Barry Truax
Riverrun

Contexts for *Riverrun*

The birth and early development of granular synthesis

The contributions of Barry Truax to the development of computer music embrace not only an extensive repertory of works extending over forty years, but also a substantive contribution to an important and now widely used digital method of generating and processing audio material, known as *granular synthesis.* An additional dimension, which will be considered further in Chapter 4 (Westerkamp), is his important role in the development of the World Soundscape Project (WSP), concerned with exploring the relationships between people and the environment. The WSP was pioneered by R. Murray Schafer, who directed the project between 1965 and 1975 at Simon Fraser University, Vancouver, drawing together a team of co-researchers, including Truax. On Schafer's departure, responsibility for continuing SFU's role in the project passed to Truax, embracing both the preservation of the existing library of source recordings and also new additions to this important resource, which has been creatively explored by several composers, including Truax and Westerkamp. This in turn led to a powerful intersection between the aesthetics of soundscape composition and Truax's pioneering techniques of synthesis and signal processing that are the central focus of this chapter.

Today the technique of granular synthesis is encountered in a variety of software tools developed for computer music composition, both commercial and noncommercial, such as Csound, Max/MSP (since 2013 marketed simply as Max), Pd, and Reaktor. These various examples range from implementations of a relatively rudimentary nature to those of considerable sophistication, albeit only fully achieved in a creative context when employed by those who understand the inherent complexities of the underlying algorithms and how best to apply them.

The concept of coding and manipulating sound information in a granular format can be traced back to the theory of communication first put forward by Dennis Gabor in 1946[1] and further expanded by Claude Shannon and Warren Weaver in *The Mathematical Theory of Communication*, published in 1949.[2] In his paper Gabor noted that traditional approaches to wave analysis involve two distinct methodologies, the first considering the evolution of composite sound waves as a function of time, and the second concerned with their spectral content, measured in terms of the constituent frequency components and amplitudes at specific instants during the time continuum. Whereas each mode of analysis involves mapping just two variables at a time, combining the results to achieve a combined perspective requires a three-dimensional matrix, which creates several practical challenges when the resulting parameters are used for further study.

Gabor sought to overcome these difficulties by proposing an alternative method of wave analysis, based on the mathematical principles associated with quantum theory. Put in simple terms, this involves digitizing an audio wave function with respect to time as a series of grains, providing successive snapshots of the associated spectra, each lasting just a fraction of a second. In analysis terms, the primary variables are the duration and frequency of each grain, the nature of the amplitude envelope that is applied to minimize any discontinuities each side of the grain, and the density of grains in terms of the degree or otherwise of overlap between successively sampled grains, or conversely the time interval between grains. The significance of these variables and the ways in which they can be used to control the processes of granular synthesis will be studied more closely in due course.

The first composer to explore the possibilities of granular principles in the composition of music was Iannis Xenakis in the context of his pursuit of the possibilities of free stochastic music. Writing in 1954 in rejection of traditional notions of linear polyphony, he observed:

> When linear combinations and their polyphonic superpositions no longer operate, what will count will be the statistical means of isolated states and of transformations of the sonic components at a given moment. The macroscopic effect can then be controlled by the mean of the movements of elements we select.[3]

Although at this early stage of study it was not yet possible for Xenakis directly to explore the creative possibilities of such principles in a digital context, the foundations for pursuing such an aesthetic were established. The writings of Xenakis on his emerging ideas during the next decade can be studied further in his book *Formalized Music: Thought and Mathematics in Composition*, first published in French as *Musiques Formelles* in 1963, and subsequently in an English translation in 1971.[4]

The following extract from the latter provides a useful starting point for a study of the characteristics of granular synthesis:

> All sound is an integration of grains, of elementary sonic particles, of sound quanta. Each of these elementary grains has a threefold nature: duration, frequency, and intensity. All sound, even all continuous sonic variation, is conceived as an assemblage of a large number of elementary particles adequately disposed in time. So, any sound complex can be analyzed as a series of pure sinusoidal sounds, even if the variations of these sinusoidal sounds are infinitely close, short, and complex. In the attack, body, and decline of a complex sound, thousands of pure sounds appear in a more or less interval of time, Δt. Hecatombs of pure sounds are necessary for the creation of a complex sound. A complex sound may be imagined as a multi-colored firework in which each point of light appears and instantaneously disappears against a black sky. But in this firework there would be such a quantity of light organized in such a way that their rapid and teeming succession would create forms and spirals, slowly unfolding, or conversely, brief explosions setting the whole sky aflame. A line of light would be created by a sufficiently large multitude of points appearing and disappearing instantaneously.[5]

These visual analogies provide a useful gateway into a deeper understanding of the nature and function of acoustic grains. Without access to computing resources during this formative period of study, Xenakis had to devise purely analog means of generating such quanta. Indeed, the expected correlations here are not always immediately evident, since in many works it is the theories of granulation that inform the musical style of writing rather a more direct algorithmic implementation. The works from this early period that perhaps become closest to such an implementation are *Analogique A* (1958), scored for nine string instruments, and *Analogique B* (1959), an entirely electronic work. In the case of *A* the acoustic quanta comprise a blend of pizzicato and very short bowed sounds, whereas in the case of *B* they are composed of very short synthesized sine waves, extracted and juxtaposed using tape editing techniques.

A somewhat more refined and perhaps better-known electronic work from this period that makes extensive use of granulated materials is *Concret PH* (1958), which he composed for the Philips Pavilion constructed for the World's Fair held in Brussels the same year. Here the sound source is that of burning charcoal, where by extraction the individual crackles are manipulated in a granular fashion. The interest stimulated by this imaginative project led to the granting of access to an IBM 7090 computer, courtesy of IBM-France, providing him finally with the resources necessary to develop digital algorithms for generating performance data, and a resulting series of stochastic works, notably *ST/10-1, 080262* and *ST/10-3, 060962* (*Atrées*), both composed in 1962. In so doing he laid important foundations for the all-digital implementations of granular synthesis that were to follow.

