LEONARDO MONOGRAPH NUMBER 1

EXTENDED MUSICAL INTERFACE WITH THE HUMAN NERVOUS SYSTEM

ASSESSMENT AND PROSPECTUS

David Rosenboom

LEONARDO MONOGRAPH SERIES
International Society for the Arts, Sciences and Technology
(ISAST)
San Francisco, California, U.S.A.
Original 1990, Revisions to Parts 5 and 6 and additional Appendixes, 1997

LEONARDO MONOGRAPH SERIES
MONOGRAPH NUMBER 1, 1990 & 1997

Coordinating Editor
Pamela Grant-Ryan

Manuscript Preparation for 1997 Version
Karen Beardsley

THE AUTHOR

David Rosenboom (b. 1947), composer, performer, conductor, interdisciplinary artist, author and educator, has explored ideas in his work about the spontaneous evolution of forms, languages for improvisation, new techniques and notation for ensembles, cross-cultural collaborations, performance art, computer music systems, interactive multi-media, compositional algorithms, and extended musical interface with the human nervous system since the 1960's. His work is widely distributed and presented and he is known as a pioneer in American experimental music. Rosenboom has been Dean of the School of Music, Co-Director of the Center for Experiments in Art, Information and Technology, and Conductor of the New Century Players at the California Institute of the Arts since 1990. He taught at Mills College from 1979 to 1990, was Professor of Music, Head of the Music Department, Director of the Center for Contemporary Music, and held the Darius Milhaud Chair from 1987 to 1990. He studied at the University of Illinois with Salvatore Martirano, Kenneth Gaburo, Lejaren Hiller, Soulima Stravinsky, Paul Roland, and Gordon Binkerd, among others, and has worked and taught in innovative institutions, such as the Center for Creative and Performing Arts at SUNY in Buffalo, New York's Electric Circus, York University in Toronto, where he was Professor of Music and Interdisciplinary Studies, the University of Illinois, where he was recently awarded the George A. Miller Professorship, New York University, the Banff Center for the Arts, Simon Fraser University, the Aesthetic Research Centre of Canada, the San Francisco Art Institute, and the California College of Arts and Crafts.

HISTORICAL NOTE

The original version of this monograph was written in 1989 and published in 1990. It's purpose was to document work that had taken place since the publication of my earlier book, *Biofeedback and the Arts, Results of Early Experiments* in the mid-1970's and the time of the original publication of this monograph. Subsequently, after all copies of the first edition had been distributed, the publication materials were lost in the tragic events of the Oakland fire. I have continued to receive many requests for copies of the monograph that I have not been able to fulfill. Now, with the advent of the World Wide Web and the assistance of ISAST, I am able to make a version of this document recovered from old computer files available again. Thanks are due to my assistant, Karen Beardsley, for conforming these old files to the published version as closely as possible. Hopefully, the ideas contained herein are sufficiently powerful to stimulate new ideas and inspirations.

International Society for the Arts, Sciences and Technology (ISAST)
425 Market Street, 3rd Floor, San Francisco, CA 94105, U.S.A.
1990 Version Copyright © 1990 ISAST. All rights reserved.
ISBN 0-9625355-0-8
1997 Version Copyright © 1997 David Rosenboom and ISAST
Printed in the United States of America

CONTENTS

Preface

The purpose of this monograph is severalfold: (1) to give a detailed description of some work done in the mid- to late-1970s in which I was able to achieve the spontaneous generation of formal musical architectures with a computer music system by using a detailed analysis of evoked responses to features in those architectures recorded from a performer's brain; (2) to provide an overview of some historical events related to the development of artistic works that are in some way responsive to bioelectrically derived signals; (3) to describe briefly the emergence of the biofeedback paradigm and to discuss biofeedback modeling; (4) to survey accumulated knowledge regarding interpretation of electroencephalographic phenomena with particular emphasis on event-related potentials (ERPs) and their relation to aspects of selective attention and cognitive information processing; (5) to present a speculative model for the general interpretation of electroencephalographic waveforms; (6) to discuss some inferences and speculations relating these phenomena to musical experience; (7) to provide an assessment of some methods and techniques that have been applied to realizing works of art with these phenomena; (8) to describe some specific algorithms for generating self-organizing musical structures in a feedback system that relates a limited model of perception to the occurrence of event-related potentials in a performer's brain; and (9) to discuss the potential of new and emerging technologies and conceptual paradigms for the future evolution of this work. Finally, an actual score containing a conceptual scheme for a biofeedback work involving electroencephalographic phenomena and electronic orchestrations is provided in an appendix to stimulate further thinking and ideas for applications in the arts.

The writing is addressed to those with an interdisciplinary interest in the arts (particularly music) and the sciences (particularly those of the brain, psychology and perception, and the study of self-organizing systems). However, readers whose backgrounds are in the arts or sciences alone, or even other areas such as cognition, philosophy, computer science or musical instrument design, are encouraged to read on as well. Many references are provided with which the reader may enhance her or his knowledge in a particular sub-discipline. Those who may find some of the technical descriptions difficult should first skim through the entire document and then return to individual sections for further study.

It is hoped that the ideas presented herein may contribute in some way toward increasing our breadth of understanding concerning dynamic processes in the arts and sciences.

Part 1—Historical Background

In the history of this youthful world, the best product that human beings can boast of is probably Beethoven; but, maybe, even his art is as nothing in comparison with the future product of some coal-miner's soul in the forty-first century.
-Charles Ives, 1920 [1]

Since the discovery of electrical pulsations arising from within the human brain, imaginative souls have speculated that through a direct connection of the brain to devices for sound production and visual display, internal realities would eventually be made externally, materially manifest. In turn, these would become enfolded through the senses into an evolving interplay among the fabricated models of cognition, the passages of consciousness, and the energetic, though capricious, environment. A global music, reflecting the morphodynamic holarchies of existence, might come into being.

My God! What has sound got to do with music!
-Charles Ives, 1920 [2]

EARLY SUGGESTIONS AND EXPERIMENTS

In a now-famous 1934 paper, the pioneering physiologists E. D. Adrian and B. H. C. Mathews reported on experiencing a translation of the human electroencephalogram (EEG) into audio signals. While listening to his own 'alpha rhythm' (large-amplitude smooth waves of 8-13 Hz) presented through a loud speaker, Adrian tried to correlate his subjective impression of hearing the alpha come and go with the activity of looking or not looking with his eyes [3]. Interestingly, this alpha rhythm, as it came to be called, was originally known as the 'Berger rhythm'. An inherently rhythmic quality had been observed in waves of electromotive force first detected on the scalp in the 1870s. However, it was Hans Berger who in 1929 provided the first comprehensive description of the human electroencephalogram [4]. Using crude, early instrumentation, Berger recognized spontaneous oscillations in the EEG, particularly when detected over the occipital area of the cortex (back of the head), where they are manifest with relatively large amplitudes. Since the occipital area is most heavily involved in the processing of visual information from the eyes, this Berger rhythm was subject to variation associated simply with opening and closing one's eyes. Most likely, the significance of the term 'alpha' lies in nothing other than that this 'rhythm', a relatively repetitive and coherent waveform, was just the first to be detected with early instrumentation. These instruments were relatively insensitive to other EEG components of much lower amplitude. During succeeding decades, numerous other scientists reported various methods of generating, in the auditory domain, stimulus frequencies that in some way followed brain rhythm frequencies. Usually this involved some aspect of alpha rhythm [5-9]. These methods were considered an aid in analyzing what was obviously a complex electronic manifestation of brain activity.

This use of auditory translations of EEG patterns allowed observers and investigators to employ the considerable integrative powers of auditory perception and feature extraction to guide them toward some insight into the form of these signals. Today we might listen to the sonic translations of complex number sequences to divine mathematical patterns that might otherwise go undetected. I have personally experienced the phenomenon of discovering and squashing elusive and subtle 'bugs' in large computer programs simply by listening to an audio translation of the raw machine code of these programs [10]. Stories are sometimes told of similar events achieved by programmers in the early days of computing, when we used to sit around and watch running program code displayed on old-fashioned panel lights. Other programmers have also been observed indulging in this kind of sensory analysis when all other systematic methods of axiomatic computer science have failed. The powers of the sensory-feature extracting mechanisms and the integrative powers of high-level image synthesizers in the brain are awesome. They simply await our conscious probing and comprehension through open focus, discipline and practice. We also live in a fantastically rich contemporary music milieu in which, as musicians, our ears are evolving even greater powers to help us manage sometimes immense and deep formal architectures. What we may yet discover by listening to our own biological computer, if I may invoke such a simple analogy, is practically unfathomable.

FIRST APPLICATION IN THE ARTS

The observation can be made that, throughout the history of advances in science and technology, artists have always been ready to experiment with applications of each new breakthrough or development, almost as soon as it is conceived or realized. Brain science proves no exception. In the past 25 years, composers like Alvin Lucier, Richard Teitelbaum, myself and others have produced major works of music with EEGs 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 [11-13]. Greatly amplified alpha signals were used to activate, either acoustically or mechanically, an array of otherwise performerless percussion instruments. This produced the startling effect of a percussion ensemble seeming to activate itself, almost invisibly, but somehow following activities inside the solo performer's mind. Teitelbaum's *Organ Music* and *In Tune*, both realized in 1968, added heart beat and breath sounds, sensed with contact microphones, to EEG signals in the creation of an electronic music texture [14]. 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 1968-1969 in the laboratory of Les Fehmi, an early biofeedback researcher at the State University of New York at Stony Brook, after a suggestion by E. E. 'Ted' Coons of New York University. I developed an environmental demonstration-participation-performance event entitled *Ecology of the Skin* in 1970-1971. It involved biofeedback monitoring of brainwaves and heart signals from performers and audience members and their translation into a musical texture, along with synchronous electronic stimulation of visual phosphenes (colored patterns often seen with eyes closed) 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 10 participants according to a simple statistical measure, the amount of time spent per minute producing alpha waves (see Fig. 1).

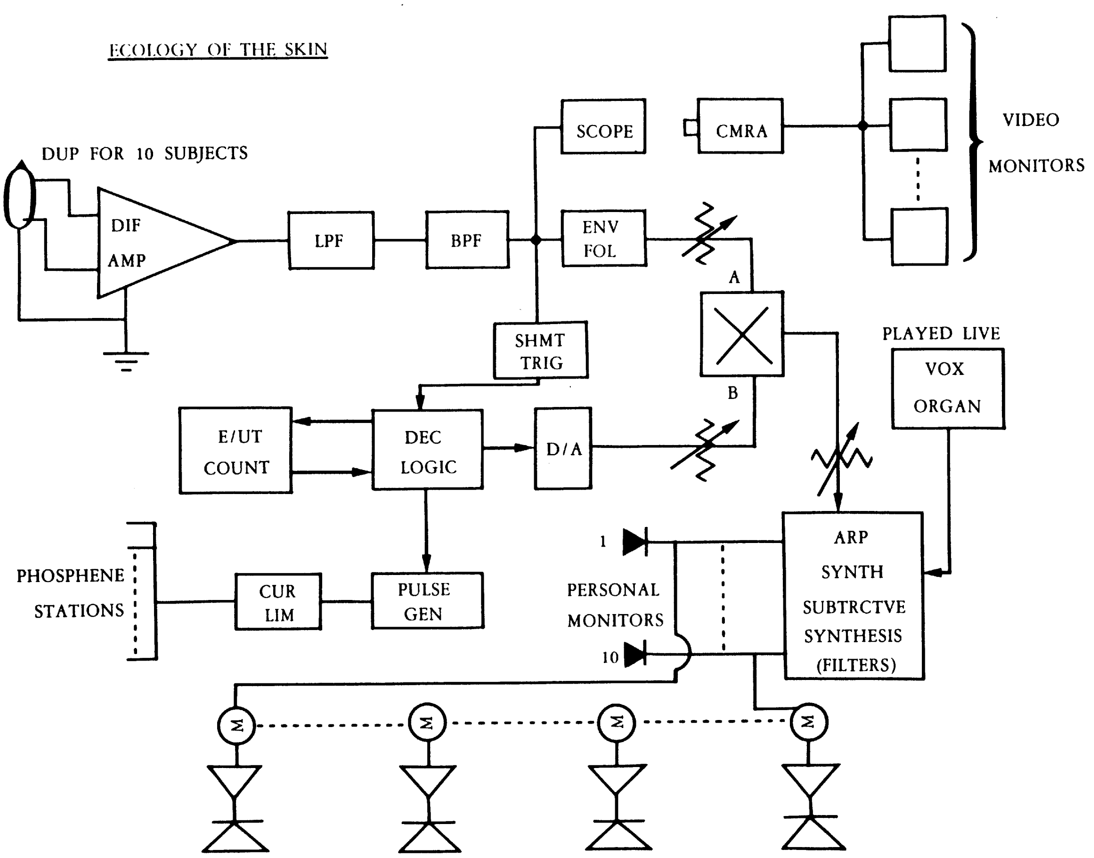

Ecology of the Skin (1970)

Fig. 1. *Ecology of the Skin* (1970). Setup diagram for 1970 biofeedback installation/performance/participation work, *Ecology of the Skin:* DIF AMP = differential brainwave amplifier; LPF = low pass filter; BPF = band pass 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 converters; CUR LIM = current limiting circuit; CMRA = video camera.

Eventually, my work led to the creation of a laboratory—the Laboratory of Experimental Aesthetics at York University in Toronto with the intent to study information-processing modalities of the nervous system in relation to aesthetic experience and states of consciousness while surrounded by an environment of artistic production. Under the sponsorship of York University, the Canada Council Explorations Programme and the Aesthetic Research Centre of Canada, many individuals carried out experiments and produced artworks there over a 7-year period. These works—representing such art forms as music, visual art, kinetic art and dance—are documented in my book, *Biofeedback and the Arts, Results of Early Experiments* [15].

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 [16-18]. These signals resulted from measuring the EEG, pulse rate, respiration, galvanic skin response (GSR), blood flow volume, and the electrocardiogram (EKG). A variety of multi-sensory 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.

EMERGENCE OF THE BIOFEEDBACK PARADIGM

The biofeedback paradigm began to be clearly articulated when the work of Neal Miller, of Rockefeller University in New York, became widely known in the 1960s. Miller had devoted over 30 years to the study of animal and human learning and had gained the respect of the majority of investigators in his field. His work at this time undermined the long-held image of the nervous system as being divided into two inflexibly separated parts—the voluntary and involuntary—one subject to learning and conscious control, the other capable of executing only automatic, built-in programs. Miller, along with his associate, Leo DiCara, demonstrated that, indeed, animals and human beings could learn voluntarily to influence the behavior of bodily functions such as heart rate, blood pressure, vasomotor responses, electrodermal activity, salivation, urine formation, gastric motility and metabolic processes previously thought immune from conscious influence. This was applied to a host of visceral phenomena and was dubbed, 'visceral learning' [19].

In animals, these learning phenomena were often achieved by means of providing some form of reward for producing the desired response. In humans, the learning paradigm was dependent on presenting information to the subject about the states and trends of change occurring in these visceral phenomena, through some form of sensory feedback, while relying on the subject's own internal motivation to effect the change. Thus, the term 'biofeedback' was coined.

Biofeedback—Definition and Modeling

The term 'biofeedback' will be used herein to refer to the presentation to an organism, through sensory input channels, of information about the state and/or course of change of a biological process in that organism, for the purpose of achieving some measure of regulation or performance control over that process, or simply for the purpose of internal exploration and enhanced self-awareness. Normally, this information will be of a type not otherwise available to that organism. It does not presuppose, however, that such an external indicator could not, through disciplined practice, be replaced by an internal mechanism of which the subject can achieve awareness without the aid of an artificial monitoring system.

Over a decade ago, the feedback paradigm, as understood from cybernetics, was being considered to offer alternatives to the previously dominant behaviorist school of psychology [20]. Behaviorism could be described as an open-loop paradigm relating behaviors (effects) to causes (stimuli). The feedback paradigm is a closed-loop one in which behaviors (effects) are also treated as one of several possible classes of causes (stimuli) of the same behaviors. Cause and effect can be traced all around the closed, feedback loop, creating a 'chicken-and-egg' problem as to which came first. This problem is not always resolvable, or even a meaningful one to pose. A *process* -oriented approach to modeling the self-organizing dynamics of a biofeedback system may prove more fecund.

It may be useful to recall the basic feedback paradigm as we have inherited it from cybernetics (see Fig. 2). In the simplest case, a defined goal state is assumed for a process, the output of which is measured and compared against the desired end (Fig. 2a). An error correction signal is then generated and combined with the input in such a way as to direct or tune the process so that it eventually arrives at the goal state. The cases of both negative feedback, which can lead to stabilization of particular behavioral states, and positive feedback, which normally leads to instability and/or system energy saturation, are shown (Fig. 2b and 2c). A partial view of learning may be described as the tuning of [K] such that the desired output is maintained with minimum expenditure of energy, while making [e], the error signal, tend toward 0. This is homeostatic behavior. In a biofeedback system we may conceive of a statistical distribution of behaviors along some relevant parametric axis, as shown in Fig. 3a, in which the desired goal is represented by the mean parametric value and the achievement of the goal by the variance around the mean. If in such a closed-loop system the probability of occurrence of a particular behavior is also increased as a result of the execution of that behavior, then we have a resonant system. In such a system we have negative feedback, which attenuates behaviors away from the mean goal state by means of an error signal, but we also have amplification of the behavior associated with the goal state. A system like this will be driven into saturation and possible instability unless the amplification of the goal state behavior in the feedback loop is mediated by other concomitant goal states or input from the environment outside the system. Lack of such mediation of positive feedback amplification may underlie such phenomena as addictive or obsessive behavior, the tendency of societies to self-destruct through degradation of their support environment or through unrestrained conflicts associated with out-of-control resonances in their political and economic systems.

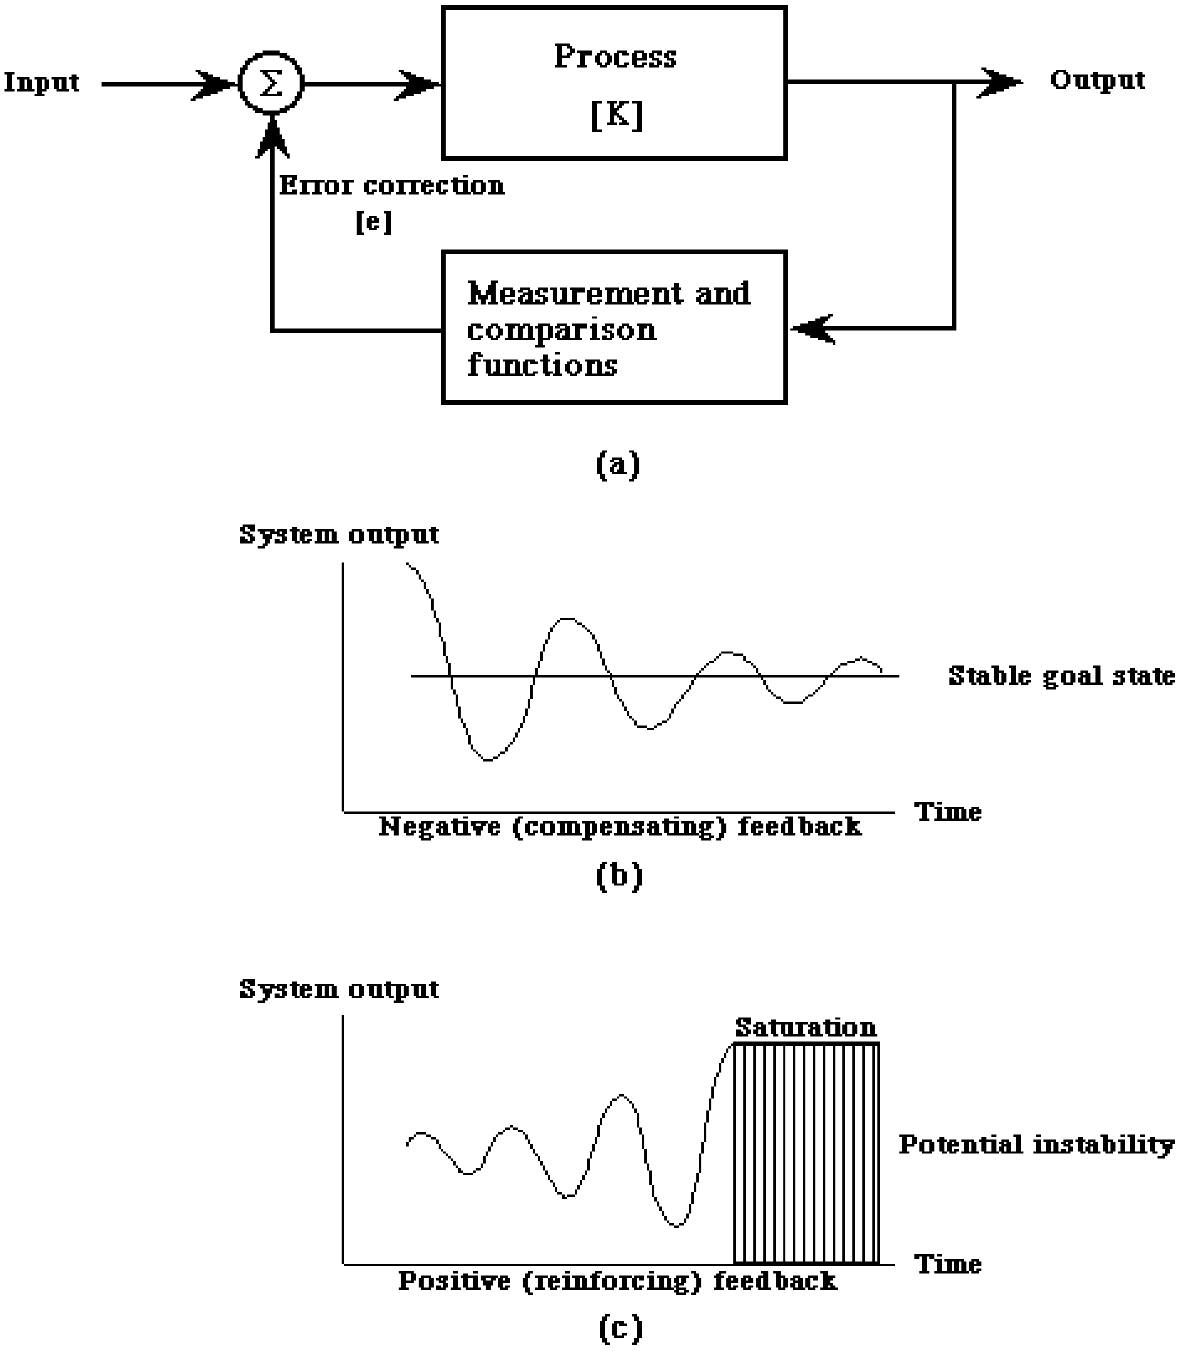

Basic Feedback Paradigms

Figure 2

Fig. 2. Basic Feedback Paradigms. (a) Feedback control mechanism in which an error signal is generated and added to the input in order to direct the system to a desired goal state. (b) Effect of negative feedback on system output. (c) Effect of unmediated positive feedback.

Today we have the language of dynamical systems theory to add to our lexicon of paradigms with which to probe behavioral patterns [21]. In a dynamics characterization of biofeedback, we would describe particular goal states as *attractors* in a behavioral phase diagram [22]. Figure 3b shows a modified statistical distribution in which the mean goal state is represented as an attractor bounded on either side by two *repellers* characterizing the state boundaries. This kind of lateral inhibition, which sharpens the definition of the desirable state, is seen in Fig. 3a to be the result of summing two Gaussian functions—the original, facilitative, behavioral distribution one, and a broad, negative, inhibitory one [23]. Behavioral patterns would be characterized by collective variables known as *order parameters*. The nature of the order parameters would be specific to the types of functions and tasks involved. The success with which this method provides a useful description of the system under scrutiny will depend, in part, on determination of the right order parameters to use. This is an area where creative solutions must often be found. Well-defined behavior patterns are attractors in the phase diagram, corresponding to stable collective states of the order parameter dynamics. Figure 4 shows a hypothetical behavioral phase space with three attractors. Behaviors occupying neighboring points in the phase space will eventually converge to the attractor region (basin of attraction). All types of attractors, *static* (fixed point, a set of dimension zero), *periodic* (limit-cycle, closed curve, a set of dimension one) and *chaotic* (where small uncertainties in initial conditions lead to large uncertainties in future behavior), as well as multiple attractors (multi-stable states), may reside in the phase diagram. Chaotic attractors may also have non-integer, *fractal*, dimensions and may be referred to as strange attractors. The system may be perturbed by environmental input or noise and shifted away from an attractor. Its characterization will then include various time constants, such as local or global relaxation time, the time required to settle toward an attractor state again. Probability distributions may be used to describe fluctuations around an attractor state (goal or mean), wherein variance corresponds to degree of fluctuation. The more stable an attractor is, the smaller the average deviation from the attractor state will be when a perturbing force of a given strength impinges upon it. It is outside of the scope of this paper to outline the analysis techniques of dynamical systems theory. However, its methods and language may lead us to a deeper understanding of the spontaneous generation of patterns, changes of state, and the self-organizing nature of human behavior, cognition and consciousness. I will return to the use of this language numerous times.

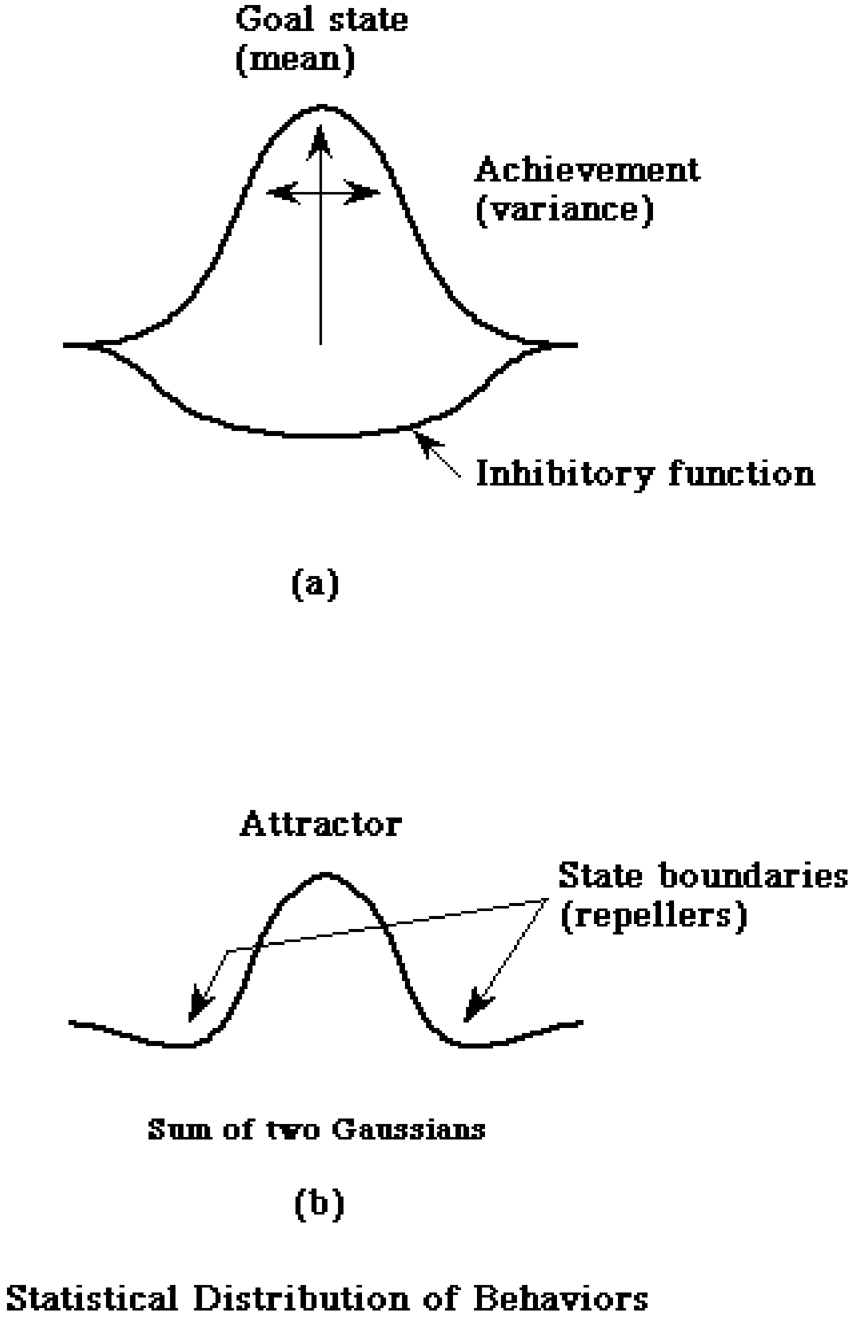

Figure 3

Fig. 3. Statistical Distribution of Behaviors. (a) Description for the behavior of a feedback control system using the language of statistics. The desired goal state corresponds to the mean of a Gaussian distribution of behaviors and achievement of the goal to the variance. A negative inhibitory function is also shown. (b) When the two Gaussians of (a) are summed, a function results with a positive peak at the mean or goal state surrounded by negative, inhibitory boundaries.

A biofeedback system may be mistakenly viewed as simply a method for stabilizing particular behaviors and, thus, as a static equilibrium system. Indeed, a surface understanding of the therapeutic application of biofeedback techniques as mediators of runaway processes deleterious to an organism, such as epileptiform brainwave patterns, heart arrhythmias, high blood pressure or uneven vasoconstriction associated with migraine headaches, may reinforce this wrong view. In fact, these processes all involve the self-organization of dynamical regimes within the organism, aided by the additional information feedback loops of the biofeedback mechanism, in such a way that the evolution of these regimes will tend toward a dynamic that promotes the self-renewal of the organism. The biofeedback system may qualify for description as an autopoietic system, capable of self-renewal merely by employing a process of self-reference. A biofeedback mechanism participates in the interaction dynamics of an autopoietic organism with its environment. Since this organism is in itself a complex, dynamical system capable of a certain degree of self-determination regarding how the information circulating in a biofeedback loop is used, a simple cybernetic view of biofeedback as a positive feedback loop is insufficient. Both positive and negative feedback functions may occur in the structure of the system's interaction matrix. Consequently, biofeedback, in its contemporary, sociocultural context, must be viewed as our participation in the 'process of self-reference' of dissipative, autopoietic organisms in interaction with an environment from which they import energy and to which they export entropy. In this light, the effect of positive feedback can be seen to be canalized (directed along chreods, lines that guide an organism's ontological development by the functions of the healthy organism in such a way as to lead to bifurcations in the evolution of the structure of behavior. The information circulating in a biofeedback process loop, therefore, can never be viewed as a static entity. The information, too, is subject to its own self-organization. Consequently, I will always refer to 'formation' as something that is 'in formation', that is, 'in the process of becoming formed'.

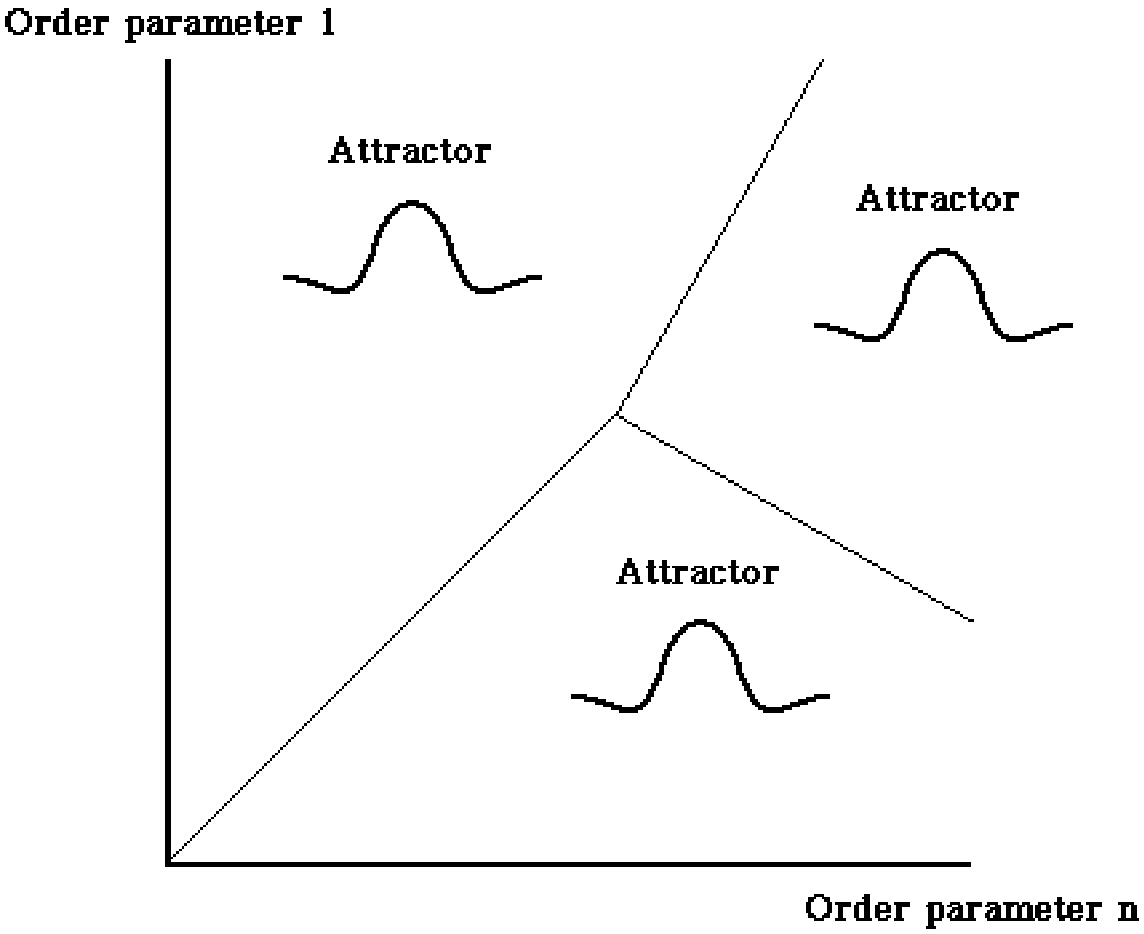

Hypothetical Behavioral Phase Space

Figure 4

Fig. 4. Hypothetical Behavioral Phase Space. Depiction of a multi-dimensional, behavioral phase space in which the order parameters correspond to measurement scales needed to characterize the system behaviors. Three attractors are shown separated by boundaries, implying that system behaviors will tend to fall into one of the three behavioral regions. Behavior trajectories—describing movement among attractor regions—may be complex and subject to description in the language of dynamical systems theory.

Eric Jantsch has provided one of the most eloquent characterizations available regarding the nature of self-organizing systems [24]. However, even such a brilliant author as he succumbs to a major blunder by falling into a common trap of popular misunderstanding (which many scientists have) about the biofeedback paradigm. This is quite understandable, since biofeedback was interpreted in an almost science fiction-like manner during its faddish period of the early

1970s. A misrepresentation developed on which numerous short-lived careers of academic and commercial charlatans were based. Biofeedback was commonly seen as the extension of a conscious control hierarchy over processes of the body, achieving "Total control of rational thinking over the body!" [25]. Jantsch was horrified by such a prospect. This fearful misrepresentation must be debunked. In fact, it has been amply demonstrated that in many examples of biofeedback processes, such as with alpha wave production in the brain, control cannot be achieved by means of rational processes. Rather, the subject must find a way to *allow* alpha wave production to evolve rather than *make* it appear. I suggest that the Zen-like state associated with achievement of what we may wrongly associate with the word 'control', is, in fact, a striking example of the quality of subjective experience associated with true conscious participation in autopoietic self-organization, including feedback with the environment. Jantsch himself provides an appropriate characterization of this experiential quality. "To live in an evolutionary spirit means to engage with full ambition and without any reserve in the structure of the present, and yet to let go and flow into a new structure when the right time has come" [26]. This strikes at the essence of experience associated with my ongoing musical work, *On Being Invisible*, described at length later in this paper. We must therefore modify and evolve our characterization of biofeedback. It should be seen as the circulation of information about functions within an organism in ever widening feedback loops involving the consciousness of that organism, to serve its creative extension, beyond the structure of its own prior self-definition, in the natural meta-evolution of its self-organization dynamics—and not merely as the extension of a conscious control hierarchy. It is, rather, the holarchic extension of dynamical processes of the self in which the conscious mentality of the individual may responsibly choose to participate as a manager or catalyst. Thus, conscious human self-management is seen as an autocatalytic agent in its own evolution. (See Part 8, *A Note on Musical Holarchies,* for a brief discussion of the meaning of *holarchy*.)

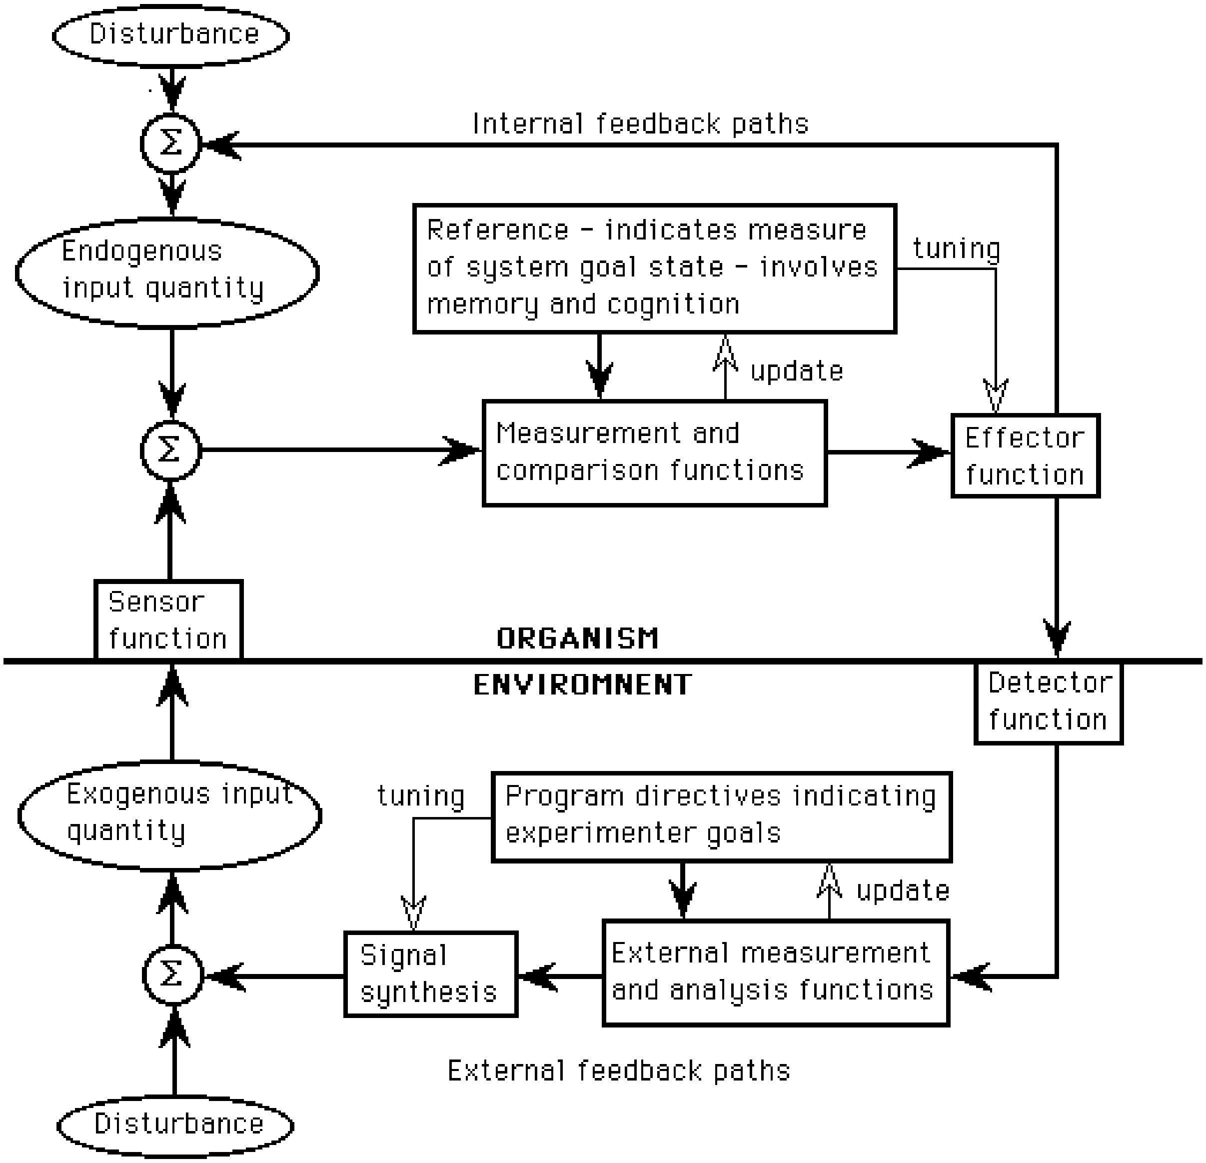

Self-organizing Biofeedback Mechanism

Figure 5

Fig. 5. Self-organizing Biofeedback Mechanism. Schema for the flow of information among components comprising a self-organizing biofeedback mechanism, including those internal to the organism and those typically present in an experimental environment. Many setups described in the text for artistic production or research include some or all of these components.

