Cloud Triptych: an exploration of stochastic movement between discrete musical behaviours

by

Simon Cummings

A thesis submitted in partial fulfilment
of the requirements of Birmingham City University
for the degree of Doctor of Philosophy

Birmingham Conservatoire
Birmingham City University

November 2016

Abstract

The primary objective of my compositional research is to create perceptible, large-scale musical transformations. To this end, my research project was to compose a series of works employing algorithmic techniques of my own devising, specifically designed to map stochastic processes of movement between epicentres of discrete musical behaviour. The processes are codified in and realised by bespoke computer programs that generate the resultant musical material. This material is then interpreted, modified as needed and finally notated as a score. The implications and validity of these processes are thereby tested in the compositions.

The specific focus of my research is music for acoustic instruments. The portfolio includes a principal work, *Cloud Triptych*, and five preliminary pieces for smaller numbers of players, exploring processes that generate portions of material to be incorporated into larger structures as well as complete compositions. An important part of this exploration was the integration of both computational and intuitive materials, and the extent to which computational models can effect an embodiment of my compositional practice, thereby simulating musical creativity. In each of the compositions the different aspects that make up any musical behaviour – such as rhythm, pitch, articulation and dynamic – are treated as independent parameters that together comprise the whole. As such, they are scrutinised and developed both on their own terms and in their role interacting with each other within larger musical entities. The central role of behaviour, and the way it can serve as a discrete epicentre of stable musical focus, is explored in several ways, abruptly alternating between contrasting behavioural states as well as various kinds of transition and transformation.

The goal and outcome of this research was the composition of the extended large orchestral work *Cloud Triptych*, which lasts approximately 28 minutes. This was created using an extensive piece of software of my own design, CloudCube, that both encapsulates and greatly expands the computational outlook of the preliminary works, in so doing approximating closer than ever to an embodiment of the diverse aspects of my compositional practice.

As well as the scores of the six compositions, the portfolio contains this 25,000-word commentary explaining the creative and technical processes I have developed, and discussing their implementation in the compositions, along with two accompanying CDs.

Table of Contents

Acknowledgements

I would like to extend my thanks to the following people and groups, for their insights and every kind of help and support:

Christopher Redgate, Rowland Sutherland, Alex Surman, Carla Rees, B.C.M.G., Richard Baker, The Curious Chamber Players, Rei Munakata, Thallein Ensemble, Edwin Roxburgh, Schubert Ensemble, Chi Hoe Mak, Laura Kox, Leeds Flute Choir, Ruth Hopkins, Thumb, Interrobang, Torsten Anders, Marc Yeats, Kenneth Hesketh, John Wall, Nikolaos Fountoulakis, Stuart Bowes, Aaron Cassidy, Matt Sergeant, Ken Kirschner, Geraint Wiggins, Murphy McCaleb, Sebastiano Dessanay, Veleka Algar, Edmund Hunt, Stuart Stephens, Thomas Simaku, Carrie Churnside, Chris Dingle, Joe Cutler, Francis Firth, Stephanie Fowler, Sanshia Bedford and Sandy Price.

Special thanks to my supervisors:

Howard Skempton, for having the loosest possible interpretation of a two-hour tutorial, the uncanny ability always to locate salient points for discussion, and for sharing quantities of wisdom that will take me years to unpack and understand.

Peter Johnson, for convincing me that my work had the potential for further investigation, and for being the best academic mentor and proof-reader I've ever had.

Extra special thanks to Herb, my beloved little cat, whose omnipresent company kept me from ever being alone.

I would also like to express my sincere gratitude to the following people:

Christopher Best, for being there at the beginning (the actual beginning) and for ongoing support and friendship. Paul Berg, for sowing the seeds of method and technique that would prove the most important of my compositional life. Stephen Thiselton, for saving the day more than once, and for helping me to apprehend and solve various elements within CloudCube. Michael Finnissy, for embodying an impossible combination of challenge and optimism, and for consistently asking the questions no-one else ever asked. Richard Causton, for countless, infinitely patient and encouraging conversations, and above all for helping me – finally! – to learn how to play.

I dedicate this Ph.D. to my wife Anna, whose love and support has been and continues to be the one constant parameter in a landscape of perpetual transition and change. Without her this research could simply never have happened.

SIMON CUMMINGS
November 2016

Contents of Portfolio

Preliminary Works

'unredeemed' self-)portrait (in the form of an eagle (2008) for solo flute

Blesi (Five Transitions of the Soul) (2009) for six players

HELP/ME: the soul-machine of the cosmology of grief (2010) for 16 players

Intense quick dream of sentimental groups with people of all possible characters amidst all possible appearances (2010) for string sextet

Four Seasons for flute choir (2016)

Main Work

Cloud Triptych for large orchestra (2016)

Contents of CDs

Disc 1: Audio

track	work	duration
1.	***'unredeemed' self-)portrait (in the form of an eagle***	6:24
	Carla Rees (flute)	
2.	***HELP/ME: the soul-machine of the cosmology of grief***	5:39
	Thallein Ensemble, Edwin Roxburgh (Birmingham Conservatoire student ensemble)	
3.	***Intense quick dream of sentimental groups with people of all possible characters amidst all possible appearances***	4:47
	Birmingham Contemporary Music Group, Richard Baker	
	Four Seasons	
4.	I. prime-temps	8:08
5.	II. samā	4:05
6.	III. fall	10:08
7.	IV. vindo	6:08
	synthetic recording*	
	Cloud Triptych	
8.	I. (complete)	5:35
9.	II. *Fig. F to end***	9:13
10.	III. (complete)	7:11
	synthetic recording*	
	TOTAL	67:18

*The synthetic recordings of *Four Seasons* and *Cloud Triptych* were created using a combination of Sibelius and Ableton Live. For obvious reasons, they do not represent the scores with complete accuracy or fidelity and are for illustrative purposes only.

**Due to the use of extended glissandi and air noise during the early sections of this movement, neither of which can be satisfactorily rendered by Sibelius, only the latter half of this movement is included.

Disc 2: Data

This CD contains the CloudCube software. Please refer to the file **README.pdf** which contains instructions for use.

Chapter 1. Introduction and Method

Introduction

This commentary is a presentation and analysis of the submitted works and an account of the research process that led to their composition. The main objective throughout has been to extend and explore the possibilities, consequences and implications of a variety of algorithmic processes and techniques. My particular interest is in perceptible processes that gradually transition and evolve over time between discrete musical behaviours.

In recent decades, the formalisation of compositional processes as computational models has enabled considerable expansion of these processes' complexity and subtlety. A fundamental consideration with regard to this expansion – and more widely, to the product of all music resulting from processes, however managed – is the issue of perceptibility. Xenakis dismissed integral serialism in the mid-1950s as an "auditory and ideological nonsense".[1] His solution lay in the use of stochastic techniques, where "what will count will be the statistical mean of the isolated states and of the processes of change which the components undergo at any given moment. Thus the macroscopic effect can be controlled by means of the movements of objects selected by us."[2] This verdict was confirmed by Gottfried Michael Koenig twenty years later: "it appears that the trouble taken by the composer with series and their permutations has been in vain; in the end it is the statistic distribution that determines the features of a composition."[3] Steve Reich sought to emphasise the audibility of formalised processes in the 1960s, stressing their intelligibility: "I am interested in perceptible processes. I want to be able to hear the process happening throughout the sounding music."[4] This perceptibility of an aural process has been argued as inherently congruous with the human experience of music: Michael Edwards likens it to the very way that the mind engages with composed sound, involving "both the enjoyment of the sensual sonic experience and the setting up of expectations and possibilities of what is to come."[5]

My own approach to composition has long been concerned with both the formalisation and aural clarity of complex processes. This has led to a focus on behaviour as the key compositional element, being a mode of action – exhibited by one or many instruments – sufficiently discrete and characteristic that subsequent alterations, due either to intrinsic states of change or extrinsic processes of transition towards differing behaviours, are audibly apparent. Initially informed by

---

[1] Xenakis 1966, p. 78.
[2] Ibid.
[3] Quoted in Berg 2009, p. 78.
[4] Reich 2004, p. 34.
[5] Edwards 2011, p. 58.

my postgraduate studies at the Royal Conservatory in the Netherlands (particularly algorithmic composition with Paul Berg, at the Institute of Sonology), this focus on behaviour led to the introduction of computer programming into my hitherto largely intuitive methodology, a skill I had learned as a child but not used since. While at first the computer acted as little more than a time-saving computational tool, I soon developed an environment where complex experiments could be rapidly tested and assessed. The computer became a kind of creative 'familiar' allowing me fundamentally to reappraise my musical understanding and compositional method. This echoes the experiences of Koenig who, noting that whereas "[e]very rule of composition that can be formulated can also be programmed and carried out by a computer [...] one does not programme known rules of composition but also tries to find out whether events not yet expressed in the form of rules are feasible. The computer thus has a stimulating effect on research in composition theory".[6] For the last fifteen years or so, my compositional methodology has been rooted in a computational, stochastic, algorithmic process-based approach. Using an updated version of BBC BASIC,[7] programs were at first created from scratch for each new composition. Increasingly, though, similarities and overlaps in method and technique have meant that new programs borrow from existing ones. The extent to which these programs incorporate aspects of my creative thinking has grown, reducing the amount of subsequent work required to be applied to their generated output, effectively emerging from the program as almost-ready compositions. The programs have thereby gradually increased the extent to which they embody my compositional practice, approaching a complex realisation of Steve Reich's idea of "a compositional process and a sounding music that are one and the same thing."[8]

These developments are discussed in Chapters 1 and 2, which respectively present the theoretical background and various preliminary works submitted in the portfolio that utilise algorithmic processes in diverse ways, generating both portions of material as well as entire compositions. Chapter 3 elaborates on the considerable development that led to the main work, *Cloud Triptych*. It describes the necessity for and approach taken to create a synthesis of the diverse elements described in the preliminary works, leading to the development of a single, large-scale, multifaceted computer program unifying and embodying them all (named, for reasons that will become clear, CloudCube).

[6] Koenig 1970, p. 3.
[7] Specifically, BBC BASIC for Windows, a re-creation by Richard Russell of the programming language originally written in 1981 for the BBC Microcomputer.
[8] Reich 2004, p. 35.

Method

In this section I outline the basic principles of the compositional method developed in the preliminary works, and brought to fruition in *Cloud Triptych*. This method has the aim of utilising perceptible algorithmic processes. It has three primary aspects:

- *Discrete musical behaviours* with probabilistic epicentres;
- *Processes of transition* between these behaviours;
- *Stochastic generation of material* using a modular approach to musical parameters.

Described in an abstract form, the method is very straightforward.

Discrete Musical Behaviours

A *behaviour* is here defined as a collection of rules that determine all relevant aspects of the desired musical material, such as tempo, rhythm, quantisation (how time/beats are subdivided), harmony, dynamics and so on. These parameters are either fixed or determined randomly either from a list or from within predefined upper/lower limits.[9] Table 1.1 illustrates examples of both types of parameter, fixed and limited.

Table 1.1. Fixed and limited parameters within a typical behaviour.

parameter	rule	type
tempo	♩ = 70	fixed
quantisation	simple/triplet subdivisions	limited
rhythm	from ♪ to 𝅗𝅥	limited
pitch	from C1 to C3	limited
articulation	staccato	fixed
dynamic	from ***ppp*** to ***p***	limited

Parameters with limits are determined by use of a random number, the range of which is either the same as the number of possible outcomes (for probabilities, between 0 and 100) or, more usually in my work, between 0 and 99. This number is then mathematically converted to an equivalent value between the upper and lower limits. For example, using the dynamic example from Table 1.1:

Table 1.2. Conversion of random number into limited value.

random number	number of options	converted value
41 {0,99)	3 {***ppp***, ***pp***, ***p***}	2 = ***pp***

[9] In practice, it is often more helpful to regard 'fixed' parameters as having upper/lower limits that are the same. Random decisions between limits are always inclusive.

It is important that the range from which the random numbers are drawn is relatively large, otherwise this has a negative impact on the probabilities of the possible options, rendering them unequal.[10] Parametric limits can themselves change over time; these changes are according to a predefined, time-dependent algorithmic process, of the kind which Koenig described as a 'tendency mask'.[11]

Once its rules have been defined, the behaviour is assigned a fixed point with respect to the work's timeline, a point referred to as its *node*. As an imaginary point or *cursor* moves along the timeline towards the node, the probability of the behaviour increases, reaching a probability of 100% at the point on the timeline where the cursor coincides with the node, as shown in Fig. 1.1.