Two composers are to be credited with pioneering the techniques of granular synthesis as we know them today, both inspired by the work of Xenakis; Curtis Roads and Barry Truax. Roads became interested in the possibilities of microsound during the early 1970s, his interests in the electronic medium being stimulated by an opportunity to access the studio resources at the University of Illinois Experimental Music Studio via a graduate student friend. At the time (1970–1971) this was a purely analog studio, based on a large Moog III modular synthesizer, an API mixing console, and a bank of analog tape recorders. Although he became aware of the functional characteristics of Music V, and managed to obtain a printout of the associated code, he was unable to access suitable computing facilities for exploring its creative possibilities.[6] As a first step in this direction, having commenced studies at Cal Arts in Los Angeles in 1972, he was able to install a version of this program on a Data General Nova 1200 computer. However, this computer had no audio converter and thus could not produce sound.

His creative interests at this time were strongly stimulated by the characteristics of rock-based styles of composition, reflecting his essentially non-classical musical background. This aesthetic perspective was to prove highly influential in shaping his initial engagement with the creative possibilities of computer music, aspects of which were to become key features of his subsequent repertory of works. During the early years he was especially interested in exploring the characteristics of micromanaged rhythm, and this led him to engage directly with a new generation of artists furthering the techniques of electronica. In due course these developments influenced key genres such as intelligent dance music (IDM) during the 1990s, associated in turn with works by artists such as Aphex Twin, Autechre, and The Orb. Roads indeed collaborated with the British duo Autechre for two concerts, the first in Los Angeles in 2001, and the second in London in 2002.

Roads first became aware of the possibilities of granular synthesis in May 1972 when attending a course on formalized music at Indiana University, conducted by Iannis Xenakis, marking the start of an association that was to prove materially significant for him. While subsequently studying at the University of California at San Diego (UCSD) in 1974–1975, he was able to access a more substantial Burroughs B6700 mainframe computer. This machine provided the facilities necessary for composing computer music via Music V, the audio data files generated by the latter being subsequently realized acoustically via a twelve-bit digital-to-analog converter operating at 20 kHz. Using the programming language ALGOL, Roads developed a free-standing granular synthesis program known as PLFKLANG, used to generate note/event data, passed in turn to a Music V program for audio synthesis. The converter, however, could only operate in mono, requiring the subsequent use of analog mixing facilities to mix the component tracks necessary to create a stereo image. This facility resulted in his first granular study, *Prototype*, completed in 1975. The functional characteristics of his program merit further scrutiny, since they illustrate well the key principles of this synthesis technique.

The waveforms used by Roads for producing grains ranged from a simple sine wave to a band-limited pulse, thus establishing a palette of sources that embraced different degrees of sonic complexity. Each of the grains was of a fixed duration (twenty milliseconds), in turn subject to a simple amplitude envelope with a sustained peak. The permissible frequency range was between 40.4 Hz and 9,900 Hz, and the permissible amplitude range was between 30 dB and 70 dB. Control of the synthesis processes was regulated by the following set of variables:[7]

1. Beginning time and duration of an event
2. Initial waveform and waveform rate-of-change (slope)
3. Initial center frequency and rate-of-change of center frequency
4. Initial bandwidth (frequency dispersion) and bandwidth rate-of-change
5. Initial grain density and grain density rate-of-change
6. Initial amplitude and amplitude rate-of-change

The basic data mapping technique can be illustrated in a graphic format, as shown in Figure 2.1.

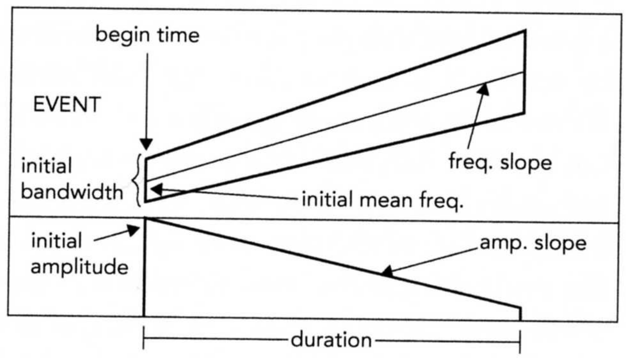

FIGURE 2.1

Variables for the control of PLFKLANG granular synthesis.[8]

Aside from the possible limitations of a fixed duration for each grain, the one drawback to this early implementation was that the system was non-real-time, requiring all the performance data to be provided in advance. Furthermore, this restrictive working environment, well known to the early pioneers of computer music, further isolated the composer from the processes of sound production. All programs and associated data had thus to be prepared using punched cards, subsequently placed in a queue for batch processing, where tasks that were least demanding in terms of computer resources were routinely given priority. Consequently, as noted in the previous chapter, sound synthesis applications would invariably be shunted to the back of the queue for overnight processing. Notwithstanding these challenging conditions, especially in the context of such a novel synthesis technique being explored for the very first time, Roads was able to make further progress with his pioneering software, producing a more sophisticated version of his program in 1981 that took advantage of a significantly improved array of digital-to-analog converters, offering up to four channels of audio with sixteen-bit resolution at a sampling rate of 40 kHz.

With the new decade the development of more affordable minicomputers opened the possibility of developing dedicated facilities for computer music, allowing for a much faster turnaround of synthesis tasks. However, as noted in the previous chapter, it was to be several years before such resources and their desktop-based successors could provide the processing power necessary to support real-time implementations of programs such as Music 11 and Csound. It was thus left to pioneers such as Roads and Truax to seek other ways of achieving such goals in the context of generating and manipulating granular sounds.