A more complex and detailed view of a biofeedback mechanism capable of self-organization is shown in Fig. 5. A large, global feedback loop, in which information circulates across the boundary separating the internal functions of the organism from the external environment, is shown. This boundary is not conceived as one generating any significant degree of autonomy of the organism from its environment. Rather, it is merely an artificial construct made relevant by the need to manifest relationships between that which is sensed as information impinging upon the organism and actions synthesized by the organism. In reality, however, the organism and its environment are better conceived as a global system of mutually coupled transformation processes, inside each of which further feedback loops are active. The organism's *sensor function* transforms energy/information patterns from the environment. The results are combined with an *endogenous input quantity* that, in turn, results from the sum of information circulating in *internal feedback paths* plus *disturbances* arising from within the organism. This data is subject to measurement and comparison functions involving memory and cognition. The results direct and tune the organism's *effector function* and, possibly, update memory. The environment in Fig. 4 is assumed to contain an intelligent process, possibly an experimenter with programmable apparatus. A *detector function* transforms information resulting from the organism's behavior. This is subject to further measurement and analysis. The result is compared with directives arising from programs or goals within the intelligent system. As a result, a *signal* intended for the organism is synthesized and, possibly, the program updated. The signal, summed with environmental disturbances, forms the *exogenous quantity* input to the organism.

Subsequent Developments Outside the Arts

Subsequent to Miller's work on visceral learning, the biofeedback paradigm was applied to functions of the organism that manifest observable effects in the electro-chemical operation of the human brain. The most obvious first candidate for exploration was the old Berger rhythm, now known as the alpha wave. Joseph Kamiya, Les Fehmi and Thomas Mulholland were but three prominent names among the scores of early pioneers. They were followed by many more. The literature expanded at a phenomenal rate throughout the 1970s. Books and articles flooded publishers' desks, the international press and media took a keen interest, careers were made and broken, and biofeedback became a fad. Underneath all this flurry, however, a steady stream of solid research and thinking flowed, though sometimes masked by opportunistic hyperbole. Some important repositories of documentation of this research include the Aldine Atherton annuals as well as *Biofeedback and Self Control* (begun in 1970 by Aldine Atherton, Chicago), the journal *Biofeedback and Self-Regulation* (begun in 1976 by Plenum Press, New York), and the excellent series *Consciousness and Self-Regulation* (begun in 1976 by Plenum Press, New York).

During the later part of the 1970s, things settled down somewhat. The chaff fell away, leaving the more solid and resilient research programs to produce of reliable results and a methodology.

In this paper I will not discuss the field of visceral learning, nor will I cover the medical or therapeutic aspects of biofeedback. I will concentrate primarily on the use of information that can be extracted from the brainwave EEG for artistic or musical purposes.

Some Applications and Cultural Implications

To date, considerable success has been achieved in medical, therapeutic and self-improvement applications. Biofeedback has been used as an aid in controlling blood pressure abnormalities and heart arrhythmias, as a treatment for migraine headaches, and for suppression of epileptiform patterns in brainwaves. It is used extensively as a treatment for hypertension and as an aid in reducing stress. It has become a basis for certain kinds of self-improvement training programs, such as those for achieving control of mental states, focus, attention, relaxation and self-integration. Complementary exercises in tension release, visualization, meditation and open-focus training also are often employed [27]. Biofeedback with EEG parameters has added to our knowledge about consciousness, states of awareness and cognitive processes.

Originally touted by the press as a panacea for all that ails and the key to self-transformation, biofeedback is now perceived in a more sober light. However, biofeedback raises issues of self-consciousness that do not fit neatly into Western culture. The achievement of success with biofeedback requires discipline, intense and regular practice, and often meditative skills. These were consistent with views held in the 1960s of transcendence and the idealism of cultural transformation. These ideals faded with the rise of 'yuppi-dom' in the 1970s, as disillusionment grew when earlier hopes for change were seen to fail or to be forgotten and, in the 1980s, as self-realization was replaced by the necessity of socio-economic self-validation. In such a climate, lack of further substantive progress in applications can only be blamed on an unwillingness to pay the price of personal hard work to achieve transformation.

EXPANSION OF APPLICATIONS IN THE ARTS

In the arts, it is not difficult to find individuals willing to apply themselves to the serious exploration of such phenomena. Transformation and personal explorations are a mainstay for experimentalists in the arts, the food of progress.

In Music

At first, the greatest expansion of artistic activities involving biofeedback occurred in the field of music. Teitelbaum's *T'ai Chi Alpha Tala* (1974), developed in our Laboratory for Experimental Aesthetics, involved transmission of alpha signals from an artist, Barbara Mayfield, engaged in the practice of T'ai chi Chu'an, by means of a tiny, encapsulated brainwave amplifier and FM transmitter attached to the artist's head. During the performance, alpha signals were extracted and used to trigger electronically synthesized melodies tuned to an Indian mode, while South Indian mrdangam master Trichy Sankaran instantly analyzed and embellished the resultant rhythmic patterns. In a later version, synchronized video processing of the performers' images was added in collaboration with video artists Dan Sandin and Jim Wiseman. Since that time, musical experiments and performances have been generated by many others in virtually all stylistic arenas, from contemporary popular music to the avant garde. Most involved the relatively direct following of some EEG phenomenon, usually amplitude of alpha waves, by a musical parameter, such as melodic shape or tone color. These processes were imbedded in a great variety of cleverly conceived, creative performance styles. Lucier's *Clocker* (1978) presents audiences with an intriguing image in which the progress of clock time seems to be warped by a performer whose GSR is used to change the settings on a time-delay device being fed the sounds of clock ticks [28]. A connection is established between changes in emotional states reflected by the GSR and the passage of subjective time. My own work in this area evolved throughout the 1970s and is documented [29-40] in several books, recordings, articles, television and video productions and films, as well as later in this document.

In Kinetic and Performance Arts

Applications in kinetic arts often involve installations or applications in performance. Pieces involving brainwave manipulated visual displays and interactive environments have been created by artists Jacqueline Humbert and C. Mark Nunn as well as myself [41]. Biotelemetry has been used in theater environments. One notable experiment, organized by Richard Lowenberg and associates in California, involved brainwave, muscle and accelerometer signal measurements from two groups of dancers, one located on the West Coast and one on the East Coast of the U.S.A. The signals were translated into sound and video displays and transmitted between the two groups by means of satellite communications provided courtesy of NASA. At the 1986 New Music America Festival in Houston, a dramatic show piece by artist Stelarc was created involving a large array of bio-sensors, a robot-like prosthetic arm, interaction with flashing lights and laser beams, and massively amplified sound [42].

In Dance

Biotelemetry techniques for monitoring and transmitting muscle signals (EMG) and brainwaves (EEG) from dancers have been developed and applied in performance. Some of these have also been applied in the kinetic arts. One recent example is a system called MADDM (myoelectronically activated, dance-directed music system), in which techniques have been developed for detecting and processing signals from several muscle groups and transmitting them to a computer music system. In this way, dance movements are given control over the playing of a musical score [43].

EVOLUTION OF THE AUTHOR’S WORK

My own work in biofeedback and the arts, begun 20 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 EEG signals. These were previously achievable only with cumbersome, non-real-time, laboratory-bound methods. Consequently, ideas which were impractical when they were proposed many years ago are now practical.

My earliest experiments involved the kind of simple parametric following referred to earlier. Soon, however, it became apparent that data from deeper statistical analyses of EEG trends would provide more meaningful signals with which to control musical forms. In *Portable Gold and Philosophers' Stones* (1972) [44, 45] 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 was 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 band-pass filters, in a device known as a 'Holophone'. This helped determine the spectral composition of the music (see Fig. 6). Measures of body temperature and GSR were also used to direct the tonality of the musical texture.

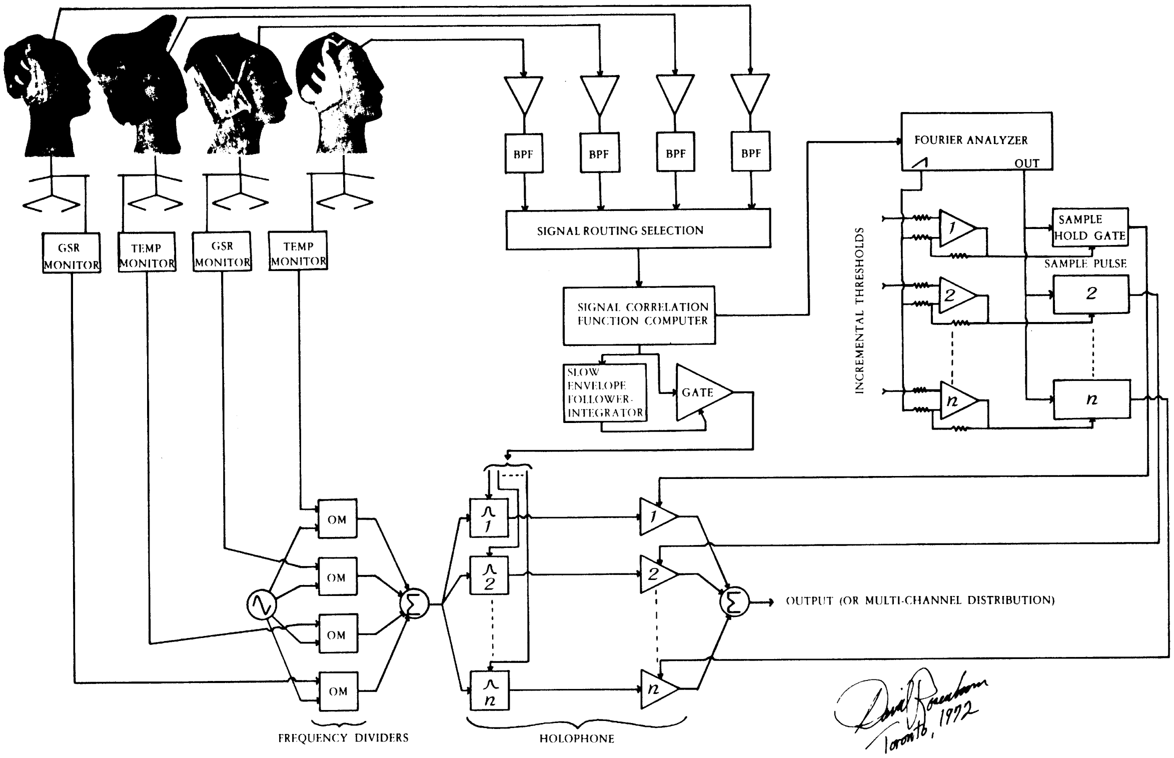

System Diagram for Rosenboom's *Portable Gold and Philosophers' Stones (Music from Brains in Fours) (1972)*

Figure 6

Fig. 6. System diagram for *Portable Gold and Philosophers' Stones (Music from Brains in Fours) (1972)*. A diagram from the score of a musical composition that includes measurement and analysis of EEG signals, GSR and body temperature changes from a quartet of performers. Equipment for signal analysis is operated by a fifth performer, who also applies the results to a sound synthesis system. The frequency dividers and holophone produce a harmonically ordered musical texture, which is made to evolve according to changes in the bioelectrical information. See the text and references [11] and [31] for detailed explanations.

Vancouver Piece (1973) involved an exploration of subtle, visual thresholds in a biofeedback installation work [46, 47]. Pairs of participants facing a two-way mirror system observed the images of their own faces becoming superimposed on each other's bodies, producing a kind of identity exchange, when they were able to produce EEG waves, such as alpha, that were in phase with each other. Further coherent wave production caused faint wisps of light to trace out horizon lines at intensity levels near the threshold of perception, around an otherwise darkened, light- and sound-isolated room. Musical textures were also produced in response to in-phase alpha wave production from the two participants.

Many other art works and research programs were carried out during the 1970s. (See references listed in the section *In Music* in Part 1 for more examples.) The culmination of these musical applications was the production of *On Being Invisible* (1976-1977). In this work, detailed at length below, complete musical forms are constructed as a result of the self-organizing dynamics of a system in which both ongoing EEG parameters and event related potentials (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.

It is possible now to imagine large-scale, musical theater or operatic works involving biotelemetric presentation by human and even non-human 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,* which awaits adequate time and support for its full realization.

Part 2—Some Bioelectromagnetic Phenomena of Significance to Paradigms of Feedback-Based Self-Organization

Here is a partial list of some interesting phenomena to explore with feedback, which can be detected relatively easily and for which the measurement techniques are practical: EEG (electroencephalogram, brainwaves), EMG (electromyogram, muscle signals), EKG (electrocardiogram, heart muscle signals), EOG (electro-oculogram, eye muscle movements), GSR (galvanic skin response, electrodermal, electrical skin resistance), ERG (electroretinogram), respiratory rate, body heat sensing at particular locations with thermistors, infra-red mapping (body heat profiles), and body movement tracking with video image analysis. All of these offer significant potential for probing aesthetic processes or for artistic production. In the remainder of this paper, however, I will concentrate primarily on the EEG, since it is packed with information about functions in the nervous system.

ELECTROCARDIOGRAM—A SPECIAL NOTE

Research has shown that attention to an external stimulus is reflected by specific profiles in heart rate change. Generally, such attention allocation is associated with a mean drop in heart rate. However, two differentiated conditions for such attention, one involving recognition of a 'clue' in an ongoing stimulus stream and one involving later recall of the 'clue' showed clear signatures in heart rate change. Both involved a brief acceleration upon detection of the clue, followed by a significant deceleration and an acceleration rebound. However, the profiles were distinctly different for the two conditions [48].

THE ELECTROENCEPHALOGRAM IN DETAIL—CATEGORIES OF USEFUL MEASURES

The human brain produces a complex, multi-dimensional, pulsating, electromagnetic field, resulting from the electro-chemical behavior of masses of neurons acting in small to very large groups. Electrical currents are set up by means of the transport of ions across the cell membranes of these neurons. Changes are produced in the electrical potential across the membranes, measured inside the cell with respect to the outside. Current inside the cell is referred to as 'source current', current outside the cell as 'volume current'. Ion transport across a cell membrane usually begins at the dendrite end and proceeds throughout the cell body, possibly producing an action potential, which travels down the axon. Among a group of cells, this creates a population of current dipoles.

Outside the cranium, voltage differentials can be measured between any two points on the surface of the scalp or between a set of scalp points and some common, neutral reference on the body relatively distant from the scalp (such as the ear lobe). The voltage gradients that can be recorded from one of these points provide a rather one-dimensional look at the complex topology of surface pulsations that reach the outside of the skull. The transmission of internally generated currents to the extra-cranial surface is confounded by a number of factors. For example, the interior environment consists of a region of relatively high conductivity, the cerebral-spino fluid, which is surrounded by a container of low conductivity, the skull, which in turn is surrounded by a skin of relatively high conductivity, the scalp. The effect of this is to impose a low-pass limitation on the EEG bandwidth and to topographically smear out the waveforms observed, such that localization of the current source-sinks that originate the pulsations becomes extremely difficult. We could call this *slow-pass filtration and spatial diffusion*. Nevertheless, considerable information can be extracted about ongoing internal activities, just as the tracking of slow seismic waves can reveal much about the makeup of the earth's interior.

The EEG waveforms can be decomposed into a set of useful categories. I proposed a general schemata for this categorization in 1976 [49] and will attempt to give a more technical explanation here. In this schemata, the brain is viewed as a generator of locally singular events (see Fig. 7). It is considered a manifestation of process rather than an object. It functions as an open system, a dissipative structure, within a non-localized environment, containing universally distributed processes and a background, whether random-seeming or partially ordered. The categorization represents that which can be observed from what is termed *point consciousness*. I use this term to refer to a state of consciousness activated to produce cognizance of clear, localized, spatio-temporal constructs acting as a frame of reference in which to locate precisely the objects of experience as focal points of attention and awareness. It is assumed there are other states of consciousness, less involved with such spatio-temporal localization and the identification of singularities of experience. The brain's own ontology and all of its animate processes unfold within a continuum that extends from its environment inside to outside. The brain, then, is considered a point concentration of a relatively universally distributed process within its environment. This, in more concrete terms, also applies directly to the electrical manifestation of brain processes. These are electromagnetic field phenomena occurring within and interacting with other fields of varying extent and central strength or density. Indeed, we may define a universe, or some subsection of it we wish to address, as a single field, albeit one full of holes, and our objects of attention as point concentrations within that field. The components of the EEG we choose for study can be considered kinds of resonances with varying amplitudes, band widths, coherence and persistence. Their relation to mental events or other aspects of experience is the object of our exploration.

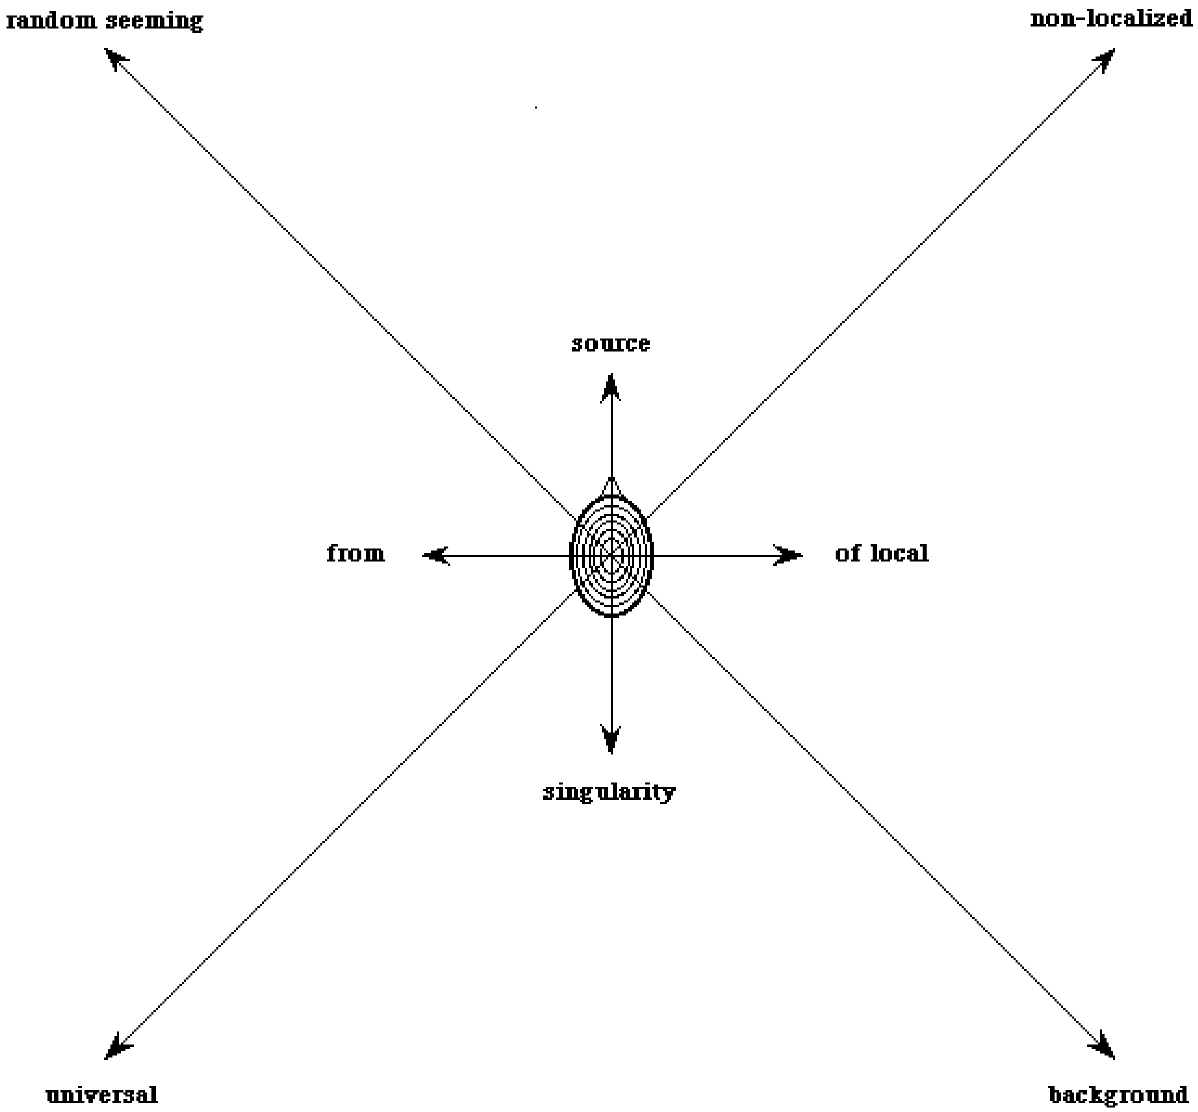

The Brain as Point Concentration
of Universally Distributed Process

Figure 7

Fig. 7. The Brain as Point Concentration of Universally Distributed Process. The brain viewed as both an originator of singular events or experiences and as a particular, local concentration of processes requiring a larger environmental context in which to be expressed. A seemingly ubiquitous background noise—similar to that thought to arise from residual activity left over from the origin of the universe—is included.

Random-Seeming Background

After all conceivable methods of EEG waveform decomposition are exhausted, there remains a background of random-seeming phenomena, about which little is understood. An analogy can be drawn with the background noise permeating the universe. It may be that the mass action of neural circuits exhibits behavior like that of certain intrinsically random, dynamical systems [50]. From this may come irreversibility and instability, both necessary for ontological evolution to occur.

Long-Term Coherent Waves

The next category is designated *long-term coherent waves*. It includes the traditional brain rhythms-alpha, beta, delta and theta. Emphasis is placed on the statistical property of *coherence*, in order to distinguish this idea from the mere presence of wave energy in a particular frequency band. Coherent waves are relatively stable, exhibiting high autocorrelation. There is a potentially serious pitfall associated with the use of simple band-pass filters to detect traditional brain rhythms. One may detect the presence of signal power in a particular filter pass-band, or even by means of Fourier analysis, in relatively incoherent, noisy signals with low autocorrelation. Since these coherent waves may result from widespread cortical synchrony among populations of neural circuits, it is critical to be able to distinguish this from pass-band energy resulting from broad-band noise. When EEG biofeedback information is derived from use of band-pass filters alone, experimental results can be confounded. This is particularly true for feedback of beta wave information, as discussed below.

Coherence is measured by the standard autocorrelation technique, in which a time-slice of the ongoing signal is compared with a series of successive time-delayed versions of itself, (see Fig. 8a). A curve is obtained, plotted as a function of time delay. Stable, repetitive patterns in the ongoing signal will periodically reinforce themselves during the processes of delay and compare. These will be revealed in the autocorrelation function, even if the original patterns are of very low amplitude and are buried in noise. Unstable, quasi-repetitive patterns will show up but will decay in amplitude over the delay axis, (see Fig. 8b). The rate of decay gives a measure of *coherence time*. Random signals will produce very rapidly decaying functions.

A form of the standard autocorrelation function for an input, time function, $f_A(t)$, is as follows:

$$C_{AA}(\tau) = \lim_{T_i \to \infty} \frac{1}{2T_i} \int_{-T_i}^{T_i} f_A(\tau) f_A(t - \tau) dt \quad ,$$

where,

$C_{AA}(\tau) =$ the computed autocorrelation function,
$T_i =$ the time period over which the integral is computed, and
$\tau =$ the time delay parameter.

Note that cross-correlation, $C_{AB}(\tau)$, between two signals, $f_A(t)$ and $f_B(t)$, merely requires a substitution of $f_B(t - \tau)$ for $f_A(t - \tau)$ in the above expression [51].

It is relatively easy to convert the autocorrelation function of a wave into a power spectrum, showing amplitudes and phases over the delay period for the major frequency components comprising the overall waveform. One way is to multiply the correlation function by sine and cosine waves being swept across the frequency range of interest. The resultant amplitudes are plotted as a function of frequency, giving the power spectrum. One can test for particular frequencies by using sine waves at or near those frequencies. The individual results for sine and cosine waves can combined in a way that gives a phase plot for various frequency components. The combination of these measures gives a much better picture of those long-term, coherent components that may actually be present in the ongoing EEG. Better data on which to base feedback indications for a subject can be provided than that obtained by simple band-pass techniques, which may also introduce their own resonances into the measurements. When autocorrelation and power spectrum measures are not feasible, it is still useful to employ band-pass filters. However, the experimenter must be aware of the nature of her or his instrumentation and proceed with caution. An unfiltered, raw tracing of the EEG on an oscilloscope screen or oscillograph should always be available for visual inspection and comparison with the behavior of the filters used. 'Long-term' in this categorization refers to coherence times of 1 or 2 seconds or more.

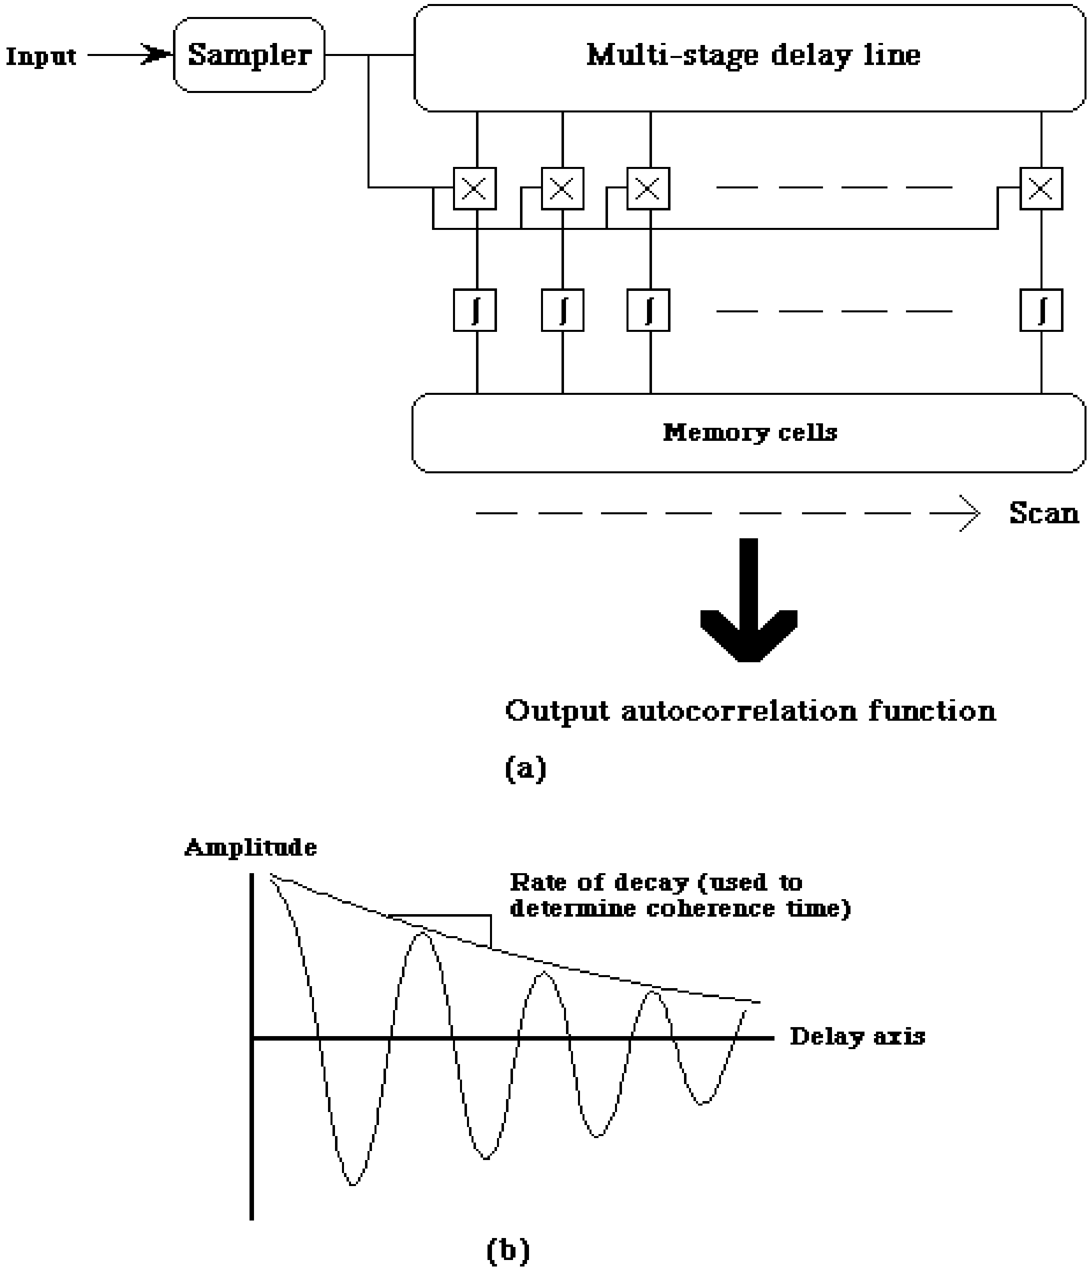

Autocorrelation

Figure 8

Fig. 8. Autocorrelation. (a) Diagram showing how the autocorrelation function of a signal may be calculated by means of sampling, successive time delays, multiplication and weighted time averaging (integration). (b) A typical autocorrelation function for a relatively coherent wave with a degree of instability. This instability is shown by how the function decays over the delay axis.

Arbitrary Historical Categorizations

Coherent waves have been classified traditionally in four categories. The definition of these categories is, in my opinion, rather arbitrary and based primarily on a sequence of historical events. For example, alpha waves are called 'alpha' simply because they were the first to receive serious experimental attention and the easiest to detect with early, primitive instrumentation. The others followed in due course.

Alpha.

Alpha waves are defined as coherent waves in an 8- to 12-Hz band. They are relatively large in amplitude over certain areas, most notably over the occipital cortex. When reinforced with biofeedback, they tend to be associated with a Zen-like state of high attention without a locally specific focus, an object of attention, or the subject being engaged in making differentiations or abstractions. They are associated with open focus, clarity and alertness. One cannot learn to produce alpha voluntarily by effortful trying. Such conscious effort merely interrupts alpha. One must learn to allow alpha to occur rather than make it occur. Such a state has numerous parallels in the practice of the arts. One cannot force the process of creativity. It must be allowed to evolve within its own subtle conditions for occurrence.

Beta.

Beta waves are poorly defined in the literature as nearly any coherent wave energy above alpha frequencies. Various researchers list different definitions. Most commonly, frequencies from about 12 to 20 Hz are accepted as beta waves.

In my experience, beta wave reinforcement is associated with a highly vigilant state of unfocused attention, a state in which one seems right on the edge of making rapid and complex abstractions, logical conclusions, calculations, observations or insights. It can be highly creative and productive.

Some researchers have reported strongly negative affective experiences associated with beta feedback. In my opinion, this is not due to any deleterious psychological implications, though it can be unsettling to some to be in such a vigilant state of attention. Rather, it is due to improper procedure and instrumentation. In my experience, when feedback is not given solely on the basis of the presence of signal power from beta-range, band-pass filters, but includes the use of autocorrelation analysis to measure waveform coherence, the experience is nearly always a positive one, sometimes even leading to elation and ecstasis. If, however, just a band-pass filter and amplitude follower are used, it is easy to give feedback inadvertently on the basis of highly erratic, incoherent, high frequency EEG activity. This can indeed be associated with negative feelings, feelings of dissociation, disconnectedness, jangling nerves, tension, anxiety, and a host of other bad experiences.

Possibly, this is experienced in conjunction with beta waves because the beta bandwidth is wide, compared to the other categories, and poorly defined in the literature. The beta bandwidth can be from 10 to 20 Hz wide in some cases. The alpha bandwidth, by comparison, is never more than 4 or 5 Hz wide. Consequently, the greater selectivity of an alpha filter makes it less susceptible to activation by irregular EEG waveforms, even though this is by no means the preferred method.

Theta

Theta is defined as relatively slow, coherent waves in a 4- to 8- Hz band. When reinforced with biofeedback, they tend to be associated with Yoga-like states of deep relaxation, or perhaps daydreaming—relatively unfocused. This state has more passive, but attentive, qualities.

Delta
Delta waves are those from about 1 to 4 Hz or below. Normally, they are present only in states of very deep sleep, states of anesthesia, or other relatively unconscious states.

Coherent Waves as an Integrated Continuum
These traditional categories of brainwave analysis have led, I believe, to a fragmented view of the coherent waves. I prefer to view them as occupying a continuum (see Fig. 9). In general they represent what might be termed *idling states* of the brain; that is, during the presence of strong coherent waves, the brain is not engaged in making spatio-temporally localized perceptions, differentiations among objects of perception, or cognitive categorizations. The brain is on the verge of doing so and is in a state of readiness to process information, whether of exogenous or endogenous origin, but is not engaged in such activity at that point. There may be a profound and strong awareness of everything in the environment, such as in the alpha state, but no abstractions are being made. One may be aware of the environment of a room, for example, but not engaged in abstracting the idea of squareness.

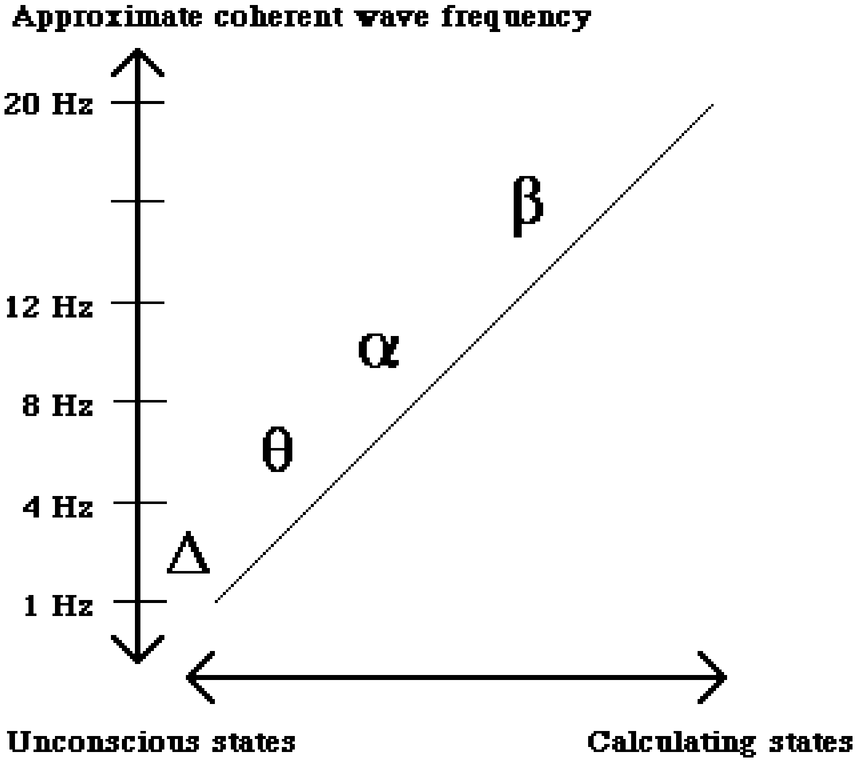

Coherent Waves as Integrated Continuum

Figure 9

Fig. 9. Coherent Waves as an Integrated Continuum. Chart showing the relationship of coherent, EEG wave frequencies to qualities of states of consciousness, ranging from the relatively unconscious to the calculating

states. (See the text for descriptions of these states.) It is proposed that viewing the coherent waves as part of a continuum may be better than using the rigid, traditional classifications of delta, theta, alpha and beta.

Qualitative differences in these waves are associated with different states of vigilance or alertness, as well as the propensity to process perceptions of changes in the external or internal environment in a particular mode, should they occur. At the low end of the continuum are what I term *relatively unconscious states*, like that represented by delta waves. Perceptions are not likely to be processed at all during the delta state. Above that we move into the theta state, during which perceptions or abstractions may be processed in a way that is semi-conscious, subliminal or dream-like. In the alpha state, of course, attention may become more highly activated, but this state is usually interrupted by the act of perceptual differentiation or cognitive categorization. During high beta production, one moves to the opposite end of the scale from the unconscious states into what I term the *calculating states*. Here the brain is vigilant and prepared to engage in quick, logical thinking with great speed. It races toward complex abstraction and differentiation when beta is interrupted. Unlike alpha, however, there is less emphasis on openness and on the free ranging consciousness associated with open focus. During both alpha and coherent beta production, there is a feeling analogous to that of being deeply engaged in something. Only the 'something' is not identified. During the alpha state, it is as if attention is high and equally distributed over the entire external and internal environment. By contrast, during beta, the feeling is more as if one were about to jump off a high diving board of consciousness into a particular logical or combinatoric peregrination of the mind. All these states have important analogs in the practice of music and other arts.

These long-term coherent waves may be summarized as regular-tending cycles of non-singular animation. The brain system is in a state of preparedness to process point-concentrated experience in a particular mode, but consciousness is neither ahead of the present (making predictions) nor behind the present (extracting time-bound associations); low frequencies tend toward the unconscious states, high frequencies toward the calculating states.

Sometimes the EEG will exhibit coherent energy distributed over a mixture of these band definitions. Relative levels of alpha, beta and theta bands may reflect intermediary states along the continuum described above. Usually, focusing attention will interrupt coherent waves, particularly alpha. Relations among ongoing, coherent wave categories and cognitive processing activities are not clearly established in the literature at this point. It has been suggested that complex topographic distributions of theta, beta and, to a lesser extent, alpha band intensities can reflect observable distinctions generated by a subject engaging in different types of cognitive tasks. For instance, discriminations between verbal and spatial tasks [52] and processes demanding internal mental focus and rejection of sensory input versus those requiring intake of sensory information [53] have been studied in this way. There is conflicting evidence regarding whether interhemispheric asymmetries in the coherent waves reflect engagement in different types of cognitive activity or differences in sensory intake versus sensory rejection modes. A tendency toward balanced topographic distribution of coherent waves, possibly reflecting a spread of cortical synchrony through resonant coupling of dynamical regimes, has been observed to accompany practices in meditation. Relative levels of theta, alpha, and beta amplitudes have been used in the composition *Chilean Drought* (which I composed with Jacqueline Humbert) to determine the audio mix played for an audience of three different vocal settings based on a text from a news account of the 1968 drought in Chile. A solo brainwave performer listens. The performer's states of consciousness, reflected by her or his relative brainwave amplitudes, influence how the news is heard in this musical setting [54, 55].

Some things can be observed in the ongoing EEG without the aid of spectral band decomposition. Detection of specific waveform signatures may also be of interest. For example, Roy John and others have demonstrated voluminous evidence for the occurrence of new patterns in the ongoing EEG that accompany learning [56, 57]. Changes in the EEG accompanying stimulus-response conditioning were observed early on in the history of this research. When an unconditioned stimulus is paired with a previously conditioned stimulus, a change from relatively low-frequency, high-voltage, synchronous (coherent) waves to high-frequency, low-voltage, desynchronized activity occurs. As training progresses, this change becomes localized over cortical areas relevant to the learning task. Today we may be able to decompose this desynchronized activity by means of pattern analysis, extraction of ERPs, or techniques for probing extremely complex patterns from dynamical systems theory. Further work involved the use of what are termed 'tracer- conditioned stimuli' (TCS), an idea originally introduced in the Soviet Union in the 1940s and the United States in the 1950s. As an unconditioned stimulus becomes a conditioned stimulus, new patterns appear in response to that stimulus. The TCS leave a kind of modulation imprint on the EEG from certain anatomical regions of the brain and, as learning progresses, this imprint spreads to other regions originally not involved in the processing of the TCS. After the introduction of signal-averaging techniques, however, most of this kind of research became focused on ERPs.