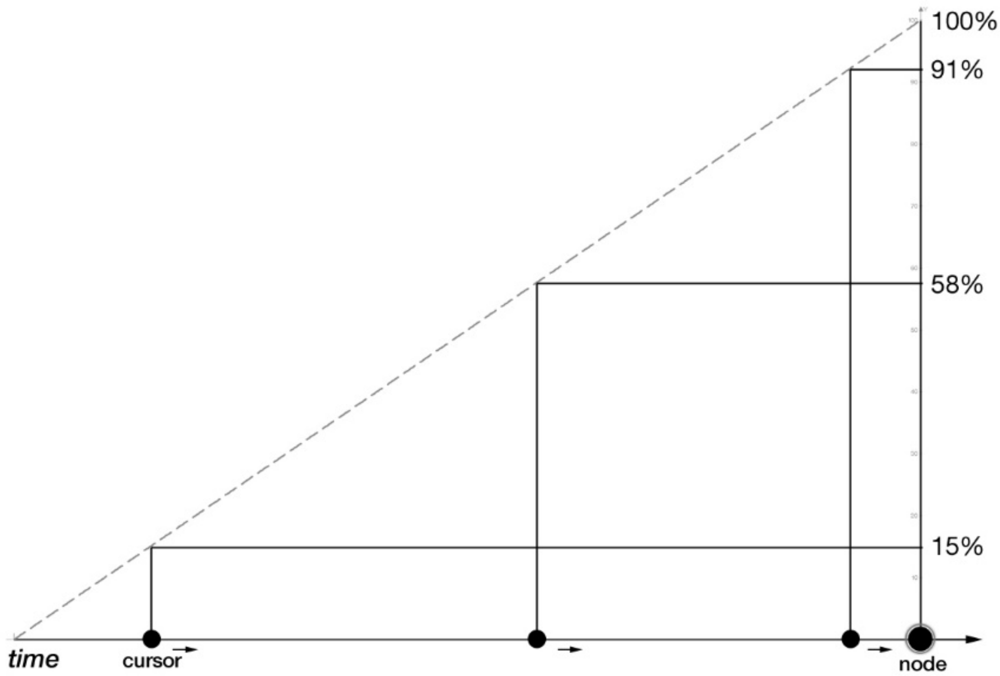

Fig. 1.1. Increasing probabilities as the cursor approaches the time position of the behavioural node.

The probability decreases accordingly as the cursor moves beyond the node further along the timeline. Proximity and probability are therefore directly proportional; in this way, the node represents the epicentre of the behaviour.

[10] For example, if the range of random numbers comprised only 0–9, and this was used to choose between four options, these options do not have an equal probability due to the fact that the range of random numbers (10) is not commensurate with the range of options (4). In this instance, two of the options will have 20% probability and two will have 30% probability (which options have which probability is determined by how the calculations are rounded). Using a range from 0–99, in conjunction with carefully-calculated rounding thresholds, produces results in which all options have an essentially equal probability. Even more accurate results would be produced using a larger range, e.g. 0–999.

[11] Koenig 1970, p. 8. Among the procedures Koenig created for *Project 2* (given names in capital letters), "TENDENCY works with a mask, the edges of which move [...] independently of each other." (Ibid.) A second procedure, ALEA, then selected randomly from within the limits created by the tendency mask.

Processes of Transition

The purpose in devising this method is to facilitate transitional processes, and to that end a minimum of two behaviours are required. While these behaviours have a maximum probability of 100%, for this to be meaningful their *range of influence* – the extent of their probability, from their epicentral node to the position(s) on the timeline where the behaviour's probability falls to 0% (or some other minimum) – must also be defined. In a system of two behaviours, positioned at the start and end of the timeline and where each has a linear range of influence extending for the entire duration, the respective probabilities for each behaviour are as shown in Fig. 1.2. Behaviour A descends from 100% as the cursor moves through time away from it; behaviour B does the reverse, ascending to the same.[12]

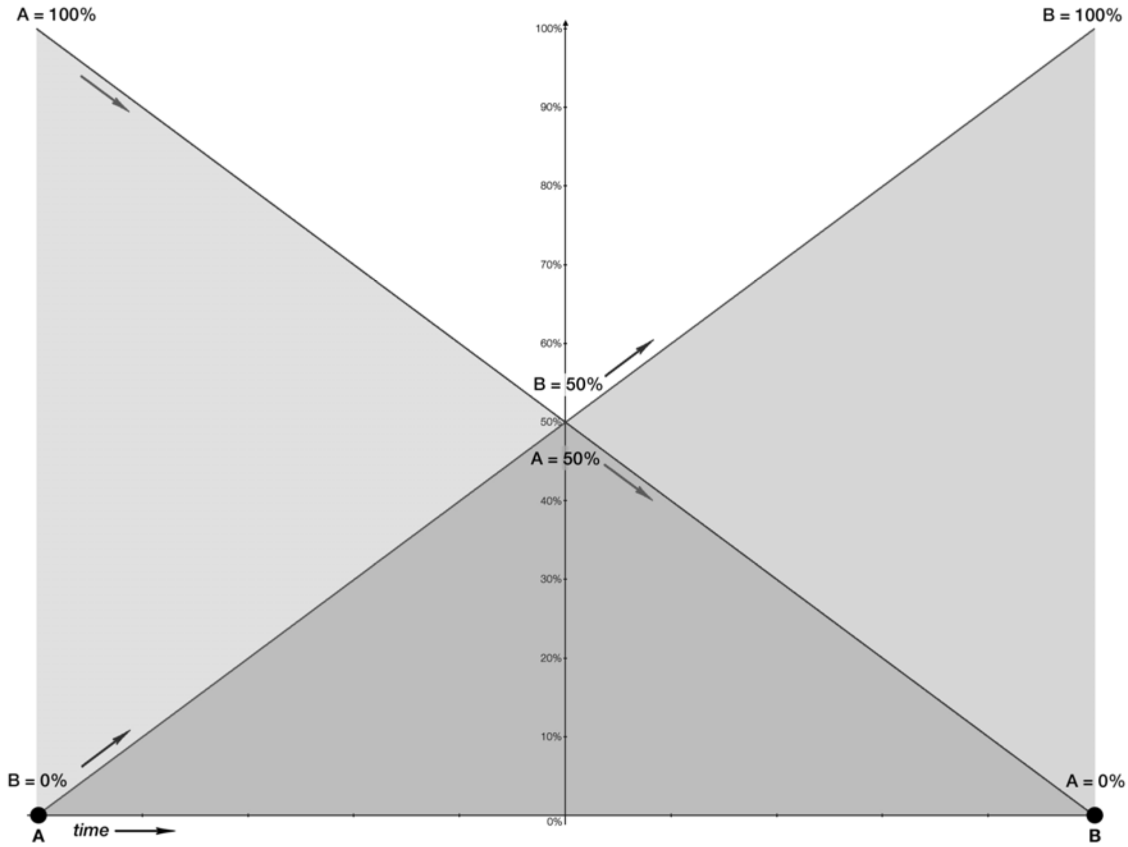

Fig. 1.2. Probability distribution in a two-behaviour linear system.

In the above system, the transition between behaviours is linear, as the probability distribution changes at a constant rate. If the rate of change is variable, following a curved rather than a straight path, then the result is a non-linear transition, an example of which is shown in Fig. 1.3.

[12] For clarity, probabilities are referred to as percentages, but may equally be regarded as equivalent values between 0 and 1. The total probability of the system at any moment is always equal to 100% (or 1).

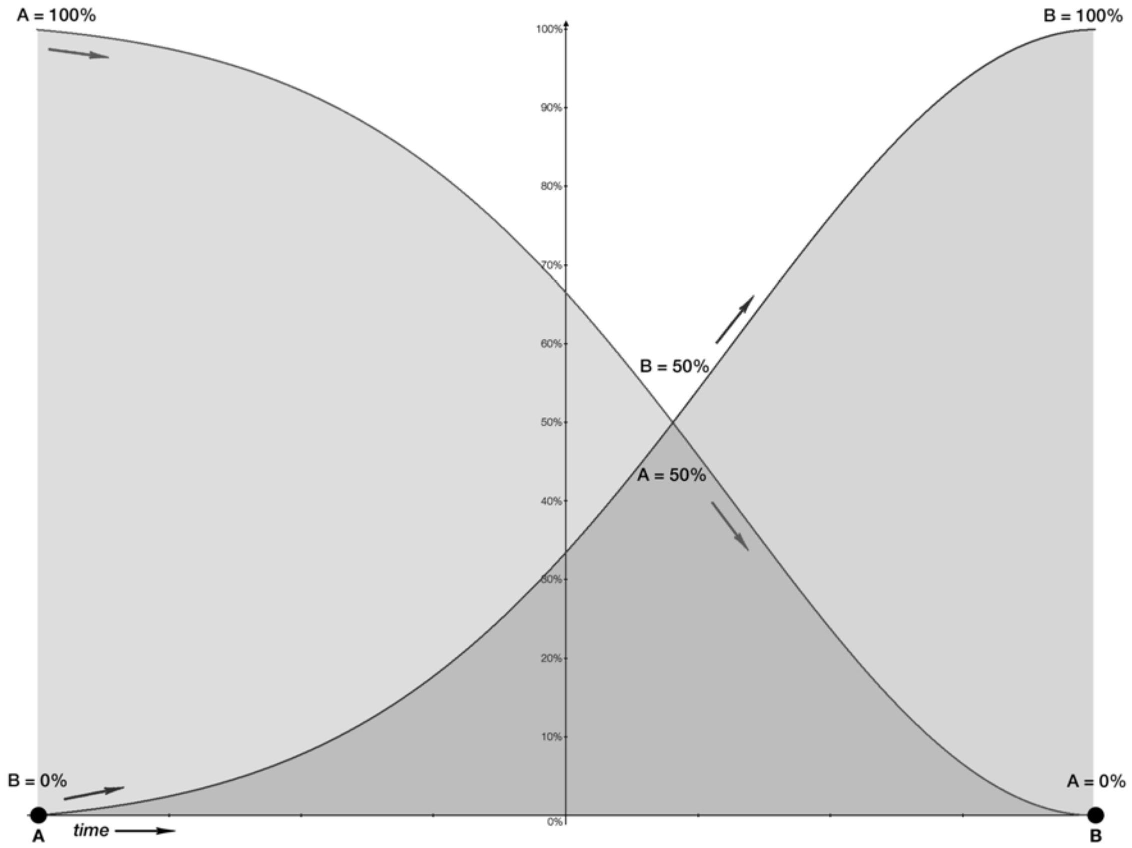

Fig. 1.3. Example probability distribution in a two-behaviour non-linear system.

Of paramount importance with respect to these transitions is their aural perceptibility. In part, I subscribe to Steve Reich's broad view, expressed in the late 1960s, that "[t]o facilitate closely detailed listening a musical process should happen extremely gradually. [...] We all listen to the process together since it's quite audible, and one of the reasons that it's quite audible is because it's happening extremely gradually."[13] In tandem with this is the distinctiveness of the behaviours between which the transitions are taking place. The intended result is not unlike the visual process in the *Metamorphosis* woodcuts of M. C. Escher, a continual transition alighting on instances of distinct behaviour.[14]

[13] Reich 2004, p. 35.
[14] cf. Locher and Veldhuysen 2000, pp. 1–11.

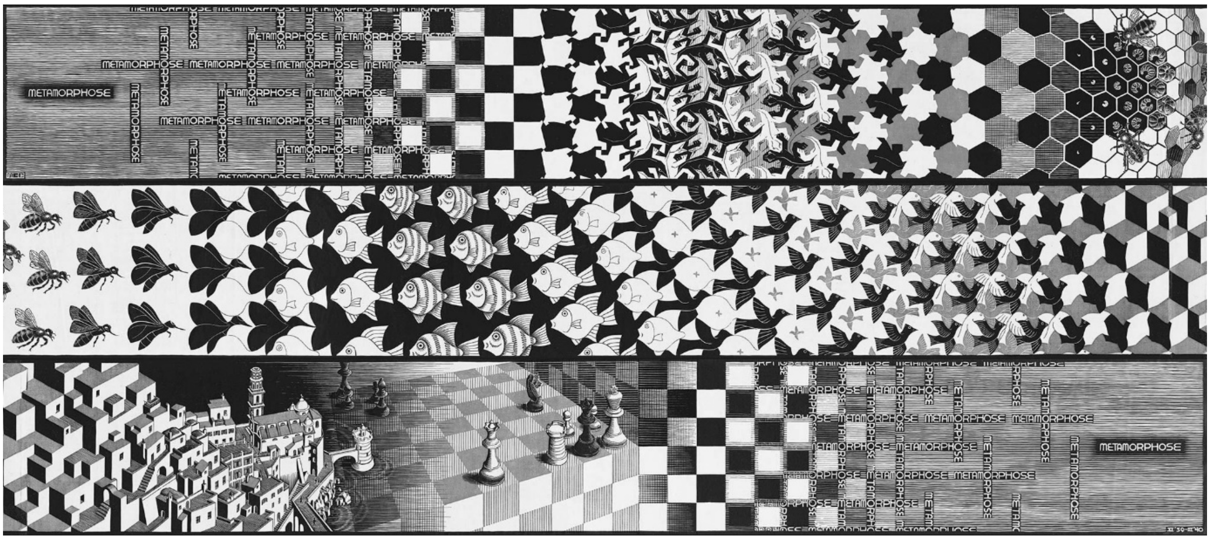

Illus. 1.1. M. C. Escher, *Metamorphosis II*, 1940.