Further work developing the techniques of microsound led Roads to Paris in the early 1990s, and in 1993 he took up an appointment as director of pedagogy at the center established by Xenakis in 1966 for research and development in the field of stochastic music. Known originally as Centre d'Études de Mathématique et Automatique Musicales (CEMAMu), and subsequently as Les Ateliers UPIC, it was finally renamed Center for the Composition of Music Iannis Xenakis (CCMIX) in 2002. This direct association with the work of Xenakis acted as a major stimulus for the development of Cloud Generator, a granular synthesis program for the Apple Macintosh microcomputer, which he developed in collaboration with John Alexander and completed in 1995. Written in the programming language C, and with the enhancement of a graphical user interface, this versatile resource was to provide the kernel for Roads and others to extend the possibilities of the technique, in due course embracing a fully interactive real-time working environment.[9] The initiative for crossing this threshold, however, is to be credited in the first instance to Barry Truax, whose innovative work is the central focus of this chapter.

Truax's initial training was in physics and mathematics, studying at Queen's University, Ontario. However, childhood interests in music were rekindled during his undergraduate studies by coming in contact for the first time with synthesizers developed by Hugh Le Caine at Expo '67, Montreal.[10] After graduation, in 1969, matters came to a head for him, and instead of pursuing a purely scientific career, he chose instead to explore the musical applications of technology, transferring to the University of British Columbia, Vancouver, as a postgraduate student to study electronic composition. In 1971 Truax left his native Canada for a two-year period of study at the Institute of Sonology in Utrecht, bringing him into direct contact with the work of Gottfried Michael Koenig. During the 1960s Koenig had become interested in the creative possibilities of algorithmic methods of music production, developing ideas that were closely related to those of Xenakis.

Toward the middle of the decade, he started work on the first of two composing programs, PROJECT 1 (1964), soon to be followed by PROJECT 2 (1966). PROJECT 1 was inspired by his interests in the techniques of serial composition and their possible extension through the processes of automation, the associated software producing data for the construction of instrumental scores. The data parameters for each note thus produced are the choice of instrument, entry delay (= metric duration), pitch, octave register, and the dynamics. In the case of the pitch parameter, three-note groupings are generated from two source intervals provided by the composer, subsequently transposed to create a twelve-note row.

A key component of the data generation process is the use of random probability in association with a rule-based algorithm, the latter being chosen from a repertory of seven possibilities of varying orders of complexity. As was to be expected at the time, the development of the program required the services of the University of Utrecht's mainframe computer, an IBM 7090, written initially in ALGOL 60 and subsequently rewritten in an early version of the FORTRAN programming language, FORTRAN II. Tasks (or "jobs," as they were commonly called) had to be submitted for batch processing in the manner described above, and although as data generation programs they were not subject to the same order of delays as those producing fully synthesized audio data, the waiting times were still significant.

The same conditions applied to the early development of PROJECT 2. The latter program took these generative processes a stage further, expanding the choice of parameters to include modes of articulation and an expanded set of rhythmic components, including rests and tempi. The format of output data was also expanded to embrace groups of data of variable content in terms of the note/events and ensembles, created from selections of one or more of these groups. From a compositional point of view, the most significant difference is the change from an environment where score data is generated section by section to one where the processes are combined to create a generative model that is controlled by the composer in a "top-down" manner, changing parameters that influence the overall evolution of the associated algorithms but not directly determining the specific details of the resulting data.

The control functions provide different types of weighting to the selection procedures that determine the content of each group. The function called ALEA makes entirely random selections from a table of values, each element once chosen remaining eligible for reselection, whereas the function called SERIES removes values from the table after selection to prevent repetitions occurring. When all the available elements have been used up, this function may then be reprogrammed using a fresh copy of the table. Alternatively, the function RATIO makes weighted selections from an associated table. A different type of selection procedure is provided by a function called TENDENCY. This allows the composer to apply "masks" to the random-number generator, dynamically adjusting the range of values from which it may make choices.

In 1970 the Institute of Sonology purchased a PDP-15/20 computer with funds provided by Utrecht University. This was a notable development in the European history of computer music, matched at the time only by the acquisition of a PDP-15/40 computer for the Elektronmusikstudion (EMS Stockholm), funded by Swedish Radio in the same year, and two small PDP-8 computers acquired by Peter Zinovieff a year earlier for his private studio in London, to support his commercial company, EMS London.[11] In all three cases the primary purpose of these computers was to control synthesis hardware, leading to their classification as hybrid systems. In the case of the Institute of Sonology, however, the acquisition of a dedicated computer also stimulated further work on new versions of both PROJECT 1 and PROJECT 2, rewritten in FORTRAN IV. With the elimination of the delays routinely associated with batch processing, the time between submitting tasks and receiving the output data on a printer was dramatically shortened from several hours to a matter of minutes. This in turn raised the possibility of directly generating sound output from the computer via suitably designed digital-to-analog converters.

During 1972 Koenig and Stan Tempelaars developed the first version of the Sound Synthesis Program (SSP). In so doing they came directly into contact with the constraints of computers in terms of such demanding applications. Although the PDP-15 was a very new design, taking full advantage of the most recent advances in digital engineering, as one of the first generation of smaller computers it had only limited workspace—just 12K of eighteen-bit core memory.[12] Furthermore, the Institute of Sonology computer had no mass data storage facilities, significantly limiting the use of dynamic memory management techniques to expand this workspace. SSP was thus limited to the production of simple fixed format waveforms, producing a mono output, restricted in turn to the parameters of time, pitch, and amplitude supplied from associated data tables.[13] Although the software was capable of generating sound output in real time (that is, without an intermediate delay once a synthesis task was launched) it did not allow any direct composer interaction during execution. Any notion of performance control therefore seemed as remote as ever. With the arrival of Barry Truax, this Rubicon was finally to be crossed.

