About *Changes: Sixty-four Studies for Six Harps*

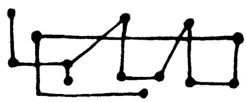

James Tenney

A. Introduction

My intentions in this work were both exploratory and didactic. That is, I wanted to investigate the new harmonic resources that have become available through the concept of "harmonic space" much more thoroughly than I had in any earlier work. At the same time I wanted to explore these harmonic resources within a formal context which would clearly demonstrate certain theoretical ideas and compositional methods already developed in my computer music of the early 1960s, including the use of stochastic (or constrained-random) processes applied to several holarchical perceptual levels, both monophonically and polyphonically. The references to the *I Ching* or "Book of Changes" (in the titles of the individual Studies) derive from correlations which were made partly for poetic/philosophical reasons, but also—and perhaps more importantly—as a means of ensuring that all possible combinations of parametric states would be included in the work as a whole. I must confess that I frequently thought of the twenty-four Preludes and Fugues of J. S. Bach's *Well-tempered Clavier* as a kind of model for what I wanted to do with the work, although it seems highly unlikely that these Studies themselves will ever betray that fact to a listener. A large mainframe computer was used in the composition process, generating coded numerical output which was then transcribed into standard musical notation. Two separate FORTRAN IV programs were involved, the first dealing with characteristics of the set of sixty-four Studies as a whole, the second determining the details of each individual Study.

B. General Features

The harps are tuned a sixth of a semitone (16.66 . . . cents) apart, so the ensemble is capable of producing a tempered microtonal set of seventy-two pitches in each octave. This tuning system (which I call the 72-set) provides very good approximations of most of the important just intervals within the 11-limit, with the worst case being the three-cent error for the 5/4 major third (383¢ instead of 386¢). The relations between some of these just intervals and their nearest approximations in the 72-set are shown in Table 1 (where interval sizes are rounded off to the nearest cent).

ratio	size	pc number in 72-set	size	error
8/7	231¢	14	233¢	+2¢
7/6	267¢	16	267¢	±0¢
6/5	316¢	19	317¢	+1¢
11/9	347¢	21	350¢	+3¢
5/4	386¢	23	383¢	−3¢
9/7	435¢	26	433¢	−2¢
4/3	498¢	30	500¢	+2¢
11/8	551¢	33	550¢	−1¢
7/5	583¢	35	583¢	±0¢

(etc.—larger intervals which are octave-complements of these have the same absolute values for error)

TABLE 1: A COMPARISON OF SOME IMPORTANT JUST INTERVALS WITH THEIR APPROXIMATIONS IN THE 72-SET

Each of the Studies is correlated with (and named after) one of the sixty-four hexagrams in the *I Ching* or "Book of Changes." This correlation is based on the configuration of adjacent digrams in the hexagram, as follows: of the three disjunct digrams in each hexagram, the lower one is associated with pitch, the middle one with temporal density, and the upper one with dynamic level. Each digram may take one of four different forms, and each of these is interpreted to mean one of four possible "states" in a parameter—low (⚍), medium (⚏), high (⚎), and full (⚌). Thus, for example, the hexagram associated with the fifth Study is number 59 ("Dispersion"), which has the following form:

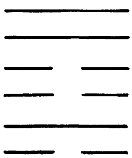

} upper digram: dynamic state = full

} middle digram: temporal density state = medium

} lower digram: pitch state = high

Relative means and ranges corresponding to the four different states are shown in Example 1.

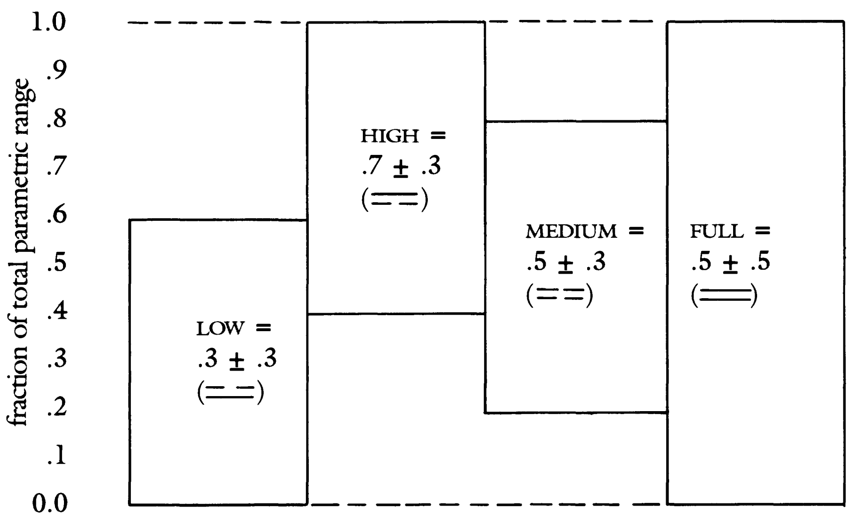

EXAMPLE 1: RELATIVE MEANS AND RANGES CORRESPONDING TO THE FOUR DIGRAM STATES

Actually, the parametric states of each Study are determined by *two* hexagrams—the first one (for which the Study is named) corresponding to the parametric states at the beginning of the Study, the second to those at the end. Where these terminal states differ in a given parameter, a gradual transition from one to the other is produced by the program using a half-cosine interpolation function. At lower holarchical levels, linear interpolation is also used for such changes of state during the course of a temporal gestalt-unit (or TG). In both cases, two mean-values are used for each TG—an initial and a final one, and these terminal values are connected by the interpolation function. For this purpose, the following formulae are used:

linear interpolation:

$$v_t = v_1 + (v_2 - v_1) * (t - t_1)/(t_2 - t_1)$$

half-cosine interpolation:

$$v_t = \frac{v_1 + v_2}{2} + \frac{v_1 - v_2}{2} \cos\left(\pi * \frac{t - t_1}{t_2 - t_1}\right)$$

where v_t is the value in the parameter at time t, v_1 the initial value (at time t_1) and v_2 the final value (at time t_2).

The first program generates two non-repeating random sequences of hexagram numbers, one for initial states, the other for final states, so every possible combination of parametric states occurs once at the beginning of one of the Studies, and once at the end of (usually) a different one. "Changing lines" for the initial hexagram are then inferred—such as would effect its transformation into the final hexagram. Because of this indirect way of deriving changing lines, they occur more often than they do when the *I Ching* hexagrams are obtained in the traditional ways, where the probability of a changing line is one in four, or 25%; here approximately 50% of the lines are changing.