Short-Term Transient Waves

Short-term transient waves, as contrasted with coherent waves, embody the epitome of singular experience. They are known in the literature as event related potentials (ERPs) or evoked responses (ERs). They are transient, non-repetitive waveforms associated with events that take place in the brain during the first 1,500 milliseconds (msec) or so after the onset of a clearly differentiated stimulus. Consequently, ERPs are associated with the making of a discrimination at some level in the sensory or cognitive information processing hierarchy. They represent highly singular, spatio-temporally localized phenomena, both from the point of view of the subject having the experience and from the point of view of the observed waveform. They could also be viewed as short-term resonance phenomena coupling the brain, as an electrochemical information processor, with events in the environment.

These ERPs, as recorded with surface electrodes on the scalp, are typically of very small amplitude, around 25 microvolts or so, and are buried deep inside the ongoing, complex EEG waveform. Thus, sophisticated measurement and statistical analysis techniques must be brought to bear on the task of detecting ERPs and extracting them from background noise. They are mutually exclusive with the long-term coherent waves discussed above. Since they are associated with opposing kinds of experience, ERPs always interrupt any coherent waves that may be present. Thus, ERPs reflect an important principle of perception and cognition; that is, the apprehension, perception or recognition of an entity is always achieved at the cost of exclusion of competing entities, at least temporarily.

A kind of quantum exclusionary principal of cognition reveals itself here. Recognition, as we normally conceive of it, is a discrete process. This is not to say that other modes of consciousness might not exist, in which awareness is distributed more evenly across the entities of perception or cognition. However, these are fundamentally different, requiring special terminology for their description. Thus, the long-term coherent and short-term transient waveforms are associated with polar aspects of conscious experience: non-singular awareness and highly singular differentiation. Both are fundamental to the maintenance of the viable human organism.

A large portion of the technical discussion in this paper will focus on use of ERPs in biofeedback paradigms that probe musical experiences. Feedback paradigms involving the coherent waves alone have been well described in the literature and explored extensively in artistic applications [58]. ERPs, however, have not been explored extensively in this context [59]. They offer considerable potential for applications in music, particularly as they relate to the perception and comprehension of formal musical architectures.

Complex, Ongoing Waves

In this category I have placed what has not been analyzed in the preceding three categories. The EEG no doubt contains a complex, ongoing background component that is not random but is patterned with a degree of complexity to make any analysis seem an insuperable task. It is likely that as life experience progresses, aspects of this patterning evolve in a manner reflecting the self-organization of the information of experience. Complex patterns of baseline activation must build up in neuronal masses over time and in response to learning. Focused experience may interrupt these and superimpose temporary patterns necessary to deal with the localized, spatio-temporal constructs of singular events. The results of these discriminations, however, must, after an appropriate period of reverberation and widespread diffusion in the brain, submerge into the ongoing patterns. Each such experience may leave a minute tracing on, or in some way contribute to, the evolution of the electromagnetic field pulsations maintained by the organism as an integral part of itself. In some sense this electromagnetic field is an emergent, global property of the sufficiently complex self-organization of a critical number of individual, electro-chemical processing units. In discussing the representation of perceptions, John describes projection pathways of average, coordinated, spatio-temporal firings of neural ensembles from intracranial recordings, their spread of activation, and the establishment of complex gradients of ionic charge [60]. This is envisaged as a complex, three-dimensional volume of isopotential contours or convoluted charge surfaces, termed a *hyperneuron*. He hypothesizes that every representational system has its own unique hyperneuron, embodied as a particular distribution of energy established by the statistical properties of local ensembles contributing to the coherent spatio-temporal patterns within a volume of neural tissue. It is assumed that the emergence of a hyperneuron is dependent on the existence of a group of critical size, containing sufficiently organized processing units.

John lists several levels of organization of these processing units and their emergent properties, as follows: First-order, emergent processes are termed *sensations*. These result from the statistics of spatio-temporal patterns in stimulus-bound neuronal ensembles. Second-order processes are *perceptions*. This is the interpretation of the meaning of sensations by means of an interaction between sensations and memories. Third-order processes produce *consciousness*. Here a multi-dimensional representation of the state of the organism and its environment is developed. This is integrated with memories and the needs of the organism, generating emotions and programs of behavior addressing those needs. There can, of course, be many levels of consciousness. *Subjective experience* constitutes fourth-order information. It is derived by organizing information about the contents of consciousness and multi-sensory perceptions, memories, emotions and actions into episodes of differentiable experience. Fifth-order information results in the *self*. It is derived from the accumulation of personal history in long-term memory and its integration into episodes of subjective experience. Finally, sixth-order information comprises *self-awareness*, the interpretation of current subjective experience in the context of salient features derived from analyses of the accumulated history that constitutes the self.

Representations of multiple items on some lower levels often exhibit invariances that share a common informational feature. These invariances constitute the representation of information on higher levels. This is known as *paradigmatic bootstrapping*. Diverse and differentiable qualities are distinguished on higher levels from uniformly represented information at lower levels. This reflects the fundamental polarities of behavior that underlie everything the brain does: fragmentation and reduction versus synthesis and integration. The world is torn apart in order to be resynthesized and its image stored inside the accumulating history of the self in an individually unique way. There could not be individual beings without both of these processes operating in balanced concert. Access to that history is mediated by memory retrieval processes, which operate through independent mechanisms on every level of information analysis. This view of memory is distinct from the history that constitutes the self. Information retrieved from memory is added to and compared with incoming information from sensory channels to feed integrative processes on each level. The resultant organization of the organism is, on the other hand, the history that constitutes the self. As these orders of information develop, 'feedback' from higher levels down to the lower levels takes place, refining and tuning their evolving processes, while 'feedforward' from the lower levels up through the hierarchy, informing integrative processes on higher levels, takes place as well. All of this must be reflected in the ongoing life of the evolving electromagnetic entity that inhabits the body, the minute glimpses of which we record, analyze and call 'brainwaves'.

THE EVENT-RELATED POTENTIAL (ERP) IN DETAIL

ERP research offers significant potential for probing detailed aspects of perception, cognition and conscious experience, as distinct from the more global aspects represented by changes in the long-term coherent waves [61]. In this section, I concentrate on details of the human ERP, as recorded from the surface of the scalp with traditional EEG monitoring techniques.

Highly localized ERPs can be recorded from electrodes implanted in specific neural tissue sites inside the brain. Such ERPs can often reveal information about the operation of neuron groups of limited size and topographic extent. Recent research has shown significant correlation between some of these intracranial events and conventional ERPs recorded from surface electrodes [62]. Sometimes they respond to more specialized processes [63].

How ERPs are Detected and Measured

ERPs normally are extracted form the ongoing EEG by means of an on-line, signal averaging computer. An epoch of the EEG lasting approximately 1,000 to 1,500 msec is sampled, beginning with the onset of a stimulus of interest or sometimes slightly before this onset, in order to capture features that may relate to expectancy, such as the contingent negative variation (CNV). (See Fig. 10 for a depiction of an idealized, auditory ERP.) In a typical experiment, a number of these EEG epochs (in the literature this number may range from a dozen or so to thousands) are arithmetically averaged to produce the ERP. The primary purpose of signal averaging is to reduce the presence of large-amplitude, background noise signals. These are presumed to vary randomly and thus will tend to be canceled out over the course of a time average. A single-trial ERP is difficult to obtain due to its exceedingly small amplitude (approximately 25 microvolts peak-to-peak) and because it is buried in large-amplitude noise. This constitutes a primary technical limitation. Averaging techniques can alleviate the signal-to-noise problem. The resultant signal is sometimes referred to as the *averaged evoked potential* (AEP).

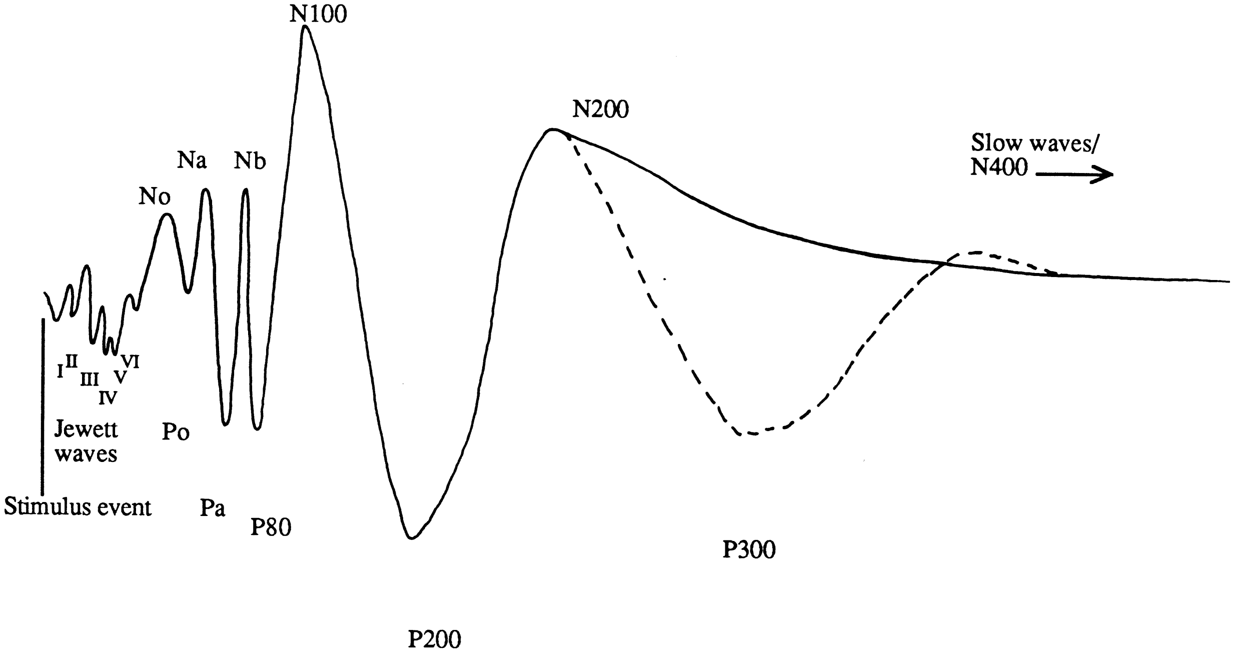

Note: Drawing follows convention of showing positive polarity in downward direction.

Form of an Idealized Auditory Event Related Potential (ERP)

Figure 10

Fig. 10. Form of an Idealized Auditory Event-Related Potential (ERP).
Approximately 450 milliseconds from an auditory ERP showing its most important waveform peaks. These may be analyzed using principal component analysis (PCA). Such ERPs are usually extracted from the ongoing EEG by averaging many epochs of the signal recorded synchronously with presentations of a stimulus event. All components are not always as clearly evident. P300, which may or may not occur, reflects processes of selective attention and may have endogenous origins.

Given an experimental situation in which the stimuli may be repetitive, it is possible that ERPs associated with the first few presentations of the stimulus contain the most important information. The remaining ones may merely indicate that no significant changes have taken place. Alternatively, some setups may involve significant changes over the course of trials due to learning, adaptation, stimulus generalization or categorization. An unweighted, averaged ERP measurement will be relatively insensitive to these phenomena. Consequently, various schemes intended to bias the averaged ERP so that it more strongly reflects the most recent trials are sometimes used. The calculation of an average can be weighted so that ERPs from the more distant past contribute less and less to the ongoing computation. Such weighted averaging schemes cause the ERP to reflect events from the recent past more strongly but incur the cost of degrading the signal-to-noise ratio. One simple scheme is to multiply the existing average by a factor, ß, and multiply the incoming signal by a factor, 1-ß, prior to adding them together. Another scheme involves applying a weighting factor, which decreases exponentially as a function of the time elapsed since the associated stimulus, to each member ERP contributing to the average. Eventually, this factor will decay to a negligible value and the associated ERP can be dropped. It has thus faded from the system's memory.

Sometimes it is necessary to examine ERPs from individual experimental trials. The extremely poor signal-to-noise ratio involved in such cases would make this all but impossible. However, C. D. Woody, some time ago, devised a more involved scheme for detecting single-trial ERPs [64]. It was designed to enable research into the variability of response latencies of individual waveform peaks in single-trial ERPs. Since such response latencies can vary quite considerably in different situations, across different subjects, in different task situations, or even within a specific experimental setup, such variations become a significant factor in the measurement statistics involved. This adaptive filter algorithm involves recording an ERP epoch, identifying the particular waveform peak of interest, and then recording its temporal position. This peak is then stored for later use as a template. After the next epoch is recorded, it is cross-correlated with the stored template several times in different time-shifted positions. When the largest crosscorrelation value is detected, the peaks are assumed to have lined up and the target peak is identified. The new waveform is averaged with the template, and its latency and amplitude recorded. This new average then becomes the new template for the next test. In this way, important waveform peaks, along with their latencies and amplitudes, are identified by comparison with a template being continuously refined. Woody also suggested making a rank ordering of trials on the basis of correlation with a template and separately averaging high- and low-correlation epochs to search for systematic differences among trials. The process is computation intensive, however, and only recently could be applied effectively in real-time situations with inexpensive computing resources.

There are other major limitations imposed by the signal-averaging technique. Since a number of epochs from the ongoing EEG must be sampled by the computer and combined into the deduced average, the timing of the epochs sampled must be consistent and meaningful. The computer cannot simply detect an arbitrary ERP buried in the ongoing EEG tracing. It must be told where to look in time to extract meaningful data. Consequently, the experimenter can only look for ERPs where they are predicted to occur. The most clearly defined and logical periods to look for these, of course, are during the first second or so after a clearly differentiable stimulus event. If there are meaningful ERPs to be found at points that are not as clearly defined, they will be missed. The detection of ERPs, then, can often be the result of a self-fulfilling prophecy. Nevertheless, by careful experimental design, much useful information has been culled from ERP experiments. Searching for ERPs in the ongoing EEG would require continuous matching to a repertoire of waveform templates. The computations are time consuming, the templates we have are rather imprecise in their definition, and ERPs in normal situations are subject to considerable, natural variability. This seems an obvious place to apply artificial intelligence techniques that can deal with imprecise set membership with calculated confidence ratings.

Contingent Negative Variation (CNV)

Contingent Negative Variation (CNV) is a general biasing of the EEG in the negative polarity direction. It has been theorized to accompany expectancy or anticipation, particularly when a response such as a motor action is intended to follow the expected stimulus. It may show anticipation of a cue for the orienting response. It can be specific to time, space, message content or a combination of these. Thus, CNV may be more generally described as a negative shift concomitant with the subject's anticipation of establishing a locally valid causal construct. It often precedes ERPs for expected stimuli.

CNV is easy to identify by visual inspection of ongoing EEG tracings, appearing as topographic differences. For instance, in an experiment in which a warning stimulus preceded a subsequent imperative stimulus requiring a non-discriminant motor response, a CNV, maximal over the frontal cortex, was elicited by the warning stimulus followed by a later CNV over the motor cortex preceding the imperative stimulus [65]. CNVs preceding voluntary motor actions are referred to as readiness potentials. In a later experiment, artificial CNVs, like those normally seen in the warning-imperative stimulus design, were successfully synthesized by adding components of voluntary readiness potentials and ERPs to the individual stimuli containing a negative after-wave that may reflect orienting or activation processes [66].

Biofeedback with ERPs

Operant control of components of the ERP by means of biofeedback in human beings has been explored only minimally, though there are some examples using animals [67]. During the early 1970s, at the Experimental Aesthetics Laboratory at York University in Toronto, I began to investigate that possibility and to explore applications of the results in a model for an adaptive, interactive, electronic music instrument that would be sensitive to different information-processing modalities of the nervous system [68]. In view of the relevance of various ERP components to the significance and meaning of stimuli, along with their discrimination and cognitive processing by a subject, this seemed an intriguing prospect. At that time, results were limited by available instrumentation, most importantly the lack of portable, inexpensive computing power with the speed necessary to provide meaningful results to a subject in a real-time feedback paradigm. Preliminary explorations, however, were encouraging. C.M. Nunn, then a graduate student and technical assistant in this lab, created and published an elaborate design for a feedback instrument based on detection of specific, principal component criteria in ERPs [69], along with an excellent survey of the underlying theory involved. His system included built-in capabilities for statistical processing of experimental data, as well. His intent was to compare a subject's performance in a signal-detection experiment without feedback with a situation in which feedback is provided for production of ERP components associated with correct detection of target stimuli. In this way, processes of sensory fine tuning, such as for learned pitch discriminations, could be studied.

Categories of Event Significance

Investigations into the nature of human ERPs have been focused on unraveling the details of their form with the hope of finding electrophysiological correlates of higher cognitive processing. The ERP can be described as a short-term, transient, slow wave, the principal components of which are spread out over roughly a 1-second time frame following the onset of a highly discriminant stimulus. It is important to note that recent research indicates this stimulus may be of external origin and thus may elicit *exogenous* components of the ERP. On the other hand, the ERP may accompany certain internally generated events that will elicit *endogenous* components. Furthermore, the ERP elicited by an external stimulus usually contains both exogenous and endogenous events. Since the ERP is not a repetitive waveform, it is not appropriate to use frequency-domain analysis methods, such as the Fourier transform, to extract important features. Rather, a time-domain-based analysis of the amplitudes and latencies of prominent peaks that are contained in the ERP and can be correlated with neurological or cognitive events is the method used to decompose the ERP. This is referred to as *principal component analysis* (PCA). Changes in the PCA analysis are primarily related to two broad categories of factors: the *form* of a physical stimulus and the *significance* of either an externally or internally arising stimulus.

A voluminous literature exists that details research into the significance of events observed in the PCA and their classification into meaningful categories [70]. I will attempt a summary of some of those important in the auditory ERP, since they may bear the greatest relevance for music. Many of the general principles and observations contained herein, however, have analogous principles and observations with respect to visually evoked responses. Their potential for investigation in fields such as kinetic arts, video and computer graphics are assumed to be quite significant.

The complete, idealized, auditory ERP consists of 20 or so peaks of interest (see Fig. 10). Normally, however, detection algorithms focus on one or more peaks for a given experimental situation. A polarity/latency nomenclature normally is used to identify peaks. For example, N200 would indicate a negative polarity peak occurring *approximately* 200 msec after the onset of measurement. I emphasize 'approximately' here to stress the fact that these latencies can vary widely in different situations and are not well defined or consistently indicated in the literature. During the first 10 msec, a series of small peaks, the so-called *Jewett waves*, are seen. These relate to the propagation of auditory nerve volleys on their way toward the central nervous system. Various other neurogenic or myogenic peaks follow, until about P80, which is probably generated in or around the auditory cortex in the temporal lobe. Most of these are relatively insensitive to factors of attention or any kind of conscious or unconscious cognitive processing. They are associated with propagation of raw sense-organ signals through the neural distribution network.

The effects of greatest interest for our purposes begin to occur at about N100. At this point, effects seemingly due to factors of attention come into play.

N100

N100 is sometimes called the *attention wave*, although this description also applies somewhat to P200. It's peak amplitude increases when attention is directed toward stimuli in the sensory channel relevant to the situation at hand (for example, auditory or visual); it therefore represents a first stage of selective attention [71]. It can be differentially sensitive to the directing of attention toward signals from either ear [72]. It can also reflect hemispheric asymmetries, particularly in response to speech stimuli for which N100 amplitudes are greater in the dominant hemisphere [73]. It is interesting to note that the latency for N100 (100 msec) is too short for it possibly to reflect processes of attention that depend on the complete presentation of a typical complex stimulus, such as a word or musical phrase. It occurs long before the articulation of the word or phrase is complete. Consequently, it is thought that N100 may reflect the activation or directing of a *response mode*, rather than any finer discrimination as to what the word or musical phrase is or what its meaning might be. Examples of such a response-mode direction might be, 'respond to a verbal, language stimulus' or 'listen analytically to a musical rhythm grouping'. This kind of discrimination can be made before obtaining any knowledge of the content of the stimulus. The N100 peak is sensitive to signal detection, but cannot be used to discriminate among specific stimuli in the attended channel [74]. It increases in amplitude particularly when attention is directed toward detection of occasional, weak signals [75]. In summary, we may say that N100 reflects the action of an *attentional gate* and, possibly, a discrimination regarding which further sensory detection and feature extraction mechanisms to activate or which type of response mode to alert. It may reflect modality or location-specific attention. It appears when there is selective attention to a subset of stimuli and, therefore reflects the stimulus set. N100 is an exogenous component of the ERP, locked to the occurrence of a physical stimulus.

P200

Large-amplitude P200 peaks also reflect selective attention processes, though of a type somewhat different from N100. Like N100, P200 peaks are large for stimuli that are attended to and enhanced by directing attention to relevant sensory channels, especially in the detection of weak signals [76]. P200 amplitudes do not show differentiation, however, in experiments involving attending to one or the other ear [77] or interhemispheric asymmetries in processing verbal stimuli [78]. P200 may have more to do with gating or activating a response-generating mechanism, generating a 'go' versus 'no-go' stimulus for a response process, after an attention-based sensory channel discrimination is made. P200 is also an exogenous component, always time-locked to a physical stimulus.

N200

The N200 peak is difficult to observe because it overlaps with P200 and is often obscured by it. N200 events also differentiate relevant from irrelevant stimulus modalities for situations involving both auditory and visual sensory channels. They do not always provide differentiation among relevant and irrelevant signals within the relevant modality, however [79]. The latency of N200 has been shown to co-vary with reaction times in a vigilance task, such as detecting and responding to targets [80]. This supports the hypothesis that N200 is related to decision processes that control behavioral responses to sensory stimuli in a rather discriminant way, and must therefore represent some very early stages of cognitive processing. Other ERP component latencies do not co-vary as predictably with reaction times as N200 does, at least as shown in the literature so far. This study also relates N200 with the detection of infrequent targets. N200 is considered an endogenous component of the ERP, in that it is thought to be generated internally as opposed to being locked to external, physical stimuli. In the same vigilance study cited above, N200 peaks were observed when within a train of visual or auditory stimuli some stimuli were randomly omitted. In such a situation, N200 peaks are easier to observe because the obscuring, earlier exogenous peaks are absent. N200 may therefore be summarized as an *endogenously synthesized component reflecting detection of stimuli that may or may not be physically present*. However, it does not necessarily reflect stimulus identification or classification. It does act as a stimulus for response generation and further cognitive processing (see the section P300, below). It may reflect the making of preliminary decisions based on modality parameters rather than on stimulus specifics. N200 is the first peak in the ERP sequence to be observed in response to expected-but -absent physical stimuli.

P300

The P300 peak has received by far the most attention from researchers interested in the physiological correlates of cognitive processes and conscious experience. It has a relatively large amplitude and is responsive to highly discriminant situations. It not only is responsive to relevant sensory modalities and to stimulus subsets but is elicited by detection and recognition of specific stimuli within the subset [81]. It is highly sensitive to the occurrence of infrequent events, that is, events with low probability of occurrence. In the vigilance study, P300 was associated with detection of targets, as was N200, but originated from different recording sites [82]. It is therefore considered to represent different aspects of brain function. Detection of signals is related to the occurrence of both N100 and P300. Recognition of signals is related only to P300. Clearly, recognition is only partially contingent on detection. While the earlier components, N100 and P200, are associated with detection by peripheral, sensory mechanisms of real physical signals that are being attended to, N200 and P300 may accompany the incorrect detection of absent stimuli. P300 amplitude has been shown to vary with the physical similarity between target stimuli and presented stimuli in certain situations—for example, the pitch of sounds [83]. N200 may represent the *process of detection leading to response*, while P300 may represent the *process of classification leading to image formation and memory updating*. In a signal detection experiment, P300 amplitude was shown to increase with the strictness of the criteria imposed on the classification task, independent of arousal [84]. Consequently, P300 may also reflect the quality of information received, relative to some internal measure. The progress of N200 and P300 may reflect the course of brain mechanisms building up probability descriptions of incoming events leading to response, on the one hand, and to cognitive classification and evaluation of stimulus significance, on the other. P300 is sensitive to the manipulation of psychological variables, independent from the physical aspects of the stimulus. P300 may be triggered by a definitive match between a sensory event and a stored neural template from memory. P300 is coupled with decision making in cognitive processes. The late William Chase summarized P300 by saying, "We'll call P300 the neural event associated with attention allocation for the memory control processing involved in memory reorganization" [85]. I referred to this as the allocation of energy and neural resources to the synthesis of the image, the *idiolog*, of an event that will be stored in memory along with all its attendant classifications and association pointers to other memory engrams [86, 87]. Donchin generally summarizes P300 as the manifestation of a routine activated whenever there is a need to update one's model of the environment in working memory [88]. Thus, P300 may represent the *onset of control for the updating of memory*. It may reflect the making of a decision to bring or not to bring mental processes to bear on the situation at hand.

P300 may be elicited in response to detected, omitted stimuli in a regular, stimulus train [89, 90]. John shows ERPs that allow discrimination between circumstances in which a constant, ambiguous stimulus is interpreted in different ways—for example, a vertical line being interpreted as a number or as a letter in different tasks [91].

A significant observation about this evidence is that it implies that auditory attention is not mediated by a peripheral gating mechanism but rather by a complex matrix of phenomena associated with internally activated selection mechanisms, image synthesis, retrieval and template matching processes, and stimulus-independent, perceptual decision making.

Slow Waves/N400

Some researchers refer to long, slow waves accompanying aspects of cognitive processing, learning, memory consolidation and pattern analysis. P300 appears with latencies too short for it to be involved in such higher cognitive processing. Consequently, one is stimulated to look beyond P300. The identification of components with very long latencies is technically difficult, however, because their effects become submerged in the ongoing EEG field complex. Researchers do not know what to look for and, therefore, how to guide their computers in the search for patterns. In an experiment designed to investigate the association of ERPs with semantically anomalous information, subjects were presented sentences in which the last word made sense, did not make sense or had some physically incongruous aspect, such as a sudden increase in amplitude [92]. An example from an experiment involving semantically incongruous endings is, "I take coffee with cream and *dog*," (my emphasis). ERPs were extracted for each word in the sentences. ERPs for ending words that made sense were normal. ERPs for semantically incongruous ending words exhibited a large amplitude, long latency, slow wave, termed N400. ERPs for endings with altered physical characteristics elicited a longer wave of opposite polarity, termed P560. N400 may relate to postulated, long, slow waves associated with higher processing. It could represent a *phasic augmenting* of the CNV, which may build up in anticipation of reaching the end of the sentence. In this case it is significant as a marker of complex linguistic processing. It may also represent a kind of cognitive 'double-take', a reprocessing of the final word in the sentence, which does not match the most highly probable predictions built up by the unfolding grammar and meaning of the preceding sentence content. P560, which accompanied the physically deviant ending words, is not well understood. However, the fact that it differs so markedly from N400 underscores N400's significance as an indicator of the cognitive activity required to resolve the semantic anomaly.

Emitted Potentials (EPs)

Those components of the ERP that are dependent on the meaning attached to a stimulus by a subject, that are elicited by absent (i.e. imagined) stimuli, and that are representative of *memory readout* are often termed *emitted potentials* (EPs). The endogenous components of ERPs, (i.e. N200, P300, and longer wave phenomena), are emitted potentials. In a well- known experiment, a graphic symbol that could be interpreted either as the letter B or as the number 13 was imbedded in number and letter sequences projected in subjects' visual fields [93]. The subjects were informed that the purpose of the experiment was to determine the speed with which they could name numbers or letters. ERP components—recorded from the frontal lobe, starting about 160 msec after stimulus presentation—showed significant differences, depending on whether the ambiguous stimulus was interpreted as a number or as a letter. Other experiments show that neutral or ambiguous stimuli can elicit ERPs that are typical of those elicited by stimuli that are expected to occur but in reality do not occur [94]. John has produced voluminous evidence for the occurrence of endogenous components in the ERP that accompany conditioning, stimulus generalization and learning [95, 96].

Relation of Peak Latencies to the Information-Processing Hierarchy

With these data, one can make a mapping of events on a timeline that begins shortly before the onset of a stimulus event and continues for a brief period afterward. Prior to the event onset, one may see EEG phenomena associated with *anticipation* or *expectancy* (CNV). Subsequent to the event onset, a relation can be seen between *peak latency*, or time after event onset, and the position occupied by internal events associated with specific ERP peaks in the information-processing hierarchy of the brain and its cognitive processing apparatus. The further downstream we are from the event onset, the more high-level the observed effects must be (see Fig. 11). This is obvious from the assumption that patterns of activation in neural populations will project widely, stimulating a spread of activation throughout large volumes of brain matter involved with broader and broader levels of classification and more and more refined aspects of response synthesis.

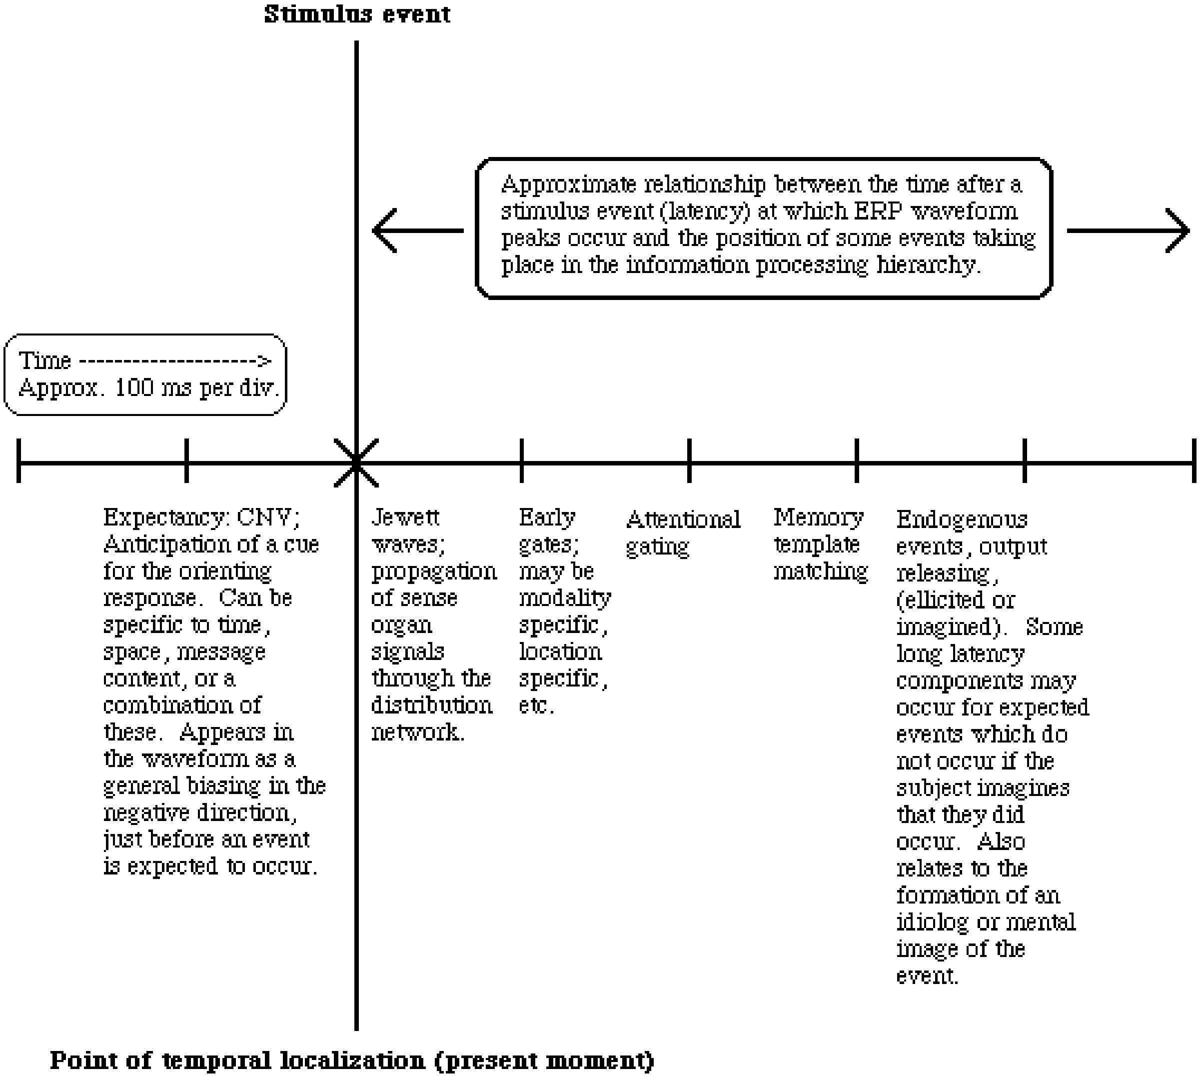

ERP Peak Latencies and the Information Processing Hierarchy

Figure 11

Fig. 11. ERP Peak Latencies and the Information-Processing Hierarchy. A mapping of some of the kinds of information processing taking place in the nervous system—thought to be associated with particular peaks in the event-related potential (ERP)—according to their latency or the time after a stimulus event at which they occur. Events associated with expectancy or anticipation of a stimulus—evidenced principally by a contingent negative variation (CNV) in the EEG—are also shown.

During the first 80 msec or so, we see signals concomitant with propagation of information from the peripheral sensory organ (ears), through the distribution network, and on to the auditory cortex. Then we see the effects of early gating, which may be modality specific, location specific, etc. Next, factors of attention and signal detection are observed at around 100 msec. From this point, information processing most likely branches out in several directions of endogenous pattern synthesis. Internally generated resonances, which have qualities of conscious decision making associated with stimulus detection and therefore involve memory activation, are likely to be seen at around 200 msec. (P300 comes too late to represent memory activation.) From here, information probably projects in two major directions simultaneously (see Fig. 12). One leads to the mechanisms that will assemble control signals required for the organism to generate a physical response; another leads to the mechanisms that will correlate the incoming patterns with those activated from memory tracings as well as formulate classifications and synthesize the idiolog or mental image of the event. The effects of this are observed at around 300 msec and, after, in late components and slow waves.

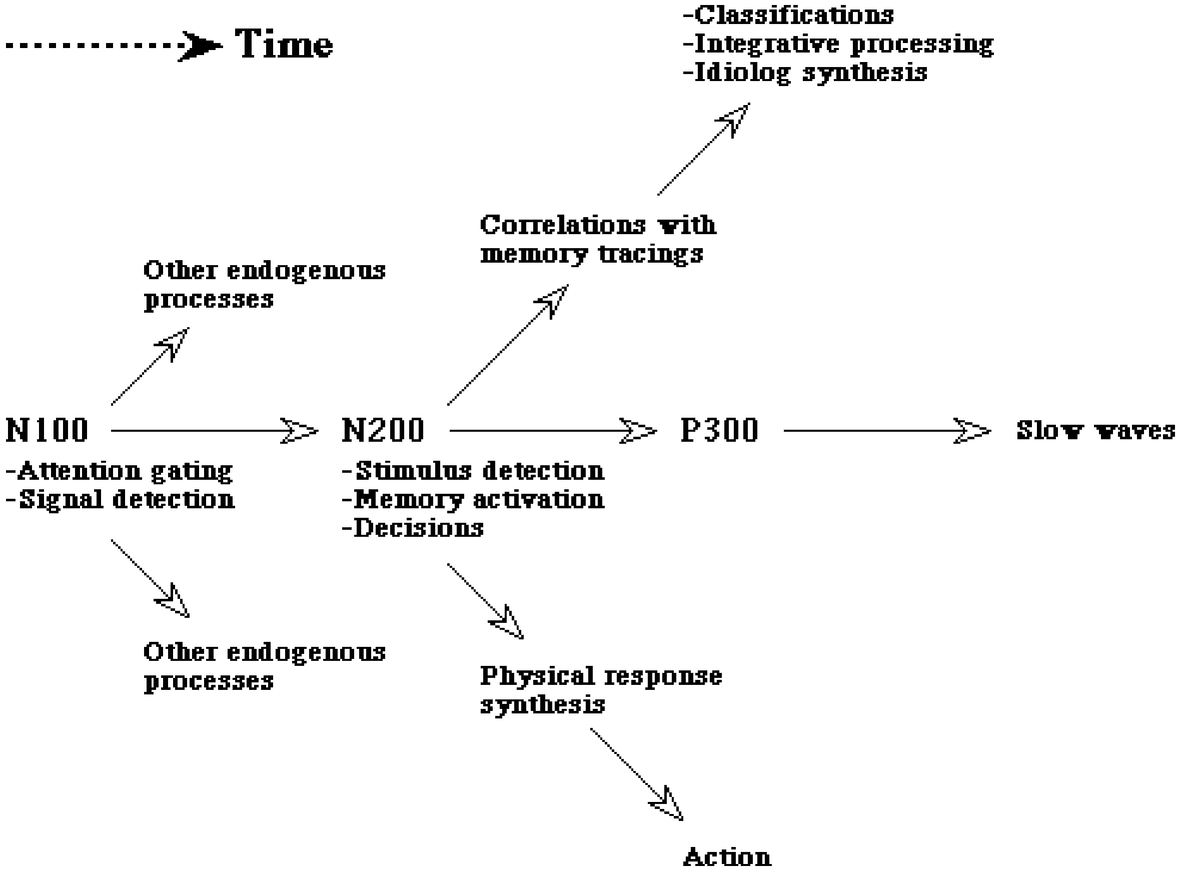

Branching of Processes Associated with ERP Components

Figure 12

Fig. 12. Branching of Processes Associated with ERP Components. As the results of particular stages of sensory information processing become available, multiple, subsequent processes may be activated simultaneously. These may use the resultant data in different ways. The significance of ERP components

may be affected when a stimulus acquires different meanings, appears in different contexts or is associated with particular tasks.

EEG Topography

Many of the experiments heretofore cited involve topographic differentiations in the locus of important ERP or coherent wave recordings. I have outlined very little about EEG topography so far. However, distributions of differences in ERP components from different anatomical regions of the brain are often quite significant. The contribution of exogenous versus endogenous processes to the overall ERP varies widely for different brain regions. They can be used to track the spread of activation of processes through different regions of neural tissue and they can show different degrees of engagement of separate parts of the cortex in a given task. Topographic mapping requires, of course, that recordings be taken from many sites simultaneously. This demands multi-channel EEG recording and analysis equipment, which is often expensive and cumbersome. Furthermore, in artistic situations, it is often unacceptably restricting to encumber a subject or performer with many electrode attachments. It may be necessary for some types of measurements, however. Convenient methods of attachment and less expensive, high- quality hardware need to be developed.

An important type of data for biofeedback research involves the degree to which the cortical synchrony observed in coherent waves is anatomically widespread. A simple but useful method of showing EEG coupling between cortical potentials from different areas has been available for some time and could be applied in real-time feedback with the current generation of microcomputers [97]. This method simply classifies the EEG waveforms from several different recording sights as to polarity and direction, i.e. positive-rising, positive-falling, negative-rising and negative-falling. Using a standard information theoretic model of uncertainty reduction, coefficients of information transmission are computed to show the degree of coupling between recording sites. This method is computationally simple, fast and effective. In the study cited, coupling between sites was displayed during a variety of tasks, such as rapid, silent reading, examining the details of a picture, listening to Mozart with eyes closed, and composing a letter mentally with eyes open. Listening to Mozart was associated with the most widespread coupling.

Musicologists might recognize this method of classifying waveforms as strikingly similar to methods for classifying melodic and other musical parametric contours according to sequences of ups and downs—a technique pioneered, I believe, by Charles Seeger many decades ago [98]. I have used a related system for correlating melodic and rhythmic contours in my recent composition for percussion and computer music system, *Zones of Influence*, in which distances in a *concept space* are calculated, showing relative morphological similarities and differences among these contours [99-101]. This method of comparison reflects the degree to which the sequences of ups and downs of two shapes are similar. Recently, Polansky has developed a formal taxonomy of distance functions on shapes [102]. In this classification, morphological distances that are based on direction and that preserve order are called *ordered linear direction* (OLD) metrics. I suggest that such morphological metrics may prove useful in tracking the spread of activation patterns, as seen in EEG waveforms, across topological and temporal dimensions.

