musical instrument, capable of high-level pattern generation in response to several types of input signal analysis or gesture capture. Feedback directing the ongoing evolution of the system's hierarchical structure-generating capability can come from a variety of sources. We have focused on feedback from aspects of attention shift. These can also come from simple performance actions, deterministically given by the performer in order to push the system in one direction or another. All of these are legitimate applications with rich musical potential.

Current Technical Issues and Future Prospects

The architecture, speed, and memory capacity of affordable computers is now approaching that required to realize the entire system described here efficiently. Even with modern high-speed microprocessors, however, a certain degree of parallelism is desirable. The computing required by the systems I have described falls mainly into three categories: (1) signal analysis, (2) musical structure generation and, (3) sound synthesis. Each of these is quite complex and ideally should be performed by independent parallel processors.

The EEG and MIDI

The synthesis equipment described in some of the 1970s examples above could be considered somewhat old-fashioned by today's standards. However, I want to stress the indisputable fact that some of the expressive power achieved with these older machines is yet to be matched with modern digital equivalents, even though these newer instruments have vastly greater potential in terms of numbers of voices and the ability to store many patches, programs, and waveforms. The proper realization of a work like *On Being Invisible II* requires independent addressing and continuous updating of all synthesis parameters in real-time. This is very difficult to achieve in a MIDI environment. To further exacerbate the problem, the EEG analysis system must have precise knowledge of exactly when changes in multiple sound parameters occur in order to coordinate its analytical procedures with musical structure-generating mechanisms. Without this information, the resulting data will have little relevance to actual musical perception and cognition. Transmission delays imposed by MIDI, and the inability to interrogate many MIDI synthesizers as to when certain synthesis processes occur, present serious hurdles to overcome.

The idea of an EEG-to-MIDI interface is a titillating one to be sure. This has been achieved both in our laboratory at the Mills College Center for Contemporary Music and in my private studio, as well as by others. See Knapp and Lusted (1989) for a description of their Biomuse system. In and of itself, this is a rather trivial development. A small personal computer equipped with a low-speed, low-resolution, analog-to-digital converter, a MIDI interface, and some simple software, along with a good EEG preamplifier, is the simplest way to accomplish it. The difficult part lies in how to extract truly meaningful data that bears a direct relationship to the production of musical sound from EEG signals. Though brainwave control of MIDI devices can certainly be fun, my experience leads me to issue a strong cautionary message to those who wish to use this method and who expect to obtain results with the precision required to produce data upon which conclusions about musical information processing can be based.

The EEG and DSP

DSP co-processors, which can perform nearly instantaneous updating of digital synthesis parameters such as those found on the Digidesign Sound Accelerator card or in the NeXT computer, offer strong potential. However, it is still very difficult to implement all three of the major computation tasks listed above within one typical commercially accessible personal computer and still expect to keep the tight timing constraints under control. At the Center for Contemporary Music, plans are being made to construct a general purpose DSP input peripheral device capable of implementing some of the EEG signal preconditioning and analysis algorithms. This will further relieve some of the timing constraints that result from trying to run both the EEG analysis and musical structure-generating programs on a single computer. My experiments are presently being conducted with the aid of the Apple Macintosh II computing platform, Digidesign Sound Accelerator cards using the Motorola 56001 DSP chip as a co-processor, a data acquisition system from GW Instruments, and various MIDI devices. There is often simply no substitute for an analog synthesis module capable of instantaneous response upon receipt of a computer-generated voltage step or trigger.

Techniques for Brain Imaging

Other technical developments on the near and far horizons are very exciting indeed. These include advances in multichannel, topographic mapping of electrical activity in the brain, often referred to as brain imaging, and the application of superconducting quantum interference devices (SQUIDs) to the detection of localized neuromagnetic fields. Other techniques being used in research and medicine include:

- Positron Emission Tomography (PET), which maps the flow of radioactively tagged blood in the brain as different regions become activated
- Magnetic Resonance Imaging (MRI), which records radio waves resulting from the realigning of molecules in the body when it is placed in a strong magnetic field
- Computerized Tomography (CT), which uses X-rays

EEG Expert Analyst and Musical Inference Engine

Advances in software, particularly in the field of expert systems, may prove useful for the types of applications described here. The detailed knowledge we now have about the dissection and categorization of EEG phenomena could be coded into the knowledge base of an EEG analysis expert system. This could then be linked to a musical inference engine containing rules for musical production and a model of perception. The musical rules would include provision for describing compositional intent, aspects of what constitutes a multilevel formal architecture, and a set of transformation behaviors with which to implement morphological dynamics. The model of perception would require rules for musical feature extraction, detection of the boundaries separating temporal gestalts, and perceiving holarchic form.

Such an integrated, feedback-based, self-organizing system could become a powerful tool for explorations in composition, performance, and perception, as it includes an expert signal analyst (herein focused on the EEG), a musical inference engine, a synthesis mechanism, and an intelligent performance input structure. Figure 9 shows a potential organizing scheme for it. This goal is furthermore imminently achievable with existing affordable technology.

Fig. 9. Potential organization of an intelligent input analyst and musical expert system/inference engine for live performance.