On the basis of the initial and final parametric states of each Study, the first program also determines (1) whether it is to be monophonic or polyphonic, and then (2) the average vertical density of its elements, (3) the overall duration of the Study, (4) its average clang-duration, and (5) the initial and final tonic locations for the Study, as described more fully below. To determine whether a Study was to be monophonic or polyphonic, it was first considered *potentially* polyphonic if at least one parameter was in the "full" state, either at the beginning or at the end. When this was the case, a weighted random decision was made, with the weighting adjusted in such a way that approximately half of the sixty-four Studies would be polyphonic, the other half monophonic.

Both temporal density and vertical density vary exponentially in the Studies—i.e. the probable distribution of values in these two parameters will be uniform on a logarithmic scale. Thus, for example, the average temporal density mTd of a

TG will be computed as mTd = 2^{mS}, where *S* is the stochastically controlled variable, and *mS* its average value. Similarly for vertical density: mVd = 2^{mZ}. But while the mean values for temporal density depend directly on input data, those for vertical density are determined by a formula which relates them to pitch-range, average temporal density, and the number of polyphonic strata, as follows:

$$mZ = .5 + (1 - \sqrt{mS/1.6}) * \sqrt{nP/195} \; / \; \sqrt{Nst}$$

where *mS* is the average value of the temporal density exponent and 1.6 is the maximum value it can have in any Study; *nP* is one-half of the number of pitches in the range (always $\leq$ 195); and *Nst* is the number of polyphonic strata in the Study (either 1 or 2). The average vertical density of any Study thus varies directly with the pitch-range, and inversely with the average temporal density and the number of strata.

The total duration of each Study varies directly with the average "volume" of the three-dimensional space outlined by the ranges in the three basic parameters (pitch, temporal density, and dynamic level), and inversely with the average density of events within this space. This volume is proportional to the product of the average ranges in the three parameters, and the "density of events" to the product of (average) temporal density, vertical density, and the number of strata, as:

$$\text{Dur} \propto \frac{\text{Volume}}{\text{Density}} = \frac{nP * nS * nL}{mTd * mVd * Nst}$$

where all variables (except *Nst*) are arithmetic averages of the corresponding variables in the initial and final states of the Study. The results of this computation are later re-scaled to yield a minimum duration of 1′20″ and a maximum of 2′40″, so the average duration for the Studies in the set is about two minutes.

Each Study is organized into TGs at two holarchical levels between those of individual elements and the Study as a whole—clangs and segments. Here I have deliberately avoided TG-articulations at both the sequence- and section-levels, in an effort to enhance the sense of continuity and the perceptibility of contour at the segment-level and over the whole Study. The average clang-durations in individual Studies were made to depend (inversely) on their average densities (as defined above), and scaled to yield a minimum duration of $2.4/\sqrt{2} = 1.697$, and a maximum duration of $2.4 * \sqrt{2} = 3.394$ seconds.

The harmonic organization of the Studies will be described in more detail later, but a brief summary here may help clarify certain other operations carried out by this first of the two programs. The pitch classes (pcs) available within a given clang constitute a "mode" of (usually) seven different pcs, one of which is treated as a local or temporary tonic or root. In monophonic Studies, a new root and a new mode are chosen for each new clang. In polyphonic Studies—whose clang-boundaries are not, in general, synchronous—a new root and mode are chosen whenever the starting time of a new clang in one stratum is later than half-way through the duration of the concurrent clang in the other stratum—so pcs in the two strata are drawn from the same set more than half of the time. In both monophonic and polyphonic Studies, the series of root-progressions is controlled in such a way that each Study ends with a dominant-to-tonic "cadence" on the same root (the "global tonic") with which it began. Initial tonic pcs are ordered in a way which distributes the seventy-two pcs given by the tuning system over the sixty-four Studies as uniformly as possible, by simply omitting every ninth pc in the series from 0 to 71. The final tonic location is determined in a way which will be explained later.

The output of this first program consists of sixty-four blocks of data, each of which is used as input to the second program to generate the details of one of the Studies. Each block includes the following data: the numbers of the hexagrams defining initial and final parametric states for the Study, its total duration and average clang-duration, its initial tonic pc and the number of unit steps (in harmonic space) to the dominant of the "target tonic," the number of polyphonic strata, and the initial and final mean values and ranges for pitch, temporal density, dynamic level, and vertical density.

C. Individual Studies

In generating the output data for an individual Study, the second program works "from the top, down." That is, it first determines the duration and other parametric state values for the first segment, then for the first clang in that segment, and then for successive elements in that clang. When all the elements in the first clang have been generated, it determines the state values for the second clang and for its elements. After the last element of the last clang in this first segment has been generated, the program proceeds to the second segment, its first clang and the latter's successive elements, and so on. In the case of polyphonic Studies, these operations are carried out "in parallel," in such a way that successive elements' parametric values are generated alternately from the two polyphonic strata. This was necessary to maintain harmonic coherence between the two strata, since pitches in the two strata were to be drawn from the same set of available pitch classes at any given moment, whenever this was possible.

The number of segments in a Study is approximately equal to the average number of clangs in a segment, and the average segment-duration approximates the geometric mean of clang- and Study-durations, although individual segment durations vary randomly within a range of $\pm 25\%$ of this average duration. For each segment, an initial and final mean value in each of the other parameters—pitch, temporal density, dynamic level and vertical density—are chosen within the available range around the current "global" mean for the Study, which is determined—as explained earlier—by a half-cosine interpolation between the initial and final mean values for that parameter given by the input data for the Study. Each of the terminal mean-values for the *segment* is computed as the arithmetic average of two random values, which results in a tendency toward a "triangular" frequency distribution, rather than a uniform one—peaking at the current global mean and decreasing linearly toward the upper and lower boundaries of the current range in that parameter. This was done to lower the probability of extreme mean values at the segment level, which would have resulted in overly narrow ranges at the clang level.

The average clang-duration for each Study is given in the input data for that Study, but—as with segment-durations—the durations of individual clangs were made to vary randomly within a range of $\pm 25\%$ around the average value. Parametric means for each clang are chosen within segment-means in relation to the current mean of the segment—as with segment-means in relation to the current global mean of the Study—except that here (a) the current segment-mean is determined by linear (rather than half-cosine) interpolation between the terminal values, (b) only a single value in the parameter is used for a clang (that is, its parametric mean will be constant throughout the clang), (c) this value is determined by a single random number (so the frequency distribution of clang-means would tend to be uniform), and (d) the clang-mean for temporal density is made equal to the current segment-mean itself, rather than being allowed to vary randomly around that mean, in order to ensure a sufficient range of element-durations within each clang.