This approach also bears an aesthetic resemblance to Reich's 'phase' works of the late 1960s and early 1970s, in particular *Piano Phase* (1967) and *Drumming* (1971). In these pieces Reich articulates a compositional process in which the music moves in and out of rhythmic alignment, illustrated in Ex. 1.1. Piano 1 (upper stave) continually repeats a rhythmic cell at a fixed tempo, unchanging for the duration of the entire piece. Piano 2 (lower stave) varies between sections where its rhythms are aligned with piano 1 – the fully notated bars – and episodes where it slowly accelerates (marked in the score as 'a.v.s.': accelerate very slightly), notated with dashed lines to indicate the lack of rhythmic alignment.

Ex. 1.1. Steve Reich, *Piano Phase*, sections 4–6.

By this means Reich establishes centres of clarity – akin, in the context of my work, to the nodal points of a behaviour – contrasted with episodes of complex metric asynchronicity that are clearly heard as transitioning between these centres. My approach is additionally informed by the kind of drifting coalescence found in large-scale works of ambient electronic music, such as Kenneth Kirschner's *July 17, 2010* (2010) and The Hafler Trio's *scissors cut arrow* (2004). The relationship between stability and flux operates in these works over long time-spans (each piece lasts around two hours), creating a fascinating disorientation in which the clarity of musical identity –

perceived as a kind of behavioural stasis – is continually being questioned by a sense of transition toward something else. In my own work, using the method outlined above, it is essential that the behaviours are designed to be sufficiently distinct and characteristic in their own right, else a clearly perceptible sense of transition between discrete musical identities will be weakened or lost entirely. However, I am interested in a nuanced engagement with this sense, of the kind demonstrated in these works, allowing for elements of uncertainty rather than a strict 'black and white' approach to the use of behaviours and transitions.

It will be clear from the above methodical outline that the composition process begins with the definition of musical behaviours and design of the processes of transition between them. This closely corresponds to the second and third of Xenakis' "Fundamental phases of a musical work" where, following an initial phase of preparation, the composer engages in "Definition of the sonic entities" (behaviours) and "Definition of the transformations which these sonic entities must undergo in the course of the composition"[15] (transitions).

Stochastic Generation of Material

To achieve these ends a computer program is written to carry out the process in its entirety.[16] In relation to a (linear or non-linear) two-behaviour system as just described, the program's actions can be briefly summarised:

1. Calculate the probabilities of behaviours A and B at the present timeline position.
2. Compare a random number between 0 and 100 with these probabilities: if less than behaviour A's probability, behaviour A is selected, otherwise behaviour B is selected.
3. Generate material according to the rules of the selected behaviour.

Having chosen a behaviour as outlined above, the amount of material to be generated – a quantity hereafter referred to as *structural duration*, usually specified as a number of beats – is one of the first parameters to be defined, and can range widely from just a single note or gesture to a lengthy portion of music. The program then proceeds to generating material, dealing with one musical parameter at a time. This is done via a sequential chain of modules, each concerned with a single parameter, such as rhythm, pitch, articulation and dynamic. Each module generates material according to its behavioural rules and this is then passed to the next module, down the chain until the music is completed, a process illustrated in Fig. 1.4.

---

[15] Xenakis 1992, p. 22.
[16] This in turn corresponds to the fourth, fifth and sixth of Xenakis' phases: "4. Microcomposition (choice and detailed fixing of the function or stochastic relations of the [sonic identities ...] 5. Sequential programming [...] 6. Implementation of calculations", Xenakis 1992, ibid.

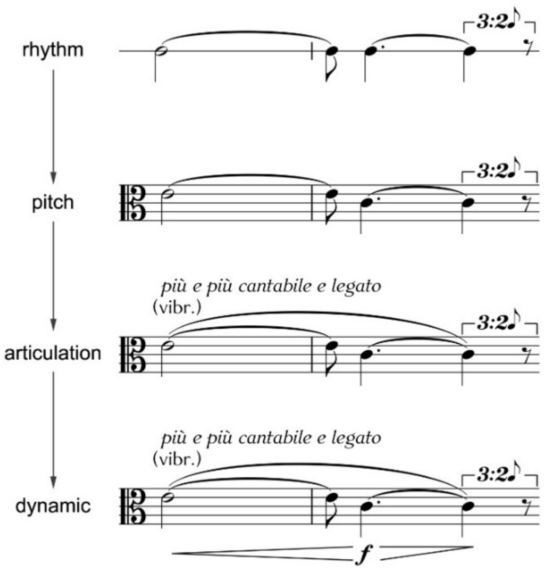

Fig. 1.4. Modular generation of musical material.[17]

The music is thereby put together from a large collection of small, stochastically-determined pieces of material which, at the macro-level, conform to a predesigned probability distribution (as in Fig. 1.2 and Fig. 1.3). This bears relation to the approach taken by Xenakis in some of his early major works, particularly *Metastaseis* (1954) and *Pithoprakta* (1956), with regard both to the independent treatment of musical parameters (a development of serialism's focus – from Xenakis' perspective – on frequency, intensity and timbre)[18] as well as the generation of material according to functions of probability (*Pithoprakta* translates as "actions through probability"). But my approach is much more closely connected to the more formalised modular method, also developed from serial considerations, devised by Gottfried Michael Koenig for his *Project 2*, the program of which "involves seven parameters which are named: Instrument, Harmony, Register, Entry Delay, Duration, Rest, Dynamics. [...] The composer determines the compositional rules on the basis of these selection principles and combinations of them, but also by indicating the order in which the parameters are to be computed, for the parameters depend on one another in such a way that, for instance, no tones may be selected which a given instrument cannot play, and vice-versa, according to the order determined by the composer. At all levels of decision, chance is given opportunities to operate to a greater or lesser extent, so that any number of variants can be composed according to the data".[19] For Koenig the order in which these parameters were addressed was somewhat fluid,

[17] The example is taken from my string sextet *Intense quick dream of sentimental groups with people all possible characters amidst all possible appearances*, bb. 32–33, viola 1; see Chapter 2, p. 32.
[18] cf. Matossian 1986, pp. 65, 85.
[19] Koenig 1970, p. 6.

but I have opted to fix the modular order as I believe it to be the most logical way of constructing musical material.[20]

In contrast to the smooth, gradual transitions of a linear system, non-linear systems produce more abrupt transitions, akin to the scientific phenomenon of 'phase transitions',[21] with a clear demarcation between the two behaviours. Non-linear transitions therefore have structural implications at the macro-level, producing a sectional behavioural distribution with elements of anticipation, overlap and bleeding between the sections.

When applied to the composition of works for more than one player, material for each instrument is created individually, in order to utilise the variety inherent in the stochastic nature of this process. Each part thereby undergoes the same fundamental processes at the same general rate but in an entirely individual way, producing an overall texture of multiple, unique, simultaneous transitions between behaviours. This resembles what Michael Parsons has described as "[t]he idea of one and the same activity being done simultaneously by a number of people, so that everyone does it slightly differently, 'unity' becoming 'multiplicity'".[22] Michael Nyman calls this 'people processes'. However, whereas for Parsons the realisation of the process is performative, and a means to considerable economy of notation (utilising the "variety ... from the way everyone does it differently" due to their "natural individual differences"[23]), in my own work the realisation is compositional, exploiting the variety from processes with stochastic parameters.[24]

My approach has obvious similarities to the 'micropolyphony' of György Ligeti's *Atmosphères* (1959) and *Lontano* (1967)[25], works in which the role of individual instruments often becomes that of a representative constituent of a larger textural mass. In such instances – usually in the form of dense canons (Ex. 1.2) – the orchestra as a whole exhibits a coherent generalised behaviour, being the accumulative product of all the players' actions, which are individually unique yet very closely related to all the others. As such, while each instrument is essentially insignificant in isolation it is nonetheless vital to the perceived cohesion of the collective sonic identity.

---

[20] Furthermore, it echoes the way in which I used to compose prior to adopting an algorithmic approach, which usually consisted of sketching out rhythms, which were then elaborated with indications of pitch contour and finally assigned specific pitches, articulations and dynamics.

[21] Phase transitions are abrupt shifts in state within thermodynamic systems, such as when water suddenly freezes into ice or boils/evaporates into gas.

[22] Quoted in Nyman 1999, p. 6.

[23] Ibid.

[24] This does not, however, preclude the possibility of 'real-time' interactions, both deliberate and accidental, of a kind not dissimilar to those Parsons is exploring. These can be actively incorporated into behavioural definitions, plus behaviours shared between multiple players will inevitably produce a concomitant effect of implied interactions.

[25] cf. Steinitz 2003, pp. 103–110.

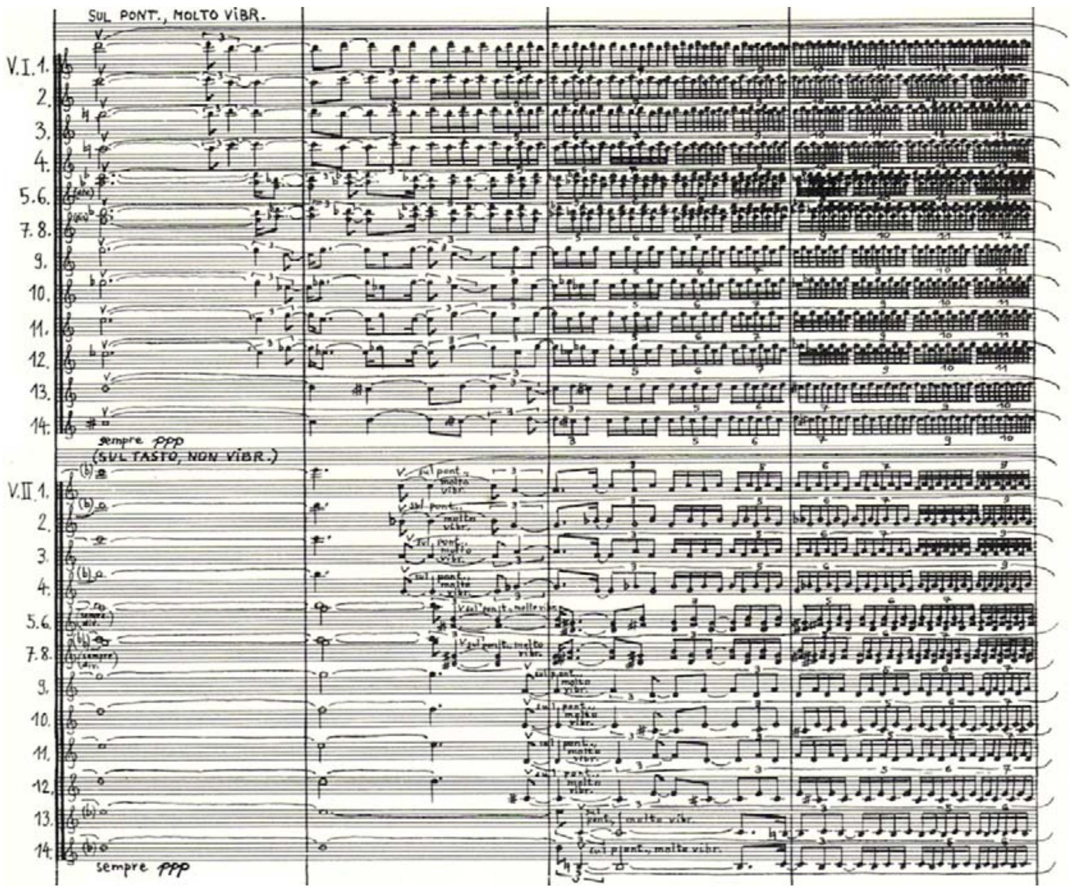

Ex. 1.2. György Ligeti, *Atmosphères*, bb. 23–26, violins I and II.

Closely related to this is the sonorist approach displayed in Krzysztof Penderecki's *String Quartet No. 1* (1960) and *Threnody to the Victims of Hiroshima* (1960).[26] Penderecki here puts special emphasis on discrete behavioural compartmentalisation, again constructed by utilising unique operations that together comprise a unified mode of action. In these works, this tends to take the form either of a shifting sound mass made up of a clustered agglomeration of individual pitches (Ex. 1.3) or a texture formed from the players rapidly executing groups of gestures so arranged as to be the aural equivalent of a random selection from a small list of options (Ex. 1.4).