It is hypothesized that, during the execution of purposive behaviors, large masses of the brain are in functional communication with each other and that an increase in correlations among low-frequency macropotentials recorded from these separate areas may be seen as a result. This correlation is independent of the voltage of the macropotentials and may reflect coordination of mass neuronal processes. Pattern-recognition techniques can be applied to analyze correlations between these macropotentials and have been used to differentiate between types of mental judgments required in a visuo-motor task design [103]. This method revealed differentiations that were not evident from visual inspection of conventional, averaged ERPs.

A DYNAMICAL SYSTEMS PERSPECTIVE ON BIOELECTROMAGNETIC PHENOMENA

The techniques of dynamical systems analysis are increasingly being applied to a host of bioelectromagnetic phenomena, including physiological rhythms, patterns in the neuro-musculature, the acquisition of learned physical movements, the generation of normal and arrhythmic heart-beat patterns, and systems of coupled, biological oscillators [104]. One primary technical advantage of this lies in the ability to describe complex behavior with a relatively small number of order parameters. Viewing the EEG as the manifestation of a dynamical system leads to interesting characterizations of movements among the subcomponents of the total electromagnetic pulsation pattern. (See the description of dynamics terms in the section *Biofeedback—Definition and Modeling* in Part 1.) If the coherent waves tend to settle into delta, theta, alpha and beta bands (and this is not absolutely certain), these bands could be viewed as basins of attraction. We could then examine the dynamics of transitions from band to band. What is the settling time under various conditions? What is the switching time from band to band? Can we more clearly characterize the perturbations that disrupt coherent waves?

Jantsch [105], in referring to the work of Freeman [106], makes the observation that the waveforms in the EEG can be interpreted as *limit-cycles* in the amplitude oscillations of extracellular potentials among groups of neurons. Freeman's methods emphasize analysis by means of phase-space plots, in which the dimensions map changes among EEG potentials recorded from different regions of the brain [107]. Jantsch further suggests that this limit-cycle behavior can be characterized as a primitive form of gestalt and that gestalt qualities may emerge at several levels leading up to thought and higher levels in the organization of life experience. The mind is described as the self-organizing dynamics of these phenomena on organismic, reflexive and self-reflexive levels [108]. I suggest that the EEG should be regarded as a low-dimensional projection of this complex, multi-dimensional activity. Low-level events—'limit-cycles' or 'primitive gestalts', to use Jantsch's terminology—become enfolded in the ongoing ordering dynamics of the brain on successively higher levels of organization. Thus, the results of sensations, experiences, thoughts and memories become widely distributed in the totality of the ongoing neural pulsation patterns. In part, then, the EEG is the *holomovement* of the electrochemical behavior of the brain. The concept of the holomovement has been explored by physicist David Bohm as a description of a primary ordering principle in the universe [109]. In this paradigm, the various energies and information associated with particular experiences are continuously being enfolded into the implicate order of the EEG holomovement. Objects of attention, either by the whole brain or by smaller neural groups, are made explicit as they are unfolded from the holomovement.

EEG State Transitions

In order to highlight significant features of this holomovement, I proposed beginning with the decomposition of the EEG into the categories *long-term coherent waves* (low-dimensional), *short-term transient waves* and the *complex, ongoing waves of hyperneurons* (high-dimensional)-all combined with a *random-seeming background.* (See the section *The Electroencephalogram in Detail—Categories of Useful Measures* in Part 2 and Ref. [110].) This categorization is further depicted in Fig. 13. If the first three of these categories can be imagined to represent states of the EEG holomovement, a series of state transitions can be described, as shown in Fig. 14. One is naturally tempted to construct a mapping of these states in a behavioral phase space along the lines of that shown in Figure 4. The existence of low-dimensional collective variables that can be used to describe observations on macroscopic neural phenomena has been pointed out [111]. However, a determination of acceptable order parameters for each dimension is difficult and tentative at present. The use of dimensions that scale correlation values between waveform phenomena as well as coupling of topographically or temporally distributed processes has been explored and may prove valuable [112]. Interestingly, sudden changes in spreading activation patterns, such as state changes in a phase space, have also been observed and studied in artificial and highly parallel networks of computer processors, such as the Connection Machine and the Hypercube [113]. A key determinant for such phenomena is a well-ordered network topology and a critical mass of processing elements. A region of network nodes (defined by an *event horizon*) that are actively engaged in a process can, under certain conditions, explosively and unpredictably expand their horizon or suddenly change their activation patterns. Groups of neurons or hyperneurons can similarly explosively spread their activation patterns or undergo sudden flips in their processing modes when certain thresholds are crossed.

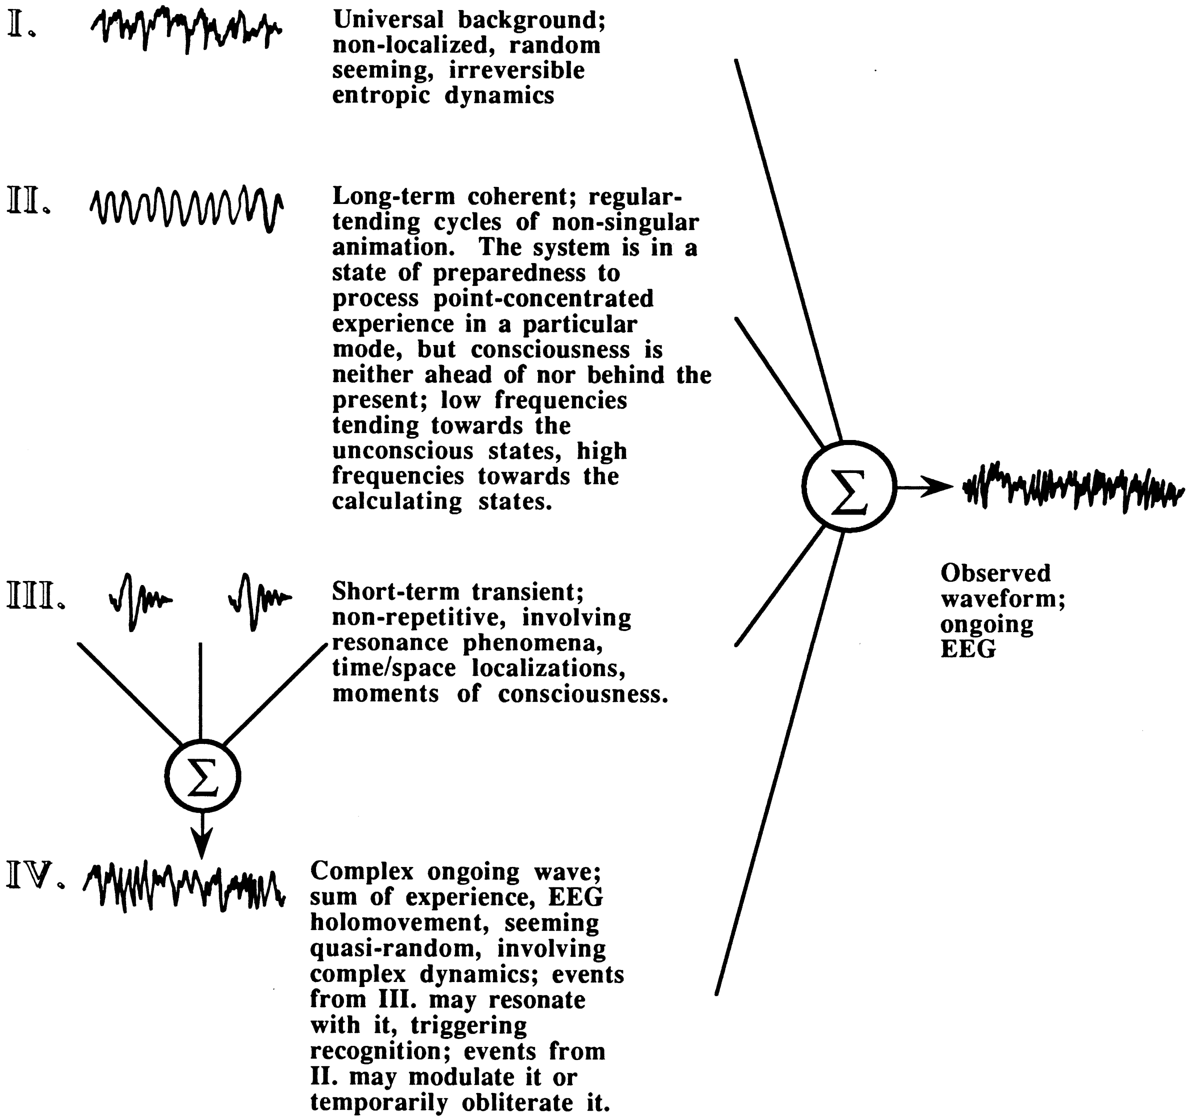

First Decomposition of the EEG Holomovement

Figure 13

Fig. 13. First Decomposition of the EEG Holomovement. A scheme for a first-stage decomposition of the EEG into four primary categories: I. the random-seeming background, II. long-term coherent waves, III. short-term transient waves, and IV. a complex, though ordered, ongoing wave. Our normal observations include only monodimensional EEG recordings. In fact, the EEG is comprised of multi-dimensional electrochemical pulsations.

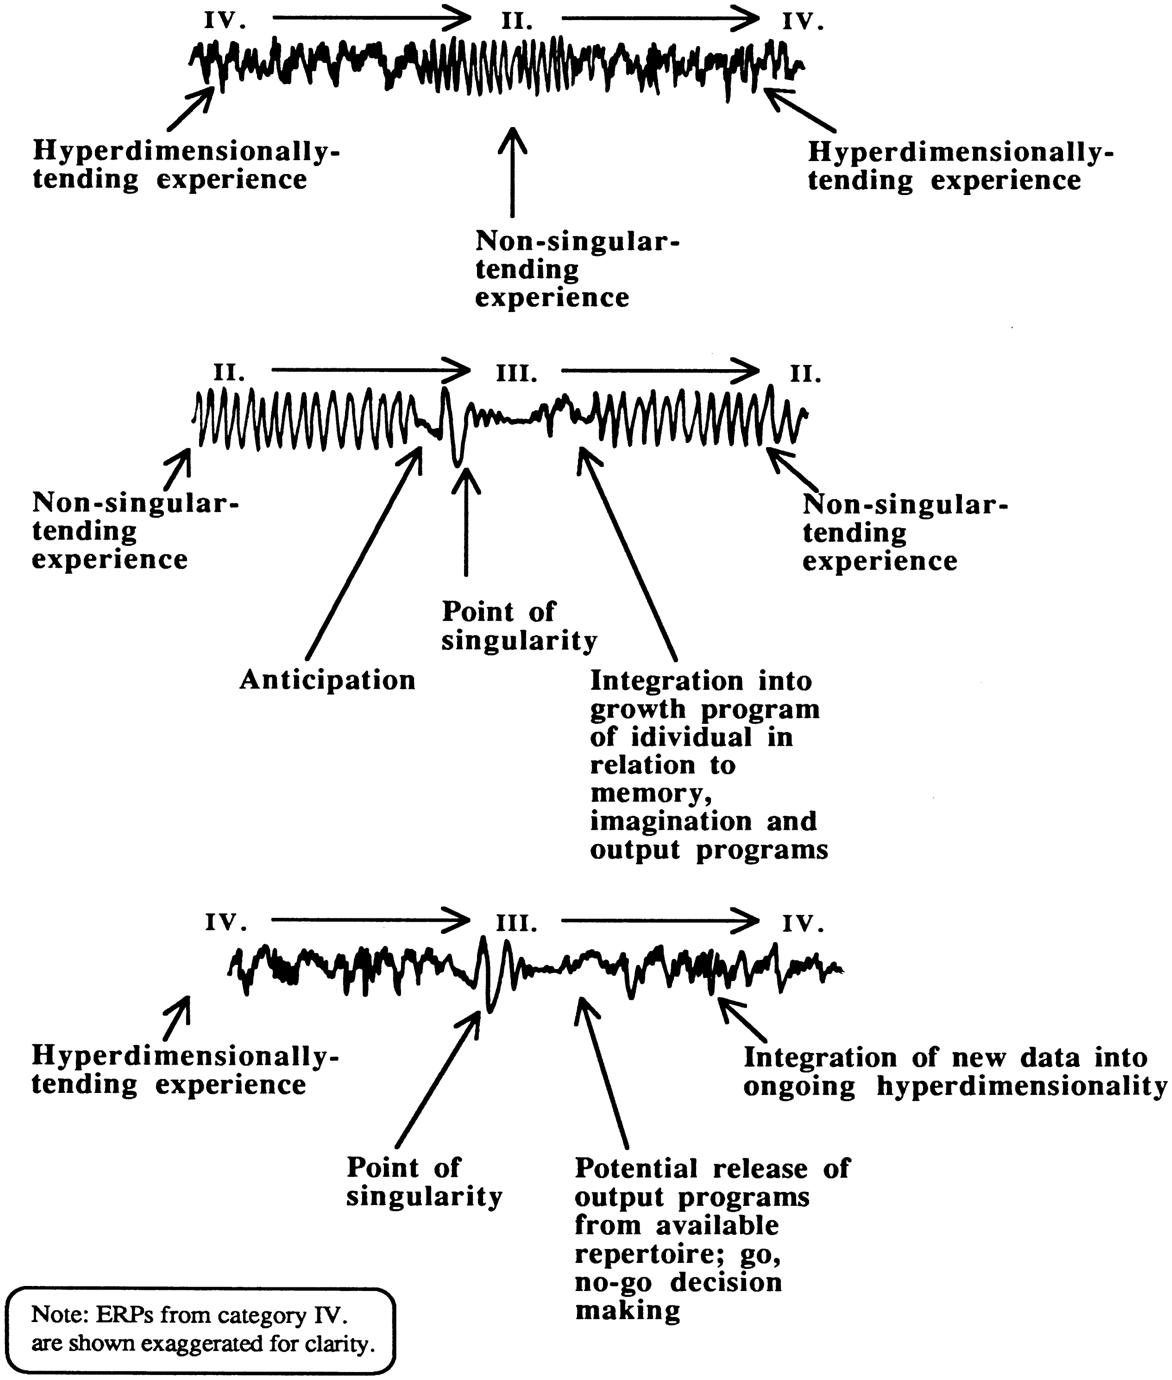

Note: ERPs from category IV. are shown exaggerated for clarity.

EEG State Transitions

Figure 14

Fig. 14. EEG State Transitions. Three depictions of the ongoing EEG changing states. These states are constituted when a dominant wave pattern is easily identified with one of the four primary categories for EEG decomposition. Each is associated with particular kinds of experience.

It has occurred to me that a reason it is so difficult to devise an appropriate set of dimensions for an EEG phase space with which to describe state transitions may be that the necessary dimensionality for various EEG states may not be constant. In fact, in considering the various states of consciousness associated with EEG waveform features, I have been led to the notion that the *dimensionality* of experience associated with these phenomena may itself be a variable mappable on some kind of axis. A tentative gesture toward visualizing such an idea is presented in Fig. 15, in which the expanding radius of approximately concentric circles represents the increasing dimensionality of experience. The diagram represents one possible projection of the concept onto two dimensions. At the center we have a dimensionless *origin* of the EEG holomovement. The origin of any space of this type is considered merely a conceptual anchor point. As one moves out from this origin, a region of space-time becomes defined. As the radius of the encompassed region of space-time increases, so does the order of dimensionality of the system. The inner circles contain the *coherent waves*, which are assumed to be associated with non-singular-tending experience of low-dimensionality. Such experience, though sometimes associated with intense awareness, does not normally produce highly differentiated perceptions subject to clear categorization through the normal cognitive mechanisms. Proceeding from the relatively unconscious states associated with delta and theta waves through alpha and the calculating states associated with coherent beta waves, we finally arrive at the outer region of complex, hyper-dimensionally tending experience (i.e. high-dimensionality), D^H. This is where most of the time of daily life is spent. An ongoing stream of highly differentiated experiences requiring many dimensions for their description, measurement and categorization are routinely associated with this kind of experience. These experiences may be grouped in complex orderings, though the extreme complexity may not permit immediate apprehension of these orderings. The ongoing complexity is interrupted by points of focus on singular experience, S_n, associated with transient waves, T_{S_n}, such as ERPs. These are of moderate dimensionality and are strong enough to engage most of the apparatus of attention and consciousness temporarily. They produce manufactured conceptions of space-time with just enough complexity (dimensionality) to allow a relatively harmonious enfolding of the experience into the ongoing holomovement. They may be analogous to what Bohm calls 'moments' in his description of the implicate order of consciousness [114]. Those singular events that are incongruous or tension generating or that somehow produce cognitive dissonance may require further conscious processing in their enfoldment. These may also be positive affective experiences that stimulate curiosity or explicate new relationships and a concomitant reorganization of memory through the activity of thought. In any case, they conclude with a *transformation process*, T_{S_n}, on the EEG holomovement, the results of which are widely distributed among the activities of hierarchically enfolded hyperneurons.

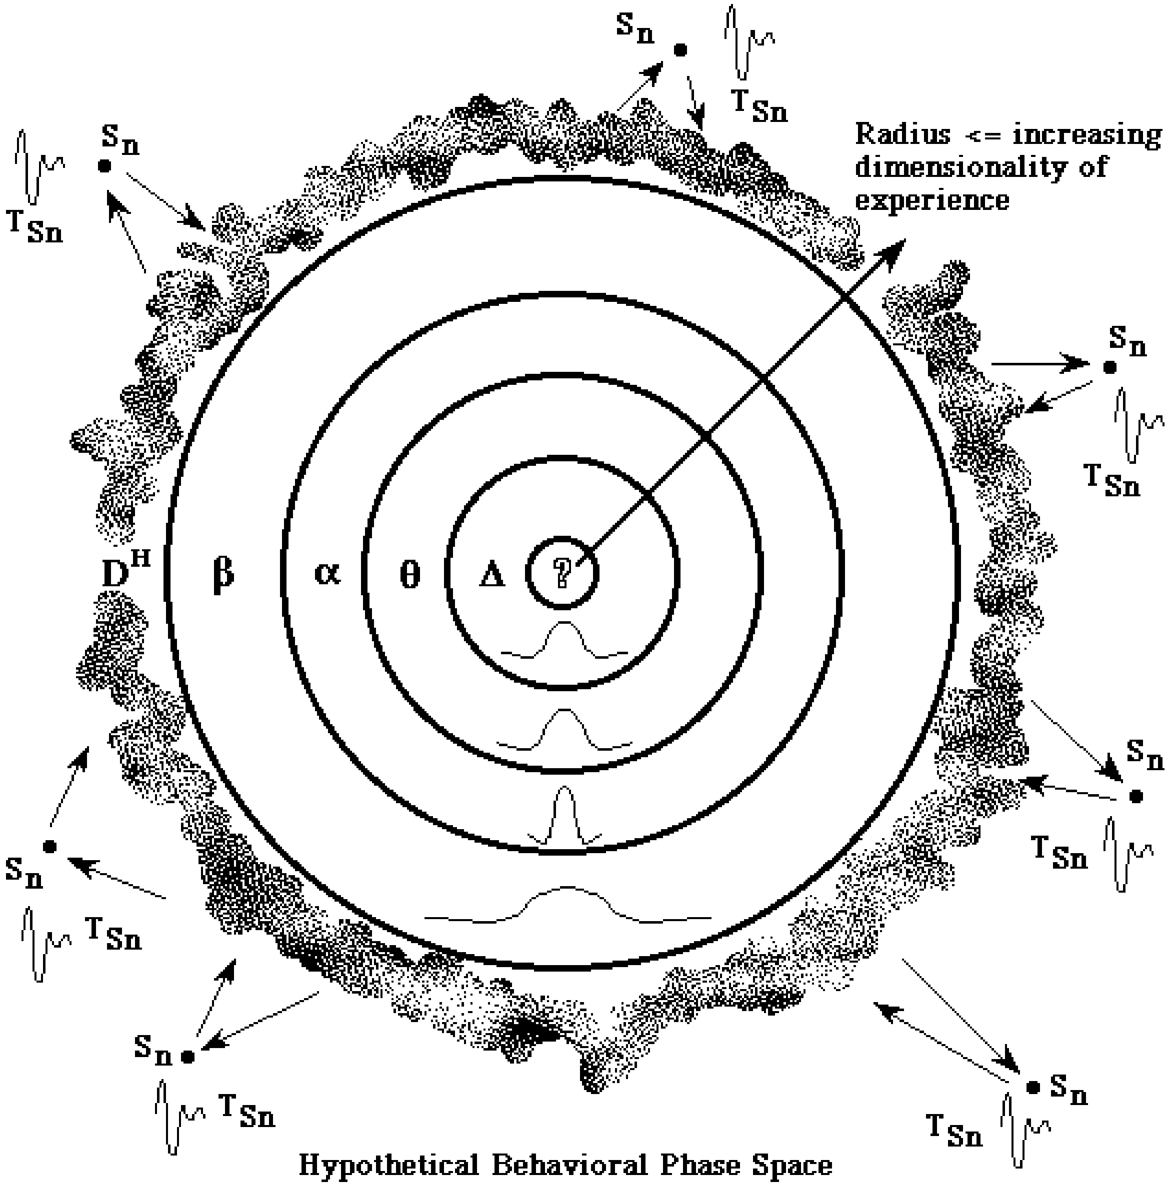

Figure 15

Fig. 15. Hypothetical Behavioral Phase Space. A suggested method for visualizing relationships among characteristics of the EEG, behavior and states of consciousness as a function of complexity. As one moves outward from a dimensionless origin, the EEG components depicted here become associated with experiences of greater and greater complexity, requiring more dimensions for their description. Temporary focus on singular experiences associated with transient waves causes these to emerge from the ongoing background complexity and to be submerged again as attention shifts.

The relative tuning strength of the coherent wave regions as attractors is also indicated using the sum-of-two-Gaussians method shown in Fig. 3. Each of the coherent waves could be described as a kind of coupling of oscillations among neural ensembles. These attractors could be revealed in spatial, (i.e. topological) and temporal frequency distributions of the ongoing EEG. Initial explorations into the possibility of achieving very widespread coupling of coherent waves like alpha, across brain hemispheres and even across individuals, have been carried out [115]. That such phase synchrony can be achieved has been demonstrated. Further explication of its meaning, anatomically or experientially, awaits more experimentation.

Part 3—Some Specific Inferences and Implications Relevant to Musical Experience

Why can't music go out in the same way it comes in to a man,
without having to crawl over a fence of sounds, thoraxes, catguts,
wire, wood, and brass?
-Charles Ives, 1920 [116]

Robert Frances reported on a set of early experiments, carried out in 1954-1955 at the University of Paris, in which an attempt was made to measure physiological concomitants of formal musical perception [117]. He and his associates took polygraph readings, including EEG, GSR, heart rate and respiration rate while subjects listened to carefully selected musical examples. His results were inconclusive as to whether specific physiological events always correlated directly with the perception of particular structural events in music. However, many interesting observations were made in this pioneering attempt. Frances was able to use clearer variations in polygraph readings—most importantly those associated with alpha brainwave blocking—to differentiate among subjects with more or less musical training, subjects instructed to listen actively versus passively, and subjects instructed to listen with an analytical versus a spontaneous attitude. He characterized the experiments as the meeting of two structures: one, the temporal structure of a musical work and two, the temporal structure of the subject brought into the presence of the work. For those subjects relatively inexperienced in music, the correlations between these two structures were inconsistent. In experienced musicians, however, the subject's physiological variations tended to accompany the events of musical structure. These musical events included important changes in melodic lines, textures and harmonies, boundaries between sections of a form, detection of themes, and the apprehension of similarity between two entities, such as melodic shapes.

Since the time of these experiments, techniques for extracting more finely grained information from signals such as the EEG have been developed. This, along with the evolving conceptual powers of theorists and experimenters, suggests that, though much has been accomplished, much more fertile territory lies in wait for tenacious explorers.

MUSIC COGNITION AND STIMULUS-BOUND ERPS

General Observations on Performance

P300 studies have been used to map the reciprocity of neural resources allocated to competing cognitive or performance tasks. P300 amplitudes can be recorded for stimuli associated with human beings performing more than one task at a time. When one of the tasks increases in difficulty, P300 responses to its concomitant stimulus increase, while responses to the stimulus associated with the secondary task decrease [118]. This is assumed to reflect a reallocation of neural resources to the more difficult task. There are obvious implications in this for the execution of the multiple, refined and complex tasks involved in musical performance. One could use related methods to study the mechanisms associated with decisions performers routinely make as to which performance functions to dedicate to continuous, finely controlled processes mediated by consciously applied volitional direction and which functions to allocate to previously learned, automatic responding and execution mechanisms.

There are also implications here for parametric masking phenomena in the perception of musical form (see various following sections).

Music and Emitted Potentials

In a series of previously unpublished experiments carried out in the Experimental Aesthetics Laboratory at York University in the late 1970s, emitted potentials were successfully detected in some musically interesting situations. Through informal communications, word had arrived of some interesting investigations by biofeedback researchers regarding the volitional distribution of attention by highly disciplined meditators such as Zen monks. The relation of this to alpha wave production had, of course, been a subject of interest for some time. It had been noticed, however, that ERPs observed in repetitive-stimulus situations differed radically between normal and highly disciplined subjects. For example, when a repetitive stimulus—such as 100 rings on a bell spaced at equal intervals several seconds apart—were presented to normal subjects, large amplitude attention waves and late ERP components were routinely observed for the first few bell strikes. However, after a certain number of such presentations, these ERP amplitudes were seen to diminish markedly, falling off roughly exponentially over time. In fact, in cases where subjects reported extreme boredom, late ERP components seemed almost to disappear. Sometimes this would accompany the subject being distracted either by something in the environment or by some internally generated dialog of thought. The bell was being ignored, and ERP late components showed this. When highly disciplined subjects were placed in the same experimental condition, and were engaged in maintaining a particularly Zen—like and alert but unfocused state, something else was observed. Recorded ERPs for repetitive bell strokes showed moderate-amplitude, late components that did not increase or decrease substantially over the whole course of an experimental session. One is tempted to conclude that this is indicative of the volitional distribution of attention evenly over all aspects of the subject's environment—in this case, a temporal environment. The popularity of gradual processes and cyclic pattern styles in contemporary music composition was showing a peak at this time. Consequently, interest in this ERP phenomenon was high.

Subsequently, I engaged in a number of further, repetitive-stimulus experiments. In an attempt to separate exogenous from endogenous ERP components, I employed the following method. First, an auditory perception threshold was determined for each subject in pre-experimental trials during which subjects were asked to press a button indicating the detection of tones presented with random inter-stimulus intervals. The amplitude of the tones was varied until the subjects accurately detected 50% and missed 50% of the tone presentations. The tones, generated on a synthesizer, were triangle waves with nearly instantaneous attacks and exponentially decaying amplitude and spectral envelopes. They were presented in a quiet room on very small speakers placed about 8 inches to either side of the subject's head. In the first experiment, tones were presented at the predetermined auditory threshold with random inter-stimulus intervals. Subjects were asked to press a button when they heard a tone. Later, ERPs were averaged separately for the stimulus presentations in three categories: (1) subjects accurately detected tones, (2) subjects missed tones that were presented and (3) subjects inaccurately detected tones when none were presented. For category (1), *correct detections*, normal ERPs were observed. For category (2), *misses*, early components of the ERP, corresponding to propagation of the signal through the sense organs and very early processing stages, were present, but late components were absent. For category (3), *incorrect detections*, early components were absent, but late, endogenous components corresponding to the *imagined* events were present. John and others have also reported recording endogenous components from expected-but-absent stimuli [119].

In a second experiment, tones slightly above auditory threshold (the lowest amplitude for which subjects could accurately detect 100% of tones played) were presented at regular intervals. In different trials this interval was varied from approximately 500 msec to 2 sec. The subjects, most of whom were musicians, sometimes reported informally that they tended to group these constant-amplitude stimuli by subjectively superimposing an imaginary accent, sometimes over every fourth and sometimes over every third tone. There seemed a natural tendency to organize the constant stimuli into chunks corresponding to 4/4 or 3/4-meter groupings. After some time, and by random determination, tones that corresponded to the first beat in these groupings were dropped out. ERPs were extracted for these absent-tone cases, in which the subject reported hearing tones by pressing the button. Again, late but not early components of the ERP were observed for these imaginary events. This experiment calls to mind the practice of good rock-and-roll drummers who occasionally avoid playing on the down beat, allowing listeners and dancers to fill in the missing beat with their own emphatic, visceral response. If better on-line techniques for recording and analyzing single-trial ERPs were available, perhaps this perceptual grouping phenomenon of superimposed, imaginary accents could be observed as it takes place. Tones could then be eliminated according to the actual grouping patterns being applied by a given subject.

I also carried out further explorations relating to the perception of form in gradual-process music. In one experiment, subjects listened to a recording of the composition *Piano Phase* by Steve Reich. (This version was performed on marimbas by percussionists at York University.) The form of this composition involves a gradual separation in the phase relationship between the execution of a common melodic pattern by two players. One performer very slowly increases her or his tempo relative to the other player until the patterns have moved one note out of phase. The process then continues through other phase relationships. My interest was in trying to track the process of pattern segmentation in perception. Initially, in a good performance at least, the two unison patterns will fuse into a singly perceived musical entity. However, at a critical point in their phase separation, as one performer speeds up, the single pattern is perceived to split into two distinctly separate entities. This is a highly discrete event in formal perception, perhaps related to the crossing of a threshold in phase separation. It seemed that there might be observable events in the EEG concomitant with the crossing of this threshold.

Since it would be difficult, if not impossible, to trigger a signal-averaging computer to begin analysis of an EEG epoch with such a subtle shift in mode of perception, and since the recording of a single-trial ERP concomitant with this event would be subject to serious uncertainties due to noise and other technical limitations, I used a measure of EEG desynchrony instead. Bilateral recordings of temporal lobe and vertex EEGs were recorded and subjected to further analysis by an on-line correlation function computer, such that a continuously updated display of the EEG autocorrelation function was visible. Subjects were first given an initial period of 10 to 15 minutes of alpha-wave feedback practice to calm the mind and isolate the experience from ongoing, daily business. (Many subjects with considerable experience and skill at brainwave biofeedback were available at this laboratory.) Next, the subjects were instructed simply to listen attentively to the recording. Points in the music corresponding to significant EEG desynchronization, as shown in the autocorrelation function, were noted.

Later, the same subjects were presented with the recording again. This time they were instructed to press a button when they first noticed the perception of two lines, rather than one line, of music. These points were again noted and compared with the corresponding EEG desynchronization record from the first presentations. In most cases there was a temporal correspondence between the point in the music at which the button was pressed and the occurrence of a significant EEG desynchronization event. Furthermore, these appeared to occur at about the same place in the music for nearly all subjects.

There are many factors in a recording of live, acoustic music like this that could cue the perception of segmented musical lines in addition to phase separation. Attack times, along with spatial and timbral differences, could be partially responsible. Consequently, in order to achieve a rigorous quantization of these phenomena, recordings of synthesized tones would probably be required so that controls for these various factors could be devised. Nevertheless, I considered the above experience to represent at least a highly replicable and observable correspondence between an important phenomenon in the perception of musical form in gradual-process music and the manifestation of that phenomenon in the ongoing EEG. Additionally, these studies hold relevance for investigating other segmenting factors in musical form perception. Skilled musicians often work with the phenomena of controlled acoustic fusion and differentiation among instruments, along with cohesion and segmentation among parts of a formal structure. Formal fusion, or cohesion, and acoustic fusion probably involve different, hierarchical levels in the activation of our cognitive apparatus. Acoustic fusion among single tones or small groups of tones may depend primarily on relatively low-level psychoacoustic processing. It is sensitive to such things as attack times, spatial separation, and crosscorrelations among waveform characteristics. Structural cohesion, on the other hand, involves comparisons on the level of formal architecture and may take place even when the requirements for acoustic fusion may not be optimum. Investigations like these begin to suggest ways of detecting correspondences between fluctuations in a bioelectromagnetic field interacting with an unfolding musical holarchy, and our analytical understanding of the form of that musical holarchy.

The results of these and other experiments guided the formation of techniques that I later applied in modeling an interactive electronic music instrument in which such phenomena are used to guide the unfolding evolution of a musical form (See Part 4, *On Being Invisible—Using ERPs to Build Formal Musical Holarchies in Real Time.*)

These experiments were considered preliminary investigations warranting more controlled follow-up experiments. For a variety of reasons (technical, logistical, economic, etc.), these have not yet been performed. However, the preliminary results seem interesting enough to report here, if for no other purpose than to stimulate ideas for further work.

MUSIC, ORDERED TIME AND MENTAL CHRONOMETRY

Time will throw its vices away and weld its virtues into the fabric of our music.
-Charles Ives [120]

Many of these ERP studies have implications for understanding the flow of musical time, performance and temporal events involved in the perception of form. *Mental chronometry* refers to the study of the order and timing of cognitive events. Studies of reaction time, time required for accurate recognition , speed of mental or physical task execution, and other related factors comprise the domain of experimental testing. Inasmuch as the order and timing of formal discriminations in musical language comprise the stuff of much creative music making, these studies have considerable relevance to the understanding of music.

The relation of ERP peak latencies to *reaction time* and *recognition time* has received considerable study. Some things are clear and other things are not. P300 latencies are not always easy to identify because sometimes the wave is sharply defined and of short duration and sometimes the wave is more spread out. This may be significant. The area contained under the P300 curve may become a useful indicator of the qualities of the underlying neural dynamics. P300 amplitudes are not always correlated with latency, i.e. long latencies do not always occur with large amplitudes. Short, large waves are associated with signal detection. This may reflect the minimal cognitive processing demands required by simply directing attention and reporting the detection of a signal. Long, large waves are associated with processing low probability and infrequent stimuli, increasing the difficulty of a task associated with a stimulus, and are often also accompanied by longer reaction times. When a subject is required to make a sensory discrimination between multiple stimuli with multiple confidence levels, highly confirming feedback stimuli (indicating correct choice) elicit short-latency P300s, while highly disconfirming feedback stimuli elicit long-latency P300s. However, P300 does not always correlate with reaction time. N200 is a better ERP concomitant of reaction time. Possibly, the endogenous processing paths branch apart after N200, as depicted in Fig. 12. This is supported by the fact that P300 latency seems sensitive to the duration of the stimulus evaluation and categorization process but rather insensitive to the response selection and execution process [121]. Reaction time, on the other hand, is strongly influenced by both. For music, this might mean that the time we require to select and execute responses in performance situations—for example, in interactive improvisation—is independent of the time we require to evaluate and classify a musical event. Consequently, the mechanisms by which we learn to recognize objects may be quite different from those we use in order to learn to react in performance situations. Perhaps ERP component studies can be used to investigate the rehearsal of control processes brought to bear on improving memory and the proprioceptive mechanisms associated with performance.

Reaction time and object recognition paradigms are particularly interesting in relation to improvisation. In developing improvisational skill, one learns a store-house of knowledge about musical objects and builds up a large file of musical object templates. All of these are linked to knowledge about musical structures and forms and the probabilities of occurrence of various objects in various situations. Together, these may describe a *musical language* or *meta-language syntax* on relatively high levels of structure, as well as on low levels of note-by-note sequences. P300 amplitude is inversely related to the probability or frequency of occurrence of a given stimulus. In an ongoing improvisation situation, one may intuit the likelihood of occurrence of events, possibly involving unconscious calculations of event probabilities and, in rapid-fire sequences, predict event occurrences along with their precise timings, while attempting to execute events synchronous with these predicted timings. One senses that a synchrony is about to occur and makes a decision to 'go for it' more rapidly than one can utilize slower, reasoning processes. One learns to maximize one's tendencies to hit musical targets correctly, achieving synchronies in complex musical language structure with other performers who are similarly skilled. Such ability requires years of practice and is difficult to describe or to teach. The study of the details of mental chronometry and concomitant electrophysiological events therefore bears considerable relevance to probing the mechanisms of this kind of high-speed musical performance and musical information processing.

Active music listening (as opposed to *passive reception* of what is essentially entertainment), in which the listener maximizes the richness of her or his musical experience by participating in an exploration of a varied stimulus field, no doubt involves similar rapid, chronometric processing. One can envisage the image of the bow of a boat racing through a sea of musical relationships, the bow spray representing a stream of rapid associations and recognitions that can be perceived in a rich musical field. It is not possible for the conscious, rational mind to catch up with this bow, or for the improvising performer to execute musical events with nearly the speed that this bow spray of perceptions can provide. Nevertheless, they are real perceptions and categorizations that can be recalled, recognized and even reproduced later. Musical cognition most likely has its own mental chronometry, separate from other frames of mind, such as that associated with linguistic processing.

P300 exhibits longer latencies to new stimuli than to old stimuli. Repeated events eventually induce automatic processing, after which new events elicit a reallocation of attention resources. Repetition in music rehearsal relegates functions to automatic processes. Some of these involve muscular learning and muscular memory. The use of ostinati and other repetitive musical processes provide, after allocation to automatic responding or recognition, constructed frameworks for the definition of other foreground events.

Music is full of sequential dependencies. The perception of these may be affected by the timing of cognitive processes, some of which may be revealed in how P300 latencies relate to response reaction times. Chase summarizes the known factors affecting response reaction time and P300 latency in the following list [122]: (1) the size of the memory set that must be scanned to classify an event; (2) stimulus frequency; (3) noise—i.e. it is difficult to distinguish events with imprecise differentiation boundaries or those that are confounded by noise, requiring more processing time; (4) fixed versus variable stimuli—i.e. pattern prediction probability; (5) the number of stimulus targets in a task; and (6) old versus new items in recognition memory. Similarly, he lists known factors affecting reaction time but not P300 latency: (1) the compatibility of a stimulus with an anticipated response; (2) sequential dependencies of stimuli on previous stimuli (particularly significant for music); and (3) comparisons among words that match in sound but not in spelling, or vice versa, such as gasp-wasp, way-weigh, cough-dough.

The last item could be referred to as the uncoupling of the orthographic from the phonologic aspects of the form of word objects. This has particular significance for music. It is a common compositional transformation device to uncouple various parametric components that together make up the form of a musical object. For example, melodic shape may be subjected to a transformation independently from its rhythm or note duration series, or vice versa. Such devices were commonly used by Stravinsky and have since been extended to nearly all conceivable parameters of musical form. Another simple example involves two different types of transposition—interval versus modal. Both maintain approximate melodic shape, but one maintains interval sizes and the other alters interval sizes. The perceptual comparisons that result become part of the language with which the music unfolds.

Such coupling and uncoupling also introduce masking phenomena. The newly transformed parametric shape becomes a *delineating* or *segmenting* parameter, which will tend to mask changes in other parameters, depending on the degree of transformation. The parameters that are held to their original contours act as the *cohesive* parameters, providing a backward reference to the original form. These processes contribute to the building of a formal gestalt hierarchy, as has been extensively treated by Tenney [123, 124].

Perception Chronometrics

I was led some years ago to construct a chronometric ordering of important stages in perceptual information processing. My interest was in understanding the subjective properties of musical attention. The intent was to construct an ordering that reflected the kinds of shifts in attentional focus that occur during the brief period immediately following a perceptual differentiation, conceived as the detection of a change of something in the environment. It is called a *priority ordering of perceptual gating* (see Fig. 16). These shifts of focus may not be rigidly discrete and, in fact, may reflect overlapping processes to some extent. I am convinced that access to these levels of perception is available to highly aware and disciplined musicians, who may use information about them to guide their functioning.