In all of my earlier stochastic music, the articulation of successive TGs was effected via the "similarity factor" only—involving differences in mean-values in various parameters. In an effort to incorporate the "proximity factor" as well, in the articulation of successive clangs, a new procedure was used here which interposes a *delay* before the beginning of each new clang (effectively prolonging the duration of the final element in the preceding clang), according to the following formula:

$$\text{Delay} = (\text{Dmax} - \text{Dur}) * (1. - \text{Pdst}/\text{Pdmx})$$

where Dur is the element-duration, Dmax the maximum element-duration possible in that clang, Pdst the pitch-distance between the two clang-means, and Pdmx the largest value this can have. The magnitude of the delay is thus determined by the relative distance between the pitch-means of the two clangs, and by the difference between the duration of the last element in the first clang and the maximum element-duration allowed for that clang (given its temporal density mean and range). The smaller the distance between the pitch-means of the two clangs (relative to the maximum value it could have, given the available range of clang pitch-means within the segment at that moment), the longer the delay is likely to be. Thus, for example, if the distance between the pitch-means of the two clangs happens to be zero (i.e. if the two clangs have the same pitch-mean, which could occur, although it's not very likely), the amount of the delay will be such that the (modified) duration of the last element in the first clang will be equal to the maximum element-duration in that clang. If, on the other hand, this distance happens to be at maximum, the delay will be zero, and the duration of that last element will remain unmodified.

The hierarchical (or holarchical), recursive character of the program, already described for segments and clangs, continues at the element-level, although element-durations are generated more simply than were clang- and segment-durations (as the reciprocal of a temporal density value for the element), and element dynamic levels are made equal to the clang-mean in that parameter (so dynamic levels remain constant throughout a clang). The value derived at this level for vertical density—truncated to the next-lower integer—determines the number of pitches in the element. As with clangs and segments, parametric values (other than dynamics) for an element are drawn from the available range around the clang-mean, but for the pitch-parameter, other, specifically harmonic procedures intervene here to determine a set of available pitch *classes* (or pcs) before the actual pitches are selected. These procedures will be described in the section that follows.

D. Harmonic Procedures

My intentions in this work, with respect to harmony, included the following:

1) that one of the pcs in every clang should function as a temporary tonic or root in relation to all the other pcs in that clang—which latter are interpreted as a kind of temporary "mode" for that clang;

2) that the root pc would change from clang to clang, by means of a root-progression chosen stochastically from a set of possible root-progressions with pre-set relative probabilities assigned to them;

3) that the pcs in a mode should tend to form relatively compact sets in harmonic space, both in relation to the other pcs in that clang and to those in the previous clang, and finally

4) that the "random walk" character of the series of root-progressions should gradually be "focused" in such a way that each Study would end with a dominant-to-tonic progression to the same root pc with which it began—and in the same mode.

To achieve these intentions required a careful analysis of the 72-set and its several possible mappings in harmonic space. For example, the pcs in the 72-set can be mapped in pitch-class projection spaces of 2, 3, or 4 (or more!) dimensions, according to the prime-limit being considered. For *Changes*, I decided to assume an 11-limit (five-dimensional) harmonic space for the modes, a 7-limit (four-dimensional) harmonic space for root-progressions, and to locate the final, "target tonic" on the same 3,5-plane as the initial tonic (which implies a 5-limit, three-dimensional space for this relation between initial and final tonic locations). Examples 2 and 3 show some of these mappings of the pcs in the 72-set, in pitch-class projection spaces of two and three dimensions (corresponding to prime-limits of 5 and 7, respectively). Note that—because the 72-set is an equal-tempered system—its lattice structure is *periodic* in harmonic space (no matter what the dimensionality may be of that space into which it is mapped). That is, it repeats itself endlessly in all directions. It was decided to use as the target tonic in each of these Studies one of the many locations of that tonic in the 3,5-plane, in a direction (in relation to the initial tonic) similar to the direction in which Bach's harmonic progressions tend to move in a mapping of the *12*-set in harmonic space—i.e. toward the left along the 3-axis (via descending fifths—e.g. V–I) and upward along the 5-axis (less quickly, and mostly via the descending minor third progression—e.g. I–vi). Example 4 shows the configuration of recurring tonics (in relation to an initial 1/1 or "0"), in an abbreviated but extended mapping on the 3,5-plane. The location used for each Study was one of the three indicated by the arrows—which one of the three depending on the estimated number of clangs (and thus, the number of root-progressions) in that Study. The numbers in parentheses give the number of unit steps, along the 3- and 5-axes, respectively, from the initial to the final tonic location.

Each of the sixty-four Studies begins (and thus ends) on a different tonic pc, and these form an ascending integer series, beginning with 0 (E↓—read: "three-sixths of a semitone below E") for Study Number 1, and ending with 71 (D♯/E♭↑) for Study Number 64, skipping every ninth pc in the series. The other pcs of the mode associated with a root are chosen from a set of alternatives—for each of six "scale degrees" (in addition to the tonic)—given as input data to the program (but common to all sixty-four Studies). These are arranged in "stacked thirds" order: prime, third, fifth, seventh, ninth, eleventh, thirteenth, and they include from three to five alternatives for each degree above the tonic. These are listed in Table 2, which gives both the pc number in the 72-set and the just ratio or ratios most closely approximated by that pc (in parentheses). The most important harmonic relationships among these various alternatives are shown in Example 5, representing their locations in harmonic space (or, more precisely, in a pitch-class projection space essentially in 7-limit form, but with the additional ratios of 11 interposed along the 3-axes (and in parentheses)). The choice of a particular pc (or ic—interval class—in