Truax's compositional imperatives on arriving at the Institute of Sonology were visionary in several important respects. Although well aware of the steady progress being made in the context of non-real-time computer music software such as the Music N series of programs in terms of an expanding repertory of synthesis and signal processing techniques, his creative imperatives were underpinned by a profound belief that the composer should be able to work with such resources in a truly interactive environment, directly influencing the processes of sound production. In an article subsequently published in the *Journal of Music Theory*, he argued the following:

> 1. That all computer music systems both explicitly and implicitly embody a model of the musical process that may be inferred from the program and data structure of the system, and from the behavior of users working with the system. The inference of this model is independent of whether the system designer(s) claim that the system reflects such a model, or is simply a tool.
>
> 2. That computer music systems form a source of data for the theorist in the form of program/system analyses, and records of user behavior (viz. protocols). In addition, such systems offer the facility to test (and therefore compare) theoretical models or concepts.
>
> 3. That real-time, interactive systems, together with composition programs of all kinds, are the most useful (i.e. observable) sources of data concerning musical activity and its modeling.[14]
>
> 4. That systems of any type will increase in their usefulness to the composer to the extent that they implement procedural musical knowledge and offer a significant learning potential for the user.[15]

The importance of the third precept to Truax, from the very outset of his career as both a composer and researcher, should not be underestimated. Nor for that matter should his concern to understand a consideration that lies at the heart of the research that has underpinned the composer studies that form the kernel of this book, which is the ways in which a composer can engage with composing software, and the influence of the latter on how it may be used:

> How much of the user's strategy remains external to the program? How much of the process must be complete before the program is used, and how much continues after it stops? To what extent, then, does the system determine the user's strategy, and to what extent does it assist the user in formulating that strategy? Important for the user will also be the time taken before the results are audible, that is, before the feedback of information allows the user to modify his strategy and come closer to his goal.[16]

In rejecting from the outset the non-real-time environment of the Music N family of software, and opting instead for the rule-based approach being advocated by Koenig and others at the Institute of Sonology, he had to engage head on with this important issue in ways that were not being considered by advocates of the former approach. In situations where the processes of composition are so far removed from those of synthesis, the role of the computer is essentially passive, merely responding to a remotely entered set of instructions. This sense of detachment is further reinforced by the requirement to specify tasks in terms of a predetermined "orchestra" of individual "instruments," realizing a succession of performance events supplied as the associated data in an accompanying "score."[17] While such a classical approach is not without its merits, and indeed served Max Mathews and his associates very well during the formative years of computer music, such a working environment was far removed from that being sought by Truax.[18]

Whereas PROJECT 1 and especially PROJECT 2 were to prove the primary starting point for his investigations, Truax was also very much influenced by the prior work of Xenakis, in particular the concepts that had been embraced in his ST (= Stochastic) series of programs, discussed earlier. The characteristic of particular interest in this context was his use of probability laws to determine the evolution of sonic entities, notably those based on the distribution theory first advocated by the French mathematician Siméon Denis Poisson in the nineteenth century. For Truax the use of Poisson-ordered distributions of discrete sound events, each of a fractional duration, provided the basis for a novel means of composing directly in sound, particularly one that could not only generate the results in real time but also be interactively controlled. In the latter context he was able to draw on some key design principles associated with PROJECT 2, notably the use of "tendency masks," to regulate the frequency boundaries of the microsound events with respect to time. Figure 2.2 illustrates the general principles involved in the design of a mask.

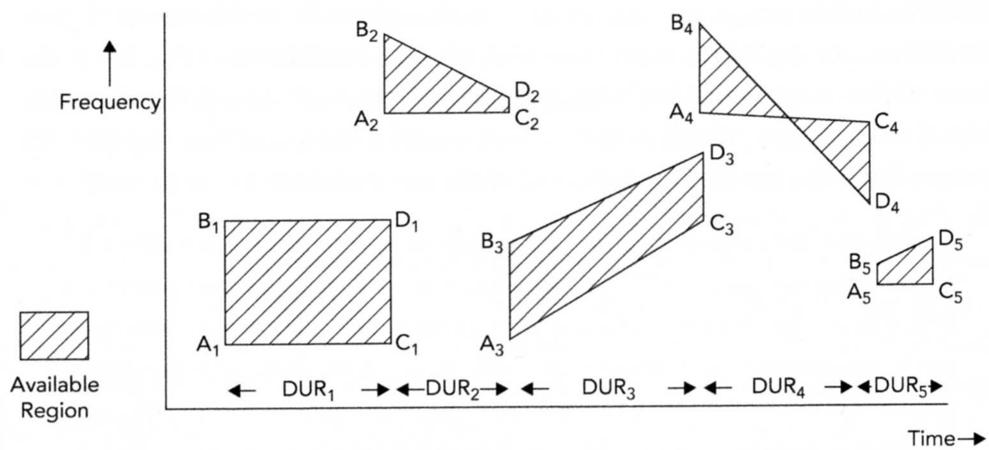

FIGURE 2.2

Tendency masks for the frequency of grains over time.[19]

The numeric data for each mask, using an integer format, consists of the following components: (i) the number of mask segments (maximum ten), (ii) the total number of events to be mapped in the distribution, (iii) the initial density in sounds per second, and (iv) the mask data for each segment (the values of A, B, C, and D in Hertz plus the duration in hundredths of a second).