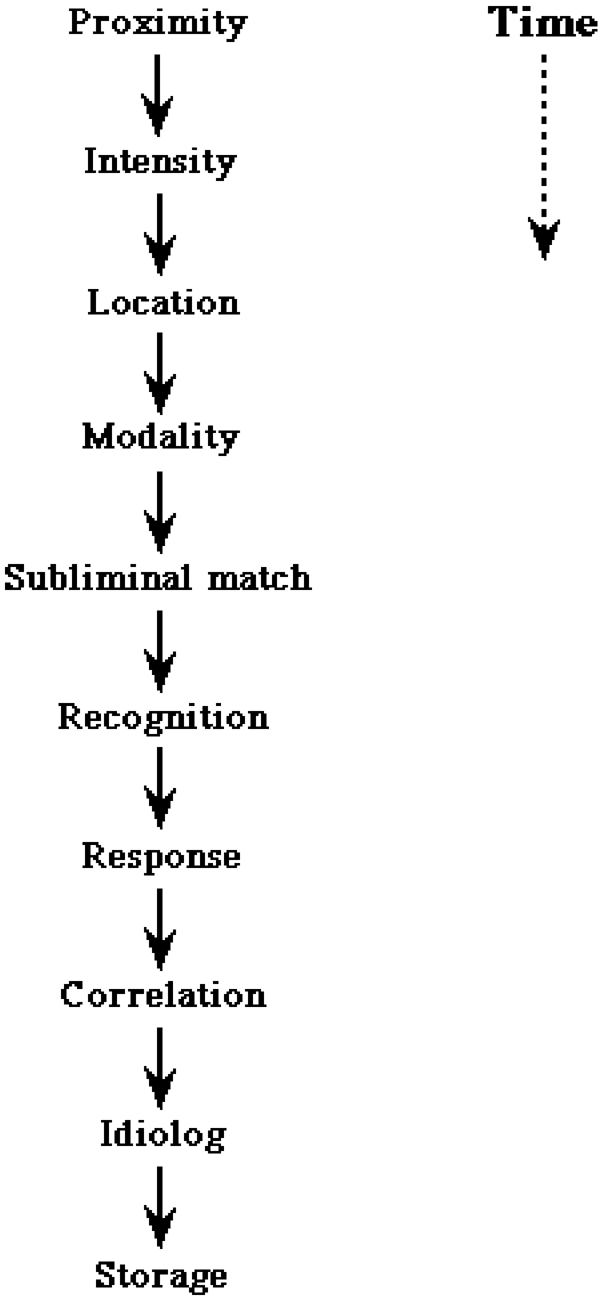

Priority Ordering of Perceptual Gating

Figure 16

Fig. 16. Priority Ordering of Perceptual Gating. A proposed ordering in time of priorities for extracting, evaluating and storing information from incoming sensory data.

Proximity to the Organism.

The highest priority for an organism detecting a change in the energy field of its immediate surroundings is to determine how close the event is to itself. This determination is made from primary sense data, is preconscious and takes place even before further information abstractions regarding location, direction or quality are determined. A decision is made about whether to activate tracking mechanisms with which to follow the event. Determinations regarding arousal level are made. Selections are made regarding which secondary reflex arcs to activate, and/or determination is made regarding which primary reflex arcs (such as spinal reflex arcs) may already be activated. These are examples of decision making at the most primitive level. The decisions are extremely rapid and represent a stage of processing barely higher that the sense organs and primary reflex arcs themselves. Conscious, cognitive processes are allowed access to this information only long after its primary biological significance to the well-being of the organism is determined.

Intensity.

Next, a measure of the intensity of the energy change is extracted, informing higher processes and feeding back to primary reflex response generators. This is still preconscious. Awareness of sense modalities or sensory qualities has not yet occurred.

Location.

A determination of the location of the source of the energy change is made, provided the information available is sufficiently complete to allow it. If the determination is made, the information is fed back to response generators and sense organs for further tuning. Of course, the information is fed forward to higher processes as well. Errors in location determination may be made at this stage, sometimes causing inappropriate response patterning. If information is not sufficient to extract spatial location with a degree of confidence that crosses a set threshold, the organism must wait for the time required by higher calculating mechanisms to give further estimates of position or for more sensory information to arrive. This can be dangerous to the organism. Consequently, confidence thresholds for events of high intensity can be set quite low to maximize the organism's opportunity to react with minimum delay.

In low-danger situations, such as musical performance, feedback from higher-volitional mechanism can override the tuning of these early detection processes so as to minimize response time to incoming events. Broad tuning shortens response times; narrow tuning requires more processing time for reliable determinations to be made. By means of controlling one's state of consciousness and therefore the level of cognitive interference with these primary processes, one can tune the system to respond nearly instantaneously to very small energy changes. Responses may be directed further by information based on probabilities of event occurrence, calculated by higher mechanisms and fed downward through the chain of processing units. Biofeedback training can be used to achieve this control. The basic effect of such training is to place individuals in control of allocating their neural resources to those processes of interest. Feedback training with high-frequency coherent waves—for example alpha, and beta—has the effect of deallocating resources from the internal dialog of discriminations and emotions and from the ongoing activity of local, spatio-temporal differentiations and abstractions emerging from the uncontrolled flood of sense data. In a coordinated way, however, it places these resources on alert and in reserve readiness for their immediate activation toward focused discrimination and response generation when required by volitional consciousness. Furthermore, such activation can, in the right circumstances, take place with minimum disruption of the coherent wave state.

In music, this relates directly to performing actions that can manifest in a special state of alertness, with great speed, and in close communication with other performers. States of preparedness, just prior to beginning the execution of very complex but highly rehearsed performances, involve some of these processes. Finally, a state of performance consciousness is sometimes achieved, especially during improvisation, in which one can be surprised by one's own actions and choices, seemingly arising from a special part of the mind, separate from that normally associated with conscious determinations. Such actions can, however, be highly logical and can provide rich results on which to build a carefully constructed musical edifice to follow in real-time performance.

Which Sense Modality

Next, and only next, comes conscious awareness of which sensory channel is responsible for the incoming information.

Subliminal Memory Match

Next, a feeling is generated that an occurrence of a match to something held in memory has taken place. However, awareness of just what the match is has not yet become possible. At this point, action synthesizing and further processing of stimulus information may split off from each other. This kind of action synthesis is of a more deliberate kind than that referred to above. The deliberation may not be highly conscious; however, the actions involved are sufficiently complex to require that significant neural resources be devoted to their construction.

Recognition or Not (of What)

At this stage, a preliminary recognition of the qualitative and quantitative content of the incoming information is determined. The determination of whether there is a match to something identifiable in memory takes place. If no match is identified, resources are activated to further analyze and classify the data and to update the model of the environment currently held in working memory. This currently relevant model of the environment may be a broad-based one or it may be relatively limited if the current focus of the individual is restricted to a particular sphere of activity. For instance, in a musical situation the environment may be the particular piece being performed, the language or musical style currently involved, the domain of improvisational material presently at hand, or the acoustic landscape of a rich sonic surround. The individual may change the scope of the focus of working memory at any time. However, this may require conscious, volitional action. One must remember to be aware of the scope of the current content of working memory in order to develop appropriate contextual evaluations of experiential events.

Message-Specific Response

At this stage, the decision is made as to whether and how to generate a response specific to the complex information received. The response construction will be relatively deliberate at this stage; however, I still refer to rather rapid, possibly learned, responses.

Correlation Connection

This is a stage of further analysis that may take place consciously or somewhat subliminally. In any case the reverberant pattern of activity elicited by the incoming information will be correlated with information scanned in memory. Bear in mind that this is not a sequential scanning, as in search procedures with sequential computers. It is more like a multi-dimensional waveform correlation process in which particular resonances brought about by wave-pattern reinforcements are noted and classified. The detection of such correlations or resonances is what is meant by *correlation connection.* Activity at this stage is composed of the occurrence of such resonances.

Formation of the Idiolog (Judgment)

This is the stage at which the idiolog is synthesized. The idiolog—as the term is used in Clynes [125]—is the total, multi-dimensional, electrochemical, electro-magnetic and anatomical morphology, *which must be synthesized* as a result of the characterization of a differentiated multi-sensory event. The scope of the event may be large or small. It is a representational entity that will be added to the accumulating history of the individual. It is not external reality but a self-constructed and self-organized interaction with physical reality, requiring judgment. It is somewhat akin to John's *hyperneuron.*

Storage

Finally, the synthesized idiolog is integrated with and enfolded into the ongoing patternings of the large-scale representational system that constitutes the brain's evolving memory.

Evolving Models for Mental Chronometry

Mental chronometry—or *scheduling theory*, as it is also called—has largely been based on the traditional concept of *additive factors* as developed by Sternberg [126]. This method predicts reaction times on the basis of a breakdown of times required to perform individual cognitive components of particular tasks and their addition in sequence. The method works to some extent but reflects only a *sequential* view of information processing. It is much more likely that in a series of *cascaded neuronal processes* it is not necessary for one process to play itself to full completion before the next process begins. The paradigm works for sequential computers carrying out discrete mathematical and logical tasks with precision. However, the brain is not optimized for such tasks. It is much more suited for rapid generalization and recognition activities. Certainly we know now that the brain is a *highly parallel, concurrent processor* involved in extracting different kinds of information from the same data items at the same time with varying degrees of certainty. For efficient allocation of brain resources, it is in fact essential that cascaded processes overlap and continuously update one another with more refined information. Furthermore, the brain is much more of an analog device, working with continuous, graded responses, quickly passing from unit to unit, than the current digital vogue tends to communicate. Both *discrete, action potential* and *continuous, ionic potential-gradient* processes are involved. McClelland addresses this in his analysis of cascaded stages [127]. In this view, each processing stage is considered to begin its work after its predecessor has begun, but before that predecessor reaches its maximum activation state.

In addition to cascade stages, we must account for *concurrence* and a widespread fanning out of projection paths from one representational or processing entity to many others. Schweickert applies *critical path scheduling theory* to Sternberg's additive factors analysis in an attempt to account for *concurrency* but does not really address *spreading activation* [128]. A hypothetical depiction of concurrency and spreading activation is shown in
Fig. 17.

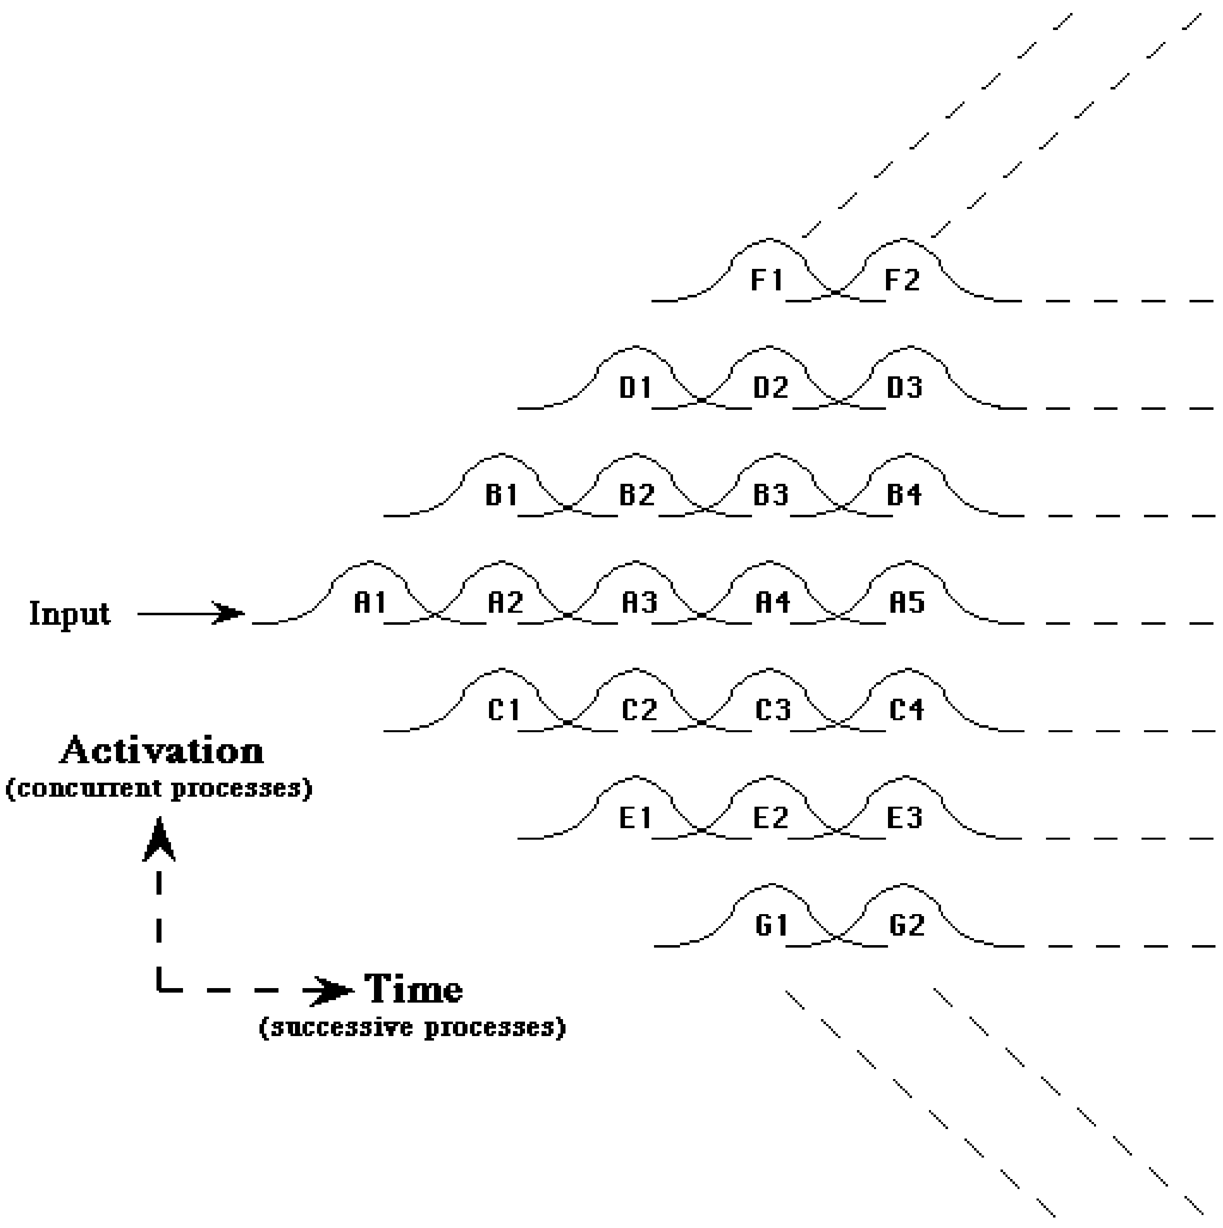

Concurrency and Spreading Activation

Figure 17

Fig. 17. Concurrency and Spreading Activation. A view of how the activity of neuronal groups may be stimulated from an input and then spread to both successively and concurrently activated groups.

MUSICAL FEATURE EXTRACTION

The neurological mechanisms underlying the extraction of what we refer to as musical *features* has, in my opinion, been the subject of only very preliminary and tentative investigations. Extraction of musical features requires more than mere recognition. It implies a drawing out of features by means of effort. It is an activity highly dependent on learning, practice, the tuning of perceptual mechanisms, rehearsal, and the development of skill in identifying phenomena that are potential candidates for musical features. It involves giving attention to similar differences and different similarities, to paraphrase Bohm [129]. Consequently, examining processes of generalization and differentiation may be the best way to begin such an investigation. The behavior of various components of the ERP, as they reflect and track mechanisms of selective attention, could be useful in discovering clues about the neural activity involved in extracting musical features.

John reports on recording ERPs to varying stimuli that a subject has perceived to be equivalent to a previously learned test stimulus. These stimuli evoked response shapes similar to those of the test stimulus [130]. Furthermore, a stimulus that is half-way between two differentiated test stimuli on some parametric axis—like frequency—may be differentially generalized to be equivalent to one or the other test stimulus. ERPs to the midpoint stimulus will correspond to whichever test stimulus it is being generalized to, possibly indicating differential readout from memory. This demonstrates an effect of experience on perception. It also has implications for the use of multi-dimensional scaling in musical perception research, such as has been extensively applied in studying timbre perception [131, 132]. Such mappings equate the distance between musical objects in a parametric space with perceived similarity or dissimilarity. The method has also been extensively applied by myself and others as a compositional aid. In this application, multi-dimensional *concept spaces* are constructed with arbitrarily chosen axes relevant to the composer's intentions. The only requirement is that the parameters chosen for these axes serve as segmenting parameters in the perception of musical gestalts. The effects of applying transformations to musical objects—for example, melodic or rhythmic shapes—can be seen in how their positions change in the chosen concept space. Regions of the space can be selected for manipulation or mapping into other higher-level spaces, as well.

Stimulus generalization can be mapped in such spaces as a region inside which differentiations tend not to be made. Shepard has developed a theory of generalization that characterizes the metric nature of these regions and the probabilistic nature of generalization processes within *psychological space* mappings [133]. He identifies stimuli in such spaces according to whether they belong to 'consequential regions'. Stimuli lying inside a consequential region are generalized because they are associated with similar consequences for the organism. He further points out that generalization results from cognitive processing and is not to be considered merely the failure of sensory discrimination. In such a space the probability that stimuli will be generalized is approximately equal to a negative exponential function of their distance apart in the space.

The computer language for composition and musical research HMSL (Hierarchical Music Specification Language) includes features directed toward facilitating this kind of exploration and musical production [134]. HMSL contains a mapping facility for the behavior of musical objects, in which such objects carry a weighting factor, often used to determine the likelihood that objects will be executed in non-deterministic sequences. These objects also contain links to other objects. The links have tendency values associated with them, which could be likened to connection strength in a neural network. Multi-dimensional mappings of these musical objects can also be made. Transformations of the objects appear as trajectories in the mapping space. The trajectories are associated with musical qualities or features, which are also mappable in higher-level spaces.

Obviously, what constitutes musical features and, further, a mapping of relationships among those features requires a deep understanding of the psychological space in which these features are imbedded for given individuals. Such a space will have a number of dimensions corresponding to the minimal set of parametric scales required to achieve reliable differentiation within the space. Physical scales, such as frequency, are almost never, by themselves, sufficient.

Perception of Musical Form

Musical form can also be considered a feature to be extracted from effortful listening. Again, the perception of different similarities and similar differences is key in finding those musical features that carry the information that articulates the musical form. Tenney has presented the concept of *temporal gestalts* (TGs) as primary units in the perception of form [135]. Musical entities are grouped into TGs according to the traditional gestalt grouping principles of *proximity* and *similarity*, applied, however, in the time domain. To these I would add the time-based principles of *common fate* and *calling for similar action in response*. For example, parameters may move in some psychological space along trajectories with similar shapes. The 'common fate' principle refers to an often-ignored, strong, cohesive force that binds objects together in perception when they move in a similar manner. The principle of 'calling for similar action in response' is a psychological one related to Shepard's idea of the consequence associated with a stimulus for an organism. The identification of TGs involves the detection of parametric differences that outline their boundaries. Numerous 'difference detectors' have been tried, and Tenney and Polansky have outlined an approach [136]. I have used a related but somewhat different approach in making predictions about the kinds of parametric changes that will be perceived as structurally significant to a subject (discussed in the section *First Few Versions* in Part 4). Most approaches involve an examination of contextually sensitive changes of rate-of-change in some musically significant parameter. The determination of the appropriate parameter, which may be multi-dimensional, is critical, of course. Further difficulties are encountered in tracking the changing thresholds of difference required of a given parameter to induce a boundary discrimination and in allocating the appropriate weight to a given parameter as a determinant of formal perception in a given musical context .

It is assumed, then, that the identified perceptual units (TGs) are enfolded in a hierarchical (or possibly, holarchical) musical form on several levels, leading up to the entire work as an entity. ERPs and other EEG components are related to detecting these perceptual units, maintaining the attentional focus appropriate for extracting musical features, and building a cognitive, holarchic entity in response to a musical experience. The difficulty in relating these EEG components to such musical perception and cognitive activity is bound up as much with the limitations on our present understanding of what musical features and musical forms can be as it is with technical difficulties in sensing these components or designing appropriate experiments with which to investigate their behavior.

Absolute Pitch and P300

An interesting item appears in the literature regarding absolute pitch and P300. Klein et al. [137] reported that music students with absolute pitch did not exhibit strong P300s in a task requiring detection of an infrequent pitch, while music students without absolute pitch did. It is suggested that persons with absolute pitch have an internal standard that allows them to retrieve the name of the pitch without engaging in a process of comparing the tone with a standard. The large P300s of persons without absolute pitch suggest that comparison and judgment processes are required for these subjects.

Part 4—*On Being Invisible*—Using ERPs to Build Formal Musical Holarchies in Real Time

In 1976 I began creating a work entitled *On Being Invisible*, which, for me, contains the richest aesthetic, symbolic, and metaphorical content arising from the import that biofeedback systems had on my work as a composer. *On Being Invisible* is a self-organizing, dynamical system, rather than a fixed musical composition. The title refers to the role of the individual within an evolving, dynamical environment, who makes decisions concerning when and how to be a conscious *initiator* of action and when simply to allow her or his individual, internal dynamics to co-evolve within the macroscopic dynamics of the system as a whole. Consequently, the work is always ongoing. Within the corpus of my music, the title serves as a label for a period of work with these ideas from about 1976 to 1979. Recently, after concentrating on other things for several years, I have begun new work with this system, calling it *On Being Invisible II*. This work is stimulated partly by advances in technology that only now make the realization of earlier concepts possible, and it is partly the result of interest in applying new knowledge within a still very rich musical paradigm.

BASIC PARADIGM—THE ATTENTION-DEPENDENT SONIC ENVIRONMENT

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.

This raised interesting questions about proprioceptive learning. To what extent can a language develop through this kind of interaction with an environment? Piaget has pointed out extensively that the acquisition of language skills necessarily involves physical interaction and feedback with one's environment. This results in what he termed 'communicative competence', (as opposed to what Chomsky has termed 'linguistic competence').

In a sense, *On Being Invisible*, a musical composition, has at the core of its structure a model for a way in which language can be acquired. Leaving open questions about physical versus cerebral interaction with an environment, this system 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 what might be called the *cybernetics of language and cognition*.

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 between neighboring goal states, which might otherwise be generalized according to Shepard's theory, is produced by the sum-of-two Gaussians method (discussed in the section *Biofeedback—Definition and Modeling* in Part 1 and depicted in Fig. 3). Indeed, Shepard suggests that his negative exponential generalization function is made to look more Gaussian as a result of 'noise' in the internal representation of a stimulus. In the case of the model implemented in *On Being Invisible*, the noise is associated with an uncertainty in the precise representation of the goal state. Thus, there is a stochastic shaping of language structure as it develops. Further work will probably produce refinements in this representation, the exact nature of the functions involved, and the proper metric space with which to compare them.

FIRST FEW VERSIONS

Each version of *On Being Invisible* employed a system 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 detected and made predictions about the perceptual effect of various phenomena in an unfolding musical structure; (3) a *perceiving, interacting entity* (human performer); (4) an *input analysis system* for detecting and analyzing bioelectromagnetic and other input signals; and (5) a *structure-controlling mechanism* that directed (1) and updated (2) in response to corresponding information from (4) and (2).

An early pre-brainwave version of *On Being Invisible* is, perhaps, worth noting for its relevance to evolving systems concepts. Its functional diagram is shown in Fig. 18. A structure-controlling feedback loop was established here that accepted occasional influence from outside the loop in the form of vocal sounds. To begin with, two Buchla Music Easel synthesizers constituted the *structure-generating mechanism* and sound synthesis system. Structure generation was 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 were combined and sent to a time-delay mechanism implemented with two tape recorders. Both the delayed and non-delayed signals were added to vocal sounds in a connection and summing matrix (mixer) and sent to the *input analysis system.* The analysis consisted of performing an ongoing autocorrelation or cross-correlation on the audio signals from this mixer, chosen in real time by the performer. (See the section *Long-Term Coherent Waves* in Part 2 and Fig. 8 for a description of these correlation functions.) The actual hardware device used was a Princeton Applied Research Correlation Function Computer. 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. A large variety of scan rates were immediately accessible, though relatively slow rates, well below the audio range, were used most often. (The Buchla Music Easel allowed for the storage of complex patches by wiring them onto plug-in circuit cards. This also allowed nearly all continuous controls and discrete selection functions in the synthesizers to be directed from currents generated by external devices, such as a computer with digital-to-analog converters.)

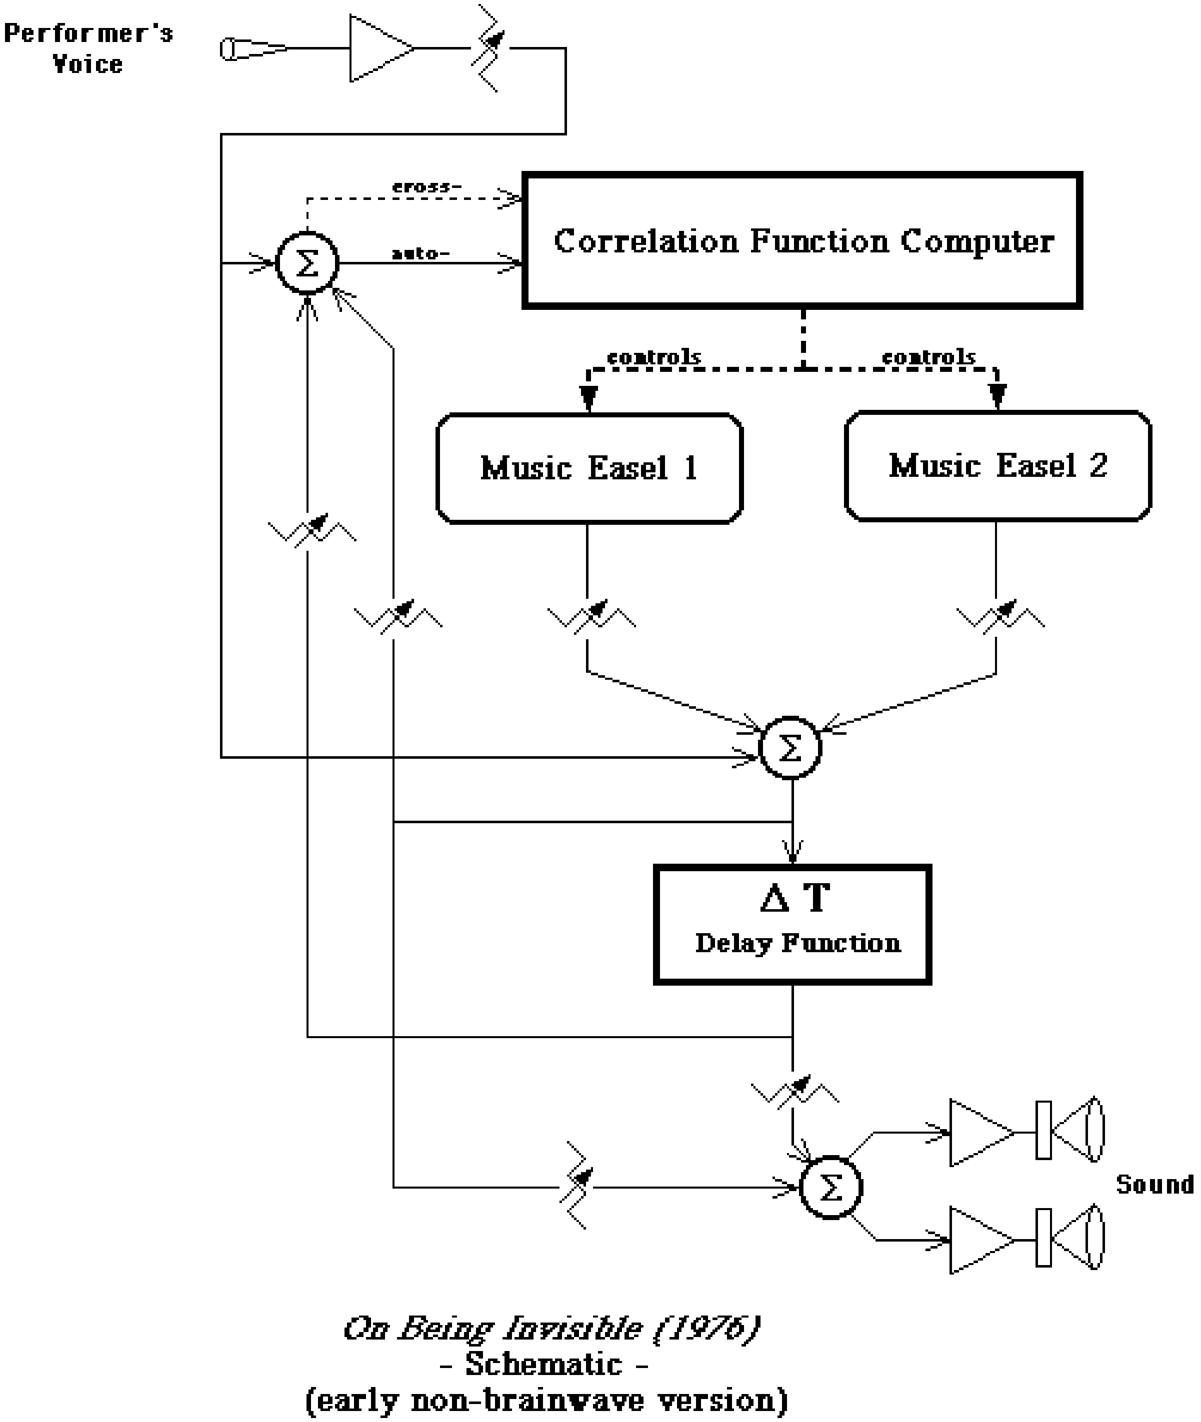

Figure 18

Fig. 18. ***On Being Invisible (1976).*** **The schematic diagram from an early, non-brainwave version of this musical work, which exhibits a certain capability for self-organization in the construction of sound patterns.**

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 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 that would in turn direct global parameters in a synthesis patch.

Some important general systems principles can be seen on closer analysis of the schematic for this piece (see Fig. 18). 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, patterns resident in sequencers. Short-term memory, which is most important for the maintenance of musical patterns, resides in two places—the tape delay system and the correlation function. Note that both of these involve time delay and show striking similarities. The tape 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 one 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 the lack of long-term coherence. To produce a presentation for an audience, activity from nearly any set of points around the loop can theoretically be tapped and projected by means of amplifiers, loudspeakers and, possibly, visual displays.

The entire configuration acts like a non-equilibrium 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 non-evolving, equilibrium systems, or at least the result of attempts to create them. These instruments are constructed to tend toward 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, non-equilibrium 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. There is, rather, a co-evolution of the performer with the performance system, the structure of which is an extension of her or his 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 19 shows a schematic for the next version of *On Being Invisible* that was created in 1976-1977 and first performed at The Music Gallery in Toronto on 12 February 1977. In this version the feedback loop encloses the *perceiving, interacting entity* (the human performer). Energy and information exchanges with the environment took place through the performer, and the sound output was processed by the performer inside the primary loop. The *model of musical perception*, the *structure-generating mechanism* and the *structure-controlling mechanism* all resided inside the software of a minicomputer (Interdata Model 74). The *input analysis system* consisted of a Correlation Function Computer and Fourier Analyzer (Princeton Applied Research), which produced continuously updated 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 could also be output as varying voltages. In this way the highly selective amplitude envelopes of any desired frequency component in the signal could be made available to the computer or synthesis system. One or two channels of brainwave inputs were derived from a combination of electrodes located at the vertex (top of the head) and sites over the occipital or temporal lobes (back or sides) of the performer's brain. Usually, the raw EEG was input to the analysis system. However, sometimes, theta-, alpha- and beta-band filters and envelope followers were used independently. The outputs of these were patched into the synthesis system, producing rhythmically synchronous triggers for sound events. The system maintained the capability to feed sounds into the analysis system from the voice or small acoustic instruments (monkey drums, Tibetan cymbals, snake charmer's pipes, etc.). A novel use of pressure-sensitive touch sensors in a keyboard array 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.*

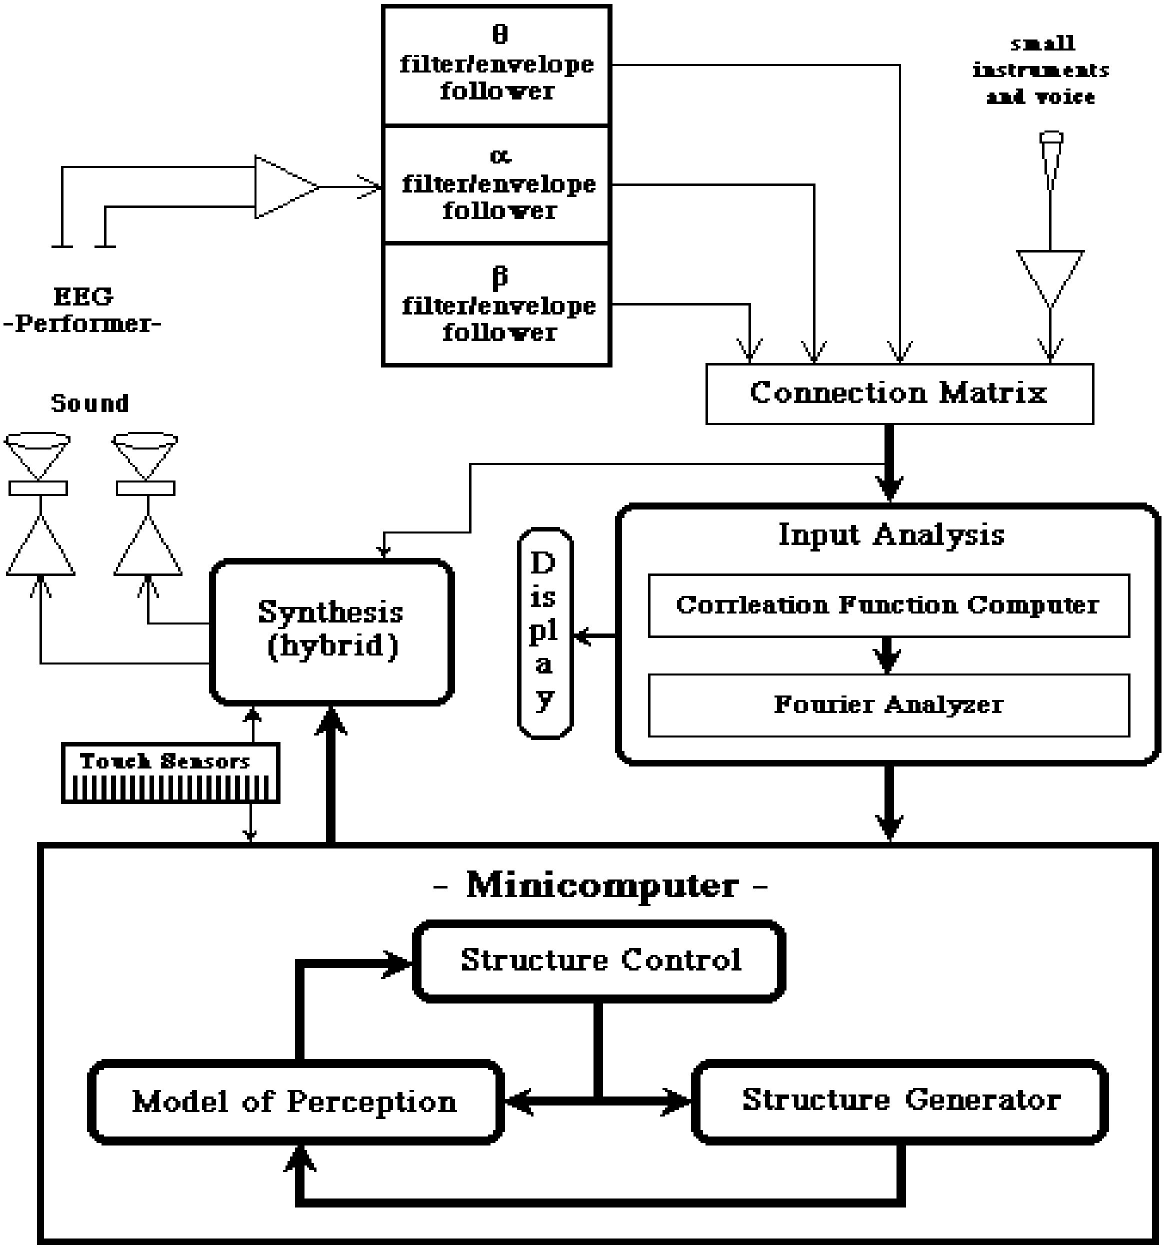

On Being Invisible (1976-7)
- Schematic -

Figure 19

Fig. 19. *On Being Invisible (1976-1977).* The schematic block diagram for a brainwave version of this composition, including all the major system components: a musical structure generating mechanism, a model of musical perception, the perceiving, interacting, human performer and an input analysis system.

The hybrid synthesis system (i.e. digitally controlled analog) consisted of Buchla 200 Series modules and a Music Easel. The control voltages for these were generated by the minicomputer. In addition, I devised a system 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, and its global variables were set by the *structure-controlling mechanism.* Gaussian distributed random values were generated and output 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 means of these parametric values were initially allowed to move according to a 'random walk' algorithm at a rate somewhat slower than that at which discrete values were output. This constituted a second level of control in the structural hierarchy. These mean values would eventually be constrained by the *structure-controlling mechanism*, however.

The *model of musical perception* represents a major addition to the system. Its purpose in this version of the work 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* (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 non-linear 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 [138]. It can be expressed as follows:

$$D_{(t)} \leftarrow \left|P'(t)\right|'_{(t)} \ \textbf{if} \ \left|P'(t)\right|'_{(t)} \geq 0 \ \textbf{else} \ D_{(t)} \leftarrow 0 \quad ,$$

a simple, discrete-time version of which was implemented by performing,

$$D_{(t)} \leftarrow \left|P_{(t)} - P_{(t-1)}\right| - \left|P_{(t-1)} - P_{(t-2)}\right| \quad ,$$

then,

$$\textbf{if} \ D_{(t)} < 0 \ \textbf{then} \ D_{(t)} \leftarrow 0 \quad ,$$

where $D_{(t)}$ is the difference function applied to successive values of an acoustic parameter, P. If $D_{(t)}$ is > 0, 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 $D_{(t)}$ crossed T, the value of $D_{(t)}$ could then be compared with corresponding values from other parametric contours to determine if the element would 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 made 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, i.e. something starting to move or starting 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 to be just 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, i.e.

$$S = D_{(t)} - T$$

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 *input analysis system* to determine if evidence of attention shift was present in the EEG. Significant EEG desynchronization, interruption of ongoing coherent waves, and EEG state changes (as discussed in the section *EEG State Transitions* in Part 2) all contributed to this determination. The *structure-controlling mechanism* was responsible for making the determination. If the prediction was confirmed, this mechanism would increase the probability that the kinds of change in musical parameters associated with the prediction would occur again. If the prediction was proven false, the probabilities associated with these kinds of change were 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 would be 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 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 an 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, or temporal gestalts.

The application of this difference detector and threshold updating mechanism to simple melodic sequences is illustrated in Figs 20, 21, and 22. 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. Note how the changes in the third-to-last note of the melody alter the placement or strength of initiators in the immediately succeeding time region.