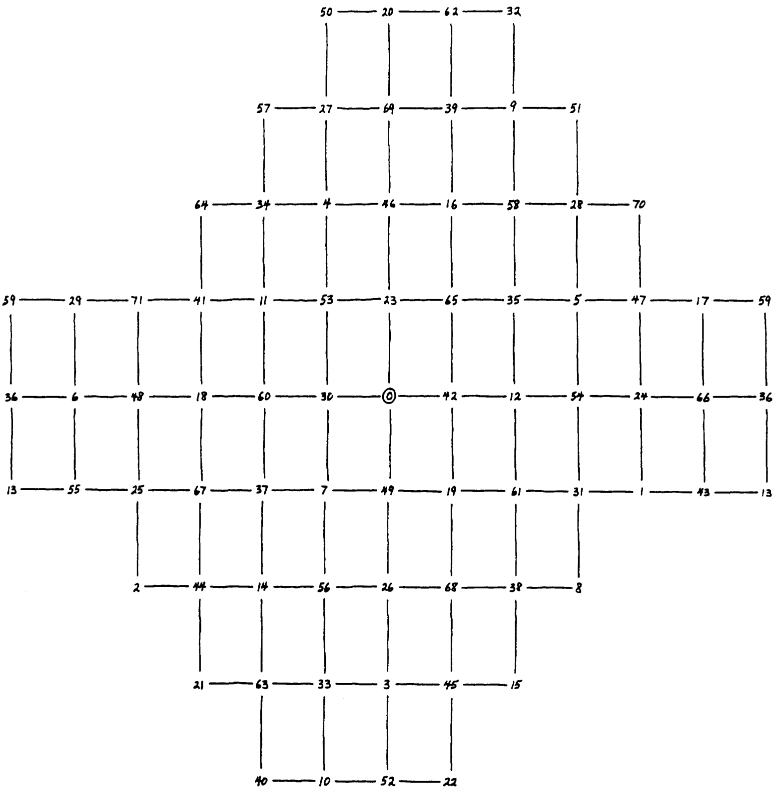

EXAMPLE 2: THE 72-SET IN THE 3,5-PLANE

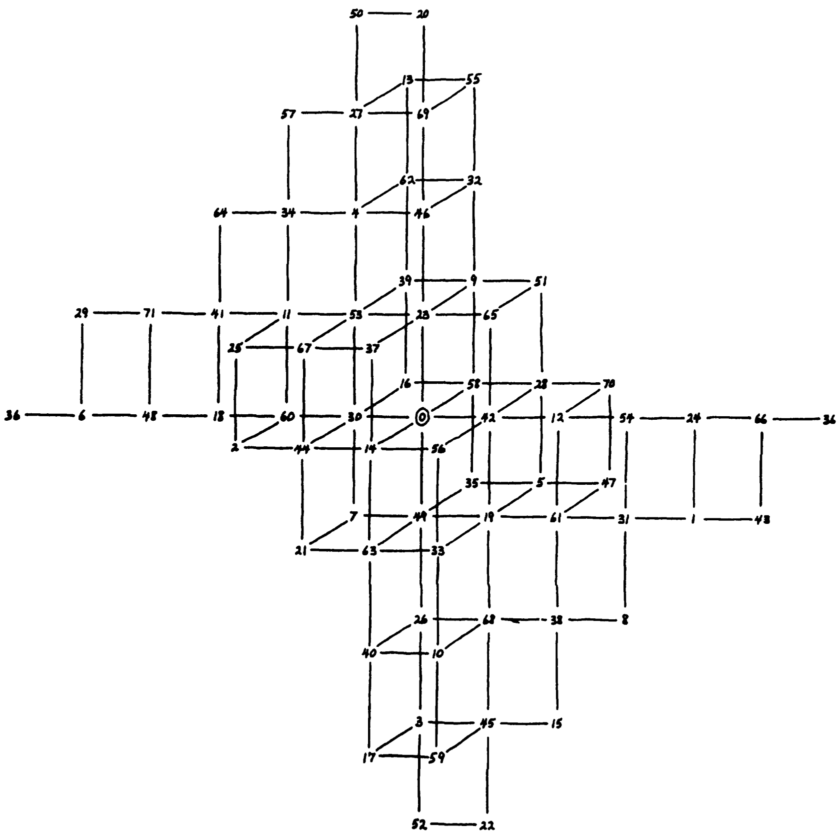

EXAMPLE 3: THE 72-SET IN 3,5,7-SPACE

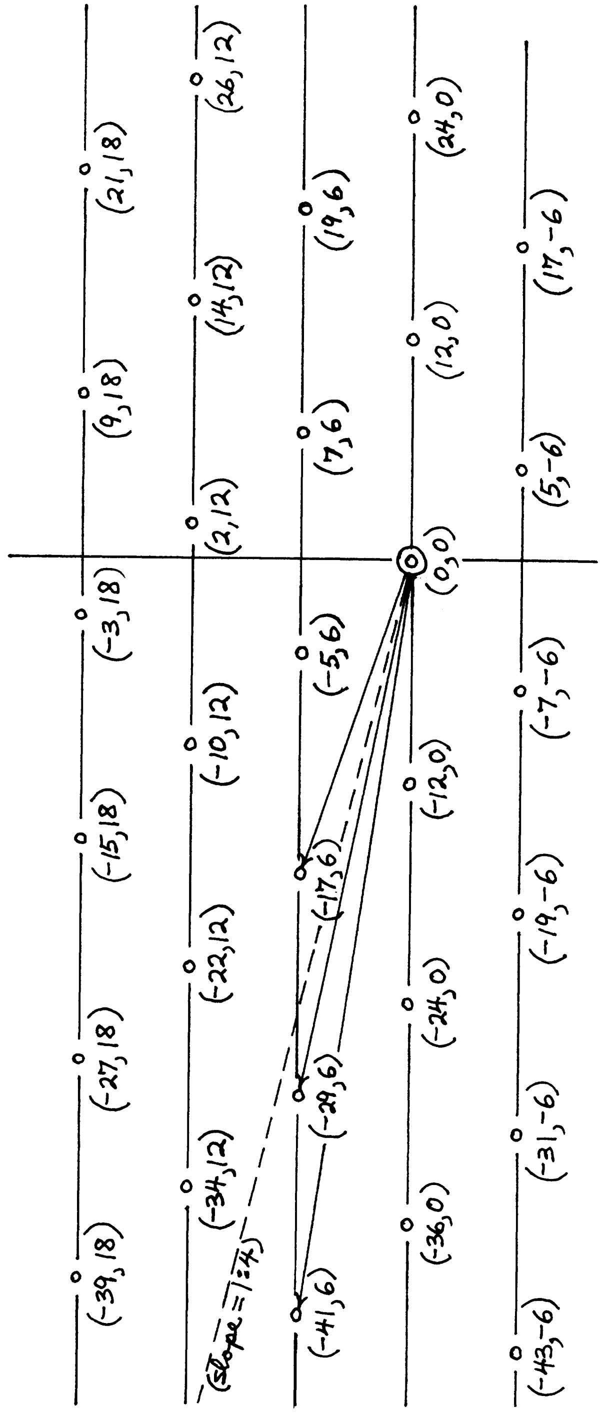

EXAMPLE 4: RECURRING TONICS IN THE 3,5-PLANE

relation to a given tonic pc) for each degree is determined by several conditions, some of which might be described as "rules," while others are more statistical in character. The rules include the following:

1) in the initial (and thus also the final) tonic set, the fifth is always made equal to 42 (3/2), and the seventh is allowed to equal 58 (7/4) only if the third (already chosen) equals 16 (7/6);