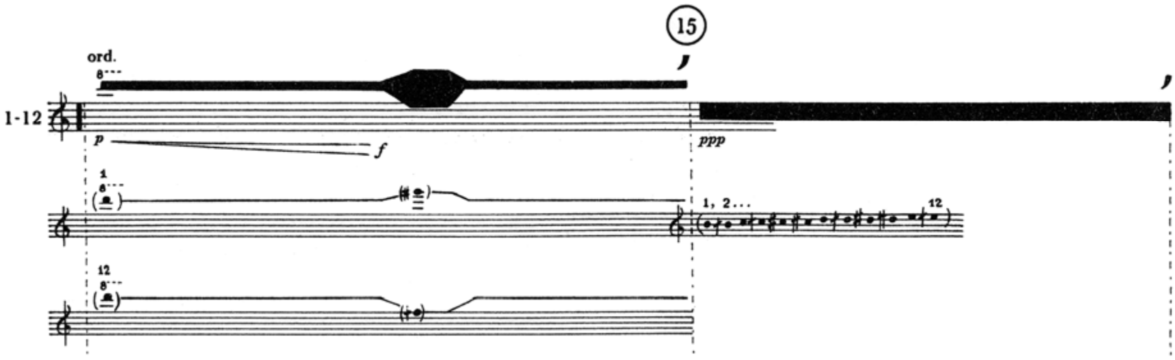

Ex. 1.3. Krzysztof Penderecki, *Threnody to the Victims of Hiroshima*, bb. 13–15, violins 1–12.

[26] cf. Schwinger 1989, pp. 124–131.

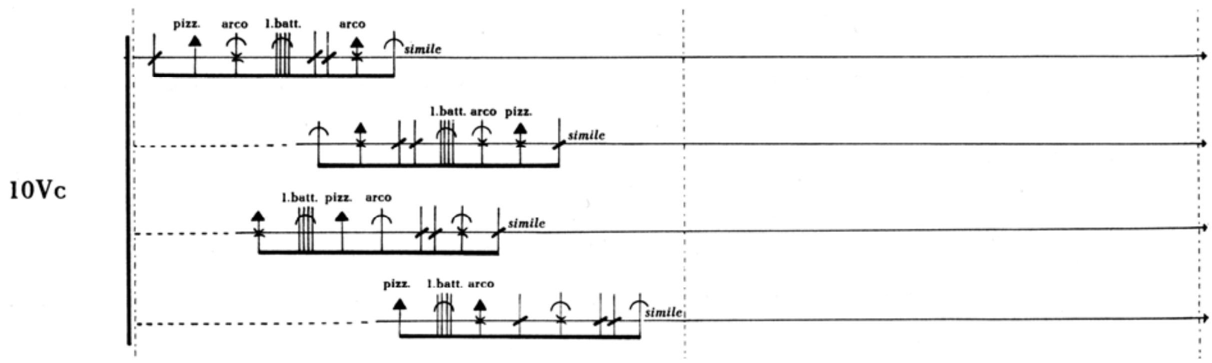

Ex. 1.4. Krzysztof Penderecki, *Threnody to the Victims of Hiroshima*, bb. 6–7, cellos.[27]

These examples again illustrate how individual players' actions are of secondary importance to the combined effect of many instruments, reinforced further by Penderecki through compartmentalising the works into discrete behavioural sections. This relationship between the individual and the mass in both the presentation and perception of distinct musical identities is of fundamental significance to the way I have developed my own approach to the design of musical behaviours.

Above all, though, my compositional method aspires to the evocative description of a new music as envisaged by Edgard Varèse in the mid-1930s:

> ...taking the place of the linear counterpoint, the movement of sound-masses, of shifting planes, will be clearly perceived. When these sound-masses collide the phenomena of penetration or repulsion will seem to occur. Certain transmutations taking place on certain planes will seem to be projected onto other planes, moving at different speeds and at different angles. There will no longer be the old conception of melody or interplay of melodies. The entire work will be a melodic totality. The entire work will flow as a river flows. [...] In the moving masses you would be conscious of their transmutations when they pass over different layers, when they penetrate certain opacities, or are dilated in certain rarefactions. [...] An entirely new magic of sound! [28]

Having explained the three aspects of my compositional method, I now discuss some of the finer points pertaining both to specific modules as well as to the method as a whole.

Quantisation and Rhythm

Rhythms are determined via a three-step process, starting with a quantisation module. Individual beats can be subdivided in various ways according to different kinds of underlying tuplet relationships (triplets divide beats into three, quintuplets into five, etc.). I have employed various

[27] Earlier editions of the score carried this explanatory footnote, omitted in later editions: "Each instrumentalist chooses one of the 4 given groups and executes it (within a fixed space of time) as rapidly as possible."
[28] Varèse and Wen-chung 1966, p. 11.

quantisation grids, each of which divides an individual beat into a fixed number of tuplet subdivisions, from three up to a maximum of seven. There is also a grid including just grace notes, the appoggiatura and acciaccatura, taken here to mean just after and just before the beat respectively.[29] The grids are referred to using the terms *grace*, *simple*, *triple*, *quint*, *sext* and *sept*.[30] All of the grids can be seen in Fig. 1.5, which consists of a single beat; the large numerals indicate the specific grid (3–7), the small numerals indicate the decimal position within the beat where the subdivisions occur. The grace grid is indicated by the infinity symbol (∞), as its durations, computationally speaking, are infinitely short.

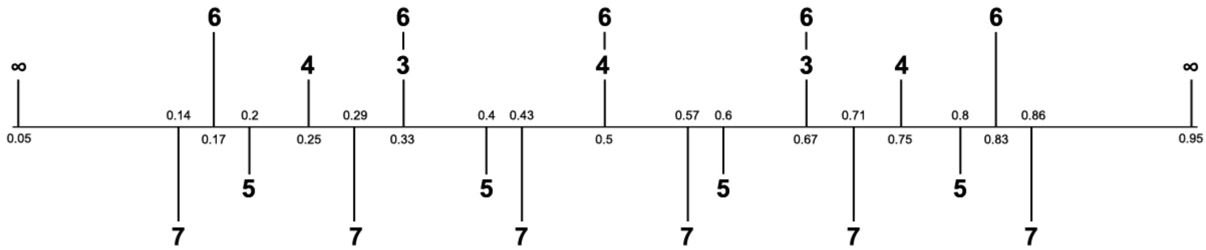

Fig. 1.5. The collection of quantisation grids.

Each behaviour has a rule concerning which grids are available to it. Having established this, the quantisation module then combines the relevant grids into a *quantisation lattice*; if all grids are included, as in Fig. 1.5, the result is twenty distinct points of beat subdivision (including beat position 0.0, occurring on the beat; this is implied in all the quantisation grids).[31] The minimum number of grids used in the lattice is zero; this would simply quantise all rhythms to the nearest beat.

Having determined the quantisation, the process moves to a rest probability module, establishing the likelihood that durations will be rests rather than notes. This is specified as a numeric value between 0 and 100. The quantisation lattice and rest probability are then passed to a rhythm module, which specifies durational limits. The limits are decimal values with respect to crotchets: therefore, with limits of a quaver and a minim, these would be specified as 0.5 and 2 respectively. The rhythm module chooses random durations between the limits; due to the random durations having decimal values they are at this stage 'floating point' rhythmic positions, which are now quantised according to the nearest points on the quantisation lattice. This process is shown in Fig. 1.6, using a quantisation lattice comprising the sext and sept grids.

[29] To clarify, events happening at the decimal beat position 0.05 occur *after* the appoggiatura, which is either tied to the previous note or constitutes a grace note rest at position 0.0.
[30] Although its name and position in this list suggests two subdivisions, the 'simple' quantisation grid is invariably used to divide a beat into four subdivisions, as shown in Fig. 1.5.
[31] Beat position 1.0 is not used; this is regarded as position 0.0 of the following beat.

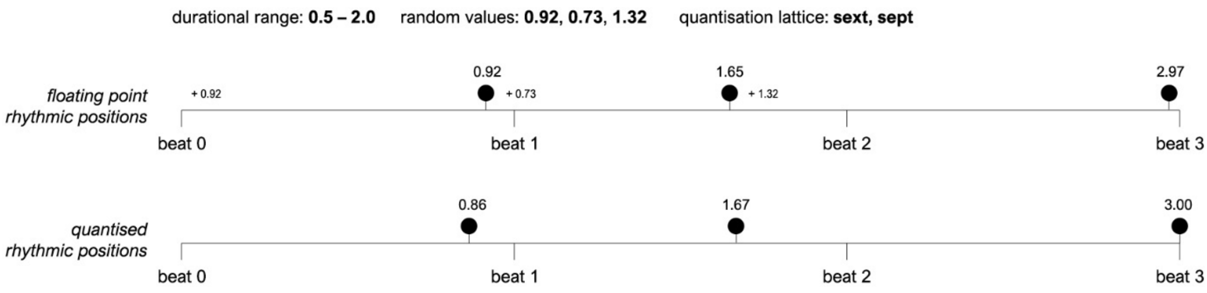

Fig. 1.6. Random durations with floating point and quantised rhythmic positions.

As each duration is quantised, the rest probability is used to determine whether it is a note or rest.

It should be noted that this process can and regularly does create instances where rhythmic subdivisions from more than one quantisation grid occupy a single beat. Where these occur, an intuitive decision is made on how best to resolve this, according to context.[32] The main criterion for these manual interventions is to make the minimum change possible to the output from the program. Fig. 1.7 illustrates a typical example, in which two notes are initially assigned beat positions from different quantisation grids (66.33 = triple, 66.80 = quint). Shifting 66.33 to a quint position (option 1) involves an addition of 0.07, whereas shifting 66.80 to a triple position (option 2) involves a subtraction of 0.13; therefore the first option is selected.

Fig. 1.7. Manual adjustment of quantisation positions.[33]

Pitch and Harmony: Weighted Tonality

I use the term 'weighted tonality' to denote an approach to pitch and harmony implemented in the majority of the pieces in the portfolio. In my work there has always tended to be a compositional imperative whereby, at any particular point and however briefly, one pitch-class is regarded as a tonic or fundamental, with all other pitch-classes heard in relation to it. This led me to construct a series whereby the eleven remaining pitch-classes are arranged in an order of increasing dissonance – or, more accurately, increasing *displacement* of a fundamental pitch-class that I refer to as the *primary*. The construction of such a series was in part inspired by Paul Hindemith's 'Series 1', expounded in his 1937 treatise *The Craft of Musical Composition*. Hindemith was concerned primarily with an arrangement that reflected Pythagorean frequency relationships, comprising "[a] single

[32] It would be straightforward to remove the possibility of such instances, forcing each beat to comprise just a single quantisation grid, but I find the need for intuitive involvement stimulating, so this particular 'flaw' has been allowed to remain.
[33] The example is taken from my orchestral piece *Cloud Triptych*, b. 92, flute 2.

tone conceived as the root of a scale; the chromatically arranged twelve-tone series born of the tensions set up by the juxtaposition of vibrating units in the proportions of the simple numbers from 1 to 6".[34] He likened this to a solar system, where the tonic "is the sun, surrounded by its descendant tones as the sun is surrounded by its planets. Series 1 shows us the distance of the planets from the central star. As the distance increases, the warmth, light, and power of the sun diminish, and the tones lose their closeness of relationship."[35] However, the emphasis on displacement rather than dissonance in my own work is crucial. For example, in conventional tonality the subdominant is highly consonant in relation to the tonic, and unsurprisingly therefore is placed third, following the octave and dominant, in Hindemith's Series 1. Yet harmonically the perfect fourth exerts a powerful displacing force (since the tonic acts as its dominant), and is therefore positioned much farther away from the primary in my own series.

Despite being qualitatively different from conventional harmony, there is usually an aim when using weighted tonality to allude, however vaguely, to either minor or major tonal traits, and for that reason there are two weighted tonality pitch series, shown in Table 1.3.

Table 1.3. Major and minor weighted tonality pitch series.

	1	2	3	4	5	6	7	8	9	10	11	12
major	PRIMARY	V	III	VI	II	VII	IV	vi	ii	vii	iii	♭V
minor	PRIMARY	V	iii	vi	ii	vii	IV	VI	II	VII	III	♭V

With weighted tonality, it is possible to restrict the range of pitch-classes that can be chosen in order to create certain harmonic colours. I use the term *chromatic index* to denote the maximum number of degrees away from the primary that can be used. A chromatic index of two, using pitches 1 to 3, produces purely triadic material, whereas a chromatic index of five, pitches 1 to 6, leads to a more rich but still essentially diatonic harmonic palette. Employing the more distant pitches 7 to 12 results in a chromatic palette that, as in conventional tonality, increasingly obfuscates the primary.[36] Likewise, an avoidance of the primary, just using pitches 2 to 5 for example, could lead either to the harmonic implication of an underlying fundamental, or indeed shift the harmonic focus away from it. This method of arranging the twelve pitch-classes as a 'displacement series' with respect to a primary, and then controlling the range of the series currently available, I refer to as 'weighted tonality' both because one is able to control the apparent 'weight' of the primary (due to how reinforced it is by non-displacing adjacent pitches) and also

[34] Hindemith 1945, p. 53.
[35] Op. cit., p. 57.
[36] Within the context of weighted tonality, the terms 'diatonic' and 'chromatic' specifically refer to these 'displacement' extents with respect to the primary: thus 'diatonic' involves pitches 1 to 6, closest to the primary, whereas 'chromatic' includes the more remote pitches, 7 to 12.

because more remote pitch-classes can be heard to have increasing 'weights' due to their ability to harmonically displace the primary. It should be noted that although the primary can be regarded as a form of tonic or fundamental, I do not deem it essential within weighted tonality for its pitch-class to be reinforced or even actually present in lower octaves, as is the case in conventional tonality. Furthermore, while the role of register with respect to the perceptible clarity (or otherwise) of a harmonic scheme is important, this is not considered within the weighted tonality method outlined here, but must of course be borne in mind when designing musical behaviours.