The articulation of the microsound events within each segment is thus a function of their density and the pitch at which they are reproduced, the latter values being determined by random selection within the frequency boundaries prescribed by the tendency mask at the corresponding instant in time. Given the already mentioned restrictions on real-time computation, the audio output was necessarily monophonic, the density characteristic being simply a measure of duration of each component event. In this context Truax was to note:

> The Poisson distribution applies in cases where the density of events is sufficiently low (less than ten to twenty events per second), such that separate events may be identified, since a large density of events merges into a continuous distribution, just as high sound density results in a fusion of events into a continuous texture.[20]

This point of fusion was clearly regarded as an important boundary in operational terms. Proceeding in the opposite direction, techniques of granular synthesis, such as those that subsequently were to form the basis of his programs GSX and GSAMX, extend from the initial production of composite spectra to recognizable grains that can be individually perceived as discrete events. There are thus important similarities between the two methods of generating microsound, and this connectivity extends to the nature and content of the sounds themselves.

Truax developed three versions of his real-time synthesis program at the Institute of Sonology, known respectively as POD4, POD5, and POD6. Given the limited amount of memory available in the computer and the lack of mass data storage facilities, significant ingenuity was required to generate the constituent waveforms. The PDP-15, however, provided a limited but nonetheless useful amount of local data storage via a proprietary system known as DECtape, consisting of a pair of small magnetic tape spools that could be used to store waveform data. In the case of POD4, ten fixed waveforms of fifty samples each were provided, with the option of constructing a further eight derivatives that were user-defined. In the latter context the ability to audition the resulting waveforms in real time allowed instant evaluation and, where appropriate, modification of the sample selections made to construct these derivatives. In the case of POD5, a slightly different approach was taken. Instead of five times fifty sample versions, a single five-hundred sample waveform was used in association with an array of sixty different amplitudes that were in turn dynamically scaled. This technique made it possible to create smooth attack and decay envelopes for the resulting microsound streams. The extended sample length also facilitated the construction of a series of harmonic derivatives by cyclically reading every other, every third, or every fourth, etc., sample. It also allowed the construction of more elaborate user-defined waveforms and the use of amplitude modulation.

The control of frequency was equally ingenious, circumventing all the calculation overheads that are necessary when producing sound at a fixed sample rate. Indeed, had this method not been possible the POD programs could not have operated in real time. The solution came with the ability directly to microprogram the operation of both the real-time computer clock and the cycle time of the processor. Via a combination of repeated cycles of both timing variables, it became possible to produce a logarithmic table of pitch values with a resolution in excess of thirty steps per octave, in the case of POD5 working with a lower limit of 50 Hz and an upper limit of 8,000 Hz.[21] Most significantly, the technique allowed the normal correlation between speed and pitch, such as that observed when speeding up or slowing down the playback of an analog tape recording, to be completely decoupled. Such a facility was way ahead of its time, anticipating processing features that were to be embraced by more advanced implementations of techniques such as granular synthesis in more conventional processing environments.

Although the methods of waveform production described above were of primary importance for both POD4 and POD5, both programs allowed the use of externally generated waveforms, sampled via an analog-to-digital converter. Although the full potential of such a facility was not to be realized until more than a decade later in the context of GSAMX, the possibilities of thus embracing the subtleties of acoustically generated sounds were not lost on Truax. His ultimate quest at this time, however, was to embrace a technique that would allow the production of sounds that could vary dynamically in real time, not merely in terms of amplitude but also timbre. The solution here was to be provided by John Chowning, following a visit to the Institute of Sonology in 1973, when he introduced Truax to the principles of FM synthesis. Consequently, the first real-time implementation of this important technique came into being as POD6, requiring simple changes to the underlying processes of computation that in turn were to lead directly to the development of GSX and the composition of *Riverrun*.

As described in Chapter 1, all that was required in this context at the most basic level was a single sine wave providing a carrier frequency "c," in turn modulated by another sine wave of frequency "m" and amplitude "d," where the modulation index "I" is the ratio of "d"/"m". In this context the sine wave was simply stored as a table function of 512 values, from which both modulation components in terms of their frequency and amplitude can be dynamically extracted in the manner described for POD4 and POD5. Up to thirty different pairs of FM-configured oscillators could be used at any one time, each controlled in addition to the values of the primary FM parameters in terms of the amplitude envelope and a maximum permitted value for the modulation index.[22] Truax discovered that significant subtleties in the resulting timbres could be unmasked by randomly distributing the sounds between two mono channels. By varying the relative amplitudes and binaural time delays fed to each loudspeaker, each component could be positioned at one of seven locations between the two extremes. Such techniques were to become a major feature of his granular synthesis programs, illustrating yet again the important links between the POD programs and both GSX and GSAMX.[23]

Overall control over the selection of sound object data that comprises all timbral values is exercised via one of four statistical processes, known respectively as ALEA, RATIO, SEQUENCE, and TENDENCY. These were originally developed for PROJECT 2 by Koenig, although in the latter context they were used to generate score data, whereas in the case of the POD programs they are used to select the individual audio events. ALEA is an equal probability function making an aleatoric selection from the waveforms that have been assigned as available. RATIO is a weighted probability function based on integer ratios provided by the composer as applicable for the selection process. SEQUENCE is a fixed sequence of choices made in accordance with an ordering provided by the composer in advance. In this case there is no random probability component to the selection process. TENDENCY is a tendency mask very similar to that used for the Poisson-ordered distribution of event frequencies and densities described earlier, combining both the waveforms available at any one point and the weighting ratios to be applied to their selection (see Figure 2.3).