2) in the dominant set preceding the final ("target") tonic, the third is always = 23 (5/4), the seventh always = 58 (7/4);

3) the various "thirds" between adjacent degrees may vary in size only within specified ranges: from a minimum of 12 (9/8) to a maximum of 26 (9/7) between prime and third or third and fifth, a minimum of 16 (7/6 or 75/64) and a maximum of 30 (4/3) between adjacent degrees above the fifth;

4) no "mistuned fifths" are allowed between non-adjacent degrees (as between the third and seventh, fifth and ninth, and so on)—i.e. any such interval must either be precisely equal to 42 (3/2) or differ from it by an interval greater than 3 (a "quarter-tone");

5) no octaves (either exact or "mistuned") are allowed between those non-adjacent degrees which share a common pc, or approximate that condition too closely (as between the third and the ninth or eleventh, the fifth and the eleventh or thirteenth)—i.e. no "seventh" larger than 68 is allowed, and no "ninth" smaller than 4. Thus, any interval between non-adjacent degrees must differ from an octave by at least 4 (two-thirds of a semitone);

6) if the third = 19 (6/5), the fifth must equal 42 (3/2)—thus disallowing both the "flat" and "raised" fifths when the third is of the ordinary minor form;

7) the raised fifth—46 (25/16)—is only allowed when the third = 23 (5/4).

prime: 0 (1/1)

3rds: 16 (7/6), 19 (6/5), 21 (11/9), 23 (5/4), 26 (9/7)

5ths: 35 (7/5 or 45/32), 42 (3/2), 46 (25/16)

7ths: 58 (7/4), 61 (9/5), 63 (11/6), 65 (15/8)

9ths: 5 (21/20), 7 (15/14, 16/15 or 77/72), 9 (12/11 or 35/32), 12 (9/8), 16 (7/6 or 75/64)

11ths: 28 (21/16), 33 (11/8), 35 (7/5 or 45/32)

13ths: 49 (8/5 or 77/48), 51 (18/11 or 105/64), 53 (5/3), 54 (27/16), 56 (12/7 or 55/32)

TABLE 2: ALTERNATIVE PCS FOR A MODE

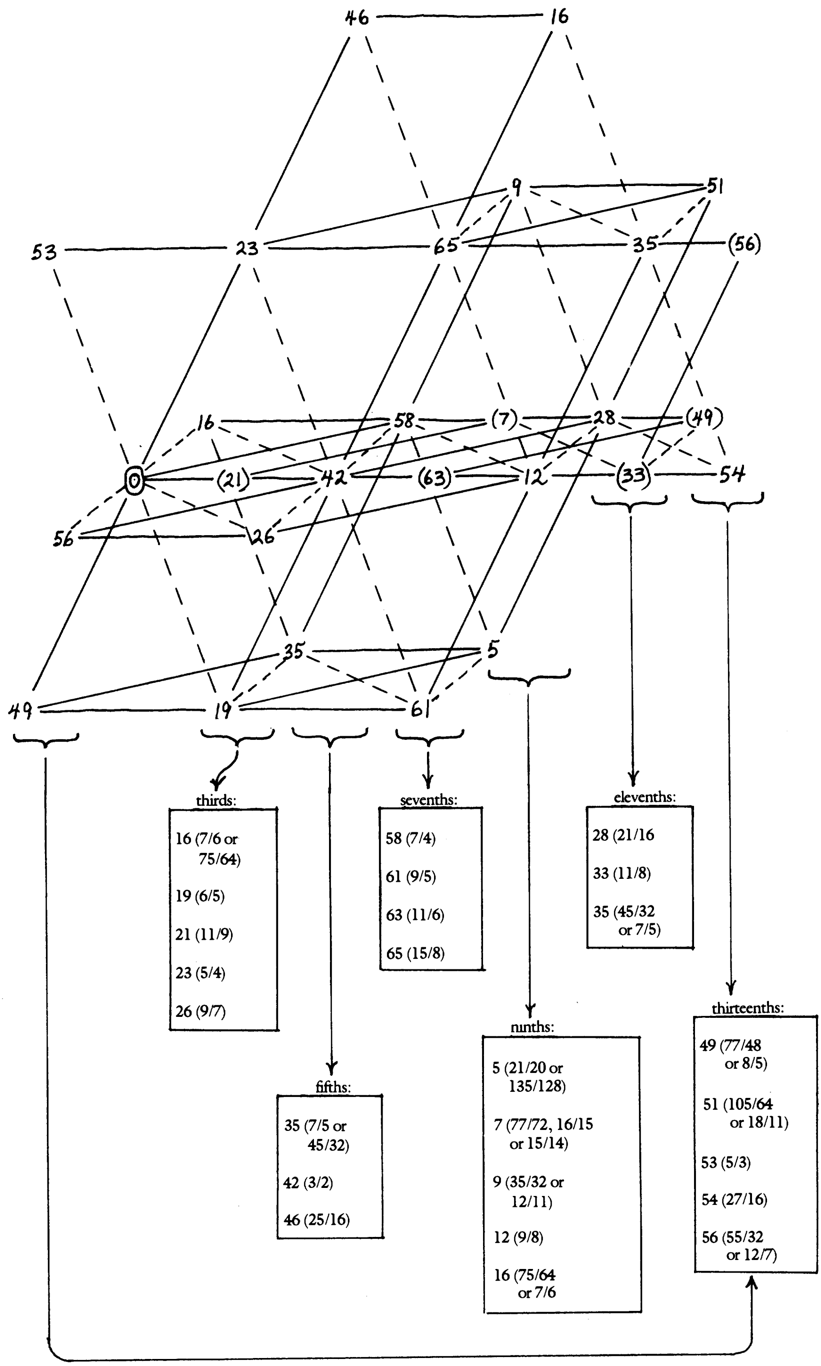

EXAMPLE 5: HARMONIC RELATIONSHIPS AMONG ALTERNATIVE PCS FOR THE MODES

Some of these rules correspond to similar rules for chord-construction in both traditional and jazz harmonic practice (and I should perhaps add here something which has not been made explicit before: the pcs of a mode are often heard simultaneously as well as successively—as "chords" as well as "melodic lines"—thus the ambivalence (which may have been noticed already) in my use of the terms "tonic" and "root"). Other rules were designed to avoid certain ambiguities and/or conflicts that might otherwise occur in the creation of these modes. Although these rules appear to be quite restrictive, a very large number of modal sets were still possible, but these were further constrained by what I have referred to (above) as "statistical" conditions, as follows:

The pcs that remain available for a given modal degree after testing against the rules just listed are assigned varying probabilities depending on the sums of their harmonic distances to pcs already chosen for that clang—and to the pcs actually occurring in the clang just preceding (I say "actually occurring" because—due to the random process involved in the selection of pitches in a clang—it is always possible that one or more of the pcs constituting the mode will not occur). The relation between these probabilities and harmonic distances varies according to the modal degree in question (the constraint is "tighter" for the higher degrees), and whether this was the first clang in the Study or not (the constraint is "looser" for the first clang), but in general that relation is an *inverse* one. That is, the lower the sum of harmonic distances between a pc and the others preceding it, the higher its probability of being chosen—and vice versa. This constraint was made stronger for higher degrees of the mode (arranged in "stacked thirds" order, remember) by raising the harmonic-distance sum to a power corresponding to the "height" of the degree, as follows:

$$\mathrm{Pr}(j) \propto 1/\mathrm{Hdsm}(j)^{n+k-1}$$

where $\mathrm{Pr}(j)$ is the relative probability of the jth pc in the set of still-available pcs for that degree, $\mathrm{Hdsm}(j)$ equals the sum of that pcs harmonic distances to preceding pcs, n is the order-number for the modal degree (i.e. $n = 2$ for the third, 3 for the fifth, 4 for the seventh, and so on), and $k = 1$ for the first clang, 2 for all later clangs. The result of all this is that there will be a tendency for pcs to form relatively *compact sets* in harmonic space, with this tendency becoming stronger for higher modal degrees—and conversely, so there is more freedom for random variation in the lower degrees.

It might be noted that the sets of alternative pcs for modal degrees yield seven different kinds of triads, only four of which are familiar in traditional western harmony (numbers 2, 4, 6 and 7, below):

1) septimal minor	0 (1/1),	16 (7/6),	42 (3/2)
2) 5-limit minor	"	19 (6/5)	"
3) 11-limit "neutral"	"	21 (11/9)	"
4) 5-limit major	"	23 (5/4)	"
5) septimal major	"	26 (9/7)	"
6) augmented	"	23 (5/4),	46 (25/16)
7) diminished	"	16 (7/6),	35 (7/5)

Another possible form of the diminished triad—0 (1/1), 19 (6/5), 35 (7/5)—was avoided because the most likely seventh degree with such a triad would have been 65 (15/8), and the perfect fourth (30 (4/3)) which this forms with 35 would have introduced an unwanted ambiguity with respect to the root. The sets of alternative pcs for scale degrees were designed to avoid pcs which might compete with the nominal root, and the perfect fifth and fourth—and even the (5-limit) major third and minor sixth, though less strongly—have very clear root-defining effects. Thus, the perfect fourth itself—30 (4/3)—was not included as a possible eleventh in a mode, and 49 (8/5 or 77/48) was only included as a possible thirteenth because of its dual character—and harmonic distance values for this interval were set to correspond to its interpretation as 77/48, rather than 8/5. (The same thing was done for the interval formed by pc 26 (9/7), to avoid its interpretation as 32/25, which—because of the way in which I calculated harmonic distances (for an explanation of which, see below)—would have given it more prominence than I thought it should have.)

The seventh-chords which can arise by way of this procedure for constructing modes include most of the traditional ones (major, minor, dominant, half-diminished, minor-major, augmented, and so forth, but not the diminished seventh), plus several others which are of interest, including the one used by Ives as the primary chord in the "Choral" of his *Three Quarter-tone Pieces*—0 (1/1), 21 (11/9), 42 (3/2), 63 (11/6). Ninth-chords include—again—all of the traditional ones, plus the blues "flat 7, sharp 9," and a very interesting group of new ones with pc 9 as the ninth of the mode. This pc—at the "quarter-tone" position between the 12-set's minor and major seconds—functions in the 72-set most frequently as the major third above the "dominant" 7th—i.e. it can be analyzed as 58 (7/4) + 23 (5/4) = 81 (mod 72) = 9 (35/32). The fact that it occurs in a "dominant"-type pc set more often than the more familiar minor or major ninth is suggestive: perhaps the latter are merely the best approximations available in the 12-set for this interval! Finally, the eleventh-chords include a good approximation of Partch's "otonality hexad"—0 (1/1 ± 0¢), 23 (5/4 − 3¢), 42 (3/2 − 2¢), 58 (7/4 − 2¢), 12 (9/8 − 4¢), and 33 (11/8 − 1¢).

The basic formula for the harmonic distance between any two pitches is $\text{Hd}(a/b) = k \log_x(ab)$, where a/b is the frequency ratio representing the interval (in its maximally reduced, "relative prime" form), and k simply determines the unit of measurement (with base-2 logarithms, if $k = 1$, Hd is in "octaves"). The form used in this program, however, is a bit different, in two respects. First, it is a measure of the harmonic distance between pitch-classes, rather than actual pitches. Second, since we are dealing here with a tempered system, a *tolerance* rule is invoked, which essentially says that we can assume the *simplest* integer ratio within the tolerance range around the tempered pitch to be the harmonically effective one (that tolerance range is here taken to be ± one-half the size of the smallest step in the tuning system—i.e. ± 1/144 of an octave or 8.33. . .¢). The first qualification means that we are concerned with a distance not between points in the full, n-dimensional harmonic space itself, but rather between points in the (n-1)-dimensional "pitch-class projection space." This, in turn, means that the formula for harmonic distance must be replaced by another of the form: $\text{Hd}\ (a'/b') = k \log_2(a'b')$, where $a' = a/2^i$, $b' = b/2^j$, and i and j are the largest integer exponents which yield integer values of a' and b'. The second qualification means that —where there are two or more relatively simple integer ratios defining intervals within the tolerance range of a pc, the one whose ratio-terms' product is smallest determines the harmonic distance value assigned to that pc. It has already been mentioned that two exceptions were made to this procedure, involving pcs 26 and 49. Pitch-class 26 (@ 433¢) approximates both 32/25 (@ 427¢) and 9/7 (@ 435¢), while pc 49 (@ 817¢) approximates 8/5 (@ 814¢), 77/48 (@ 818¢), and 45/28 (@ 821¢). While I wanted both of these pcs to be included among the available alternatives (for thirds and thirteenths, respectively), I wanted 26 to be treated as a 9/7, and 49 as a 77/48, so their minimal harmonic distance values were overridden in another part of the program with the higher values. (I see now, in studying the program again, that the value I assigned to 49 was that of 45/28 rather than 77/48—i.e. $\log_2(315) = 8.30$ rather than $\log_2(231) = 7.85$—but fortunately this error turned out to be a small one, with scarcely noticeable effect on the final results.)