The foregoing discussion perhaps indicates pitch and harmony have fundamental importance in my work. However, my compositional interest in this area is more concerned with the juxtaposition of pitch and noise and the reappraisal of what defines each (and their boundaries). This is something I have also explored significantly in my electronic music, particularly in *the Ceiling stared at me but i beheld only the Stars* (2009), the *Simulated Music* cycle (2010) and my ongoing series of *Studies* (2014–16).

Layers of Determinism

The end result of the modular process is output from the computer program, which for speed has usually been presented in the form of a neatly-ordered list, providing a breakdown of all of the material generated. The list forms the basis for notation in the Sibelius program. This basis is subject to interpretation and alteration, highlighting the fact that intuition and creative whim remain integral to the composition process. In other words, the processes of transition and stochastic generation outlined above are just one of several 'layers of determinism'. They are, in fact, the innermost layer, where automated decisions are made within the system according to the rules of the predefined behaviours and the transitional processes that comprise and control that system. This layer can be described as being *in machina*. Tempo is a parameter that, unless the music is unsynchronised, cannot be determined independently for each player, but must apply to all players the same. It therefore exists in a secondary layer where decisions are made outside the central system, again according to predefined rules and processes but operating either irrespective of activity within the system (passive) or in relation and/or response to that inner activity (active). This constitutes an *ex machina* layer of determinism. Both of these deterministic layers are subject to an ad hoc, intuitive, outermost layer of decision-making which can be appropriately described as *deus ex machina*.

A vital aspect of the *deus ex machina* layer of determinism is the composer's prerogative to reject any or all results produced by the inner layers, or to put it another way, the option of running processes repeatedly until a desired outcome is achieved. Each time the program is run the results are, of course, technically 'correct', but there may be concerns regarding the material itself and/or its juxtaposition with other material, so judgement and final decisions remain here with the composer. This prerogative was exercised to a considerable extent in the composition process of my string sextet *Intense quick dream of sentimental groups with people of all possible characters amidst all possible appearances* (see p. 36). The parts were created in order, beginning with violin 1. As the program created new material for each successive player, initial results were on numerous occasions rejected as a result of considering the material on its own terms in addition to how it would sit alongside the existing music. This process was quick and instinctive, and there was no precise criteria for rejection or acceptance, though coincidental appearances of 'imitation' were regarded positively. A simple example is the use of tremolando in bars 28–31, being passed quickly between the violins and cello.

While it may seem contradictory to design self-contained algorithms and then critique their output, I am in agreement with Xenakis:

> When I used programs to produce music [...], the output sometimes lacked interest. So I had to change. [...] Other composers [...] have acted differently [, declaring:] The machine gave me that so I have to respect it. This is totally wrong, because it was he who gave the machine the rule. [...] A rule governs just one aspect of the music. If the composer says: I've kept to the rules, everything is under control, he's wrong [...]. A composer can never boast that he has everything under control.[37]

Most fundamental of all to the *deus ex machina* layer are the preparatory composition stages, where the system as a whole – behaviours, time positions, probabilities and everything else – is designed and constructed. In this regard there is a continual, experimental back-and-forth between the conceptualisation of the system and its realisation, not only in terms of trial and error but also of mutual influence, each inspiring the other. John Bischoff has described this approach as 'sculptural':

> Starting from an initial idea or perception, I composed each piece by adding a feature at a time, making aesthetic choices by repeated close listening. As a result of this process, some features were refined, others were dropped. Finally, the identity of each piece emerged. Although this way of working within a medium is common in visual art, it is unusual within the field of computer programming. It is generally thought that programmers implement their ideas on the computer with as few unforeseen developments as possible, and that ideas flow in only one direction, from the operator to the machine. From the vantage point of an artist,

[37] Varga 1966, p. 201. The *in machina* and *ex machina* layers operate with respect to time, while the *deus ex machina* layer does not; as such, there are certain similarities between these layers of determinism and Xenakis' 'in-time' and 'outside-time' structures. (cf. Xenakis 1992, pp. 160-161).

> though, it is just as easy to see the flow going in the opposite direction: the medium reveals itself as the artist proceeds, and such revelations shape the direction in which the artist continues. This interaction between artist and medium, and the intimacy it suggests, is no different for a computer artist.[38]

Therefore, within my compositional method there is an ongoing dialogue between preconception and spontaneity, between formalised process and intuition.

[38] Bischoff 1991, p. 37.

Chapter 2. Preliminary Works

I here present five preliminary works that illustrate aspects of my exploration of stochastic transitions between musical behaviours, and their gradual evolution and sophistication.

'unredeemed' self-)portrait (in the form of an eagle (2008) for solo flute

This work is one of my earliest experiments with contrasting musical behaviours. The second in a planned series of four solo works,[39] it took as inspiration an interpretation by Michael Sadgrove of the large tapestry that hangs at the east end of Coventry Cathedral, Graham Sutherland's *Christ in Glory Surrounded by the Tetramorph*. Each of the four instruments – the others being oboe, clarinet and trombone – was assigned a unique trajectory, shown in Fig. 2.1.

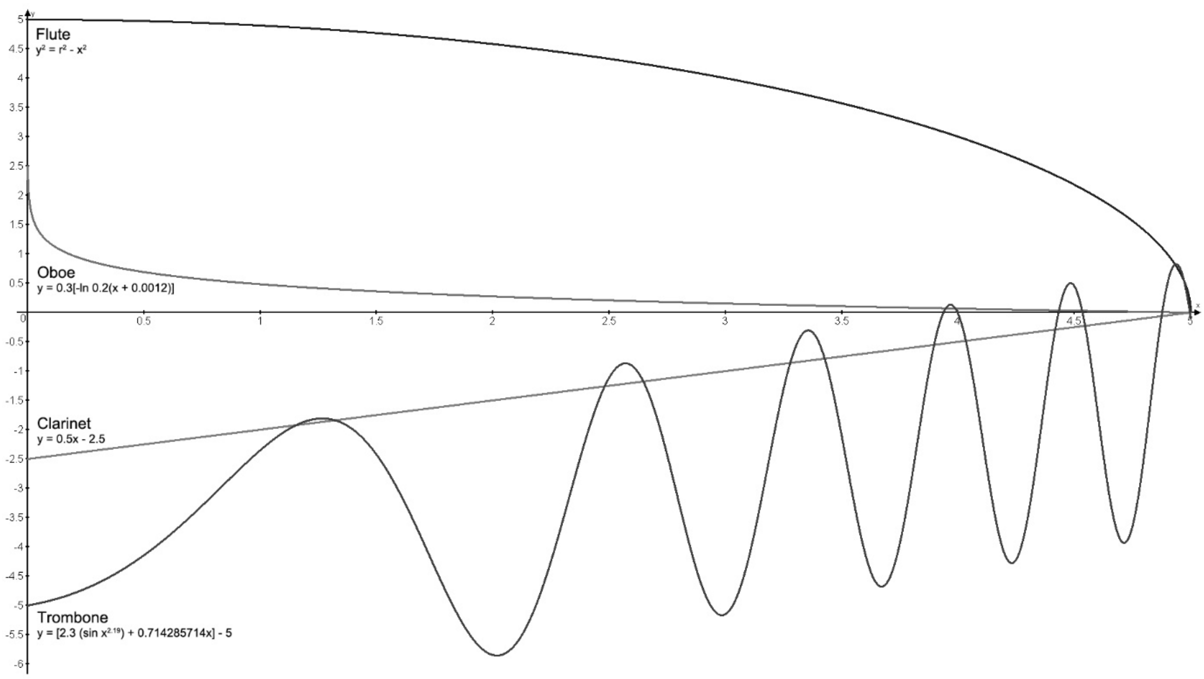

Fig. 2.1. *'unredeemed' self-)portraits*: trajectories used by the four instruments.

These trajectories are functions arranged to coincide at the same horizontal point, x=5, and their contours were designed principally for use in establishing pitch tendency (in the graph in Fig. 2.1, a vertical distance of 2.5 corresponds to an octave; hence the instruments begin in octaves). However, all four trajectories together affect each instrument's behaviour as well as the overall dramatic interplay and counterpoint.

[39] These works are designed also to be performed simultaneously as a quartet, in addition to every other combination of duo and trio.

The work for flute seeks to embody the following description from Sadgrove's book of the figure of the eagle on the Sutherland tapestry:

> The eagle, looking for prey, is my oppressive side, ready to exploit those who are weaker than I am, impatient with weakness, vulnerability and compassion.[40]

Following research into the primary behavioural characteristics of eagles in a variety of situations,[41] I defined three behaviours for the flute, together encapsulating the bird's predatory nature. The first, 'searching', although not exactly passive, is the most self-contained and reserved, sustaining long durations in various ways. The second, 'watching', is quiet and essentially still but restive, and the only one of the three to include substantial periods of silence. The third, 'attacking', consists of rapid, violent gestures, brief but brutal; the music associated with this behaviour is unique in that it was composed entirely intuitively. Table 2.1 shows a summary of all three behaviours' main characteristics.

Table 2.1. *'unredeemed' self-)portrait (in the form of an eagle*: summary of the three behaviours.

'searching'	'watching'	'attacking'
• senza vibrato • harmonics • trills/tremolandi • double-/triple-tonguing • ***pp*** – ***mp***	• pitched breathing through the instrument • flurries of notes, sharp, staccato • occasional 'calls' • rapid repeated notes • rest (falling probability throughout) • ***pp*** – ***mp***, calls = ***f***	• grace note plunge (to begin) and ascent (to end) • growling/roaring, incorporating the voice • jet whistles • percussive tonguing • high grace note figurations of approximate pitch • ***mf*** – ***ffff***

Examples 2.1 to 2.3 illustrate each behaviour.

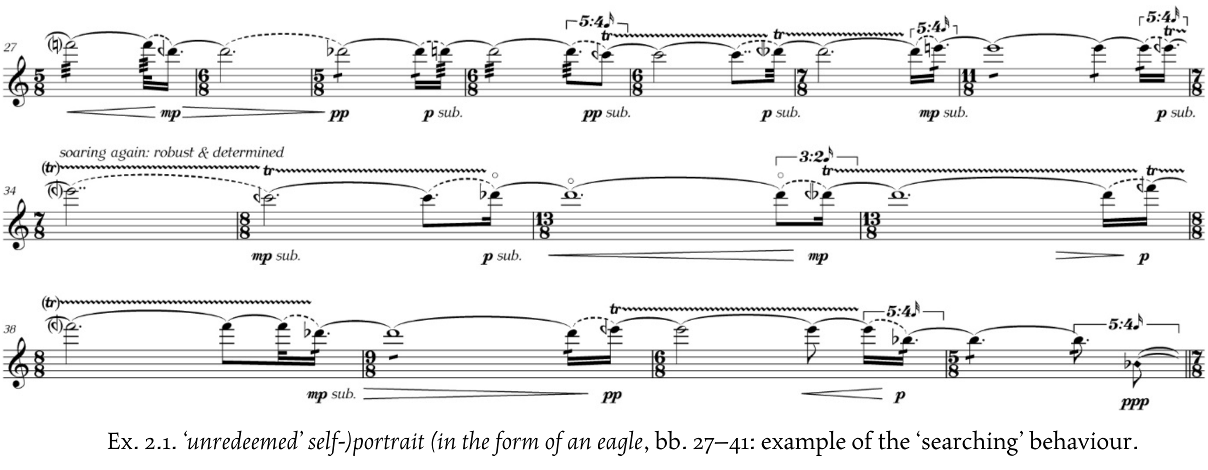

Ex. 2.1. *'unredeemed' self-)portrait (in the form of an eagle*, bb. 27–41: example of the 'searching' behaviour.

[40] Sadgrove 1995, p. 113.
[41] Principally from the 1997 David Attenborough BBC documentary *Wildlife Specials: Eagle*.

Ex. 2.2. *'unredeemed' self-)portrait (in the form of an eagle*, bb. 16–26: example of the 'watching' behaviour.

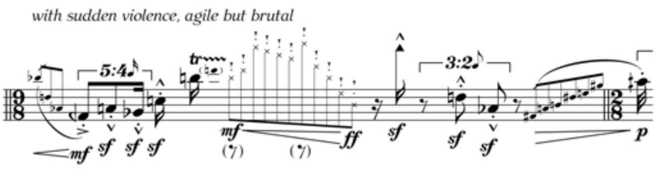

Ex. 2.3. *'unredeemed' self-)portrait (in the form of an eagle*, b. 62: example of the 'attacking' behaviour.