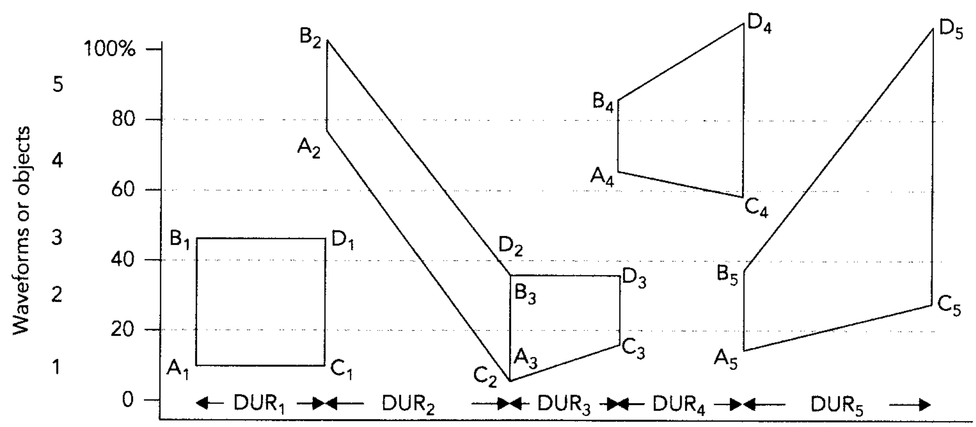

FIGURE 2.3

Tendency masks for available waveforms and their weighting ratios.[24]

In this context the vertical coordinates are expressed in percentage terms, since the total number of waveforms available will be determined in advance by the composer and vary from context to context.

In 1973 Truax returned to Vancouver, taking up a post combining teaching and research in the newly created Department of Communication Studies at Simon Fraser University.[25] In the latter context his primary interests focused around the further development of his work at the Institute of Sonology, and he was immediately faced with the challenges of modifying his POD programs to run on a different computer. Although the high-level commands were written in FORTRAN, for which a compatible compiler was available, the time-critical synthesis routines had been written in low-level assembler code, specific to a PDP-15. During 1974–1975 he modified the programs to run on a Hewlett-Packard 2116 machine, taking advantage of the slightly larger memory (16K of sixteen-bit memory rather than 12K of eighteen-bit memory) and, most importantly, both a disk unit and a mass storage magnetic tape unit. In 1978 he transferred his work to a NOVA 3 computer with non-real-time synthesis on a Varian computer located in Computing Science, and finally in 1981 to a PDP (LSI) 11/23, the first computer that could be dedicated to his personal research.[26]

A primary concern during this transitional period was the enhancement of POD with the ultimate objective of facilitating real-time generation of polyphonic textures. Unfortunately, the computer technology available at SFU still had some way to go before such a goal could be realized. A partial solution was identified from an ongoing collaboration with EMS Stockholm, where a near-identical PDP-15 computer could be used to operate a bank of digitally controlled analog FM synthesis oscillators. Truax was thus able to produce a hybrid polyphonic version of POD6 that could work in real time at EMS. However, back at SFU, his only option was to develop a non-real-time version, known as POD7. In accepting the consequential time delays as individual voices were synthesized and incrementally recorded, he developed an interactive audio file management system that allowed enhanced facilities for producing synthesis data and interactive feedback once the results were computed (see Figure 2.4).

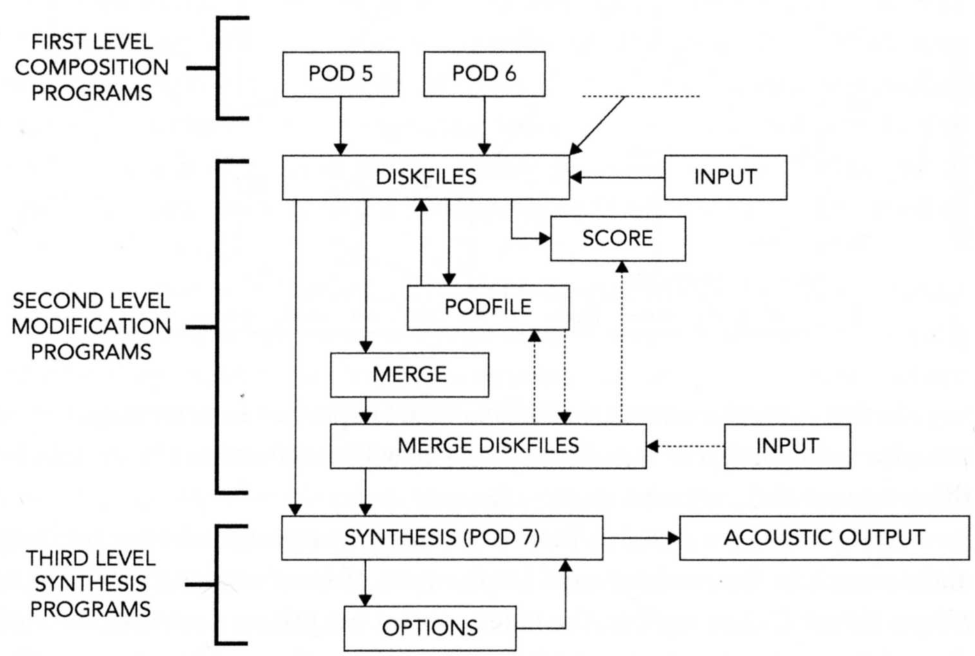

FIGURE 2.4

Program structure of the POD system.[27]

Having accepted the necessity of changing to a non-real-time synthesis environment, Truax used the opportunity to enhance the quality of the sound output, concentrating now almost exclusively on the FM synthesis technique introduced with POD6. The new program, for example, used a wave table of 8,192 discrete values, a substantial improvement on the 512 values of the earlier programs.[28] It also facilitated for the first time the direct generation of spatially distributed sounds within a stereo listening field.