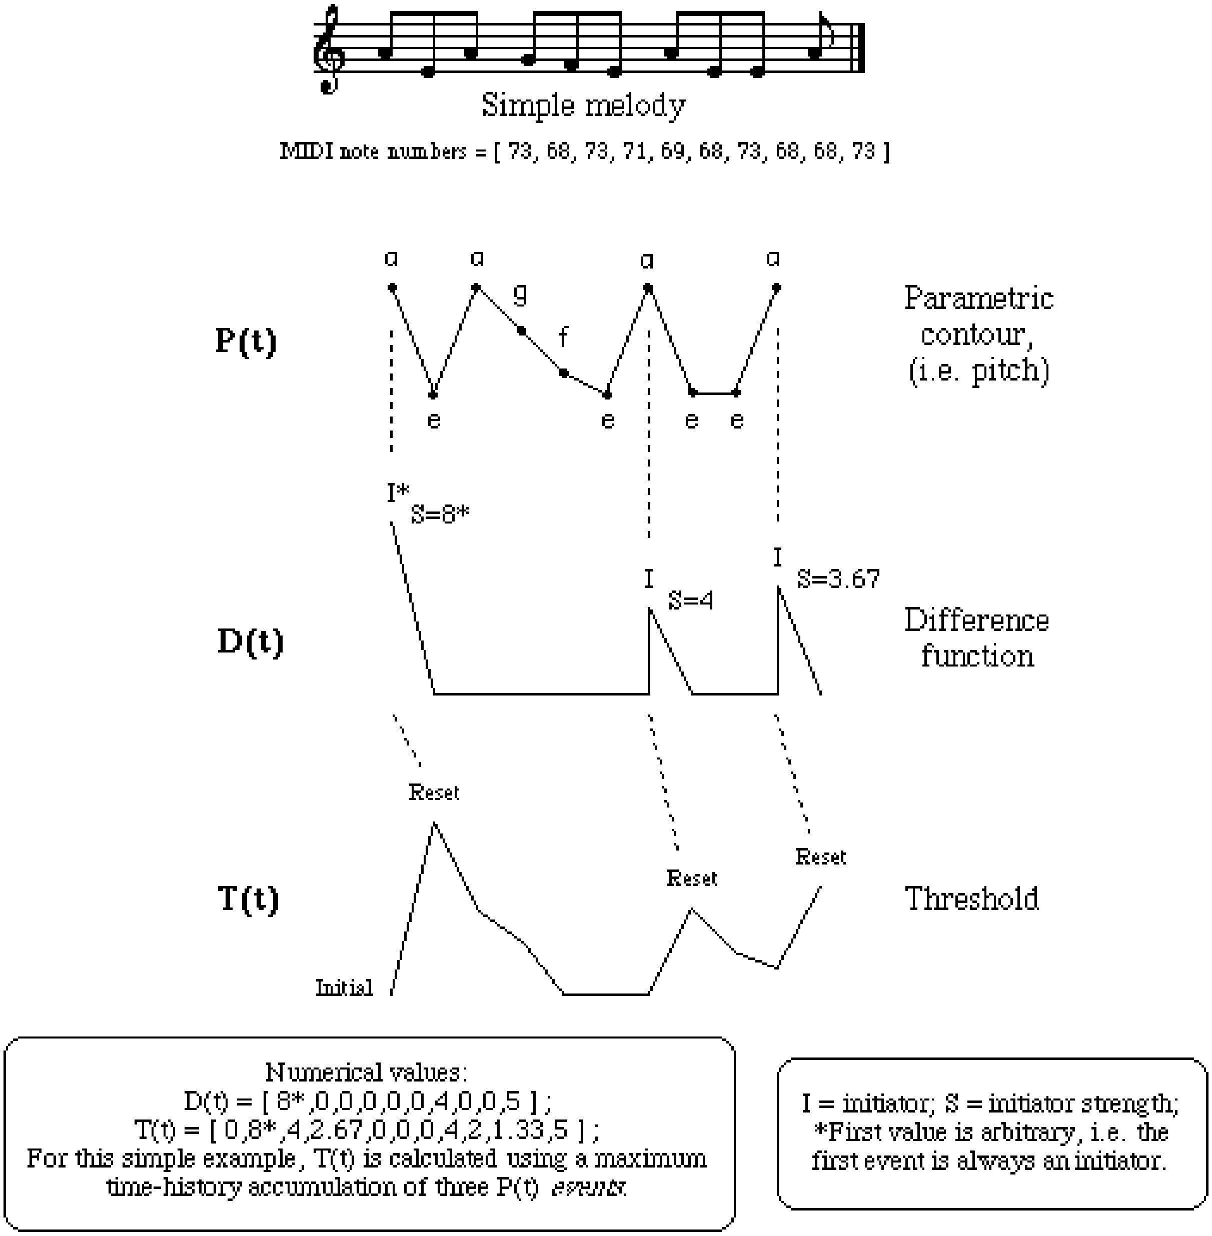

Numerical values:
D(t) = [8*,0,0,0,0,0,4,0,0,5];
T(t) = [0,8*,4,2.67,0,0,0,4,2,1.33,5];
For this simple example, T(t) is calculated using a maximum time-history accumulation of three P(t) *events.*

I = initiator; S = initiator strength;
*First value is arbitrary, i.e. the first event is always an initiator.

Application of the Difference Function to a Simple Parametric Contour

Figure 20

Fig. 20. Application of the Difference Function to a Simple Parametric Contour. An example of how the model of perception identifies important structural features in a simple parametric contour, i.e. a melody, in this case. These features, shown as initiators (I), separate potential groupings in time—or temporal gestalts—which may be associated with subtle shifts of attention by a listener.

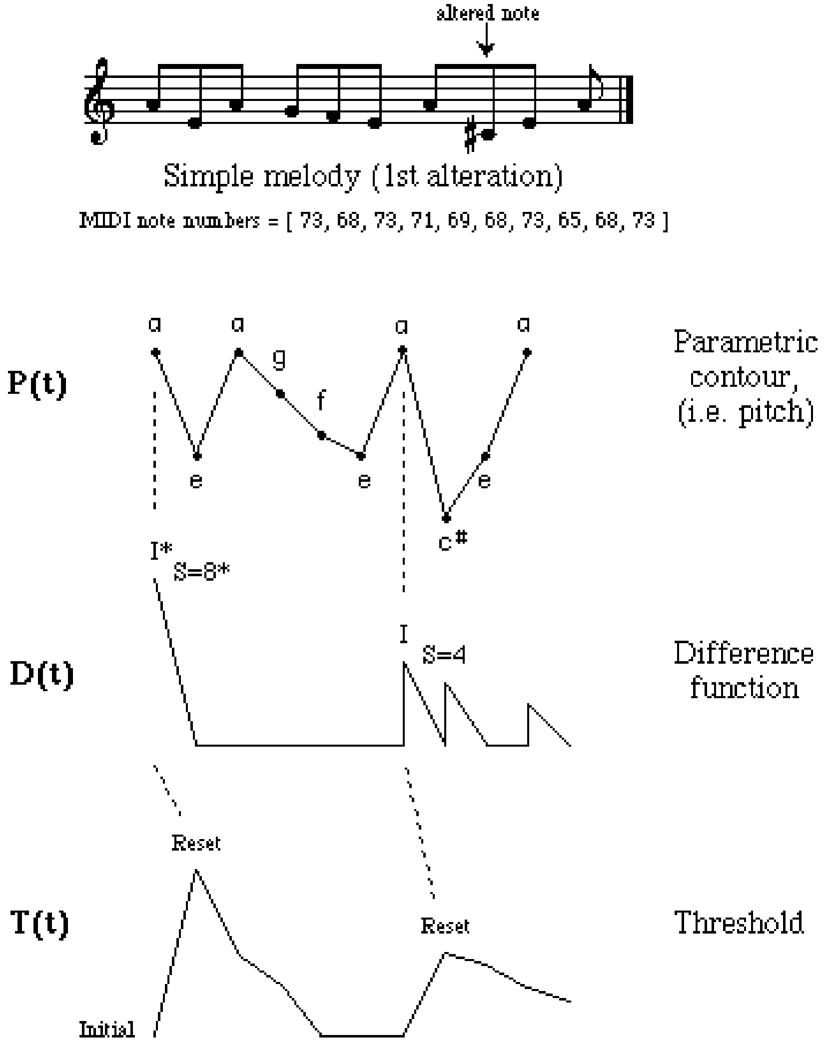

Numerical values:
D(t) = [8*,0,0,0,0,0,4,3,0,5] ;
T(t) = [0,8*,4,2.67,0,0,0,4,3.5,2.33,1.67] ;
Values are calculated as in Figure 20.

I = initiator; S = initiator strength;
*First value is arbitrary, i.e. the
first event is always an initiator.

Application of the Difference Function to First Alteration of Parametric Contour

Figure 21

Fig. 21. Application of the Difference Function to First Alteration of Parametric Contour. An example for comparison showing how lowering one note in the melodic contour affects the identification of possible temporal groupings. Note that only two initiators (I) are detected in this version.

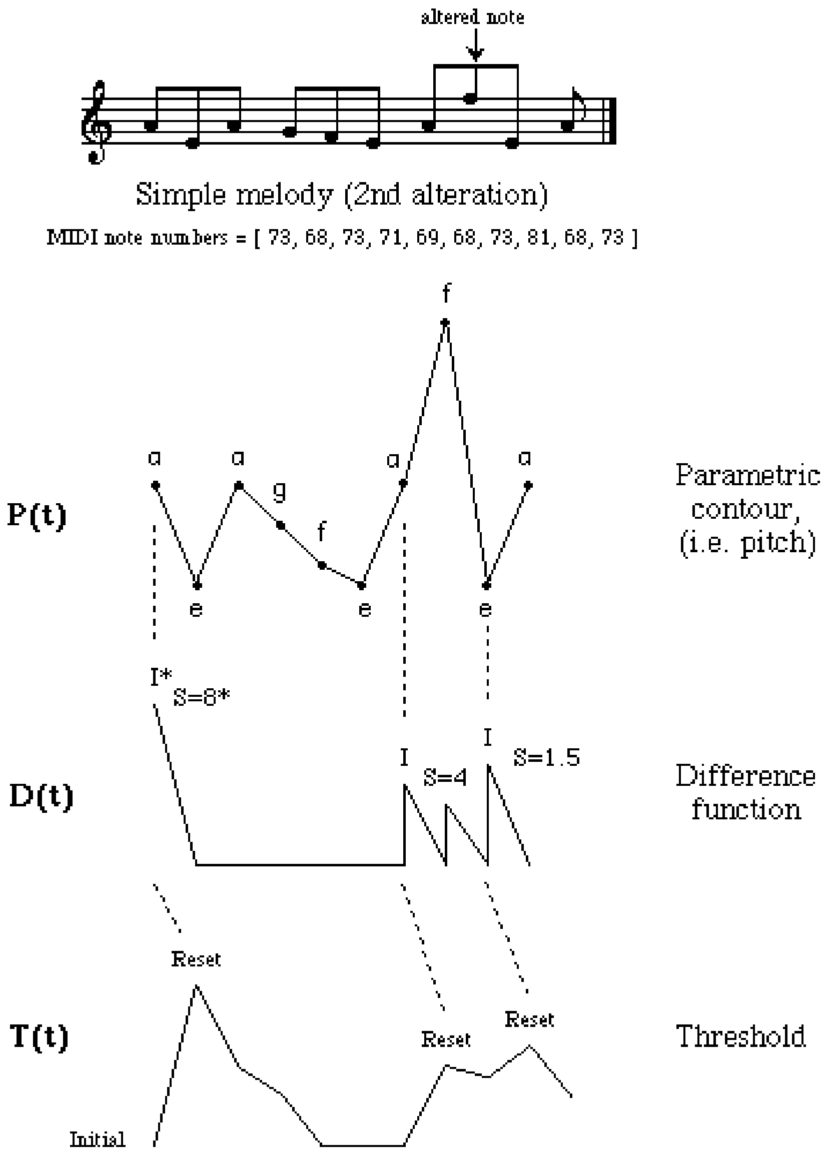

Numerical values:
D(t) = [8*,0,0,0,0,0,4,3,5,0] ;
T(t) = [0,8*,4,2.67,0,0,0,4,3.5,5,2.5] ;
Values are calculated as in Figure 20.

I = initiator; S = initiator strength;
*First value is arbitrary, i.e. the
first event is always an initiator.

Application of the Difference Function to Second Alteration of Parametric Contour

Figure 22

Fig. 22. Application of the Difference Function to Second Alteration of Parametric Contour. A second example for comparison showing the effect of raising the altered note of the melodic contour. Note that three initiators (I) are again identified, with the last one in a different position.

FURTHER REFINEMENTS—REAL-TIME, CONVERGENCE AND DIVERGENCE OF PATTERNS

Subsequent versions of *On Being Invisible* contained many refinements. First, the ability to account for the flow of real time was added. This replaced the arbitrary use of element (e.g. a 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 (see Fig. 23). Remember that an 'event', here, is defined as a unidirectional (increasing) change of rate-of-change, $D_{(t)}$, in some parameter, $P_{(t)}$, 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 measurements. 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 [139]. (See Fig. 23 for a sketch of the behavior of T with exponential decay.)

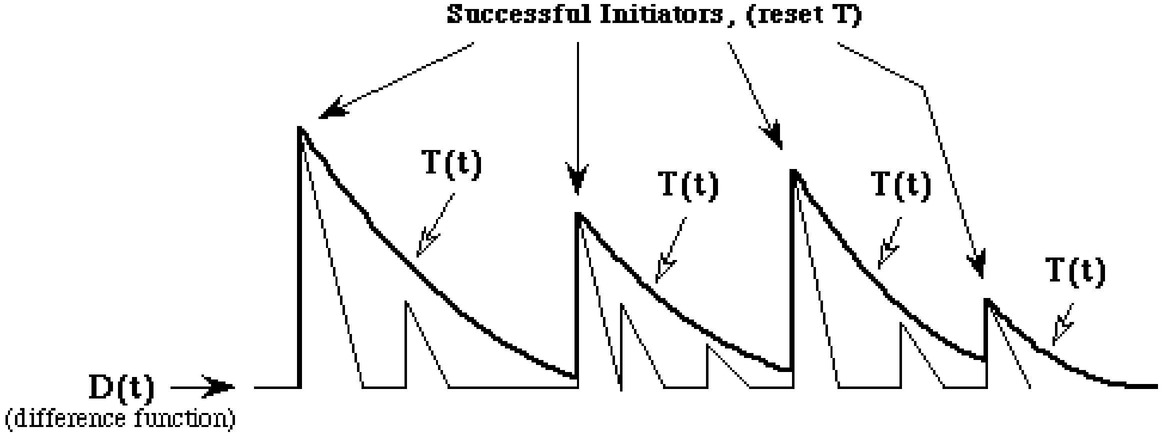

Effect of Exponential Decay of Threshold Function, T(t), on Initiators in Succeeding Time Regions

Figure 23

Fig. 23. Effect of Exponential Decay of Threshold Function, T(t), on Initiators in Succeeding Time Regions. When the threshold function, T(t), is given an exponential decay after the identification of an initiator, subsequent peaks in the difference function, D(t), must be larger or more removed in time in order to be

considered initiators. This method may be more reflective of the persistence and decay of processes evident in the nervous system than that involving a simple time average of threshold function values.

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-controlling 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, these will converge into patterns. Second, if predictions are unsuccessful, the musical structure is made more open to stochastic influences. Consequently, if successful predictions 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. Many of these are directly associated with the phenomena of attention discussed in previous sections. Repeating patterns may fail to elicit attention shifts because of the subject's boredom, volitional shifting of attention 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-controlling mechanism* used S to direct the rate of convergence or divergence of patterns. When a high-S prediction was successful, convergence to repeating patterns was more rapid than when a low-S prediction was successful. Correspondingly, a high-S, unsuccessful prediction would cause relatively rapid divergence and a low-S, unsuccessful prediction, less rapid divergence.

Convergence and divergence are achieved by adjusting variables in various stochastic canons. To create divergence with a Gaussian-distributed canon, for instance, 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 a filter window size applied to the output of some random generator. However, convergence was also dependent on processes used to build *hierarchical structures*.

HIERARCHICAL STRUCTURE BUILDING

At a certain point in an *On Being Invisible* session or performance, the performer could activate a *hierarchical structure building* part of the *structure-controlling mechanism.* Usually, the performer would use her or his 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-controlling 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; that is, 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 combinatoric possibilities. This recombination potential was simply 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—Information, Memory, Expectancy and Pattern Evolution

The *hierarchical structure builder* contained the beginnings of a simple musical inference 'engine', but with a unique quality. It had to make predictions in real time 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. An analysis of the sequence of clangs was carried out, inspired originally by concepts from information theory.

In 1958 Edgar Coons and David Kraehenbuehl published a stimulating paper in which a method for quantifying the information value of events in a sequence, with particular reference to musical structure, was presented [140]. 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 non-confirmed 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, predictions involving the occurrence of dissimilarity (maximally informed events) versus similarity (minimally informed events), along with sequences of these, were taken into account. Initially, I tried to incorporate a variant of this method into the *On Being Invisible* programs. A practical problem, however, prevented a full realization. The Coons and Kraehenbuehl analysis method 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 can soon grow out of hand. A full-scale analysis of this type requires tabulation of all the 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 what was for me an opportune time [141]. This study showed enhancements in the amplitudes of N200, 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 the hierarchical musical patterns that I 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 each time a clang is played. The form of the memory/expectancy function is adapted from Squires et al., as follows:

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

where, M_{EN} is the memory/expectancy function for event E at position N, as a function of the sequence of past events, S_i. As the algorithm runs, S_i is set equal to 0 if the event at the current index, i, is not equal to E and set to 1 if it is equal to E. An exponential decay of memory for past events results from the evaluation of α, where $0 \le \alpha \le 1$. 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 a simple example. (see Fig. 24.) Consider the following primitive sequence of three events, or musical objects, simply labeled, **A**, **B**, and **C**. I will call these clangs, following Tenney's terminology. Higher-order groupings of clangs will be called *sequences*. In the terminology of HMSL, these would be programmed as *collections*.

ABCABCABCCBAABCCBAABCABCA. . .

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 [142]. 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 [143].

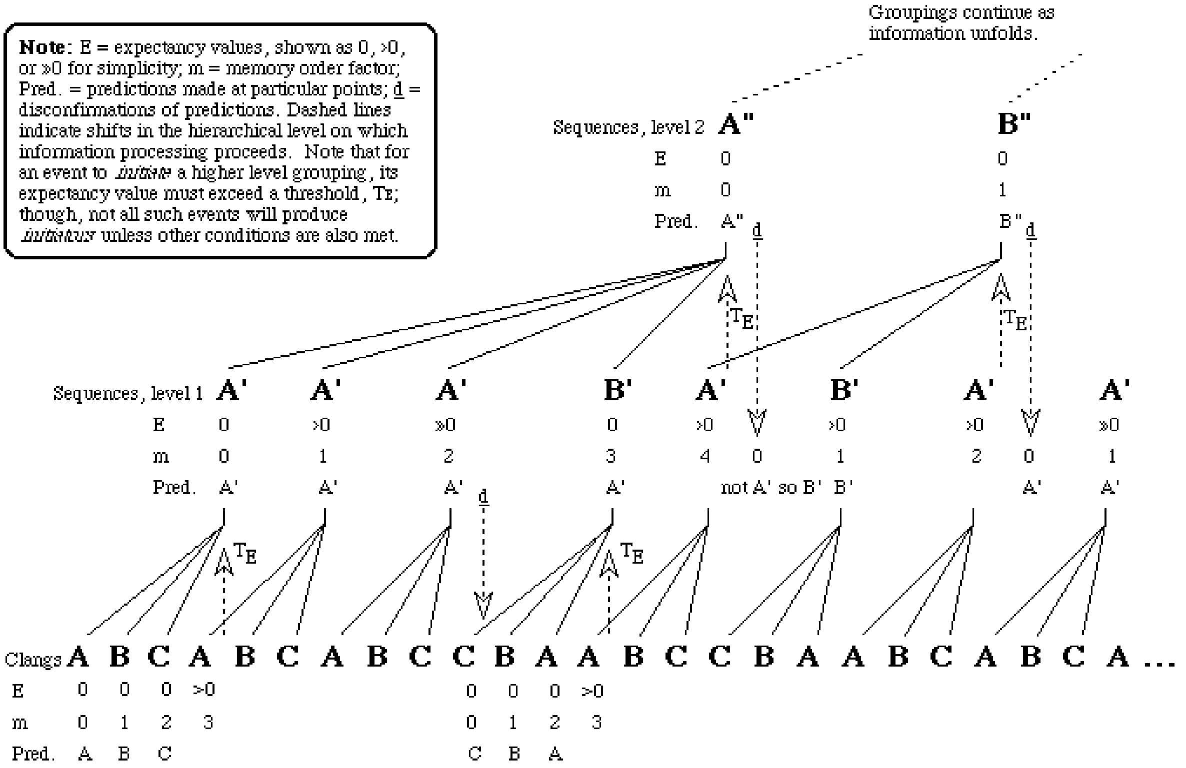

Illustration of the *Hierarchical Structure Builder* Operating on a Simple String of Events

Figure 24

Fig. 24. Illustration of the *Hierarchical Structure Builder* Operating on a Simple String of Events. The actions of the algorithm used to implement the *hierarchical structure builder* can be traced on three temporal grouping levels—temporal gestalts—in this example.

The first 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 non-zero expectancy value is obtained. It is compared with an expectancy threshold, T_E, 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 material 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 recur. 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 non-zero. 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, **A**, 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 non-zero expectancy value is obtained and, again, the preceding clangs, up to the limit of the memory-order parameter, 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 the memory-order parameter 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'** disconfirms the prediction of another **B''** but it cannot trigger a new grouping, since 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 levels of grouping, *n*th-order transition probability tables (Markov chains) can be built to reflect the likelihood that particular clangs or sequences will tend to follow one another or remain bound in high-level groupings. The probabilities in this table can be skewed, by successful detections of attention shift to increase the likelihood that particular transitions will recur. When tests for attention shifts are unsuccessful, the corresponding transition probabilities are decreased. 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 difficult, on the basis of EEG concomitants of attention shift alone, to tell 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, however, 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 does not 'know' about such things as retrogrades, inversions or transformations on parametric contours. 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), 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, however, 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.

This structure builder operates on several important principles or rules:

> *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, T_E, is considered a *potential* initiator of a new sequence grouping.
>>
>> *Rule:* A sequence grouping must contain at least two or more 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 shift 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 toward 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. However, in this case, the potential for such evolution is imbedded in the structure of a artificially intelligent musical instrument.

Global Parametric States

The analysis of parametric values via the difference function, $D_{(t)}$, continues on the second and subsequent hierarchic levels as well, but with a difference: changes in the global qualities for each parameter of a clang, what Tenney calls 'state variables', are evaluated instead of the individual element values. These are the average of the parametric values for the elements of a clang , adjusted to reflect element durations thus,

$$\sum\left(\frac{\text{parametric values} \times \text{durations}}{\sum \text{durations}}\right) \quad ,$$

following Tenney's suggestion for these calculations. 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, i.e. 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 TG analysis involves the question of *parametric weighting*. How important is a change of a given size in one parameter, such as loudness, in relation to another, such as pitch, 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 Varèse, TGs are delimited more by amplitude events than by pitch [144]. The solution, in general, is an empirical one; that is, one should adjust the weighting values until the model behaves as desired. 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 a 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, such as 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 temporal gestalt analysis in musical form perception. To date, at least some preliminary spade work has been carried out. Rich possibilities for investigation remain on the horizon. One may apply the system to the investigation of form perception in precomposed, fixed musical works simply by using only some parts of the *On Being Invisible* feedback loop, that is, the *model of musical perception* and the *input analysis system*. Instead of a spontaneously generated musical structure, a precomposed, fixed one is output. A time record 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 from a random starting point, but 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 I have described often yields inspiring surprises.

The Musical Result

It is nearly impossible in a written text to give an adequate description of the musical sound one might experience during a performance of *On Being Invisible*. When the work is realized with a powerful electronic music system, an enormous range of timbres can result. After all, it is a fundamental behavior of the algorithms implemented that a full range of possible timbres is searched, albeit by stochastic methods, when disconfirming feedback from the performer is persistent. Realizations to date have typically consisted of four to eight electronic 'parts', each having several sub-'voices', which organize themselves according to the hierarchical structure-generating scheme described above. Some qualities common to all realizations include exploration of a wide range of often-unexpected timbres, surprising pattern derivations, and a tendency for the music to converge upon patterned relationships out of random fields—as attention locks onto them—and then for the patterns to diverge again—as attention wanders—into complex variations or even disorder. When an ordered section persists long enough for hierarchical pattern structures to form, multi-layered counterpoints or rhythmic constructions can often emerge. By contrast, earlier works I produced with coherent EEG waves (delta, theta, alpha and beta) were often typified by sound textures with more static large-scale forms but richly structured inner detail.

ALGORITHMIC IMPROVISATION WITH ERPS AND OTHER INPUTS

Is it the composer's fault that man has only ten fingers?
-Charles Ives, 1920 [145]

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

Because 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 in a group context involving improvisation with an attention-dependent sonic environment can be exciting, though results often can be 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. I have focused on feedback from aspects of attention shift. Feedback 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.

Part 5—Technical Developments—Old Problems and New Possibilities

TECHNICAL PROBLEMS TO SOLVE

Technical Limitations Due to Computers

Most technical limitations, both past and present, hinge on the architecture, speed and memory capacity of computers. The types of analyses of EEG signals described herein are mostly statistical, time-dependent processes. They require time-consuming and often repetitive, computation-intensive processes like averaging, integration, filtering, template matching and prediction. Even if these could be performed instantly, constraints would still be imposed by the speed with which the incoming data arrives. A fundamental restriction on this lies in the relatively low frequencies of EEG signals. The brain apparently does not need high-speed, serial processes to achieve its profound, integrative results. Probably, massive parallel processing in the human brain compensates for the relatively slow processing speed of individual computing units (i.e. neural circuits). Electronic computers are only now beginning to show the potential of parallel architectures. The computing required by the systems I have described falls mainly into three categories: signal analysis, musical structure generation and sound synthesis. Each of these is quite complex and, ideally, should be performed by independent, parallel processors.

In the computer systems available for live performance at the time *On Being Invisible* was originally realized, memory limitations dictated a modest storage capacity for clangs and sequences. The full realization of the potential envisaged for the creation of complex musical architectures requires a dramatic expansion. The more extensive search and comparison operations implied by more memory demand increased speed as well.

The number of levels on which real-time, hierarchical analyses could be performed in 1977 was similarly limited. Fast, powerful computers may be capable of maintaining analyses on several hierarchical levels at once. Restriction on the use of algorithms requiring massive numbers of predictions, information-theoretic tallies and the like (as described in the section *Musical Inference—Information, Memory, Expectancy and Pattern Evolution* in Part 4) may be removed by faster, parallel processors as well.

Problems in EEG Analysis and Extracting ERPs

Major technical problems still lie in the techniques for extracting information from ERPs in real time for use in feedback paradigms. We need to refine our methods for obtaining ERPs from single stimulus-events. This is still confounded somewhat by the small voltage of the ERP, the presence of noise, and the difficulty of predicting the form of the ERP being measured. Donchin and McCarthy report on the use of a refined adaptive filter technique based on the Woody method (described in the section *How ERPs Are Detected and Measured* in Part 2) to extract single-trial ERPs with some success [146]. Many ERP components are slow waves. Consequently, we can use filtering to advantage. Most of the energy in P300, for instance, is below 1 Hz, while most of the energy in the ongoing EEG is in the 4- to 12-Hz range. A low-pass filter, set for 3 dB attenuation at around 5 to 6 Hz, can aid in recovering P300 information with less signal averaging than would otherwise be required. I say this with caution, however, because I am not completely confident that we know what we may be throwing out with this procedure.

The point of using signal averaging, of course, is to remove so-called, 'confounding variables'—be they unwanted electrical activity, random noise or physiological and psychological variables. These are assumed to be random, over some reasonable number of trials, and will therefore average out. It may be, however, that in certain situations these confounding variables may be the very ones that are the most interesting, particularly when explorations involving artistic production and performance are involved.

We need better ways to see the topographic distribution of EEG events, rather than being limited by visual interpretation of the conventional, multi-channel EEG tracing. The electromagnetic field pulsations of the brain are multi-dimensional. The concept of brain imaging with practical, affordable systems needs to become an area for serious technical development.

The Question of Information Bandwidth Limitation on the EEG

A criticism regarding the use of EEG information as a control source for electronic music production has focused on the limited information bandwidth of the EEG signal. It is true that the normal frequency bandwidth—around 25 Hz—cannot possibly carry enough information to update the fast parameters of timbre synthesis in reasonable, musical time. There is a great deal of delay in the transmission of electrotonic effects through tissue masses. Possibly, biomagnetic effects may require less time delay to manifest results outside the skull. (See the section *The Potential of Superconducting Sensors and Biomagnetism*, following). But, even there, what is recorded outside the mass of billions of neurons is a highly integrated, averaged summation of activities produced by many singular, low-level, neural circuit events. All of these processes take time. Furthermore, changes in EEG waveforms that represent intelligent activity do not occur fast enough to carry nearly the amount of information imparted to an instrument by the average performing musician. But then, most present implementations of MIDI (Musical Instrument Digital Interface) are not sufficient for this either! Nevertheless, clever methods for analysis and information extraction can result in useful controls. Furthermore, there is no reason to use EEG information to replace or even mimic what performing fingers or embouchures can do. After all, that is the purpose of the motor cortex and the physiognomy it controls.

The first and most obvious method of constructing a musical biofeedback loop involves the direct mapping of a measured phenomenon onto variations in a sonic pattern in a direct, one-to-one, manner. This usually involves some primary acoustic parameter such as pitch, loudness or timbre [147]. However, the best use of brainwaves may not lie in trying to couple them to low-level note production. Why not continue to use the proprioceptive, neuromuscular system—i.e.

brain, neuron and limb effectors—to do that? Some of the most interesting results we can obtain from brainwaves involve signals related to such high-level functions as information-processing modes; global operating states; the representation and manipulation of musical, cognitive processes; and states of consciousness. Therefore, *why not apply these to high levels of musical form*?

So we upgrade the level of representation for information we expect to extract from the signals we record and also upgrade the level of musical form production to which we apply the results. A solution to the information bandwidth problem can be found that mimics an approach recently applied at Sandia National Laboratories (Albuquerque, New Mexico) to solve the problem of how to make efficient use of the computing resources that reside inside large parallel-processor computers: *scale up the problem* to fit the resources available [148]! The information contained in ERPs and much of the ongoing EEG manifests from global, high-level information processing by massive neural resources. Therefore, in musical work with this data, it makes the most sense to apply the information obtained to relatively high-level aspects of musical form generation.

The human brain achieves its speed by means of *massive parallelism*, rather than through rapid processing of signals by single processing units. The speed I refer to here is the speed with which physical, sensory signals are reduced (fragmentation/reduction) and recombined (synthesis) into the mental representation of an event (idiolog) on a relatively high level of image formation. Consequently, the bandwidth achievable by direct electronic interface is low if one measures the bit rates of atomistic, low-level representations, but is quite high if one takes into account the significance of the signals that can be acquired that relate to high-level events, albeit relatively slowly.

Thus, this bandwidth question becomes something of a red herring if one thinks carefully about what can be extracted from the signals. The relevant question requiring creative solutions is simply, *What is the significance and best use of information that requires a relatively long time to derive?*

POTENTIAL OF NEW MEASUREMENT TECHNOLOGY

No-Contact Electrodes—A Dream

Many physical limits are imposed on the situations that can be designed, particularly in the arts, due to the cumbersome nature of EEG and other physiological sensing apparatuses. EEG electrodes are sensitive, require careful attachment with messy electrolyte pastes, are intolerant of physical movement and involve delicate wires. The monitoring equipment must be in close proximity to the subject so that electrode wires do not become noise antennas due to excessive length. These are just some of the physical encumbrances of most sensing situations. Consequently, the possibility that someday a sensing technology will be developed that is free from the requirement of direct physical contact has inspired dreams of exotic applications and science-fiction images alike. The first glimmer of such a technology has appeared in the form of methods for sensing magnetic fields that emanate from the nervous system.

The Potential of Superconducting Sensors and Biomagnetism

The body is a highly electrically active mass of tissue. Magnetic fields occur, of course, in the vicinity of electrical charges that are in motion. Techniques for measuring these fields associated with biological organisms have been developed over the past decade or more—the study of which is termed biomagnetism [149]. Magnetic fields associated with the discharge of action potentials from single neurons have even been measured [150]. The brain's electrochemical processes induce the transport of ions, setting up electrical current dipoles. These current dipoles, though very small, generate magnetic fields. Such tiny fields can be detected outside the head. Furthermore, the detectors used do not require physical contact with the scalp. The measurement of the fluctuation and topographic distribution of these magnetic fields is termed the magnetoencephalogram (MEG) [151, 152].

The strengths of extracranial magnetic fields comprising the MEG are but a tiny fraction of that constituting the earth's magnetic field. They measure typically less than 1 picotesla, or 0.00001 Gauss (1 tesla = 10^4 Gauss). By comparison, the magnetic field of the earth averages around 0.5 Gauss—about 50,000 times larger. The Earth's magnetic field exhibits various patterns and modes of behavior. It is subject to small, continuous pulsations in the range of 0.1 to 100 Hz, has significant components in the range of 8 to 16 Hz, and exhibits a peak around 10 Hz. Curiously, this peak corresponds to the dominant alpha-wave frequency. The possibility of some entrainment function has been suggested but not thoroughly explored [153]. The idea that some resonant coupling may be involved is hard to ignore, however.

The MEG is a manifestation of currents arising from ion movements produced by changes in the electrical potential of neuron cell membranes These ion movements begin at the dendrite end and proceed throughout the neuron's cell body, creating a current dipole. The current inside the cell is called the *source current*, and the current outside the cell is called the *volume current*. The orientation of individual dipoles with respect to the surface of the scalp is critical in order for them to be detected. Fluctuations in that part of the magnetic field that is perpendicular to the head can induce current flow in the loops of a coil wound in a plane at right angles to the magnetic field lines. Fields from current dipoles oriented tangentially to the scalp can be measured. Those oriented radially do not contribute to the component of the extracranial magnetic field that is perpendicular to the scalp; therefore these cannot be detected (see Fig. 25).

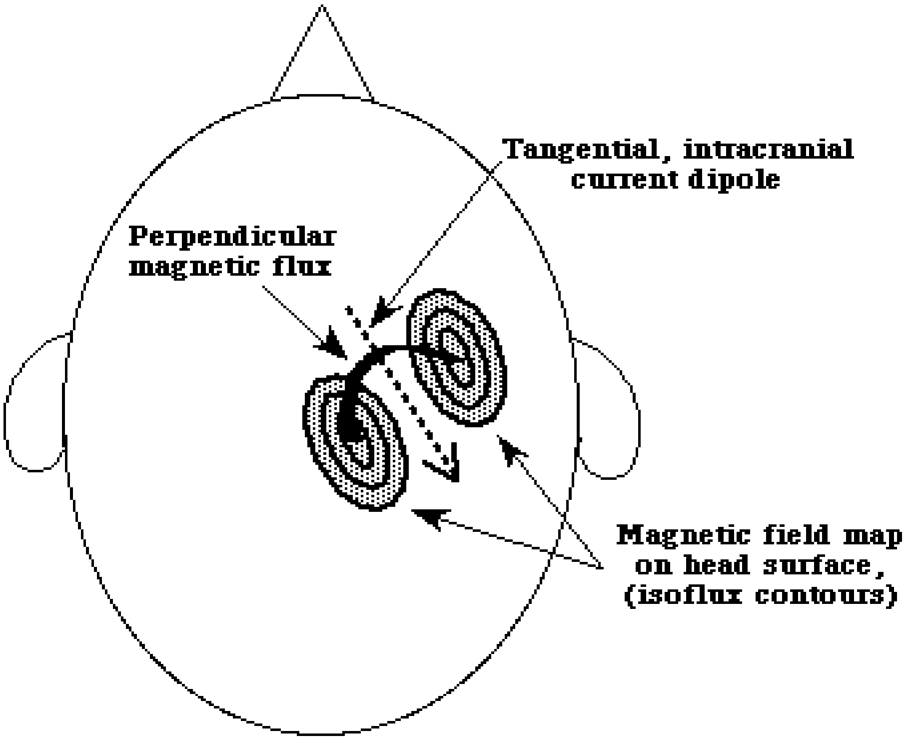

Orientation of an Idealized Extracranial Magnetic Field Contributing to the MEG

Figure 25

Fig. 25. Orientation of an Idealized Extracranial Magnetic Field Contributing to the MEG. Extracranial magnetic fields are produced by current dipoles within the brain, which run tangentially to the scalp. Measurements of the magnetoencephalogram (MEG) can be used to create a magnetic field map of brain activity on the head surface.

The current flow induced in a sensing coil by the tiny magnetic field fluctuations associated with brain currents will be correspondingly tiny. Superconducting materials must be employed to detect them. The instrument used to measure these fields is known as a SQUID (Superconducting QUantum Interference Device) [154]. A SQUID *magnetometer* has three components: (1) detection and input coils, (2) SQUID and electronics and (3) a cryogenic vessel (Dewar) containing liquid helium to maintain the cool temperatures required by superconducting materials. (See Fig. 26 for a generalized diagram of SQUID components.)

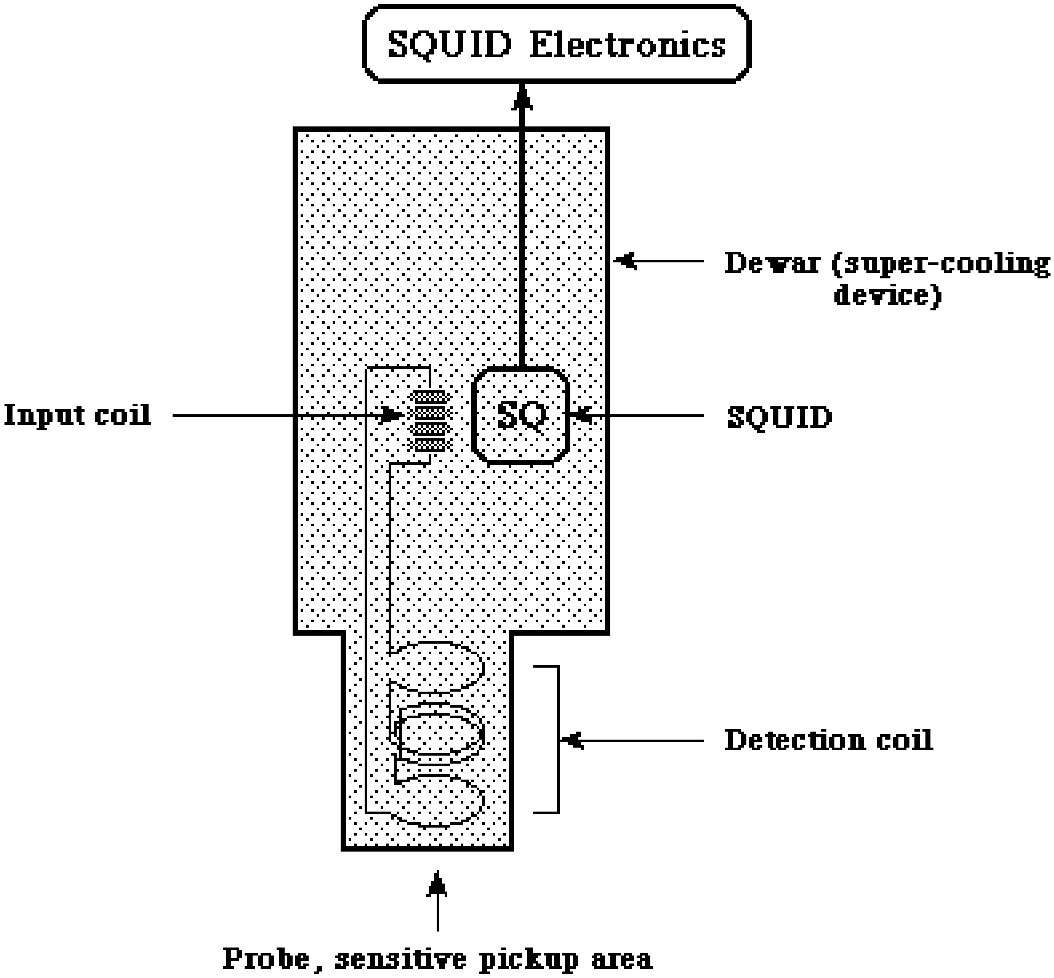

Generalized Diagram of SQUID Components

Figure 26

Fig. 26. Generalized Diagram of SQUID Components. Major components of the SQUID instrumentation used to detect tiny magnetic fields that make up the magnetoencephalogram (MEG). Detection coils comprising a magnetometer, the SQUID input coil, and the SQUID itself must be kept super-cooled in order to achieve super-conductance.

The detection coil is connected to the SQUID input coil, forming a superconducting circuit that acts as a flux transformer. When a magnetic field is present at the detection coil, current flows in the superconducting circuit, the value of which is proportional to the instantaneous value of the magnetic flux present. The current in the input coil impresses a magnetic field on the SQUID. A preamplifier measures the response of the SQUID. This comprises a highly sensitive current-to-voltage converter (greater than 107 V/A), the output of which is linearly related to the instantaneous value of the magnetic flux passing through the detection coil. Modern SQUIDs use direct current (DC) biasing techniques.