These three behaviours come into effect according to the proximity of the instrumental trajectories over time. Using a vertical threshold of five quartertones, the flute's behaviour is determined by the positions of the oboe and clarinet trajectories with respect to that of the trombone:

- searching: when the oboe or clarinet trajectories are beyond the threshold;
- watching: when either the oboe or clarinet trajectories are within the threshold;
- attacking: when both the oboe and clarinet trajectories are within the threshold.

This yields the result shown in Fig. 2.2, with 'searching' passages in white, and episodes of 'watching' and 'attacking' coloured light and dark grey respectively. In addition, the process can be seen to have created the work's macro-structure, with an emphatic phase transition appearance, abruptly shifting between behaviours.

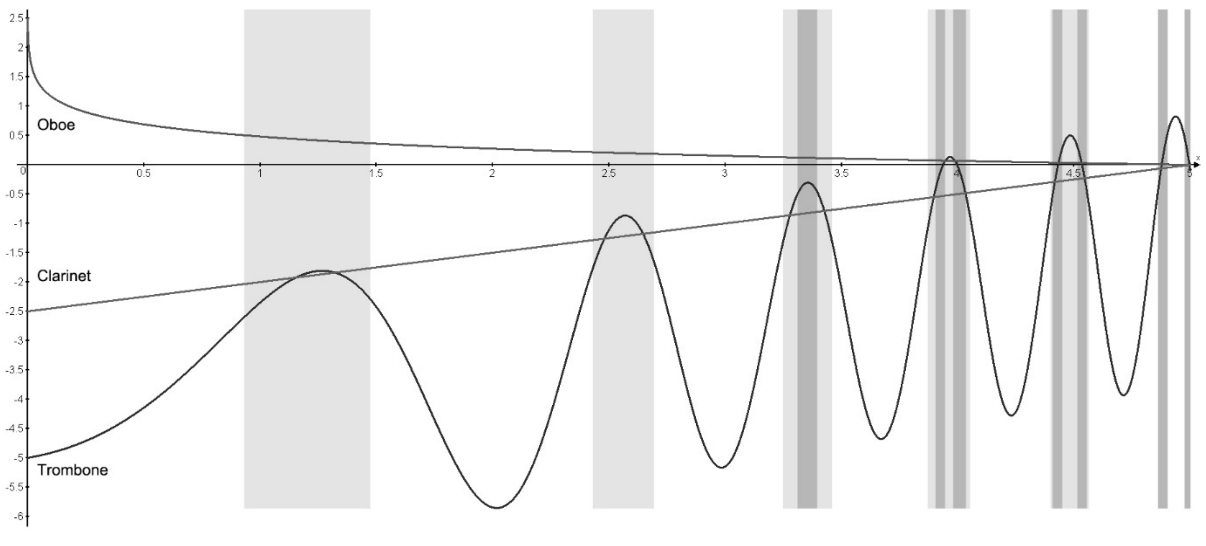

Fig. 2.2. *'unredeemed' self-)portrait (in the form of an eagle*: behavioural macro-structure.

This piece was an important compositional step forward for me in many respects, above all in my use of discrete musical behaviours as a determinant at both the micro- and macro-structural levels. It was also significant in its integration of computational and intuitive materials. As noted above, the entirety of the 'attacking' behaviour was freely-composed, not on account of any particular difficulty in realising it computationally, but in order to involve myself directly at these dramatically pivotal moments. In addition, this work was the first to implement an approach to quantisation at beat level, resulting in tuplet subdivisions that are restricted to one or two beats, avoiding more complex irrational possibilities.

Blesi (Five Transitions of the Soul) (2009) for six players

In this chamber work, composed for the Swedish ensemble Curious Chamber Players, I was keen to explore the possibilities of transitioning smoothly between behaviours, rather than 'jump-cutting' between them. The piece takes its title and inspiration from the eponymous geothermal phenomenon found within the Geysir National Park in Iceland. It comprises two deep pools of water, one of which is a light green colour, its water boiling hot and its surface shimmering and steaming (being the origin of its name, "one that blazes"). A thin connecting rivulet allows water to flow across into the second pool, which is a deep azure blue; its temperature is very warm, but nothing like the searing temperatures of its source. This struck me as an apt metaphor for the tenuous, often tepid, relationship, described by numerous theologians, between the individual and God, where the blazing source becomes drastically cooled.

This idea is explored in part through the notion of being 'out of sync', with respect to both time and pitch. The work's 'source' is characterised by the number seven, being demarcated by a series of seven fundamental pitches, rising through the scale of D♭ major, culminating on C♮, and also rooted in a fundamental metre of 7/4. The members of the ensemble gradually rebel and ultimately disconnect from this underlying order, moving into a different metre, 6/4, subsequently using a rival series of eight fundamental pitches, rising through the octatonic scale from D♭ to C♮. The result of this is two structural layers, forming a large-scale 7:6 episodic relationship, as shown in Fig. 2.3.[42]

[42] The numbers six and seven were chosen for their symbolic connotations within various religious traditions, seven being held as the 'perfect' number, whereas six is commonly associated with evil.

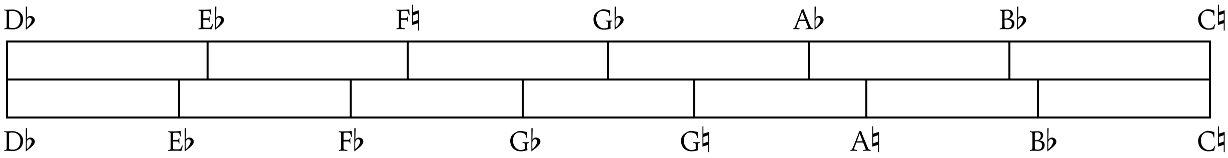

Fig. 2.3. *Blesi*: structural overview.

Use of the major scale in the 'source' structure draws on Boris Mouravieff's concept of the seven steps of the scale as "a useful framework or model to make sense of the unfolding of significant activities [...]. This evolution would take place in seven steps and move in a certain direction, determined by the initial plan or desire. [...] unless certain specific conscious impulses were added along the way, there would always be a tendency to go astray or deviate from the original plan."[43]

The instruments break away gradually from this structure, two at a time, the implications of which are explored in five episodes, each of which undergoes a different kind of transition. The inspiration for these came from a Nine Inch Nails lyric, in the song 'No You Don't':

> so many dirty little places
> in your filthy little worn out broken down see through soul [44]

Those five epithets, 'filthy', 'little', 'worn out', 'broken down' and 'see through', form the basis for the transitions. A two-behaviour setup of the kind in Fig. 1.2 (p. 11) was used. Each transition takes place between contrasting initial and terminal behavioural states, largely defined by a gradual shift in their respective musical parameters, as shown in Table 2.2. It must be stressed that, for the sake of brevity, the below summary (and every subsequent summary of behavioural parameters in this text) is greatly simplified, necessarily omitting a variety of additional aspects contained within each behaviour. Many of these parameters are not static but themselves undergo changes over time. For example, in the 'see through' terminal state, the A major weighted tonality gradually becomes chromatically reduced to the point where only the note A is heard.

[43] Le Mée 2003, p. 63.
[44] Nine Inch Nails 1999.

Table 2.2. Summary of initial and terminal behavioural states in *Blesi*'s five transitional episodes.

episode	initial behaviour	terminal behaviour
1. filthy	• *purity/cleanliness* • durations: ♩. – 𝅗𝅥. • pitches: E♭, B♭, G♮ • vibrato • ***pp – mp***	• *dirt/interference* • durations: ♪ – ♪. • pitches: E♮, D♮, A♮, G♭ • senza vibrato • increasing bursts of 'interference': multiphonics, tremolando, fluttertongue, sul pont., glissando, Bartók pizz., bow pressure • ***mf – ff***
2. broken down	• *order* • durations: 𝅗𝅥. – 𝅝. • pitches: E major (weighted tonality) • ***mf***	• *chaos* • durations: ♩. – 𝅗𝅥‿♪ • pitches: random chromatic • ***pp – fff***
3. little	• *confidence* • bold, expansive, trivial – like a collection of independent, polyphonic 'speeches'	• *insecurity/guilt* • reflective, wistful
4. worn out	• *energy* • wide intervals: 7–15 semitones • rest probability: 20% • ***mf – fff***	• *exhaustion* • small intervals: 1–2 semitones • rest probability: 90% • ***mp***
5. see through	• *substance/density* • durations: 𝅗𝅥 – 𝅝 • rest probability: 0% • pitches: random chromatic • ***ff***	• *ethereal/sparse* • durations: grace notes • rest probability: 90% • pitches: A major (weighted tonality) • ***pp***

This was the earliest manifestation of weighted tonality in my work; only seven notes of the chromatic scale are used, and the aim was to allude to major tonalities. The notes were arranged as shown in Table 2.3.

Table 2.3. *Blesi*: weighted tonality pitch series.

1	2	3	4	5	6	7
PRIMARY	V	III	VI	II	VII	IV

To implement the transitions I wrote a computer program comprising five separate procedures containing the rules for each transition. The program created most of the material used in each transition, although all involved a certain degree of freely-composed music. Fig. 2.4 shows the output from the program for the cello undergoing the 'filthy' transition, beside its corresponding music. As can be seen, various violent 'bursts' have been realised intuitively, and there are many other instances of this throughout the piece, such as in the 'broken down' transition, where the choice of register for each pitch was also determined intuitively.[45] Whereas these uncomputed moments are generally very brief, for the 'little' transition, due to its intended highly flamboyant nature, the program took a similar approach to the 'attacking' behaviour in *'unredeemed' self-)portrait (in the form of an eagle*, simply calculating points where each of the two behaviours were to be heard, and all of the music was then composed entirely freely to fit these structural divisions.

The composition process of *Blesi* successfully proved to me that gradual stochastic transitions between discrete behaviours were both musically and dramatically effective. It further reinforced my approach to integrating computational and freely-composed elements, while at the same time greatly decreasing the extent to which the output from the program required editing.

HELP/ME: the soul-machine of the cosmology of grief (2010) for 16 players

Composed for Birmingham Conservatoire's Thallein Ensemble,[46] this piece began from my considering the possibilities of a composition unable to wrench itself away from a single pitch, articulated by what I described in my initial sketches as a "series of agonised sighs/surges". The piece has a similar structure to *Blesi*, comprising an introduction and coda surrounding six episodes, each of which undergoes a small-scale transition. The aim with this piece was to explore further using these algorithmic techniques to create cohesive portions of a larger whole, focussing less on transitions between two contrasting behaviours than on either single, ongoing processes with increasing intensity or static, locked-in processes cycling for eternity. An additional aim was to examine again the integration of intuitively composed elements within an overall computational environment. The emotional tone of the piece is one of anguish and distress, and this influenced the choice of instrumentation which for the most part employs only softer timbres (i.e. no double-reed instruments) and lower registers. For inspiration I drew on the convoluted verse of Japanese cyberpunk poet Kenji Siratori, extracting six short quotations from his 2008 'noise mantra' *This Machine Kills* that capture something of the intended emotional subtext.[47]

[45] It should also be noted in Fig. 2.4 that some of the 'senza vibr' instructions (at beats 8 and 11) have been ignored until after the first 'burst'.

[46] A welcome commission, as I directed the ensemble myself as an undergraduate in the mid-1990s.

[47] Siratori's writings are also the basis of my 2012 work for voice and five players, *the octave of the grief of the clone that leapt to the remainder of night sky*. This piece forms part of a tangential compositional thread that defers using the work's processes to create notated material, instead handing them over to the players who work through them in real-time, more directly comparable with Michael Parsons' 'people processes'.