By the end of the decade the production of POD-generated works was gathering pace, including *Sonic Landscape No. 3* (1975), *Sonic Landscape No. 4*

(1977), *Androgyny* (1978), and *Aerial* (1979) composed by Truax, and works by other composers who had opportunities to use the POD software, including *La mer à l'aube* (1977) and *Heliograms* (1977–1980) by Jean Piché.[29]

Notwithstanding the additional composing facilities provided by POD7, Truax still yearned for a system that could synthesize polyphonic voices in real time. The breakthrough came with the acquisition in 1981 not only of the PDP-11/23 computer, but also, a year later, a fast front-end digital signal processor, the sixteen-bit DMX-1000, manufactured by Digital Music Systems and first released commercially in 1979. The latter device is distinctive in several respects. First, it was designed specifically for synthesizing and processing audio data. Second, it used two high-speed memories operating in parallel. One memory stored up to 256 programming instructions, the other stored the associated data, both under the management of the host control computer. Third, it was microprogrammable, allowing the use of fully optimized signal processing instructions. The basic machine instruction cycle was just 200 nanoseconds, and this speed allowed audio processing several orders of magnitude faster than that possible with a PDP-15 or PDP-11 computer.[30] The memory stored a maximum of 4K words, and this restriction had implications for the storage of wavetables. This limitation, however, was significantly offset by the superior performance of the processor.[31]

The DMX-1000 was shipped with a software synthesis program known as Music 1000, developed from Music 11, the version of Music N written by Barry Vercoe at MIT specifically for the PDP-11. Although it was possible to execute the program in real time, it still required the preparation of a traditional "orchestra" and "score" in advance, with no opportunity for live interaction. For Truax the attraction of the DMX lay in the power and versatility of its signal processing facilities, finally making it possible for POD programs to generate polyphonic textures in real time. Whereas the control environment, written using a high-level compiler, could be hosted by the PDP-11/23 computer with relative ease, the synthesis kernel had to be entirely rewritten in DMX-1000 microcode.[32] The rewards for this extra work were considerable. For example, a reworking of POD6 as POD6X facilitated the generation of six voices of FM simultaneously in real time. In a similar vein, a similar reworking of POD7 as POD7X created a highly responsive system for creating, storing, verifying, and merging sound data in real time, in the latter context using a subroutine known as CONDUC.[33]

The early 1980s was to prove an important watershed in the development of Truax's POD-based compositions, starting with his four-channel work *Arras* (1980). His growing interest at this time in exploring the timbral possibilities of polyphonic FM textures required extensive use of POD6 and POD7, building up the textures layer by layer. The first compositions to be realized using the new DMX-1000-based system were *Wave Edge* (1983), followed by *Solar Ellipse* (1984–1985), two FM works inspired by the Chinese

Book of Changes, more generally known as *I Ching*.[34] What amounted to a quantum leap in terms of the available computing power, however, was leading to a reappraisal of the techniques being used for real-time synthesis, not least in terms of the range and versatility of the tools that could now be provided for interactive composing.

It has already been noted that a key feature of the POD systems was their engagement with the creative possibilities of microsynthesis, building up textures from aggregates of individual sonic elements. Hitherto, as noted earlier, Truax had regarded the point where these elements cease to be recognizable as discrete events as a boundary in terms of their minimum duration that could not usefully be crossed. A growing awareness of the underlying significance of Gabor's "Theory of Communication," however,[35] and the ways in which it usefully underpinned the principles of granular synthesis being explored by Roads, led him to look more closely at the way in which POD generated microsounds.[36] He concluded that the POD system of generating grains was not capable of managing a truly granular environment. Hitherto the restrictions had been a matter of necessity, since in a pre-DMX-1000 era it was impossible to provide the necessary dynamic control of individual grains at densities that exceeded these boundary conditions. The aural effect was thus one of essentially static textures of limited musical value. Now this was no longer the case, and it became possible to investigate the possibilities of building up textures from aggregates of individual sonic elements lasting just a fraction of a second, first realized with *Solar Ellipse*.

The key to overcoming this hurdle lay in harnessing the significant increase in computing power provided by the DMX to control the generation of grains at much higher densities, combining and overlaying several granular streams or voices simultaneously. Furthermore, a key requirement was that the improved system also had to operate in real time, allowing the composer to interact directly with the parameters controlling the processes of synthesis. Two versions of his new granular software were developed during the mid-1980s, GSX and GSAMX, supported in due course by GRMSKX, an additional control facility allowing the implementation of tendency masks in a manner similar to those offered by POD. GSX allows two basic modes of synthesis, fixed waveform and frequency modulation, whereas GSAMX uses sampled sounds as the source material.

Although the techniques of granulation used for the GS programs share features in common with those of Roads, the control environment develop several characteristics a stage further. The minimum duration of individual grains that could be controlled by the PDP-11/23 was approximately eight milliseconds (ms) per voice (125 grains per second), the number of simultaneous voices being limited by the speed of the DMX-1000 synthesis engine. In the case of fixed waveform and sampled sound synthesis, a maximum of twenty voices can be produced, resulting in in a total grain density of 2,500 grains pe second. In the case of FM synthesis, the maximum number of voices reduces to eight, with a corresponding grain density of 1,000 grains per second. From the composer's point of view, it is the nature of the grains for individual voices that are of paramount importance, and it is from this perspective that the key operational characteristics of the programs will now be considered in more detail.[37]

The following basic parameters define the nature of each grain:

1. The waveform of the grain
2. The amplitude of the grain[38]
3. The frequency of the grain
4. The frequency range of the grain
5. The duration of the grain
6. The duration range of the grain
7. The delay between the end of the grain and the start of the next grain

The minimum duration of each grain, as noted above, is 8 ms, and the maximum is 120 ms, whereas the delay between grains can be varied between 0 ms and 120 ms. It is interesting to compare the upper limits of these two parameters with the earlier cited POD limits of approximately no more than ten to twenty events per second for each event to be individually perceived. The maximum grain duration of 120 ms corresponds to slightly more than eight events per second, well within this perception boundary. If an element of delay between grains is also introduced, this process of differentiation becomes even more marked. In theory there is no reason why these two limits should not be extended further; however, for all intents and purposes, the 120-ms limits adopted by these programs for these duration parameters are more than adequate for the purposes of granular synthesis. The ability to modulate both the frequency and the duration of each grain using an associated random probability algorithm is a powerful enhancement to the granulation processes, the significance of which will be returned to and discussed in more detail later.