Considerable background magnetic field noise exists in typical measurement environments. This is due to fluctuations in the earth's magnetic field and extraneous fields from electric motors, elevators, passing subway trains, etc. The magnetic fields associated with intracranial ion currents are much weaker than these unwanted fields, (typically, several thousand times smaller). To reduce background noise, a gradiometer is constructed consisting of two coils wound in opposite directions. One, the pickup coil, is placed close to the head. The other is placed parallel to the first, some distance away. This makes for a kind of *differential detector*. Relatively uniform background fields will induce equal and opposite signals in the two coils, which cancel each other out. Small fields nearer to the pickup coil induce unequal currents, which do not cancel. This is because the strength of a magnetic field decays with distance from the source according to the well-known, inverse- square law. An induced current, then, flows to the SQUID sensor and produces a voltage proportional to the rate of decay of the magnetic field along the axis of the gradiometer. Second- and third-order gradiometers can be constructed for greater background rejection with a consequent trade-off in reduced total-system sensitivity. Flux transporters can be wound so as to respond to differences of rate-of-change (i.e. second derivative) in the spatial gradient of the magnetic field, making them relatively insensitive to uniform background fields and to background fields with uniform spatial gradients. The choice of detection coil diameter and the separation distance between pickup and compensating coils (known as the 'baseline') determine the spatial resolution achievable with a given device. Typical coil diameters allow spatial resolutions of 2 to 3 cm, with gradiometer baselines between adjacent coils of about 3 to 4 cm. Signal-to-noise ratios (S/N) of 3 to 15 can be achieved in this way. However, background fields in unshielded settings can still cause problems. Consequently, signal averaging—employing 10 to 20 event trials—is still used to improve S/N. As with surface-recorded ERPs, this results in a loss of sensitivity to variations among individual events. Adaptive filters with pattern template matching and various prediction techniques can be applied in an effort to get meaningful single-trial data. The cryogenic vessel (Dewar), filled with liquid nitrogen, is quite cumbersome and does impose physical limits on the possible orientations of the device. Several orthogonally oriented magnetometers, placed within the same Dewar at some distance from the pickup coil, have sometimes been used to try to detect background fields and subtract them from the signal of interest. This can successfully reduce noise components below 10 Hz, but is not so effective for reducing noise above 10 Hz. Instruments with up to seven channels are available.

MEG has been extensively applied in epilepsy research [155, 156], in the study of visually evoked fields (VEF) [157], and in somatically evoked fields (SEF) magnetic fields associated with the flow of currents in the brain responding to electrical stimulation of the fingers [158].

There are some advantages to this approach. The detection device requires close proximity, but not physical contact, with the head. Typically, the pickup coil is placed about 7 mm from and tangential to the scalp. Signals arising from much more highly localized current source-sinks in the brain than those detected with traditional surface electrodes are routinely extracted by means of MEG. The detected magnetic field emerges from one small region of the scalp and enters another (see Fig. 25). An example of this is reported in a study involving detection of fields that reverse their direction in synchrony with the frequency of a repetitive somatic stimulus [159]. These show a qualitative similarity to what would be associated with a current dipole source located just under the scalp.

The SQUID as EEG Sensor

The SQUID itself may be refined by further developments in the field of high-temperature superconductors [160]. SQUIDs are made possible by the effect of magnetic fields on the operation of devices known as 'Josephson junctions'. Josephson junctions are based on the phenomenon of *quantum tunneling* (by Cooper pairs of bound electrons) across a thin oxide barrier separating two superconducting materials. This tunneling through the barrier can result in current flow without any voltage. When the junction is placed in a magnetic field, however, the device can be switched from zero to a finite voltage. When two such devices are connected in a ring, current flow around the ring becomes sensitive to the magnetic field flux passing through the ring (see Fig. 27). In this way, extremely sensitive detectors of current, voltage and magnetic fields—like those associated with the brain and the heart—can be fabricated. An added bonus results from the fact that very little background noise exists in superconducting materials due to the free flow of electrons.

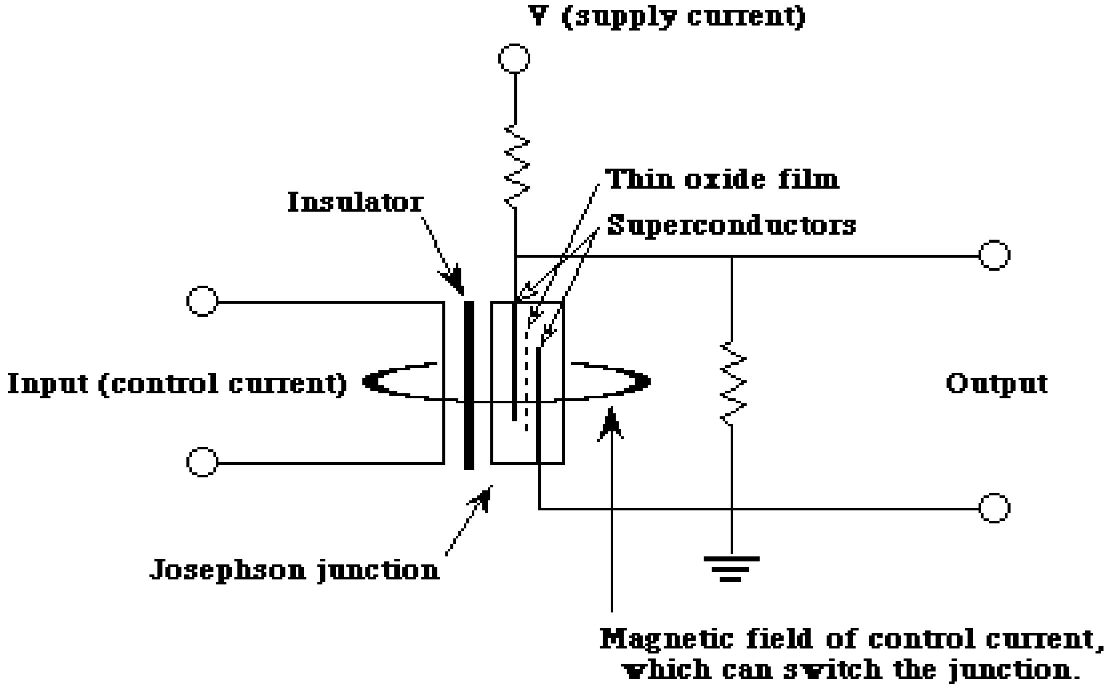

Josephson Junction Magnetic Field Sensitive Switch

Figure 27

Fig. 27. Josephson Junction Magnetic Field Sensitive Switch. A small magnetic field generated by an input current can switch the Josephson Junction—created by sandwiching a thin-film in between two superconductors—from zero to a finite voltage. If two of these are then connected in a ring, the current flowing around the ring becomes sensitive to the magnetic flux passing through the ring. Currents induced in pickup coils by magnetic fields emanating from the brain can be detected with the aid of these devices.

Present applications of SQUIDs in biomagnetism are limited by the relative immobility of devices that require super-cooling to operate at superconducting temperatures. Recently, great strides have been made in the development of high-temperature superconductors; however, additional materials that can operate at yet higher temperatures are needed in order to make practical devices. Such materials will certainly have an impact on the fabrication of junction devices and SQUID magnetometers [161]. One can imagine a myriad of applications in bioelectromagnetism if high-temperature devices become available. Many companies and agencies are currently working to develop such practical devices [162]. Significant hurdles in the fabrication of the thin-film ceramic materials required for junction devices still need to be surmounted. However, progress in fabricating both thin-films and fibers [163, 164], along with processing bulk materials and screen printing on substrates needed to manufacture useful devices [165], has recently been reported at places such as Sandia National Laboratories, Stanford University, AT&T Bell Laboratories and IBM.

NASA researchers are fabricating superconducting-insulating-superconducting (SIS) junctions from high-temperature, superconducting thin-films, which can operate without cooling in the environments of outer space [166]. These will be used to build highly sensitive magnetometers, slated for use in future, deep-space gravity surveys. Superconducting thin-films may not need cryogenic systems when deployed in the frigid environments of outer space. Natural superconducting matter states may even exist in such environments.

If high-temperature superconductors can eventually be fabricated, one can envision a kind of superconducting electrode or probe that could be positioned near the brain to record electromagnetic field gradients and flux changes associated with acts of perception and cognition. Such devices may require less cumbersome cooling apparatus, though a gradiometer-type construction would be required to eliminate uniform background fields. As presently employed, Josephson junctions are two-terminal devices that cannot impart any gain to currents flowing through them. If three-terminal devices could be constructed, these superconducting, bioelectromagnetic field probes could possibly provide the necessary amplification and signal conditioning right at the probe or electrode site. Sources of noise and signal degradation—due to background fields, large volume currents flowing in body tissue, and skin conductivity—could be eliminated before the signal gets to the analysis and display equipment.

Multiple Measurement Channels and Brain Imaging

It is abundantly clear from research data gathered thus far that information processing activity in the brain is disturbed across populations of neurons, neural groups and, depending on what level of information abstraction or synthesis is involved, large regions of the brain itself. The closest analogy in computational architectures is that of *massively parallel, distributed processing*. Consequently, we need to find clever ways of displaying a topographic mapping of large amounts of condensed data from various regions of the brain in order to discern meaningful patterns of whole-brain activity. We are confronted with two practical problems in striving toward this. The first is how to obtain simultaneous recordings with reasonable spatial resolution from many separate locations at affordable cost. Second, we must discover creative ways to display the data so that the observer (human or computer) may discern patterns within an integrated, gestalt-like picture of global activity.

Jesse Salb of the National Institute of Mental Health developed a brain-imaging device that gives a continuously updated graphic image of the distribution of potentials over the surface of the scalp. This imaging technique presents significant advantages over conventional EEG tracing techniques. A more complete picture of the distribution of activity can be absorbed at a glance. The system requires 28 electrodes and 28 channels of high-gain, differential pre-amplifiers; anti-aliasing and muscle artifact-removing filters; and analog-to-digital conversion (ADC). A microprocessor then updates a display 60 times per second with colors representing different voltage levels distributed on an image of the brain's surface. Fortunately, the sampling rates and resolution required of the ADC system are minimal, given the low frequency and limited bandwidth of EEG signals. Still, this system cost just under $20,000 when constructed in 1983 [167].

More recently, John and colleagues have developed a method for gathering and displaying computerized differential classifications of normative and dysfunctional electrical brain features in people 6 to 90 years old [168]. Part of this method, called 'neurometrics', involves a colorized, topographic display of relative signal power in delta-, theta-, alpha- and beta-bands from 19 recording sites. Actually, this display results from a battery of statistical techniques, including absolute and relative power across frequency bands, coherence, symmetry and various topographic comparisons across the whole brain. These data are then subjected to further statistical analyses, producing comparisons among normal and cognitively impaired populations from different age groups, for the purpose of neruometric diagnosis. The signal analysis and display techniques could well be applied in feedback paradigms for research or artistic production.

Jones-Leonard et. al. describe a Brainwave Neuroscience Workstation, developed by Brainwave Systems Corporation, which integrates and displays data from intracranial recordings of clusters of neuron groups [169]. The system makes use of a *stereotode*—a dual electrode probe that records neuron firings differentially, referenced to a separate electrode. Because extracellular voltage strength is inversely related to distance from the current source generator, recordings from the two tips of the stereotode can be used to map the spatial distribution of neuron firings. Such a technique is related to SQUID gradiometer construction and can be used to map the spatial gradients and differential magnetic field flux associated with the activities of closely spaced but separate neural groups.

A Brain Imager device, from Neuroscience Corporation and Darex, Inc. (San Diego, California), is currently being used for audience demonstrations at the Lawrence Hall of Science (Berkeley, CA). It contains a head cap with a full array of electrodes. Color images of side and top views of the head are displayed, depicting EEG activity in various frequency bands along with views derived from other statistical analyses. Systems for creating visual and auditory displays of a wide variety of bioelectrical signals, such as EKG and others, have also been devised, [170, 171].

Still other imaging technologies—such as PET (positron emission tomography), which maps the flow of radioactively tagged blood in the brain as different regions become activated, MRI (magnetic resonance imaging), which records radio waves resulting from the realigning of molecules in the body when it is placed in a strong magnetic field, and CT (computerized tomography), which uses X-rays—may await technical advances that will reduce their cost and stimulate creative applications.

Communication with ERPs

An experimental communication system that attempts to use the visually evoked response as a communication channel has been devised [172]. Essentially, the system identifies, by analysis of EEG data, which of several visual targets a subject is looking at. The method takes advantage of the fact that the visually evoked response is largest when gaze is fixated within the central one degree of the visual field. By means of a novel method, the fixated target can be identified with 90% reliability. The method involves testing ERPs to a number of stimuli presented in different parts of the visual field simultaneously and extracting the ERP to the target stimulus over several presentations. The response time of the subject, under ideal conditions, is about 1.5 sec. Variations due to subject differences and signal-to-noise characteristics reduce this performance in real-world situations. Nevertheless, this is an exciting result. Perhaps new developments will permit exploration of ideas like this in auditory or even musical domains.

EEG ANALYSIS EXPERT SYSTEM

Various experimental systems have been constructed in attempts to develop computers that read commands extracted from a subject's EEG or ERP components. One, at the Stanford Research Institute (Menlo Park, CA), developed by L. Pinneo, involved subjects directing the movements of a visual target on a screen by issuing mental instructions such as 'up', 'down', 'right' and 'left'. EEG patterns, hypothesized to be associated with these commands, were extracted. This was reported to have been accomplished with very limited accuracy by brute-force pattern search and matching in the computer [173].

In this document I have surveyed and to some extent cataloged many detailed items about the significance and meaning of various EEG components to some aspects of perception, cognition and neural information processing. A natural next step is to imagine that this information could be coded in a knowledge base, forming the core of an *EEG analysis expert system*. If such a system were comprehensive enough and could operate fast enough, it could be imbedded inside a feedback loop. The expert system could provide data to a musical inference engine that contained a knowledge base of musical forms and would make predictions about musical perception and generate musical results on the basis of feedback from a performer and its own self-organization directives.

The types of rules for this expert system would include those for decomposition of the EEG (along the lines proposed in Part 2, *Some Bioelectromagnetic Phenomena of Significance to*

Paradigms of Feedback-Based Self-Organization), the detection of EEG states and state transitions, and the detailed dissection of the principle components of the ERP. The system would also be required to give descriptions of various time-variant statistical features and trends for these components. A learning mechanism to continuously update and refine the rules would be necessary.

MUSICAL INFERENCE AND ERPS

A musical inference engine suitable for inclusion in this expanded system would require rules in two categories: musical rules and rules for a model of perception. The musical rules would include provision for a description of compositional intent, aspects of what constitutes a multi-leveled, 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, for detection of the boundaries separating temporal gestalts (TG segmentation), and for perceiving holarchic form.

Such an integrated, feedback-based, self-organizing system—including the expert signal analyst (herein focused on the EEG), musical inference engine, synthesis mechanism, and intelligent performance input structure—could become a powerful tool for explorations in composition, performance and perception. Figure 28 shows a potential organizing scheme for this. Furthermore, this goal is imminently achievable with existing and affordable technology. All that is required is the vision, support and time to realize it.

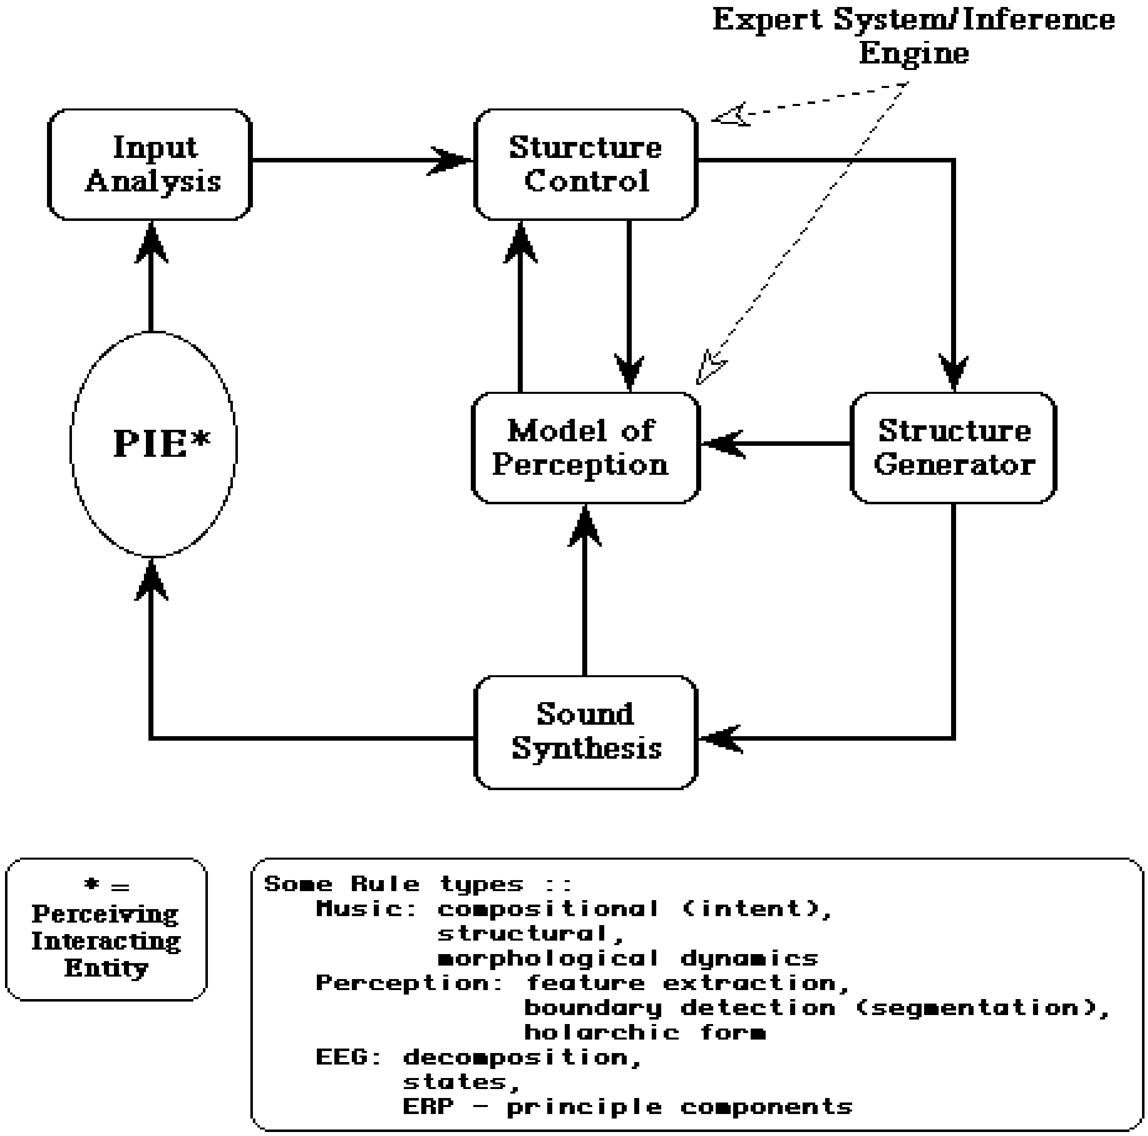

Some Rule types ::
  Music: compositional (intent),
    structural,
    morphological dynamics
  Perception: feature extraction,
    boundary detection (segmentation),
    holarchic form
  EEG: decomposition,
    states,
    ERP - principle components

Potential Organization of Intelligent Input Analyst and Musical Expert System/Inference Engine for Live Performance

Figure 28

Fig. 28. Potential Organization of Intelligent Input Analyst and Musical Expert System/Inference Engine for Live Performance. Scheme for creating an interactive music system capable of generating musical form, analyzing its own output according to a model of perception, analyzing input signals according to rules for their decomposition, and directing the evolution of a musical structure. The perceiving, interacting entity (PIE) becomes a musician acting within a new paradigm of performance practice.

Part 6—*On Being Invisible II*—A New Work: The Hardware and Software to Realize It

With all these developments in mind, I returned to this area of work in 1994 to realize a new version of *On Being Invisible*. Many of my ideas for this composition stem from research dating back 20 to 30 years. My stimulus for becoming involved in this work again came from the arrival of affordable technology necessary to realize the concepts and a renewed interest in using adaptive, intelligent instruments as an aid in exploring the musical mind. Furthermore, advances in fields such as dynamics, self-organizing systems, brain science and even, cosmology have led to new conceptual tools with which one can characterize and rethink these ideas.

MUSICAL CONCEPTION

The musical conception of *On Being Invisible* is multi-layered, with the primary feedback system described herein at the heart of it. The self-organizing dynamics of the system's behavior have also sparked ideas for compositions and schemes for organizing groups of interacting musicians skilled at improvisation. The components of the feedback system have been generalized to accommodate a broad range of signal inputs. This promotes the ability to interface the system with biosignal performers, musicians and other artists; enables it to accept acoustic inputs; and facilitates capturing gestures from a variety of sensors. Learning algorithms that guide the expert signal analysis and musical inference systems are being developed.

Finally, the major components of the feedback system have become anthropomorphized and have taken on the aspects of characters in a mythological scenario for evolution and social organization. From this, I developed a script, the intent of which is to place the work in the context of a full-scale theatrical performance. In a sense, this comprises yet another idea of opera. This self-organizing opera, first realized in 1994-1995, is called *On Being Invisible II (Hypatia Speaks to Jefferson in a Dream)*. A copy of it's program notes are presented here in an appendix.

THE EEG, MIDI AND DSP

The synthesis equipment described in some of the examples from the 1970's discussed 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, ability to store many patches, programs, waveforms, etc. In addition, the flexibility with which one could implement non-standard interfaces between sensors of unusual input signals and these older instruments, along with their complete reconfigurability through patching, is yet to be matched by the contemporary commercial instrument fare. 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 resultant data will have little relevance to actual musical perception and cognition. Serious hurdles yet to be overcome include transmission delays imposed by MIDI and the inability in many MIDI synthesizers to ascertain when certain synthesis processes occur.

The idea of an EEG-to-MIDI interface is a titillating one, to be sure. This interface was achieved in our laboratory at the Mills College Center for Contemporary Music (CCM) and also in my private studio during the 1980's and at the California Institute of the Arts Center for Experiments in Art, Information and Technology (CEAIT) in the 1990's; and others have achieved it, as well [174]. 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 from EEG signals truly meaningful data that bears a direct relationship to the production of musical sound. Although brainwave control of MIDI devices can be fun, my experience leads me to issue a strong cautionary message to those who wish to use this method and to obtain results with the precision required to produce data upon which conclusions about musical information processing can be based.

Real-time DSP (Digital Signal Processing) algorithms supporting nearly instantaneous updating of synthesis parameters offer strong potential. These can be implemented on parallel co-processors, high-level workstations and integrated media machines. However, it is still very difficult to implement all the major computation tasks required inside one typical, commercially accessible personal computer box and still expect to keep the tight timing constraints under control. My current projects at CEAIT utilize peripheral data acquisition systems capable of facilitating some of the EEG signal preconditioning and analysis algorithms along with peripheral DSP hardware for direct digital sound synthesis. This helps 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. It is hoped that it will be possible to reproduce these devices easily and inexpensively. As power and processing speed in personal computers increases, these tasks may become easier. Consequently, practice and performance with groups of individuals should be readily facilitated.

HARDWARE/SOFTWARE INTEGRATION

My own experiments are presently being conducted with a variety of hardware and software components, including Silicon Graphics Inc. (SGI) and Apple Macintosh computing platforms, DSP co-processor cards, data acquisition systems from GW instruments, computer controlled video laser disc players for visual displays and various MIDI devices. In addition, sometimes there is simply no substitute for a synthesis module capable of instantaneous response upon receipt of a computer-generated voltage step or trigger. The computers must be able to implement each of the three major computing tasks required: (1) signal analysis, (2) compositional form generation and analysis according to the model of musical perception and (3) sound synthesis. To date, however, it has not been possible to imbed all three of these demanding tasks inside one computer. Of these, signal analysis is the least demanding, because the bandwidth of the input signals involved is relatively low. Consequently, sample rates can be slow and resolution low. Programs for composition form generation and musical perception analysis have been developed using the HMSL software environment originally developed at the Center for Contemporary Music at Mills College [175] and now continued elsewhere. This software is very extensive, often consuming much of the power available in one computer. Obviously, sound synthesis could be accomplished with MIDI-based instruments and some work along these lines has taken place. However, it is very difficult to obtain enough low-level information about continuous changes in acoustic parameters of sounds in a MIDI environment to produce reliable results with the model of musical perception. The difference detector described in the section *First Few Versions* in Part 4 must operate on several acoustic parameters at once. The time-based evolution of these parameters in a MIDI instrument generally cannot be accessed by a remote computer quickly enough, if at all. I have tried two alternative approaches. One involves using widely available computer synthesis programs such as CSOUND, CMUSIC or CMIX to generate sound files in which important, predicted structural landmarks have been flagged. Later on, these can be played back and EEG parameters tested at the appropriate points. However, in order to implement dynamic evolution based on feedback, there must be many short sound files that can be selected for playback on the basis of their effects on selective attention. This produces only a simulation of continuous evolution. Nevertheless, high-quality sound can be generated this way. A second approach involves imbedding software for sound synthesis inside the compositional form-generating environment and down-loading signal-generating code to a parallel digital signal processor. This can enable real-time synthesis, which can be made to evolve continuously on the basis of feedback. A series of patchable unit generators for sound synthesis was written for the HMSL software environment by Phil Burk with a graphic user interface by Robert Marsanyi, formerly of the CCM.

My current project plans involve integrating all these processes in a multi-media environment with analysis, synthesis and multi-media resources distributed over an interactive communications network. SGI and other computing resources at CEAIT are the target platforms at the present time. In this way, the many stimulating aesthetic possibilities that emerge when individuals are able use all the techniques I have described to interact with networked systems—mirroring the concept of interacting neuronal network assemblies in the brain—may be explored.

Part 7—The Horizon from Here—Future Extensions

Many of the extensions visible on the horizon at this time involve advances in our conceptual power to characterize appropriate models of brain functions and behavior. We may apply pattern recognition techniques to the classification of EEG response waveforms. There are potential applications here for new developments in neural circuit technology and radically reconceived computer architectures. However, further progress in understanding the relationship between atomistic neural events and the global properties of cognition may ultimately have far greater significance.

John argues against the current vogue of 'connectionist' formulations of cognition on the basis of the fact that neurons exhibit considerable 'plasticity' in their behaviors during learning [176]. The connectionist view would assert that the characteristics of the 'neural patch' hold the key to behavioral description rather than that such a key lies in the behavior of the individual units that are patched together. The connectionist view may be valid in part, but overall it is too simplistic a view. In favor of the connectionist view, I would point out that a particular patch connection can have weighted probabilities associated with it, constituting a degree of behavioral plasticity in the circuit itself. HMSL contains facilities for linking musical objects together in which tendency values are associated with each link. These, along with weights assigned to the musical objects themselves, may be used in the calculation of musical behaviors, such as determining the execution order of musical items according to a variety of rules, inputs and behavioral descriptions. John posits a 'statistical view' of neuronal representation in memory [177, 178]. Statistics is a language of description, used here to characterize the emergence of global properties—in this case aspects of generalization, learning and memory. John reports on two types of neurons, identified from single-unit recordings. The first are called 'stimulus-bound' neurons, whose responses to a stimulus are invariant with respect to different perceptions that may be ascribed to that stimulus by the organism. The second type are termed 'gnostic units', the responses of which are related to perception rather than determined by sensation. These units display great variability of response for individual stimuli but show characteristic average behaviors for given perceptual modes. From this experience with recording neural units, John states an 'ergodic hypothesis': *The response to a stimulus averaged over a set of individual neural units corresponds to the response of a single unit averaged over several presentations of the stimulus.* The responses of single units are more qualitatively heterogeneous with respect to stimuli proffering different consequences for the organism than with respect to electrode position in different brain regions. With respect to the latter, the responses are relatively homogeneous, qualitatively. All of this is used to support the *statistical view*, rather than the *switchboard view* of cognition.

Further complicating the picture, the notion that the EEG arises from masses of summated action potentials is far too simplistic a picture. In fact, this idea has largely been abandoned by neurologists. The electrochemical action of the brain's tissue is far more complex. *Graded* *potentials* resulting from ion transport, electrotonic coupling, field effects and dendrodendritic synapses are a few more processes that contribute to higher cortical electromagnetic activity [179].

I am attracted to the view that the form of the laws governing brain behavior is enfolded into the behavior of neural elements or groups at every level of structural hierarchy. I believe Bohm's notion of the 'implicate order' may have direct application here [180]. There is ample evidence that the electrophysiological properties of neurons endow them with what are termed *autorhythmic oscillatory properties*, and that these develop into the network behavior of coupled oscillators and resonators, depending on the connectivity of individual neural groups [181]. The behavior of coupled oscillators and resonators can quickly produce complex, global behavior, as studied in the field of non-linear, dynamical systems. This behavior provides the context into which information arising from sensory data must be enfolded. The resulting overall behavior constitutes the EEG *holomovement*.

Part 8—A Note on Musical Holarchies

Earlier I referred to the *On Being Invisible* system in the context of building musical holarchies in real-time. I use the word 'holarchy' to refer to a structural entity that results from the simultaneous evolution of processes on both macro- and micro-scales of time, space and cognitive extent. Thus, I avoid the view of a system or structure as being the exclusive result of either top-down or bottom-up building processes. 'Holarchy' is also close to what is meant by 'coevolution' [182]. When I use the word 'hierarchy' I usually mean it to refer to a bottom-up process. This is appropriate for mechanistic systems such as computers. It is not appropriate when processes of perception or cognition are placed in the context of evolving societies or theories of learning. Learning takes place on both global and atomistic levels simultaneously. Consequently, I present in Fig. 29, a view (inspired by Jantsch, using terminology from Bohm, Tenney and myself) of the evolution and formation of the holarchic gestalt for a work of music as a function of the temporal extent of musical cognitive domain.

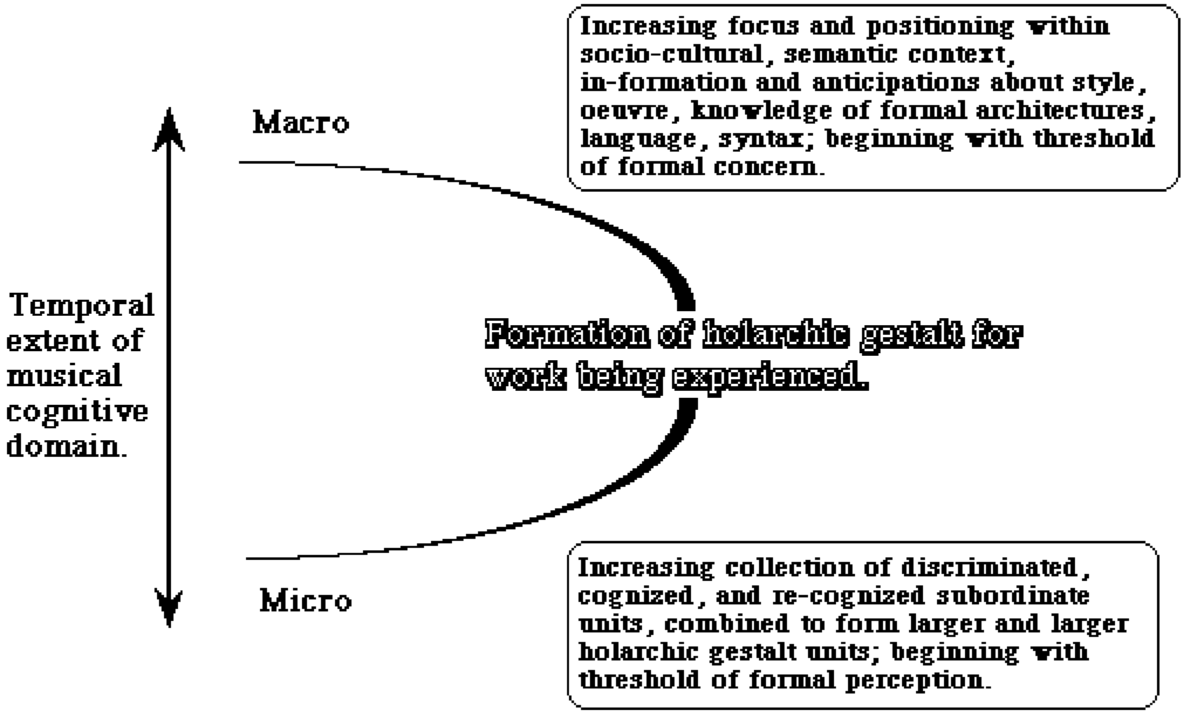

Increasing focus and positioning within socio-cultural, semantic context, in-formation and anticipations about style, oeuvre, knowledge of formal architectures, language, syntax; beginning with threshold of formal concern.

Increasing collection of discriminated, cognized, and re-cognized subordinate units, combined to form larger and larger holarchic gestalt units; beginning with threshold of formal perception.

Simultaneous evolution of macro- and micro- aspects of formal perception leading to synthesis of *holarchic temporal gestalt* for a musical experience

Figure 29

Fig. 29. Simultaneous Evolution of Macro- and Micro- Aspects of Formal Perception Leading to a Synthesis of *Holarchic Temporal Gestalt* for a Musical Experience. A view of how a total mental image for a musical work is formed. Processes of perception and cognition coevolve from both atomistic (bottom up) and global (top down) starting points at the extreme ends of a scale showing how broadly in time the musical cognitive domain extends. Macro-processes take into account a wide range of experiences in time, including history and expectations for the future. Micro-processes function on time scales appropriate for the various substructures in a work. Eventually these converge toward a final *holarchic gestalt.*

Just as the greater part of the heavens must for ever remain hidden
to the astronomer, so shall we never grasp the essence of music
completely...
-F. Busoni, 1924 [183]

Appendix 1—Study for On Being Invisible (1978):

A Conceptual Scheme for a Biofeedback Work Involving Coherent-Wave, EEG Phenomena and Electronic Orchestrations

(Composed in 1978 and slightly revised in 1989)

1. Establish the ability to monitor the theta, alpha and beta frequency bands of a performer's EEG.

2. Establish the ability to measure the amplitudes of each EEG frequency band and the stability of the EEG frequencies that occur in each band. A measurement of the variance of dominant frequencies in each band is required. The ability to measure the coherence of these dominant frequencies is desirable.

3. Establish three orchestras of electronic sound, one associated with each of the three EEG frequency bands. These orchestras should be stochastically controlled (see Fig. A-1).

Fig. A-1. Distribution of activity for three electronic orchestras, each associated with an EEG frequency band.

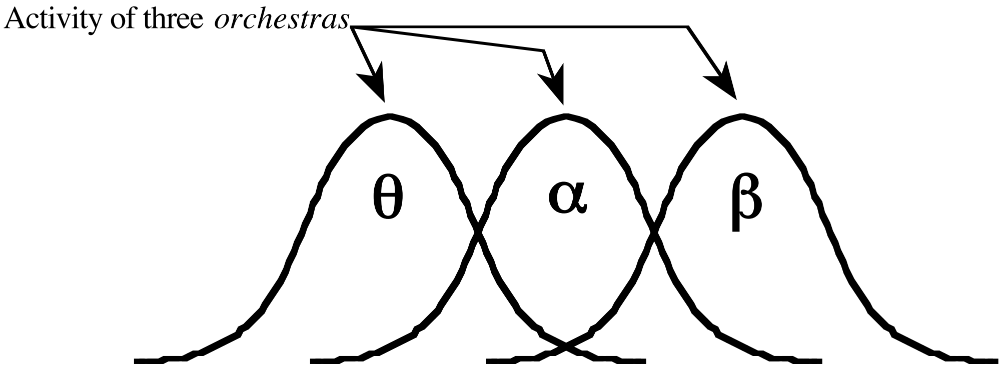

4. Make the variance allowed in the orchestra actions correspond to the variance in the associated EEG frequency band. When a stable state is achieved, i.e. coherent wave centered on the mean of the frequency distribution, the associated orchestra must trap a set of its current patterns and functions, allowing them to repeat more or less continuously. These patterns may be simple or complex (see Fig. A-2).

Fig. A-2. Relationship of electronic orchestra activity to the variance in a particular EEG frequency band.

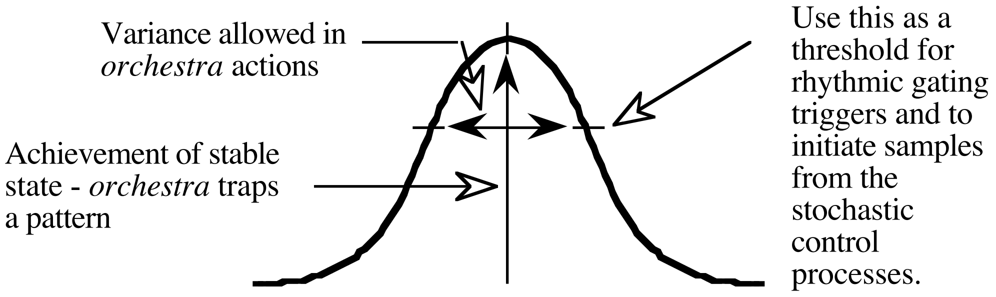

5. Establish a large table of values for control functions to be assigned to the three orchestras with the statistical characteristics shown in Fig. A-3.

Fig. A-3. Statistical characteristics for tables containing the values of functions that control the electronic orchestras.

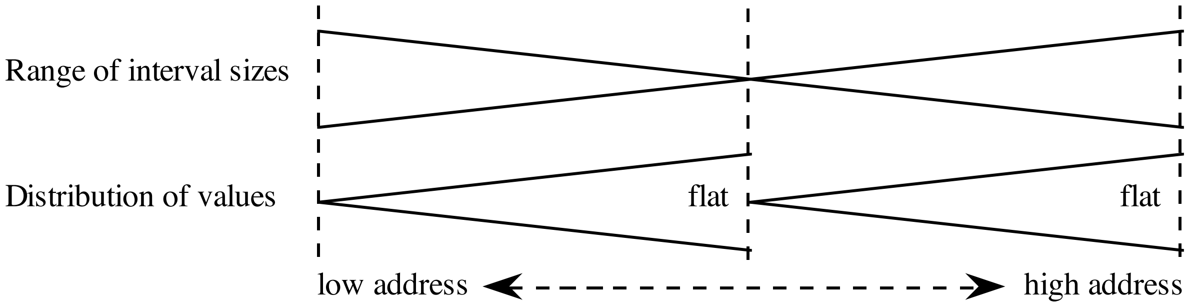

6. Assign three non-overlapping regions of the table to each of the three orchestras. Control the width, but not the placement, of random windows on these functions. Values may need to be sensed and compared with previous and adjacent values, and new events may need to be generated based on interval size.

7. Make the stability of a trapped process depend on relative as well as absolute values of activity in the EEG frequency bands. For example, the stability of the alpha orchestra is high if the EEG alpha is high, higher if theta is also low, and higher still if both theta and beta are low. To implement this, amplitudes from each band may be processed as follows:

$$\alpha \leftarrow \alpha - \frac{\theta}{2} - \frac{\beta}{2}$$

$$\beta \leftarrow \beta - \frac{\theta}{2} - \frac{\alpha}{2}$$

$$\theta \leftarrow \theta - \frac{\alpha}{2} - \frac{\beta}{2}$$

8. Proceed with a self-organizing improvisation based on relationships manifested amidst an evolving musical fabric and conscious or unconscious musical cognitive processes.

Appendix 2—Some Notes and Updates for the 1997 Re-Publication

DYNAMICAL SYSTEMS AND THE BRAIN

Some of the most exciting developments in research fields related to the topics of this monograph that have occurred since it's first publication in 1990 have emerged from the study of complex adaptive systems. The literature in this area has expanded greatly and it is far beyond the scope of this publication to provide an adequate survey. However, I want to refer the reader to the work of J.A. Scott Kelso on self-organization in the brain and behavior which I have found particularly helpful and insightful. His ideas are based on the assumption that the emergence of patterned behavior on all levels is governed by genetic processes of self-organization. Many of the results are directly applicable to the study of music perception and performance [184].

MUSIC AND THE EEG

New studies that relate brain processes observed in the EEG to musical experience have continued to be carried out. A interesting place to begin examining these is the work on musical perception and the dimensional complexity of brain activity by N. Berbaumer et. al., along with the references they site [185]. Their experiments explore what is known as the non-linear resonance hypothesis of music perception in which relationships are drawn between the complexity of music stimuli, the complexity of observed EEG waveforms, musical training and experience and the presumed activation of neural assemblies necessary to process rich associations. In these studies, complexity is given a precise mathematical definition—commonly used in studying chaotic systems—involving calculating the correlation dimension of a time series. It was determined that subjects' ratings of the subjective complexity of the music stimuli followed the mathematical complexity measures. The complexity of EEG waveforms and music stimuli are strongly related, but also strongly affected by the level of subjects' musical training and experience. An interesting and related sonification experiment with the EEG was also presented by G. Mayer-Kress [186]. In this exercise, rapid short-time synchronization events in the EEG—presumed to accompany cognitive events and perception—were mapped onto a musical texture of independent instrument sounds that were drawn into transient synchronization as concomitant events were detected in the EEG.