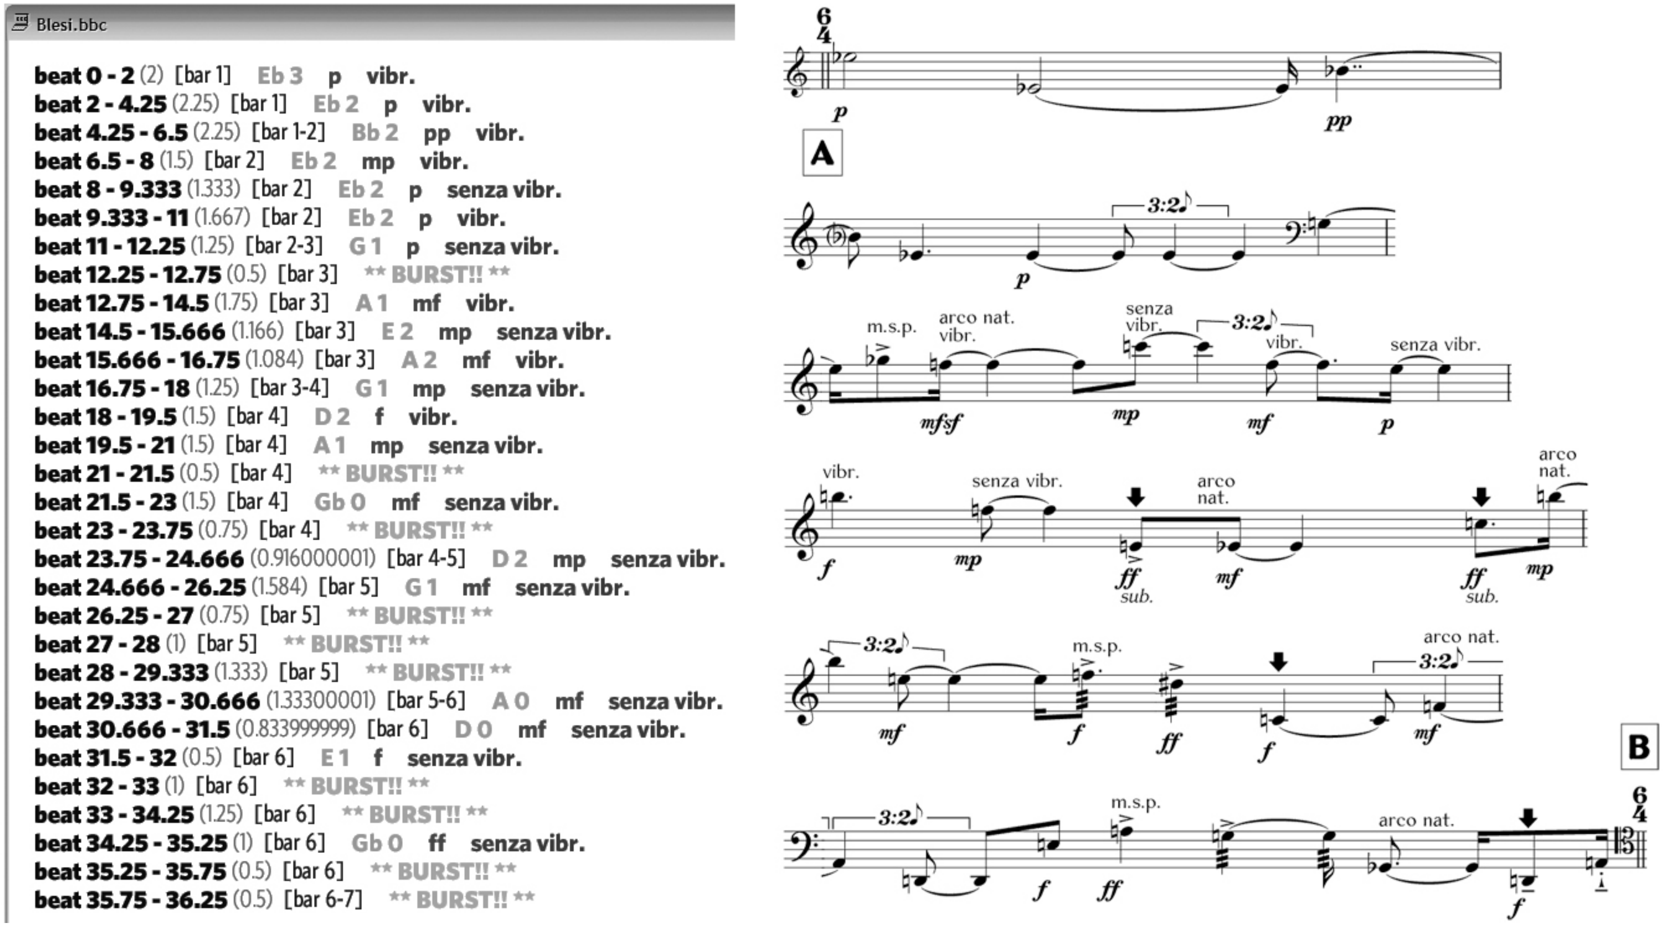
Blesi.bbc

beat 0 - 2 (2) [bar 1] Eb 3 **p** **vibr.**
beat 2 - 4.25 (2.25) [bar 1] Eb 2 **p** **vibr.**
beat 4.25 - 6.5 (2.25) [bar 1-2] Bb 2 **pp** **vibr.**
beat 6.5 - 8 (1.5) [bar 2] Eb 2 **mp** **vibr.**
beat 8 - 9.333 (1.333) [bar 2] Eb 2 **p** **senza vibr.**
beat 9.333 - 11 (1.667) [bar 2] Eb 2 **p** **vibr.**
beat 11 - 12.25 (1.25) [bar 2-3] G 1 **p** **senza vibr.**
beat 12.25 - 12.75 (0.5) [bar 3] ** BURST!! **
beat 12.75 - 14.5 (1.75) [bar 3] A 1 **mf** **vibr.**
beat 14.5 - 15.666 (1.166) [bar 3] E 2 **mp** **senza vibr.**
beat 15.666 - 16.75 (1.084) [bar 3] A 2 **mf** **vibr.**
beat 16.75 - 18 (1.25) [bar 3-4] G 1 **mp** **senza vibr.**
beat 18 - 19.5 (1.5) [bar 4] D 2 **f** **vibr.**
beat 19.5 - 21 (1.5) [bar 4] A 1 **mp** **senza vibr.**
beat 21 - 21.5 (0.5) [bar 4] ** BURST!! **
beat 21.5 - 23 (1.5) [bar 4] Gb 0 **mf** **senza vibr.**
beat 23 - 23.75 (0.75) [bar 4] ** BURST!! **
beat 23.75 - 24.666 (0.916000001) [bar 4-5] D 2 **mp** **senza vibr.**
beat 24.666 - 26.25 (1.584) [bar 5] G 1 **mf** **senza vibr.**
beat 26.25 - 27 (0.75) [bar 5] ** BURST!! **
beat 27 - 28 (1) [bar 5] ** BURST!! **
beat 28 - 29.333 (1.333) [bar 5] ** BURST!! **
beat 29.333 - 30.666 (1.33300001) [bar 5-6] A 0 **mf** **senza vibr.**
beat 30.666 - 31.5 (0.833999999) [bar 6] D 0 **mf** **senza vibr.**
beat 31.5 - 32 (0.5) [bar 6] E 1 **f** **senza vibr.**
beat 32 - 33 (1) [bar 6] ** BURST!! **
beat 33 - 34.25 (1.25) [bar 6] ** BURST!! **
beat 34.25 - 35.25 (1) [bar 6] Gb 0 **ff** **senza vibr.**
beat 35.25 - 35.75 (0.5) [bar 6] ** BURST!! **
beat 35.75 - 36.25 (0.5) [bar 6-7] ** BURST!! **

Fig. 2.4. *Blesi*, ‘filthy’ transition, cello: program output and notated music.

Each of the episodes has the same overall duration, thirty-one crotchet beats, subdivided according to a descending Fibonacci sequence of bar-lengths: 13/4, 8/4, 5/4, 3/4, 2/4; each episode audibly demarcates these bar divisions in a different way. The six episodes are summarised in Table 2.4.

Table 2.4. *HELP/ME*: summary of episodic behaviours.

<table>
<tr><th>episode/quotation</th><th>behaviour</th></tr>
<tr><td>A
“the picture of the chloroform that respires the malice of the grief of a cell is observed as if i love it”</td><td>• abject; aghast
• Fibonacci demarcations: sizzle cymbal
• violas/cellos (practice mute): utter silence, followed by weak, timid, evanescent E♭s
• pp</td></tr>
<tr><td>B
“I copy the life that the soft storage of myself receives the quickening of the replicant murder that break down the grief of the end of the world”</td><td>• massively pressurised
• Fibonacci demarcations: accents followed by portamento resetting pitch back to E♭
• violas/cellos (practice mute): deviations up to a semitone around E♭
• durations: 𝅝.. maximum (no minimum) / rests: 𝅗𝅥. maximum (no minimum)
• p</td></tr>
<tr><td>C
“the locus of the love of the clone that quiesced”</td><td>• barely animate; sluggish and vague
• Fibonacci demarcations: lion’s roar
• alto flutes/bass clarinets: deviations around E♭ –deviation decreases from 6 quartertones to nothing through each bar; senza vibr. → molto vibr in each bar.
• durations: 𝅗𝅥 – 𝅝‿♩ / rests: 0.1♩ – ♩.
• mp – mf</td></tr>
<tr><td>D
“the emotional circuit of the ADAM doll short... the horizon of her chromosome toward... the time axis of demolition line”</td><td>• bristling with nervous, increasingly frantic energy
• Fibonacci demarcations: bass drum roll
• tutti: E♭ exploded into multiple low registers; deviations around E♭ – amount of deviation decreases through each bar but range increases according to Fibonacci: 13/4 = 3 → 2 quartertones | 8/4 = 5 → 3 | 5/4 = 8 → 5 | 3/4 = 13 → 8 | 2/4 = 21 → 13
• durations: 1.2♩ – 𝅗𝅥. / rests: 0.7♩ – ♩
• random articulations: wind & brass = fluttertongue/double-/triple-tongue/trill/ cuivré/ord. | strings: trill / trem. / sul pont. / col legno trem. / pizz. trem / ord.
• reducing f → p through each bar</td></tr>
<tr><td>E
<i>“the octave of the grief of the clone that leapt to the remainder of night sky”</i></td><td>• <i>miasmic</i>
• Fibonacci demarcations: flexatone, then tam-tam
• three layers: 1) alto flutes/horn/contrabassoon: melody based on Psalm 22 plainsong
2) bass clarinets: wild descending glissandi
3) trombones/flexatone/strings: glissando deviations around E♭ – deviation rises from 7 quartertones at start to 24 from third bar
• durations: increasing through episode, roughly ♩ – 𝅗𝅥‿♪ → 𝅗𝅥. – 𝅝‿♩
• layer 1 = <i>mf</i>, layer 2 = <<i>f</i>>, layer 3 = ben <i>f</i> sempre</td></tr>
<tr><td>F
<i>“fuck junk to the angel mechanism crunch ... apoptosis season of the chromosome of yourself the scream”</i></td><td>• <i>jammed, trammeled</i>
• Fibonacci demarcations: string surges
• two layers – 1) strings: cluster around each instrument’s lowest E♭, heavy bow pressure
2) wind/brass: single crescendo around random E♭
• layer 1 = <i>mp</i><<i>ff</i>><i>mp</i>, layer 2 = 0<<i>ff</i>>0</td></tr>
</table>

While the above episodes were created almost entirely computationally, the first layer of section E (bb. 30–34) features a freely-composed melody based on a portion of the plainchant for Psalm 22 on Palm Sunday:[48]

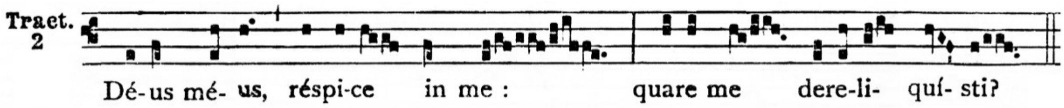

Ex. 2.4. Psalm 22 (opening), from Mass for Palm Sunday.

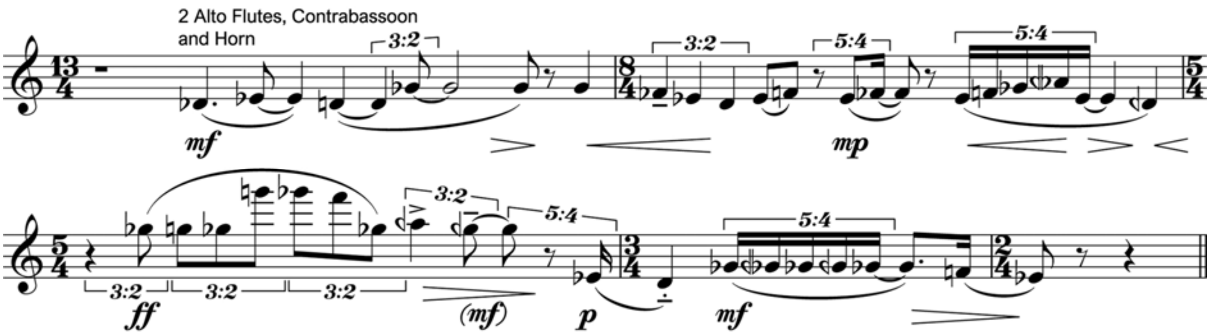

Ex. 2.5. *HELP/ME*, bb. 30–34: melody based on Psalm 22 plainchant.

The introduction and coda were also composed intuitively based on existing musical materials. The introduction uses a 17th century Scottish hymn tune, *Abbey*:[49]

Ex. 2.6. Hymn tune *Abbey*.

This is turned into a furiously overbearing, crushingly loud fanfare (bb. 1–9), providing a brutalist context for the six episodes that follow in its wake. The brief coda, occupying the final bar of the piece, is also aligned with a quotation from Kenji Siratori, hinting at a final emotional unhinging:

> the language line of artificial play human-genome yourself of the sun self-consolation in the world of the disillusionment lost the past![50]

---

[48] *The Liber Usualis*, p. 594.
[49] *Church Hymnary*, p. 39.
[50] Siratori 2008, p. 8.