One further key parameter that must be set at the outset is the linear amplitude envelope for each grain. In place of the fixed envelope used by Roads, Truax uses symmetrical attack and decay envelopes that can be varied between 1/2 the duration of each grain and 1/16 the duration of each grain, with a default value of 1/4. In practical terms, sharper envelopes may produce audible sidebands according to context, which can be usefully manipulated if so desired by using gentler envelopes. Figure 2.5 shows how Truax illustrates these basic grain characteristics.

In the case of sampled sounds, two approaches were adopted. The first required the sound extract to be loaded into the memory of the DMX-1000. This, however, could only store a maximum of 4K words, a total of 4,032 samples. Accordingly, the maximum sample length was limited to approximately

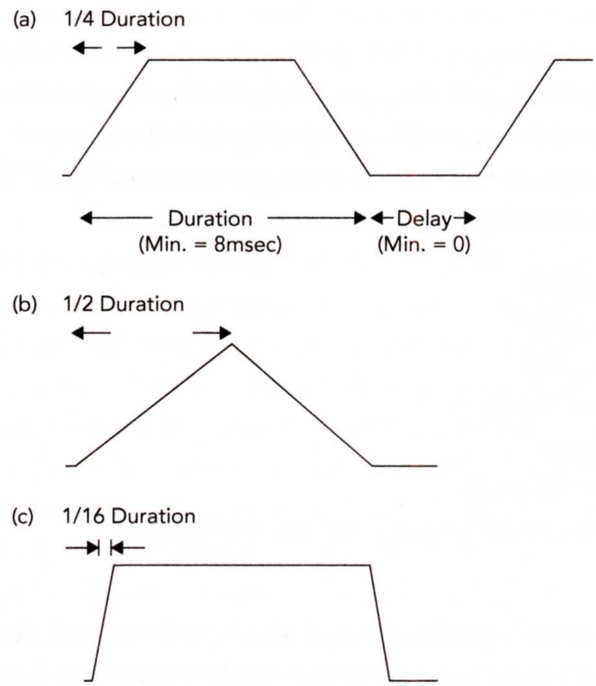

FIGURE 2.5

Grains with variable attack and decay envelopes.[39]

150–170 ms, thus allowing only very brief fragments of externally sampled material to be granulated. Notwithstanding the practical restrictions thus imposed, Truax successfully used this mode of sampling for his first work to use GSAMX, *The Wings of Nike*, completed in 1987. This was composed entirely from granulations of just two phonemes. The second approach, subsequently implemented, was altogether more satisfactory, providing the basis for the extensive repertory of works based on sampled sounds that was to follow. In place of the fixed preloaded sound extract, he was able to use an external hard disk drive for the PDP-11/23, offering a maximum capacity of five megabytes, which even at a sampling rate of 40 kHz could accommodate sounds lasting more than a couple of minutes. In this arrangement the internal memory of the DMX-1000 is constantly updated with new data from the disk, using internal buffering to maintain a smooth flow of audio for granulation.

There are some additional parameters that are specific to the mode of granular synthesis. In the case of fixed waveform synthesis these include the number of voices for each of the three possible waveforms (chosen from a library file of twenty). In the case of FM synthesis these extend to the average modulation index (the basic c/m ratio is specified separately) and index range, and the number of voices (maximum eight). In the case of sampled synthesis, in place of a frequency parameter an offset value relative to the start of the sampled sound and a speed of output parameter (default value one) that provides pitch/time transposition must be specified, along with the number of voices (maximum eight). In terms of the parameter variables that can be preprogrammed in advance using GRMSKX, these are limited in the case of tendency masks to frequency (fixed waveform synthesis or FM), offset (in the case of sampled synthesis), duration, modulation index (FM only), and, in the case of envelopes, amplitude and delay time.

In practice, given the significantly enhanced real-time control environment, Truax only made limited use of GRMSKX, choosing instead to use an alternative interactive performance control facility using ramp files to control the rate of change of parameter settings within user-specified limits.[40] The characteristics of this latter mode of control feature strongly in the work that will now be studied in depth, *Riverrun*, composed in 1986. This was the first work to be composed using GSX, and the analysis that follows provides a detailed introduction to Truax's work with granular synthesis, providing also a useful basis for further study of these techniques in later works, including those composed using GSAMX.[41]

Inside *Riverrun*

Riverrun: Technique, form, and structure

In investigating *Riverrun* in more detail, we will first look at the overall structure of the work. This is followed by an examination of the granular synthesis techniques and the GSX software that are so important to the work. Finally, we will look at how these are used to shape the music. Throughout, the text will be enhanced by the accompanying software, which contains video recordings of the composer explaining his approach and demonstrating the original GSX system in his studio, along with interactive explorers we created for this project. The explorers provide you with the opportunity to try out the techniques for yourself to deepen your understanding of their potential and the role they play in *Riverrun*. You are encouraged to download the software, as this facilitates full access and interactive engagement with these materials. Short video demonstrations of the software in use are also provided to give a quick taste of what the software has to offer.

The musical shape and structure of *Riverrun* are intimately connected to the technical means of its production—the granular synthesis techniques devised by the composer and realized using the GSX software that is employed exclusively in the composition. The work is constructed from a series of grain streams generated one at a time using the software and then superposed (using an analog multitrack tape recorder) to form sections of the work. The composer's own documentation of the work identifies five sections in *Riverrun*, and these are clearly identifiable aurally in the work as well as from examining the source materials (published by the composer on DVD and generously made available by him for use in this project).[42] There are occasionally brief overlaps from one section to another, but nonetheless the section breaks are always clearly delineated.

A distinctive feature of this work, relating both to its aesthetic and its means of production, is the fact that the work is made up out of relatively few