CONCEPTS FOR COMPOSITION

In 1991 I began developing an interactive program called *Hierarchical Form Generator (HFG)* in which a parsing algorithm automatically separates a stream of musical inputs into meaningful substructures and makes these available to performers for transformation, recall and recombination with other substructures during a performance [187]. The parsing algorithm—based on a partial model of musical form perception—is designed to detect boundaries in essentially unpredictable musical inputs and has its origins in the earlier work described in this monograph. A repertoire of software tools for composition and transformation of musical materials has been added, making *HFG* a general purpose tool for composing and improvising. To date, several musical projects have been realized with it, including the compositions *Predictions, Confirmations and Disconfirmations* (1991) for MIDI grand piano, *HFG* and automatically responding instruments, *Extended Trio* (1992), in collaboration with jazz bassist, Charlie Haden, and South Indian mrdangam and kanjira performer, Trichy Sankaran, in which *HFG* is used to link all members of the trio to responding instruments [188], and several compositions, *Lineage*, *Enactment*, *Transfiguration* and *Transference* (1992) in collaboration with Anthony Braxton, which are duets for MIDI grand piano, *HFG*, a responding playback piano and various wind instruments [189]. *HFG* is continuing to be used as a general purpose tool for composition and improvisation in new works as well.

I have experimented with various dynamical systems analysis techniques to examine the complexity of music. One of these involves calculating the correlation dimension of melodic lines and is very similar to techniques used for analyzing the complexity of brain activity in the studies cited in the section *Music and the EEG*, above [190]. Perhaps analyses similar to these could be incorporated into self-organizing feedback processes for further compositions.

Finally, the study of self-organization and emergent phenomena in the evolution of musical forms continues to be a source of inspiration and enlightenment for what I have come to call *Propositional Music*. The reader is referred to the following publications for further discussion and information about these ideas, [191], [192], [193] and [194].

BRAIN IMAGING

During the past 20 years, brain imaging techniques using positron emission tomography (PET) and later, magnetic resonance imaging (MRI) enabled researchers to map regions of the brain that are activated during various kinds of mental operations. Most of these methods are based on observing changes in blood flow as various neural groups are stimulated into increased activity. Among the findings produced with these methods is support for the idea that brain processes that are generated internally—such as from instructions or imaginations—can activate the same sensory regions that would be activated in response to corresponding, real sensory events [195]. The temporal resolution of these techniques is extremely low, however, because considerable time is required for the blood flow changes associated with the experimental events to become observable. Nevertheless, these methods have enabled important studies on voluntary attention, the convergence of endogenous and exogenous processes onto the same brain areas and changes in brain activation as a function of thinking, visualization and experience [195]. Recently, a new technique called functional magnetic resonance imaging (fMRI) that can respond to changes in brain activation more quickly has surpassed positron emission tomography (PET). fMRI is based on detecting increases in blood oxygenation in localized areas and has given rise to a new method for analyzing event related changes in the brain, known as event related fMRI. It is being used to view the activation of neural groups in regions of the brain that are involved in various cognitive processes. Knowing the precise timing of activation and synchronization events in the brain, along with phase relationships among cyclical processes, is critical to increasing our understanding of neural information processing and control systems. To add higher timing resolution to spatial mappings, researchers are combining traditional EEG and magnetoencephalography (MEG) methods with the new fMRI techniques and, sometimes, PET. The spatial resolution of fMRI and its ability to locate active neural assemblies, combined with the relatively high temporal resolution of the EEG and MEG are, thus, joined to give researchers unprecedented views of brain activity [196].

Appendix 3—On Being Invisible II (Hypatia Speaks to Jefferson in a Dream):

A Self-Organizing, Multi-Media Performance Work Utilizing Event-Related Potentials From Performers' Brains (1994-95)

PROGRAM NOTES

Performers
Player One: brainwaves (Hypatia)
Player Two: brainwaves (Jefferson)
Player Three: instrument with transient sounds (Double 1)
Player Four: instrument with sustained sounds (Double 2)
Player Five: narrator
Players Six and Seven: computer media

In the late 1960's I became fascinated with new developments in brain science as they related to musical perception and the emergence of new musical languages. Ideas from cybernetics, notably those relating to the self-regulation of systems by means of feedback, were finding their way into psycho-biological research, resulting in an explosion of interest in something popularly known as, *biofeedback*. The notion of self-regulation, that individuals may be able to achieve a degree of conscious, willful control of particular body functions formerly thought only to be regulated by unconscious, autonomic processes, captured the imaginations of many people. My own interest in biofeedback centered around the notion that self-regulation of brain functions, as could be observed through monitoring aspects of electrical brain activity, was closely related to certain processes involved in the evolution of new musical styles. Self-regulation by means of feedback is also closely related to some ideas about evolution, and models of evolution appear as a consistent, thematic referent throughout much of my musical work. Consequently, I began a long period of research in information processing modalities of the nervous system as they relate to aesthetic experience and creative activity. I produced many musical compositions and interdisciplinary, artistic pieces in which the material forms in the works were generated spontaneously by means of direct monitoring of electrical brain activity and/or other body functions. I published numerous articles about this work, two books, Biofeedback and the Arts and Extended Musical Interface with the Human Nervous System, and several recordings. This was, however, only a beginning.

In 1976, I began creating a work entitled *On Being Invisible*, which, for me, contains the richest aesthetic, symbolic and metaphorical content arising form the import that biofeedback systems had on my work as a composer. *On Being Invisible* is a self-organizing, dynamical system, rather than a fixed musical composition. The title refers to the role of the individual within an evolving, dynamical environment, who makes decisions concerning when and how to be a conscious *initiator* of action and when simply to allow her or his individual, internal dynamics to co-evolve within the macroscopic dynamics of the system as a whole. Consequently, the work is always ongoing. Within the corpus of my music, the title serves as a label for a period of work with these ideas from about 1976 to 1979. A recording of an early version was released in 1977. Recently, after concentrating on other things for several years, I have begun new work with this system, calling it, *On Being Invisible II.* This new work is stimulated partly by advances in technology that only now make the realization of earlier concepts possible, and it is partly the result of interest in applying new knowledge within a still very rich musical paradigm.

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. To accomplish this, components of the electroencephalogram (EEG) recorded from the brains of on-stage performers, known as *event-related potentials (ERPs)*, are detected, measured and analyzed. ERPs are transient waveforms in the EEG associated with the occurrence of stimulus events having a high degree of salience — particular meaningfulness — to the subject emitting these brainwaves, always in relation to a particular context of surrounding events. Next, computers are programmed to produce a stream of sonic events according to some predetermined starting point or compositional method devised by the composer. The computer software also contains a partial model of musical perception, with which it attempts to predict what events in its own, musical output might be perceived by the subject as having significance in the emerging musical structure. Usually, these correspond to boundary points, such as the end of a phrase and the beginning of a new phrase, a significant change in texture, or changes in the pattern grouping of phrases into sequences or other higher level forms. A powerful, widely-used software tool which I co-authored, known as HMSL, (Hierarchical Music Specification Language), is used to manipulate formal musical elements referred to as *morphologies*, or *morphs*, for short. ERPs from the performer-subjects are then analyzed to determine if the computer's predictions correspond to signals from the brain that should accompany important, attention-securing events. If they do not, the music generating algorithms are allowed to mutate into new forms and new predictions are tested. If the predictions are confirmed, the kinds of events reliably associated with these confirmed predictions gain prominence in the musical fabric. In this way, self-organizing, musical forms can emerge that are related to the shifts of attention experienced by the performer-subjects and that can be confirmed by brain signal measurements. In modern terminology, this system exhibits many of the characteristics of what we call, *complex adaptive systems.* Such systems are used to model the evolution of complex life forms that are often governed by simple, underlying rules. Thus, an interactive, musical system is produced that can spontaneously evolve new, emerging, musical orderings, and perhaps, even languages.

Over many years of performing, writing, producing recordings of *brainwave music*, and further thinking, the components of this feedback system began to remind me of characters in a mythological drama, the spontaneous forces of creativity, the drive to converge upon ordered relationships in society, the counterbalancing tension of divergence from order as our consciousness loses its focus on orderings from the past, and the fundamental uncertainties regarding forces in nature that are only partially knowable. Consequently, I began to think about *On Being Invisible* in theatrical or narrative terms. This raised an important question. If music combined with theater can be loosely termed, *opera*, how, then, does one go about creating a *self-organizing opera?* This question may never be fully answered, but it is far too stimulating to my imagination to stop trying.

On Being Invisible II (Hypatia Speaks to Jefferson in a Dream) is an experiment with this question. The setting is a dream in which Thomas Jefferson hears the voice of the Greek, woman, astronomer, mathematician, and philosopher, Hypatia, traversing the centuries of time and the space of continents, mingling with his own internal voices as he is writing one of his later-to-be-famous documents. The components of ideological conflict that emerge from this scene remind me of the tension associated with the individual performer in *On Being Invisible*, who must always negotiate a thin dividing line separating being part of something larger than one's self and trying to willfully direct a naturally evolving process. Hypatia, an Alexandrian who was murdered in A.D. 415 for being both Greek and a woman who dared to lecture, resided at a focal point of change in the old world, the end of Classical Greek philosophy and the beginning of the Dark Ages, the foundation of Neo-Platonism and the emergence of Plotinus, the transformation of Christianity from a moral teaching into a brutal instrument of political power, the appropriation of Plotinus and mysticism by the Christians to obscure thought and achieve totalitarian, political control, the decline of Alexandria as an intellectual center, symbolized by the destruction of the fabled library, combined with an unprecedented outpouring of romantic, multi-sexual poetry, and the labyrinthine racial-political conflicts there among Greeks, Jews, Ptolmaics, remnants of Egyptian antiquity, Copts, Islamics, Europeans, and numerous others. These are just a small sampling. Similarly, Thomas Jefferson was a figure wedged in-between the end of the Age of Enlightenment and emerging Romanticism, an American hero who espoused freedom of thought and religion but also kept slaves, a revolutionary torn between rationality and romance, who's relationships with women, from slaves to European intellectuals, symbolized the psycho-sexual dilemma of a young nation, whose brilliant inventiveness and creative genius was at once steeped in Neo-Classicism and evinced a great contempt for Plato, who was both a champion of the political avant garde and a player in the new dynamics of wealth and power, a president in the new world who was also obsessed with the mathematics of miscegenation. The *invisibility* manifest in this scenario is represented by the dream state of Jefferson in which these conflicts energize his thoughts and entreaties to wisdom are transmitted to him through warps in space-time by the reincarnated mind of Hypatia.

This realization of *On Being Invisible II* is set for two performers, representing Hypatia and Jefferson, whose brain signals are being monitored and event-related potentials analyzed. The results are used to create the forms of electronic music we hear, sequences of visual icons we see through computerized video projection, and arrangements of words spoken by electronically sampled voices. The words come from various texts by Jefferson, including selections from his letters and writings on the arts and philosophy. Hypatia's words are speculative. They come from modern authors, original words by the composer, and selections from Hypatia's contemporaries. Each of these characters has a double image on stage in the form of a musician. These are the ghost doubles of Hypatia and Jefferson, in the sense of being their personal angels and also representing human beings' propensity to make copies of themselves in nefarious forms. These musical parts are written for master improvisers to provide musical glue for the performance. Finally, a narrator represents the dream state and a neutral form of the emerging properties of a new, global consciousness.

Credits and Acknowledgments
Conception and Composition: David Rosenboom, 1994-95, based on the earlier work, On Being Invisible, (an attention-dependent sonic environment), 1976-1979.
Technical assistance and computer/video image design: Kent Clelland.
Recorded Voices: Teri DeSario and Roxanne Merryfield
Projected Slide Collages: Jacqueline Humbert
Digital video assistance: Warren Heaton
Photoshop computer assistance: Vincent Carté
Analog video assistance: Steven Kury
Media consultant: Sara Roberts
Brain science inspiration: Dr. E. E. "Ted" Coons and Dr. Lloyd Kaufman

References and Notes

1. C. Ives, *Essays Before a Sonata*, in *Essays Before a Sonata and Other Writings by Charles Ives*, H. Boatwright, ed. (New York: W.W. Norton and Co., Inc., 1964) pp. 88-89. Originally published in C. Ives, *Essays Before a Sonata*, 1920.

2. Ives [1] p. 84.

3. E.D. Adrian and B.H.C. Mathews, "The Berger Rhythm: Potential changes from the Occipital Lobes in Man", *Brain*, **57** (1934) pp. 355-385.

4. For a discussion of the early history of ideas on rhythmicity in brainwaves, see P. Andersen and S.A. Andersson, *Physiological Basis of the Alpha Rhythm* (New York: Appleton Century Crofts, 1968).

5. H.W. Shipton, "An Electronic Trigger Circuit as an Aid to Physiological Research", *Journal of the British Institute of Radio Engineers* **4** (1949) pp. 374-383.

6. V.J. Walter and W.G. Walter, "The Central Effect of Rhythmic Sensory Stimulation", *Electroencephalography and Clinical Neurophysiology* **1** (1949) pp. 57-58.

7. J. Corriol and H. Gastaut, "Le Circuit Déclancheur Électronique de H. Shipton et W. Grey Walter. Son Application en Élecroencéphalographie Clinique", *Rev. Neurol.* **82** (1950) pp. 608-613.

8. M.G.T. Hewlett, "An Electronic Trigger Mechanism", *Electroencephalography and Clinical Neurophysiology* **3** (1951) pp. 513-516.

9. E.C. Turton, "An Electronic Trigger Used to Assist in the EEG Diagnosis of Epilepsy", *Electroencephalography and Clinical Neurophysiology* **4** (1952) pp. 83-91.

10. For a more recent example of the translation of data into sound, see I. Peterson, "The Sound of Data", *Science News* **127** (1985) pp. 348-350.

11. A. Lucier, "Statement On: Music for Solo Performer" (1971), in D. Rosenboom, ed., *Biofeedback and the Arts, Results of Early Experiments,* 2nd Ed. (Vancouver: Aesthetic Research Centre of Canada Publications, 1976) pp. 60-61.

12. A. Lucier, and D. Simon, *Chambers*, (Middletown, CT: Wesleyan University Press, 1980).

13. A. Lucier, *Music for Solo Performer*, LP record, #VR 1014 (New York: Lovely Music, Ltd., 1982).

14. R. Teitelbaum, "In Tune: Some Early Experiments in Biofeedback Music (1966-74)" (1974), in Rosenboom [11] pp. 35-56.

15. Rosenboom [11].

16. M.L. Eaton, "Bio-Potentials as Control Data for Spontaneous Music", ORCUS Research Technical Publication, TP-3001 (Kansas City, MO: 1968).

17. M.L. Eaton, "Electronic Music—A Handbook of Sound Synthesis and Control", ORCUS Research Technical Publication, TP-3003 (Kansas City, MO: 1969).

18. M.L. Eaton, *Bio-Music (Biological Feedback Experiential Music Systems)* (Kansas City, MO: ORCUS Research, 1971; also Barton, VT: Something Else Press, 1974).

19. For a non-technical account of some of this history, see G. Jonas, *Visceral Learning* (New York: Viking Press, 1973).

20. W.T. Powers, "Feedback: Beyond Behaviorism", *Science* **179** (1973) pp. 351-356.

21. Dynamical systems may be thought of as those involving forms or behaviors that change over time. The study of such changing forms may also be termed 'morphodynamics'. For an excellent series of mathematical picture books from which readers of almost any discipline may gain insight into this field, see R. Abraham and C.D. Shaw, *Dynamics, the Geometry of Behavior,* 4 vols, *Part One: Periodic Behavior, Part Two: Chaotic Behavior, Part Three: Global Behavior, and Part Four: Bifurcation Behavior* (Santa Cruz, CA: Ariel Press, Inc., 1984-1985).

22. G. Schöner and J.A.S. Kelso, "Dynamic Pattern Generation in Behavioral and Neural Systems", *Science* **239** (1988) pp. 1513-1520.

23. This idea for the sum of two Gaussians was originally suggested by Heinz von Foerster. See H. von Foerster, *Observing Systems*, (Seaside, California: Intersystems Publications, 1981).

24. E. Jantsch, *The Self-Organizing Universe* (Oxford: Pergamon Press, 1980).

25. Jantsch [24] p. 248.

26. Jantsch [24] p. 255

27. L.G. Fehmi, "Open Focus Training", text of a workshop presented at a meeting of the Biofeedback Research Society and the Council Grove Conference on Voluntary Control of Internal States (1975).

28. A. Lucier, "Clocker", *Musicworks* **41** (Toronto: Music Gallery, 1988).

29. D. Rosenboom, "Method for Producing Sounds or Light Flashes with Alpha Brain Waves for Artistic Purposes", *Leonardo* **5**, No. 2, 141-145 (1972); also reprinted in F. Malina, ed., *Kinetic Art* (New York: Dover, 1974).

30. D. Rosenboom, "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, Information Processing Systems* (Urbana: Office of Continuing Education in Music, University of Illinois, 1975); also in Rosenboom, [11].

31. D. Rosenboom, *Brainwave Music*, LP record, #ST1002 (Vancouver/Toronto: Aesthetic Research Centre Publications and Recordings, 1976).

32. D. Rosenboom, "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 Pub., 1976).

33. D. Rosenboom, *On Being Invisible*, LP record, #MGE-4 (Toronto: Music Gallery Editions, 1977).

34. D. Rosenboom, *On Being Invisible*, video tape, (Vancouver: Western Front Video, 1977).

35. D. Rosenboom, "On Being Invisible", *Musicworks* **28** (Toronto: Music Gallery, 1984) pp. 10-13.

36. D. Rosenboom, "The Performing Brain", *Computer Music Journal* **14**, No. 1 (1990) pp. 48-66.

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

38. Malina [29].

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

40. Rosenboom [11].

41. Rosenboom [11].

42. S. Penny, "Charlie Chaplin, Stelarc and the Future of Humanity", *Artlink* **9**, No. 1, pp. 29-32 (1989).

43. R. Gillett, K.C. Smith and B. Pritchard, "MADDM—Dance Directed Music", in B. Truax, ed., *Proceedings of the International Computer Music Conference* (Burnaby, B.C., Canada: Simon Fraser University, 1985).

44. Rosenboom [11].

45. Rosenboom [31].

46. Rosenboom [11].

47. Grayson [37].

48. D.P. Spence, M. Lugo and R. Youdin, "Cardiac Change as a Function of Attention to and Awareness of Continuous Verbal Text", *Science* **176** (1972) pp. 1344-1346.

49. Rosenboom [32].

50. For a discussion of intrinsically random, dynamical systems, see I. Prigogine and I. Stengers, *Order out of Chaos*, (New York: Bantam Books, 1984).

51. For a canonical treatment of autocorrelation and cross-correlation functions, see G.M. Korn and T.M. Korn, *Mathematical Handbook for Scientists and Engineers* (New York: McGraw Hill, 1968).

52. A.S. Gevins, G.M. Zeitlin, J.C. Doyle, C.D. Yingling, R.E. Schaffer, E. Callaway and C.L. Yeager, "Electroencephalogram Correlates of Higher Cortical Functions", *Science* **203** (1979) pp. 665-668.

53. W.J. Ray and H.W. Cole, "EEG Alpha Activity Reflects Attentional Demands, and Beta Activity Reflects Emotional and Cognitive Processes", *Science* **228** (1985) pp. 750-752.

54. Rosenboom [11].

55. Rosenboom [31].

56. E.R. John, "Switchboard Versus Statistical Theories of Learning and Memory", *Science* **177**, No. 4052, 850-864 (1972).

57. E.R. John, "A Model of Consciousness", in G.E. Schwartz and D. Shapiro, eds., *Consciousness and Self-Regulation, Advances in Research*, Vol. 1 (New York: Plenum Press, 1976) pp. 1-50.

58. Rosenboom [11].

59. Some studies have been carried out in which ERPs to visual stimuli of an artistic nature have been recorded. Both Eaton [18] and Clynes (M. Clynes, "Towards a View of Man", in M. Clynes and J.H. Milsum, eds., *Biomedical Engineering Systems* [New York: McGraw Hill, 1972] pp. 272-358) point out that ERPs recorded for presentations of complex, visual stimuli, such as patterns with structural intricacy, random elements or deeper content (for instance, that found in paintings), tend to elicit larger peaks than do those for simple stimuli, such as light flashes.

60. John [57].

61. Because the literature in this field is so voluminous, the references cited in this article do not constitute a survey that can be considered inclusive. Rather, they represent influential stepping stones that can aid one in traversing the knowledge base of the field in a journey toward generalized observations and formation of concepts, which, it is hoped, will be applicable in developing increased understanding and new applications.

62. E. Halgren, N.K. Squires, C.L. Wilson, J.W. Rohrbaugh, T.L. Babb, and P.H. Crandall, "Endogenous Potentials Generated in the Human Hippocampal Formation and Amygdala by Infrequent Events", *Science* **210** (1980) pp. 803-805.

63. I. Fried, G.A. Ojemann and E.E. Fetz, "Language-Related Potentials Specific to Human Language Cortex", *Science* **212** (1981) pp. 353-356.

64. C.D. Woody, "Characterization of an Adaptive Filter for the Analysis of Variable Latency Neuroelectric Signals", *Medical and Biological Engineering* **5** (1967) pp. 539-553.

65. J.W. Rohrbaugh, K. Syndulko and D.B. Lindsley, "Brain Wave Components of the Contingent Negative Variation in Humans", *Science* **191** (1976) pp. 1055-1057.

66. J.W. Rohrbaugh, K. Syndulko, T.F. Sanquist and D.B. Lindsley, "Synthesis of the Contingent Negative Variation Brain Potential from Noncontingent Stimulus and Motor Elements", *Science* **208** (1980) pp. 1165-1168.

67. J.P. Rosenfeld and B.E. Hetzler, "Operant-Controlled Evoked Responses: Discrimination of Conditioned and Normally Occurring Components", *Science* **181** (1973) pp. 767-770.

68. Rosenboom [30].

69. C.M. Nunn, "Biofeedback with Cerebral Evoked Potentials and Perceptual Fine-Tuning in Humans", Master's thesis, York University (1976); and in, *Journal of Experimental Aesthetics* **1** (Vancouver: Aesthetic Research Centre of Canada Publications, 1977) pp. 11-100.

70. See E. Donchin, ed., *Cognitive Psychophysiology, Event-Related Potentials and the Study of Cognition*, Vol. 1 (Hillsdale, NJ: Lawrence Erlbaum Associates, 1984) and Refs. [30] and [57] and the many references cited therein.

71. S. A. Hillyard in Ch. 3, "Event-Related Potentials and Selective Attention", in Donchin [70].

72. S.A. Hillyard, R.F. Hink, V.L. Schwent and T.W. Picton, "Electrical Signs of Selective Attention in the Human Brain", *Science* **182** (1973) pp. 177-179.

73. Y. Matsumiya, V. Tagliasco, C.T. Lombroso and H. Goodglass, "Auditory Evoked Response: Meaningfulness of Stimuli and Interhemispheric Asymmetry", *Science* **175** (1972) pp. 790-792.

74. R. Parasuraman and J. Beatty, "Brain Events Underlying Detection and Recognition of Weak Sensory Signals", *Science* **210**, (1980) pp. 80-83.

75. T.W. Picton and S.A. Hillyard, "Human Auditory Evoked Potentials, I: Evaluation of Components, and II: Effects of Attention", *Electroencephalography and Clinical Neurophysiology* **36**, (1974) pp. 179-199.

76. Picton and Hillyard [75].

77. Hillyard et al. [72].

78. Matsumiya et al. [73].

79. J.M. Ford, W.T. Roth, S.J. Dirks and B.S. Kopell, "Evoked Potential Correlates of Signal Recognition Between and Within Modalities", *Science* **181** (1973) pp. 465-466.

80. W. Ritter, R. Simson, H.G. Vaughan, Jr., and D. Friedman, "A Brain Event Related to the Making of a Sensory Discrimination", *Science* **203** (1979) pp. 1358-1361.

81. Ford et al. [79].

82. Ritter et al. [80].

83. Parasuraman and Beatty [74].

84. D.D. Paul, and S. Sutton, "Evoked Potential Correlates of Response Criterion in Auditory Signal Detection", *Science* **177** (1972) pp. 362-364.

85. W. Chase in Ch. 10, "Report of Panel IV: Mental Chronometry", in Donchin [70].

86. Rosenboom [11].

87. Rosenboom [30].

88. Donchin [70].

89. Picton and Hillyard [75].

90. Ritter et al. [80].

91. John [57].

92. M. Kutas and S.A. Hillyard, "Reading Senseless Sentences: Brain Potentials Reflect Semantic Incongruity", *Science* **207** (1980) pp. 203-205.

93. V.S. Johnston and G.L. Chesney, "Electrophysiological Correlates of Meaning", *Science* **186**, No. 4163, 944-946 (1974).

94. M.M. Bortin, A.A. Rimm, and E.C. Saltzstein, "Evoked Potential Correlates of Expected Stimulus Intensity", *Science* **179**, No. 4075, 814-816 (1973).

95. John [56].

96. John [57].

97. E. Callaway and P.R. Harris, "Coupling Between Cortical Potentials from Different Areas", *Science* **183** (1974) pp. 873-875.

98. For some interesting examples, see C. Seeger, *Studies in Musicology 1935-1975* (Berkeley: University of California Press, 1977).

99. D. Rosenboom, *Zones of Influence*, score of musical work for computer electronics and percussion (Oakland, CA: David Rosenboom Publishing and Frog Peak Music, 1984-86).

100. D. Rosenboom, "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**, No. 1/2, 569-583 (Spring 1987).

101. L. Polansky, D. Rosenboom and P. Burk, "HMSL: Overview (Version 3.1) and Notes on Intelligent Instrument Design", in J. Beauchamp, compiler, *Proceedings of the International Computer Music Conference* (Urbana: University of Illinois, 1987) pp. 220-227.

102. L. Polansky, "Morphological Metrics: An Introduction to a Theory of Formal Distances", in J. Beauchamp compiler, *Proceedings of the International Computer Music Conference* (Urbana: University of Illinois, 1987) pp. 197-205.

103. A.S. Gevins, J.C. Doyle, B.A. Cutillo, R.E. Schaffer, R.S.Tannehill, J.H. Ghannam, V.A. Gilcrease and C.L. Yeager, "Electrical Potentials in Human Brain During Cognition: New Method Reveals Dynamic Patterns of Correlation", *Science* **213** (1981) pp. 918-922.

104. For detailed descriptions of some of the methods employed, see L. Glass and M.C. Mackey, *From Clocks to Chaos, The Rhythms of Life* (Princeton, NJ: Princeton University Press, 1988).

105. Jantsch [24] p. 161.

106. W.J. Freeman, *Mass Action in the Nervous System: Examination of the Neurophysiological Basis of Adaptive Behavior Through the EEG* (New York: Academic Press, 1975).

107. R. Pool, "Is It Healthy to Be Chaotic?", *Science* **243** (1989) pp. 604-607.

108. Jantsch [24] pp. 161-165.

109. D. Bohm, *Wholeness and the Implicate Order* (New York: Routledge and Kegan Paul, 1980).

110. Rosenboom [32].

111. Schôner and Kelso [22].

112. It is worth mentioning that I have applied the concept of 'correlation dimension' extensively in creating what are termed, 'concept spaces' for mapping relationships in musical composition. For example, a method of mapping variations on the contours or shapes of melodies and rhythmic patterns according to the degree to which each is correlated with precomposed melodies or patterns, called the 'origin' and 'target patterns', was used in the compositions *Zones of Influence*, a five-part work for percussion and computer music system, and *Champ Vital (Life field)*, a trio for violin, piano and percussion. 'Trajectories' in these concept spaces (which map transformations of the shapes, one into another) become primary compositional entities endowed with particular aesthetic qualities. These *t*rajectories may become enfolded further into structures on other levels of a 'holarchic' form. A computer language optimized for compositional experimentation along these lines, called HMSL (Hierarchical Music Specification Language), has been created. See Refs. [99], [100], [101] and D. Rosenboom, *Champ Vital (Life Field)*, score of musical work for violin, piano and percussion (Oakland, CA: David Rosenboom Publishing and Frog Peak Music, 1987) for further information.

113. J. Shrager, T. Hogg and B.A. Huberman, "Observation of Phase Transitions in Spreading Activation Networks", *Science* **236** (1987) pp. 1092-1094.

114. Bohm [109].

115. Rosenboom [11].

116. Ives [1] p. 84.

117. R. Francès, *The Perception of Music*, trans. from the 1958 edition by W.J. Dowling (Hillsdale, NJ: Lawrence Erlbaum Associates, 1988).

118. C. Wickens, A. Kramer, L. Vanasse and E. Donchin, "Performance of Concurrent Tasks: A Psychophysiological Analysis of the Reciprocity of Information-Processing Resources", *Science* **221** (1983) pp. 1080-1082.

119. John [57].

120. Ives [1] p. 94.

121. G. McCarthy and E. Donchin, "A Metric for Thought: A Comparison of P300 Latency and Reaction Time", *Science* **211** (1981) pp. 77-80.

122. Chase [85].

123. J. Tenney, *META + HODOS and META Meta + Hodos*, 2nd Ed. (Oakland, CA: Frog Peak Music, 1988).

124. J. Tenney and L. Polansky, "Hierarchical Temporal Gestalt Perception in Music: A Metric Space Model", *Journal of Music Theory* **24**, No. 2, 205-241 (Fall 1980).

125. Clynes [59].

126. S. Sternberg, "High Speed Scanning in Human Memory", *Science* **153** pp. 652-654.

127. J.L. McClelland, "On the Time Relations of Mental Processes: An Examination of Systems of Processes in Cascade", *Psychological Review* **86** (1979) pp. 287-330.

128. R. Schweickert, "Critical-Path Scheduling of Mental Processes in a Dual Task", *Science* **209** (1980) pp. 704-706.

129. Bohm [109] pp. 115-118.

130. John [57].

131. J.M. Grey, "Multidimensional Perceptual Scaling of Musical Timbres", *Journal of the Acoustical Society of America* **61**, No. 5, pp. 1270-1277 (1977).

132. D. Wessel, "Timbre Space as a Musical Control Structure", in C. Roads and J. Strawn, eds., *Foundations of Computer Music* (Cambridge, MA: MIT Press, 1985) pp. 640-657.

133. R.N. Shepard, "Toward a Universal Law of Generalization for Psychological Science", *Science* **237** (1987) pp. 1317-1323.

134. Polansky, Rosenboom and Burk [101].

135. Tenney [123].

136. Tenney and Polansky [124].

137. M. Klein, M.G.H. Cole and E. Donchin, "People with Absolute Pitch Process Tones Without Producing a P300", *Science* **223** (1984) pp. 1306-1309.

138. For a discussion of the URS model, see Clynes [59].

139. In retrospect, I see this as a first, tentative application of Shepard's generalization theory in the time domain (see the section *Musical Feature Extraction* in Part 3 of this monograph).

140. E. Coons and D. Kraehenbuehl, "Information as a Measure of Structure in Music", *Journal of Music Theory* **2**, No. 2, 127-161 (1958).

141. K.C. Squires, C. Wickens, N.K. Squires and E. Donchin, "The Effect of Stimulus Sequence on the Waveform of the Cortical Event-Related Potential", *Science* **193** (1976) pp. 1142-1145.

142. Note that simple, regular sequences—like the one 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. The latter often contain a richer fabric of information with which to produce engaging music. "Beauty in music is too often confused with something that lets the ears lie back in an easy chair" (Ives [1] p. 97).

143. Here, emphasis is placed on grouping musical events. These methods may also be applied, however, to any sequence of events considered in the abstract. Among the important concepts involved is the idea of analyzing what is known, at any particular point in an unfolding sequence that permits the formation of hierarchical groupings, and what must be known in order to detect patterns. Consequently, with these methods one could examine many kinds of sequences, such as editing schemes in film, video and television: the order in which important symbols occur in verbal, written or dramatic material; or any succession of events resulting from a time-based scanning and presentation.

144. Tenney [123].

145. Ives [1] p. 84.

146. E. Donchin and G. McCarthy in Ch. 10, "Report of Panel IV: Mental Chronometry", in Donchin [70].

147. For numerous examples, see Rosenboom [11].

148. M. Waldrop, "Hypercube Breaks a Programming Barrier", *Science* **240** (1988) p. 286.

149. S.J. Williamson, G.L. Romani, L. Kaufman and I. Modena, eds., *Biomagnetism—An Interdisciplinary Approach* (New York: Plenum, 1983).

150. J.P. Wilkswo, Jr., J.P. Barach and J.A. Freeman, "Magnetic Field of a Nerve Impulse: First Measurements", *Science* **208** (1980) pp. 53-55.

151. R. Hari and O.V. Lounasmaa, "Recording and Interpretation of Cerebral Magnetic Fields", *Science* **244** (1989) pp. 432-436.

152. R. Crease, "Biomagnetism Attracts Diverse Crowd", *Science* **245** (1989) pp. 1041-1043.

153. J.B. Beal, "The New Biotechnology", in J. White, ed., *Frontiers of Consciousness* (New York: Julian Press, 1974).

154. For applications of the SQUID in biomagnetism, see S.J. Williamson, L. Kaufman and D. Brenner, "Biomagnetism", in B.B. Schwartz and S. Foner, eds., *Superconductor Applications: SQUIDs and Machines* (New York: Plenum, 1977) pp. 355-402.

155. D.S. Barth, W. Sutherling, J. Engle Jr., and J. Beatty, "Neuromagnetic Evidence of Spatially Distributed Sources Underlying Epileptiform Spikes in the Human Brain", *Science* **223** (1984) pp. 293-296.

156. D.F. Rose, P.D. Smith and S. Sato, "Magnetoencephalography and Epilepsy Research", *Science* **238** (1987) pp. 329-335.

157. D. Brenner, S.J. Williamson and L. Kaufman, *Science* **190** (1975) p. 480.

158. D. Brenner, J. Lipton, L. Kaufman and S.J. Williamson, "Somatically Evoked Magnetic Fields of the Human Brain", *Science* **199**, No. 4324, 81-83 (1978).

159. Brenner et al. [158].

160. T.H. Geballe and J.K. Hulm, "Superconductivity—The State That Came in from the Cold", *Science* **239** (1988) pp. 367-375.

161. J. Clarke and R.H. Koch, "The Impact of High-Temperature Superconductivity on SQUID Magnetometers", *Science* **242** (1988) pp. 217-223.

162. B. Deagon, "Superconductivity Sparks Research Activity", *Computer Technology Review* **8**, No. 3, (1988) p. 4.

163. R. Pool, "Superconductors' Material Problems", *Science* **240** (1988) pp. 25-27.

164. R. Pool, "New Superconductors Come Through", *Science* **240** (1988) pp. 1613-1615.

165. D.W. Murphy, D.W. Johnson, Jr., S. Jin and R.E. Howard, "Processing Techniques for the 93 K Superconductor $Ba_2YCu_3O_7$", *Science* **241** (1988) pp. 922-930.

166. "Superconductivity: NASA Rises To The Challenge", *NASA Tech Briefs* (April 1988) p. 8.

167. M.A. Harris, "Brain Waves Appear on TV in Real Time", *Electronics* (24 February 1983) pp. 47-48.

168. E.R. John, L.S. Prichep, J. Fridman and P. Easton, "Neurometrics: Computer-Assisted Differential Diagnosis of Brain Dysfunctions", *Science* **239** (1988) pp. 162-169.

169. B. Jones-Leonard, K. Stengel, J. Zebrowski and L. Barnum, "Electrophysiology", *Infoworld Target Edition, Computers in Science*, No. 1 (1987).

170. I. Peterson, "Picture This", *Science News* **131**, No. 25, 392-395 (1987).

171. C. Pickover, "A Note on the Visualization of Heart Sounds", *Leonardo*, forthcoming.

172. E.E. Sutter, "The Visual Evoked Response as a Communication Channel", *IEEE/NSF Symposium on Biosensors*, CH2068-5 (1984) pp. 95-100.

173. "Mind Reading Computer", *Time* (1 July 1974) p. 61.

174. R.B. Knapp and H. Lusted, “A Bioelectric Controller for Computer Music Applications,” *Computer Music Journal* **14**, No. 1 (Cambridge, MA: The MIT Press) pp. 42-47, and see also, "Biomuse, In Search of the Brain-to-MIDI Link", *Keyboard* (June 1989) pp. 26-27.

175. For a description of the language, see Polansky, Rosenboom and Burk [101]. A new version of HMSL (version 3.40) with significantly enhanced features is now being used. See P. Burk, "HMSL V3.40 Beta Docs", CCM Technical Document (Oakland, CA: Mills College, 1989).

176. John [57].

177. John [56].

178. John [57].

179. F.O. Schmitt, P. Dev and B.H. Smith, "Electrotonic Processing of Information by Brain Cells", *Science* **193** (1976) pp. 114-120.

180. Bohm [109].

181. R.R. Llinás, "The Intrinsic Electrophysiological Properties of Mammalian Neurons: Insights into Central Nervous System Function", *Science* **242** (1980) pp. 1654-1664.

182. For a good discussion of the coevolution idea, see Jantsch [24] .

183. F. Busoni, "The Essence of Music: A Paving of the Way to an Understanding of the Everlasting Calendar", first published in *Melos* **1** (1924), reprinted in *The Essence of Music and Other Papers*, R. Ley, trans. (New York: Rockliff, 1957; and Dover, 1957 and 1987) p. 199.

184. J.A. Scott Kelso, *Dynamic Patterns, The Self-Organization of Brain and Behavior* (Cambridge, MA: The MIT Press, 1995).

185. N. Birbaumer, W. Lutzenberger, H. Rau, G. Mayer-Kress and C. Braun, “Perception of Music and Dimensional Complexity of Brain Activity”, (Urbana, IL: World Wide Web http://www.ccsr.uiuc.edu/~gmk/Papers/MusicEEG, 1994).

186. G. Mayer-Kress, “Sonification of Multiple Electrode Human Scalp Electroencephalogram (EEG)”, Poster presentation and sound demo at: ICAD’94, Santa Fe Institute, (Urbana, IL: World Wide Web http://www.ccsr.uiuc.edu/People/gmk/Projects/EEGSound, 1994).

187. D. Rosenboom, “Parsing Real-Time Musical Inputs and Spontaneously Generating Musical Forms: (Hierarchical Form Generator (HFG))”, *Proceedings of the 1992 International Computer Music Conference* Computer Music Association (San Francisco, 1992).

188. D. Rosenboom with T. Sankaran and C. Haden, “Extended Trio: Sampler (1992)”, on *Hallways, Eleven Musicians and HMSL*, CD recording, #FP002 (Hanover, NH: Frog Peak Music, 1993).

189. D. Rosenboom and A. Braxton, *Two Lines*, CD recording, #LCD3071, (New York: Lovely Music Ltd., 1995).

190. D. Rosenboom, “B.C.—A.D. and Two Lines, Two Ways of Making Music While Exploring Instability, In Tribute to Salvatore Martirano”, *Perspectives of New Music* **34**, No. 1 (Winter 1996) pp. 208-226.

191. D. Rosenboom, “Interactive Music with Intelligent Instruments—A New “Propositional Music?”, in I. Brooks, ed., *New Music Across America*, (Valencia and Santa Monica, CA: California Institute of the Arts and High Performance Books, 1992) pp. 66-70.

192. D. Rosenboom, “Propositional Music: On Emergent Properties in Morphogenesis and the Evolution of Music; Part I: Essays, Propositions and Commentaries”, *Leonardo* **30**, No. 4, (Cambridge, MA: The MIT Press, 1997) in press.

193. D. Rosenboom, “Propositional Music: On emergent Properties in Morphogenesis and the Evolution of Music; Part II: Imponderable Forms and Compositional Methods”, *Leonardo Music Journal* **7**, (Cambridge, MA: The MIT Press, 1997) in press.

194. D. Rosenboom, “Propositional Music: On Emergent Properties in Morphogenesis and the Evolution of Music; Essays, Propositions, Commentaries, Imponderable Forms and Compositional Method”, in L. Heninian, T. Ortiz and J. Zorn, eds., *Arcana: Musicians on Music* (New York: Hips Road) in press.

195. M.I. Posner, “Seeing the Mind”, *Science* **262** (1993) pp. 673-674.

196. M. Barinaga, “New Imaging Methods Provide a Better View Into the Brain”, *Science* **276** (1997) pp. 1974-1976

For additional information on the author’s work, see the World Wide Web site: **http://music.calarts.edu/~david**