Once the pcs of the mode for a clang have been chosen, the program is almost ready to proceed to the selection of actual pitches, within the range already determined for that clang. As at all higher levels, this involves a random process, but at this level the process is further constrained by two kinds of probability distributions, one providing some control over the rate of recurrence of each pitch, the other correlating modal degree with register. The probability of a given pitch being chosen by the random process at any moment was computed as the product of two probability "factors" stored in a two-dimensional array called PPR(N,L), where N = 1 or 2, and L is an index for pitch (L = 1 for the lowest pitch in the harps' range, L = 452 for the highest). For all values of L, PPR(1,L) was initialized at 1., so all pitches began with the same relative probabilities. Just after a pitch is chosen for an element, PPR(1,L) for that pitch is reduced to a very small value, and then increased step by step, with the generation of each succeeding element (at any *other* pitch), until it is again equal to 1.0. The result of this procedure is that the immediate recurrence of a given pitch is made highly unlikely (although not impossible, especially in long and/or dense clangs, and in a polyphonic texture), with the probability of recurrence of that pitch gradually increasing over the next several elements until it is equal to what it would have been if it had not already occurred. The other probability factor—PPR(2,L)—is used to effect a correlation between modal degree and register, as shown graphically in Example 6. Note that while the root or tonic of the mode has an equal probability of occurring anywhere within the pitch-range of the clang—and all other modal degrees are equally likely at the upper boundary of the clang's pitch-range—the higher modal degrees have low probabilities of occurring in the lower regions of the clang range (and conversely for the lower modal degrees).

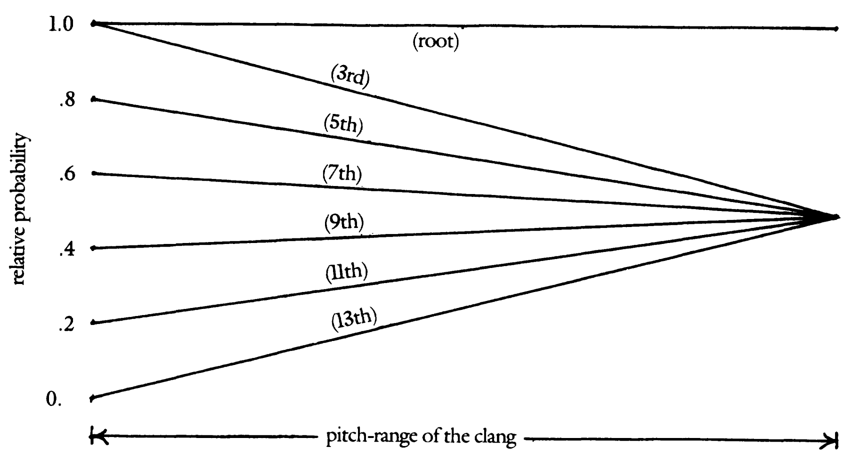

EXAMPLE 6: CORRELATION OF MODAL DEGREE WITH REGISTER

Finally, values are determined for the starting-time (or epoch), duration, pitch(es), and dynamic level of each element in the clang. Element-duration is computed as the reciprocal of a temporal density value for that element, and the epoch is given by the sum of epoch and duration values for the previous element in the stratum (plus the "delay" described earlier, when the element is the first in a new clang). These time values are initially calculated on a virtually continuous scale—as in *Bridge*—but (unlike *Bridge*) I decided in this work to quantize or "rationalize" these values so they could be represented in the standard metrical rhythmic notation in the score and parts. This was done as follows: for the epoch of each element, the program computes (and prints out with the other parametric values for that element) the absolute differences between the initially calculated value and both the nearest sixteenth-note and the nearest triplet eighth-note. It is then left up to the person transcribing the numerical output data into musical notation to decide which of the two rational approximations to use, based on the magnitude of the "error" involved, and on the epochs and errors for any other elements which may begin within the same quarter-note (since the two divisions of the quarter—by 3 and by 4—cannot generally be mixed within a given quarter in our standard system of rhythmic notation). Example 7 shows an example of a page of output data, with the values for a single element boxed, and the "error" values just described shown circled.

When the ending-time of an element equals or exceeds a predetermined ending-time for the clang, the program computes a new root pc for the next clang, and a new mode for that clang. The interval-class (ic) between this new root and the root of the previous clang thus defines a root-progression, and is determined as follows: an array is used to store an initial set of relative probabilities for allowable root-progressions (these probabilities are the same for all sixty-four Studies), as shown graphically in Example 8, and listed in Table 3. This set of probabilities is conceived as determining a smaller set of six vector components in a three-dimensional harmonic space, and these, in turn, can be reduced to a single "resultant" vector which indicates the direction and average rate of root-movement through that space—assuming, of course, that a large number of such root-progressions will be involved. The result is a kind of directed "random walk" through the harmonic space.