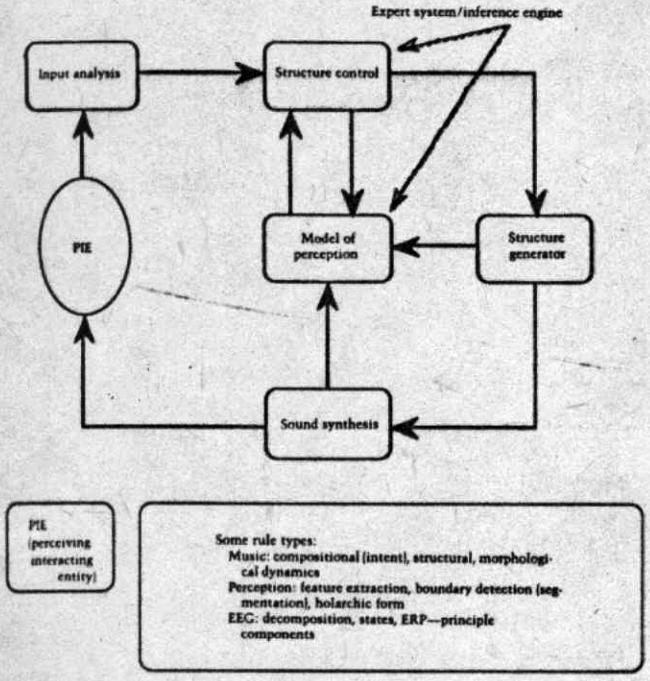

On Being Invisible II

It is now possible to imagine large-scale musical theater or operatic works involving biotelemetric presentation by human and even nonhuman performers interacting with audiences, other performers, and environments. This could create a synergistic theater, linking participants in a large-scale organism, the ontology of which could provide a script of mythical proportions. The eternal quest to understand the role of human consciousness in determining when and how to initiate action provides the essential dramatic tension. This is the grand intent of my ongoing project, currently titled *On Being Invisible II.* In this work, the major components of the feedback system shown in Fig. 9 have become anthropomorphized and are taking on the aspects of characters in a mythological scenario for evolution and social organization. From this, a script is being developed, the intent of which is to place the work in the context of a full-scale, theatrical performance. I await adequate time and support for its full realization.

References

Adrian, E. D., and B. H. C. Mathews. 1934. "The Berger Rhythm: Potential Changes from the Occipital Lobes in Man." *BRAIN* 57:355–385.

Clynes, M. 1972. "Towards a View of Man." In M. Clynes and J. H. Milsum, eds. *Biomedical Engineering Systems.* New York: McGraw-Hill.

Coons, E., and D. Kraehenbuehl. 1958. "Information as a Measure of Structure in Music." *Journal of Music Theory* II(2): 127–161.

Foerster, H. von. 1981. *Observing Systems.* Seaside, California: Intersystems Publications.

Grayson, J. 1973. *Sound Sculpture.* Vancouver: Aesthetic Research Centre of Canada Publications.

Knapp, R. B., and H. S. Lusted. 1990. "A Bioelectric Controller for Computer Music Applications." *Computer Music Journal* 14(1).

Lucier, A. 1976. "Statement On: Music for Solo Performer, 1971." in (Rosenboom 1976a).

Lucier, A. 1982. *Music for Solo Performer* (recording). New York: Lovely Music, Ltd.

Lucier, A., and D. Simon. 1980. *Chambers.* Middletown, Connecticut: Wesleyan University Press.

Malina, F., ed. 1974. *Kinetic Art.* New York: Dover Books.

Paul, D. 1986. "Biomusic and the Brain: An Interview with David Rosenboom." *Performing Arts Journal* 29 (Winter): 12–16.

Rosenboom, D. 1975. "A Model for Detection and Analysis of Information Processing Modalities of the Nervous System Through an Adaptive, Interactive, Computerized, Electronic Music Instrument." *Proceedings of the Second Annual Music Computation Conference, Part 4.* Urbana, Illinois: University of Illinois Office of Continuing Education in Music. (Reprinted in Rosenboom 1976a.)

Rosenboom, D., ed. 1976a. *Biofeedback and the Arts, Results of Early Experiments,* 2nd ed. Vancouver: Aesthetic Research Centre of Canada Publications.

Rosenboom, D. 1976b. "Prolegomenon to Extended Musical Interface with the Human Nervous System: An Outline Mandala of Electro-Cortical Forms Observable Through Point Consciousness." In M. Byron, ed. *Pieces: A Second Anthology.* Toronto: Michael Byron Publishers.

Rosenboom, D. 1976c. *Brainwave Music* (recording). Vancouver: Aesthetic Research Centre of Canada Publications.

Rosenboom, D. 1977a. *On Being Invisible* (recording). Toronto: Music Gallery Editions.

Rosenboom, D. 1977b. *On Being Invisible* (videotape). Vancouver: Western Front Video.

Rosenboom, D. 1984. "On Being Invisible." *Musicworks* 28:10–13. Toronto: Music Gallery Editions.

Rosenboom, D. 1987. "A Program for the Development of Performance Oriented Electronic Music Instrumentation in the Coming Decades: What You Conceive is What You Get. *Perspectives Of New Music* 25(1&2): 569–583.

Rosenboom, D. 1989. *Extended Musical Interface with the Human Nervous System: Assessment and Prospectus*. Leonardo Monographs. Fairview Park, Elmsford, New York: Pergamon Press.

Squires, K. C., et al. 1976: "The Effect of Stimulus Sequence on the Waveform of the Cortical Event-Related Potential." *Science* 193(9): 1142–1145.

Teitelbaum, R. 1976. "In Tune: Some Early Experiments in Biofeedback Music (1966–74)." in (Rosenboom 1976a).

Tenney, J. 1988. *META + HODOȘ and META Meta + Hodos, 2nd ed.* Oakland, California: Frog Peak Music.