In order to further ensure not only that this "random walk" will have—over the long run—the appropriate direction and rate in relation to the location of the "target tonic" (or rather, the dominant preceding this tonic), but that the movement will become gradually more "focused," and finally arrive at its goal, the set of individual root-progression probabilities is revised, for each new clang, according to the actual direction and distance remaining to the target. I won't go into more detail here about the mechanics of this process, except to note that this part of the program turned out to be more complicated than I had expected it to be, and that it didn't always work! That is, there always

```
STUDY #25. HEXAGRAMS 54 AND  8; DURATION=  86.56 SECONDS DPR=, 1.875
      231.0 195.0     1.6  1.2     4.0  1.8      0.500  0.500
      231.0 117.0     0.8  1.2     2.8  1.8      0.660  0.660
      IRT= 27  MX= -16  NCP= 41  NST=  2 RANDOM SEED=  1989965860


   INITIAL ROOT-               0 .00  30 .30  60 .01  11 .05  53 .11  16 .04  58 .02  23 .03  65 .10
   PROGRESSION PROBABILITIES:         42 .10  12 .05  61 .01  19 .01  56 .02  14 .04  49 .10   7 .01

 -0.230   0.160   0.000  -0.580   0.350  -0.130   0.290  -0.060   0.060


SEGMENT  1                0.00   12.08   12.08
108.37  72.37    1.53  1.13  4.82  0.98  0.77  0.23
238.33 183.99    0.95  0.58  3.98  1.76  0.49  0.49


                               0          27  53  69  16  39  62   9     -16   6  0       41
                                           0  26  42  61  12  35  54


CLANG  1    0.00     0.00     1.53    1.53
    122.88 112.61    1.48   1.09  3.89  0.17   0.54   0.05

 1: ST=  0.00, DR= 0.63, ET=  0.63,          CMS= 0.0

     0.2.1   0.0.0   0.0.0   0.0.0   0.0.0   0.0.0   0.0.0
       53      0       0       0       0       0       0

     M= 1, Q= 1,      0/3.000     0/4.000          DYN: 3


                                                              SEGMENT  1   0.45   12.22   12.68
                                                              196.51 160.511.62  1.18  4.45  1.35  0.44  0.44
                                                              361.99  59.952.46  0.30  3.80  1.67  0.88  0.14

                                                              CLANG  1    0.00    0.45    1.72    2.18
                                                                  268.21 129.79   1.68  1.12  3.42  0.40  0.45  0.39

                                                               1: ST=  0.45, DR= 0.27, ET=  0.72,         CMS= 0.9

                                                                   8.5.3  0.0.0  0.0.0  0.0.0  0.0.0  0.0.0  0.0.0
                                                                     27    0      0      0      0      0      0

                                                                   M= 1, Q=1,     1/3.120    2/4.047        DYN: 3


 2: ST=  0.63, DR= 0.33, ET=  0.97,          CMS= 1.2

     3.3.3   0.0.0   0.0.0   0.0.0   0.0.0   0.0.0   0.0.0
       69      0       0       0       0       0       0

     M= 1, Q= 1,     [2/3.033]   [3/4.116]         DYN: 3
```

```
3: ST=  0.97, DR= 0.25, ET=  1.22,          CMS= 1.8
   8.3.3  0.0.0  0.0.0  0.0.0  0.0.0  0.0.0  0.0.0
     27       0      0      0      0      0      0
   M= 1, G= 2,   C/3.032   0/4.032        DYN: 3

4: ST=  1.22, DR= 0.25, ET=  1.47,          CMS= 2.3
   6.2.2  0.0.0  0.0.0  0.0.0  0.0.0  0.0.0  0.0.0
     16       0      0      0      0      0      0
   M= 1, G= 2,   1/3.115   1/4.032        DYN: 3

                 0       27  53  69  16  39  62   9     -16   6  0     40
                          0  26  42  61  12  35  54

CLANG  2    0.38     1.84     2.27     4.12
   172.54  92.33     1.41  1.02  5.36  0.43  0.58  0.15

1: ST=  1.84, DR= 0.25, ET=  2.09,          CMS= 3.5
   8.2.3  0.0.0  0.0.0  0.0.0  0.0.0  0.0.0  0.0.0
     27       0      0      0      0      0      0
   M= 1, G= 2,   3/3.156   3/4.094        DYN: 5
```

```
2: ST=  0.72, DR= 0.35, ET=  1.08,          CMS= 1.4
   6.6.2  0.0.0  0.0.0  0.0.0  0.0.0  0.0.0  0.0.0
     16       0      0      0      0      0      0
   M= 1, Q=1,    2/3.056   3/4.027        DYN: 3

3: ST=  1.08, DR= 0.40, ET=  1.48,          CMS= 2.0
   5.4.3  0.0.0  0.0.0  0.0.0  0.0.0  0.0.0  0.0.0
      9       0      0      0      0      0      0
   M= 1, Q=2,    C/3.075   0/4.075        DYN: 3

4: ST=  1.48, DR= 0.25, ET=  1.73,          CMS= 2.8
   3.5.3  0.0.0  0.0.0  0.0.0  0.0.0  0.0.0  0.0.0
     69       0      0      0      0      0      0
```

EXAMPLE 7: THE FIRST TWO PAGES OF OUTPUT DATA FOR STUDY NUMBER 25

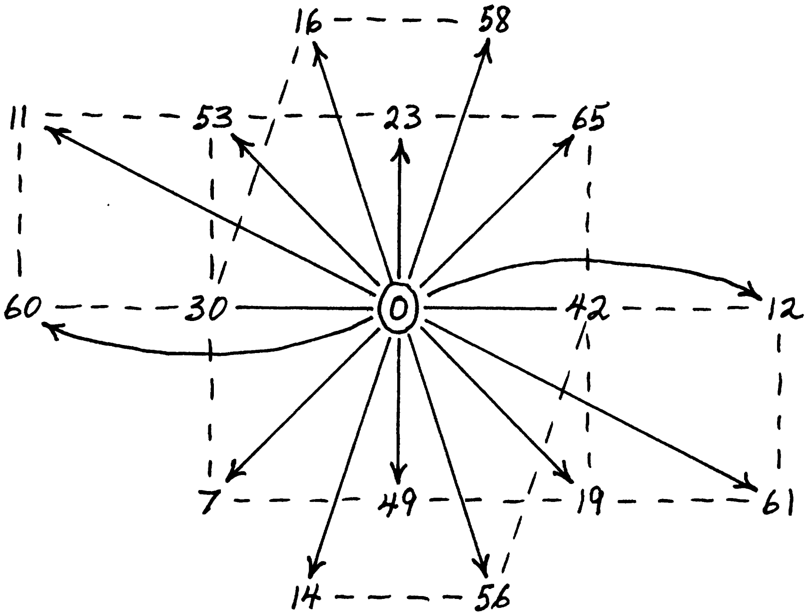

EXAMPLE 8: AVAILABLE ROOT-PROGRESSIONS

root-progression *interval-class*	*probability*	root-progression *interval-class*	*probability*
0	.00	30	.30
7	.01	42	.10
11	.05	49	.10
12	.05	53	.11
14	.04	56	.02
16	.04	58	.02
19	.01	60	.01
23	.03	65	.10

TABLE 3: ROOT-PROGRESSION PROBABILITIES

remained a certain degree of unpredictability in the final convergence toward the dominant, such that the intended target was actually missed in about one out of three runs of the program. When this happened, the output was discarded, and the Study generated again with a new random seed. Since the total duration and the average clang-duration of each Study were considered characteristic features of that Study, derived by the first program by operations on its terminal states in the three primary parameters—and not be be altered arbitrarily or contingently—the series of root-progressions was required not only to *arrive* at its target, but to arrive there *on time*. Such a percentage of "failures" is therefore not surprising—given the essentially *stochastic* nature of the process. In each Study, the program had four chances to succeed: if it arrived at the target dominant at the sixth, fifth, fourth, or third clang from the end, a "cadencing" routine was initiated, which kept it rooted on the dominant pc, and set the mode in some form of (extended) "dominant 7th," until the next to last (or in some cases, the last) clang, at which point it effected a progression to the final tonic. The similarities between this procedure and what might be inferred from many of the cadential passages in Bach's Preludes should be obvious—although profound differences will also be evident to any listener, I am sure.