The coda reworks two Easter plainchants, the melody for Flute 1 based on 'Hæc dies', that of Flute 2 on the Alleluia that follows it:[51]

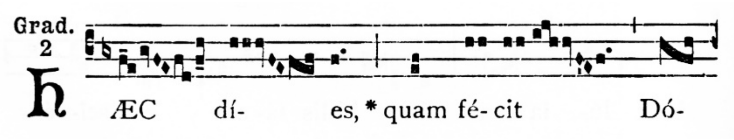

Ex. 2.7. Hæc dies (opening), from Mass for Easter day.

Ex. 2.8. Alleluia (opening), from Mass for Easter day.

These are turned into two 'delirious' bursts of melody, aspiring to freedom, but below them, in the brass, is a quotation from Variation VII of Richard Strauss' *Don Quixote*, specifically the climactic brass chords during the eponymous protagonist's imaginary ride through the air:[52]

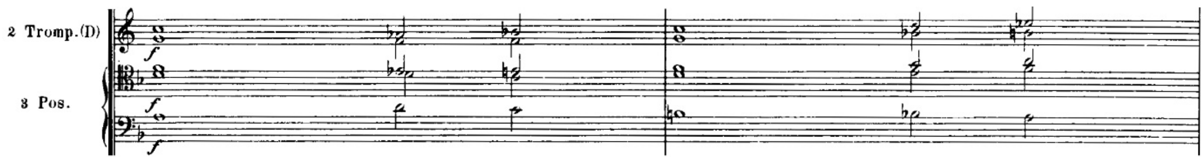

Ex. 2.9. Richard Strauss, *Don Quixote*, Fig. 58.

Transposed up a semitone (to be compatible with the present work's omnipresent E♭ primary) and performed pianissimo with mutes, it becomes a disturbing hint at the reality belied by the alto flutes' music, a reality captured in the work's closing Siratori quotation, after the final barline:

> "ADAM doll_desire_death_latency × awakening_impossibility = death"[53]

HELP/ME is an important work with regard to consolidating my thinking about rhythm. In earlier works (particularly evident in *Blesi*), my rhythmic language tended to be highly complex, often resulting in long irrational groupings. In this piece, as in *'unredeemed' self-)portrait (in the form of an eagle*, the focus on quantisation at the level of individual beats facilitates complex rhythmic counterpoint and creates a heightened moment-by-moment focus on the material.

In hindsight, *HELP/ME* has turned out to be especially pivotal, both in general terms regarding the direction my research has taken since and specifically in the way I would conceive *Cloud Triptych*.

[51] *The Liber Usualis*, p. 778–9.
[52] Strauss 1979, p. 260.
[53] Siratori 2008, p. 13.

Although I felt this piece succeeded on a technical level, it proved problematic in terms of the extra-musical considerations I had hoped to embed within it. Put simply, it has a formal elegance that, to my mind, undermines the chaotic emotional rawness of its subtext, less an embodiment than a rationalised representation. As a consequence, my compositional outlook in the remaining works of the portfolio focussed increasingly on musical processes in and of themselves, nowhere more so than in *Cloud Triptych*.

Intense quick dream of sentimental groups with people of all possible characters amidst all possible appearances (2010) for string sextet

This work, written for Birmingham Contemporary Music Group, is based on a short poem by E. E. Cummings that summarises the stages in an implied long-term process of evolution, from bestial wildness through refined sophistication to a point of indescribable transcendence.

> wild(at our first)beasts uttered human words
> —our second coming made stones sing like birds—
> but o the starhushed silence which our third's [54]

I emulated this three-stage progression by using a three-behaviour, non-linear system. The three behaviours were defined according to the characteristics described in the lines of the poem. Behaviour A, corresponding to the first line, is articulated via a collection of simple, brute gestures, behaviour B has a propensity to form melodic shapes and lines, and behaviour C emphasises harmonics and periods of silence.

Table 2.5. Summary of *Intense quick dream*'s three behaviours.

behaviour A ***wild/bestial***	**behaviour B** ***più e più cantabile e legato***	**behaviour C** ***increasingly hushed, becoming silent***
• sustained open string(s), using three types of bow pressure: normal, excessive (distorting), pitchless scraping (single-/double-/triple-stop) • shorts 'stabs', arco battuto/col legno battuto • sustained string(s), beyond the bridge • Bartók pizzicato, open string(s) • rest (duration: ♪– 𝅗𝅥)	• jeté glissando • tremolo, with initial accent (single-/double-stop) • melodic phrase • rest (duration: ♩ – 𝅝‿♩)	• harmonic • rest (duration: 𝅗𝅥 – 𝅝..)

[54] Cummings and Firmage 1994, p. 844.

The collection of gestures in behaviour A only gradually becomes fully available, so that at first the instruments are extremely restricted in their musical 'vocabulary'. The likelihood of melodic phrases in behaviour B increases over time, and the length of these phrases is directly proportional to the behaviour's probability, together engineering a clearly-articulated transition away from gestures towards passages of melody. The pitches in these melodies are determined by weighted tonality, according to an underlying harmonic sequence (which completes three cycles throughout the work's duration), the primaries of which are a falling cycle of fifths beginning on G♮. The melodies have a pitch 'ceiling' and chromatic index that both also increase over time, enabling them to soar ever higher and with greater freedom. In behaviour C, the rest probability is linked to the behavioural probability, ensuring that silence ultimately becomes unavoidable.

The three behaviours have their nodes positioned at the start, middle and end of the work's timeline respectively. Behaviours A and C have a range of influence lasting half of the work's duration, behaviour A from the beginning to the middle, behaviour C from the middle to the end. Behaviour B, by contrast, has a range of influence extending for the entire duration, with non-linear probability according to a Gaussian-type bell curve. My decision to use a bell curve was made for both musical and extra-musical reasons. Cummings' poem, although implying an evolutionary process, opts not for a diachronic description but rather a short series of synchronic 'snapshots'.[55] This abrupt effect is analogous to the phase transitions referred to previously which, as has been shown, are a corollary of using non-linear transitional processes. Furthermore, E. E. Cummings' poem presents a qualitative progression through its states of behaviour; whereas the first and last correspond to animalistic and transcendent states respectively, the middle line can be read as capturing the quintessence of human artistic expression, elegant and ecstatic. Being the most readily relatable – the most recognisably 'human' – of the three, this is the central reason why its probability is given a bell curve, since such distributions are synonymous with analysis of almost all aspects of human character and behaviour. This curve, which causes behaviour B to dominate the work's middle regions, is given a minimum probability value of 3%, ensuring its behaviour can occur at any point during the piece, resulting in the probability distribution shown in Fig. 2.5. At the start of the timeline, behaviour A has an extremely strong likelihood of being chosen. As the cursor moves closer to the work's centre, the probability of behaviour B being heard increases rapidly. Beyond the midpoint, behaviour A can no longer be heard, and as time continues, behaviour C will in turn gradually predominate and bring the piece to an end. Simultaneous with this probabilistic movement between the three behaviours, there is an *ex*

[55] The terms 'diachronic' and 'synchronic' are borrowed from Saussure's contrasting axes of study of the evolution of language; cf. Sturrock 1993, pp. 5–6.

machina process at work determining the music's dynamic tendency, approximating to an overall diminuendo from ***fff*** to ***ppp***.

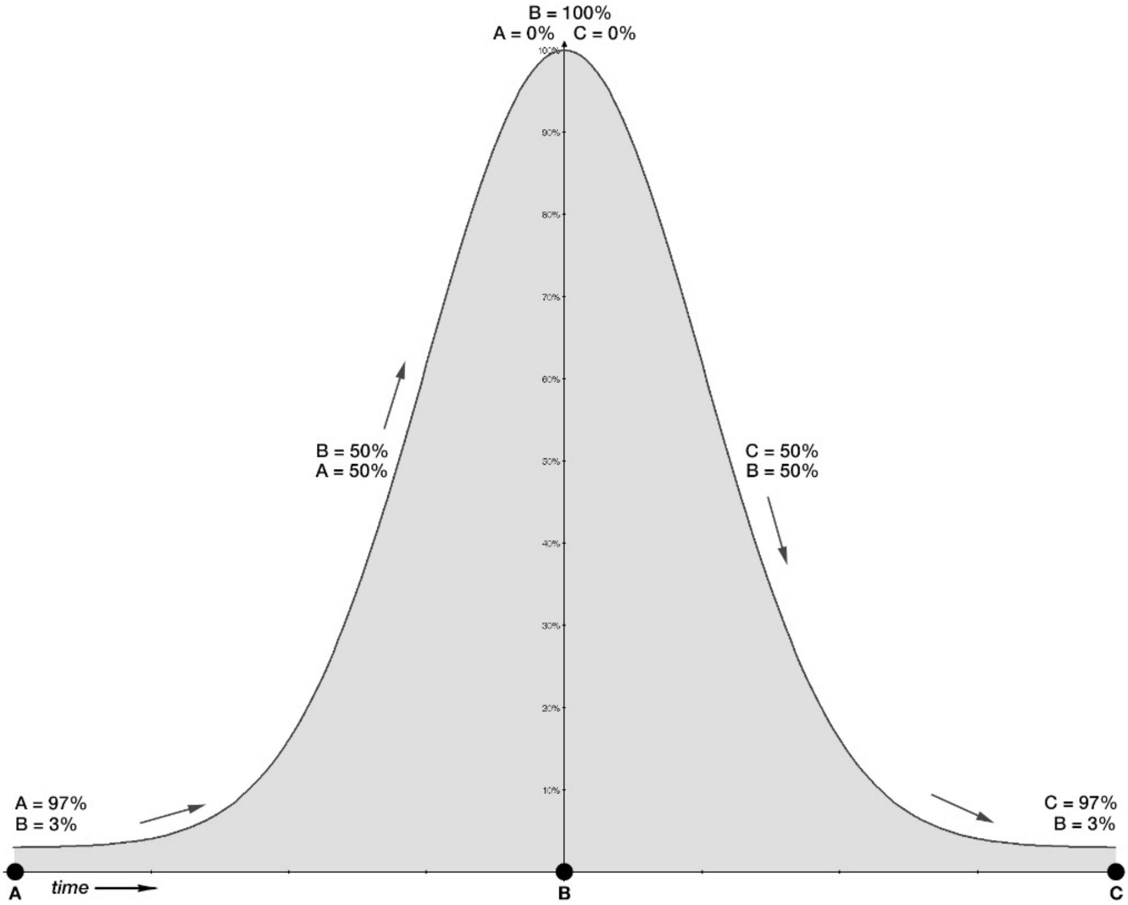

Fig. 2.5. *Intense quick dream*: probability distribution.

What makes this piece especially significant in my work is its self-contained nature, whereupon the entire composition from start to end is created by the algorithmic processes, for the most part without further necessary involvement from myself (apart from the addition of phrasing). As such, its processes for the first time become an almost complete embodiment of my compositional practice. In tandem with this is the use of an unchanging 4/4 metre and fixed tempo throughout, a combination that has become a recurring feature of my work as its computational nature has grown. I regard them as functioning like a map grid system, a fixed frame of reference within which free-flowing features are situated and coordinated.

Beyond this, although based on the E. E. Cummings poem, for the first time in my work there is no attempt here to establish or convey a specific psychological subtext. This movement away from extra-musical aspects is essential both in the remaining preliminary works and in the development of *Cloud Triptych*.

Four Seasons (2016) for flute choir

In 2015 I became composer-in-association with the Leeds-based Arts@Trinity Flute Choir, for whom *Four Seasons* was written. There were various aims with the piece, principally to explore the long-term musical potential in simple algorithmic processes – not so much transitioning between behaviours as undergoing intrinsic change, akin to a clockwork mechanism (similar to the approach taken in *HELP/ME*), and not even necessarily using stochastic calculations – and to examine them within a context where they occur at differing rates, sometimes asynchronously. An additional aim was to seek ways to compose more quickly and succinctly with algorithmic processes, with a minimum of extraneous complexity and extra-musical association, where the processes become an end in themselves rather than acting partly as metaphors for a personal subtext.

I. prime-temps

The title comes from the 15th century English term for spring.[56] This movement, for ten flutes, comprises an eight-minute process producing simultaneous, synchronised ascents. The ascents are chromatic, rising two octaves from C4, accelerating such that each successive duration is shorter by 10%; quantisation is maximal, with all grids available throughout (rest probability is zero; the players are instructed to breathe ad libitum, avoiding moments immediately before a change of pitch). Each ascent begins senza vibrato and detached; from low G vibrato and legato increase until normal levels are attained for the last few notes (from high F). The dynamics also increase, exponentially, from ***pp*** to ***f***. Each flute performs this ascent a unique number of times – flute 1 = once, flute 2 = twice, ... flute 10 = ten times – through the duration of the movement, resulting in a complex texture of ascents overlapping each other, with a generalised pitch focus (essentially but subtly determined by flute 1, undergoing a single ascent) that is also ascending.

[56] The term is derived from the Old French 'prin tans'.

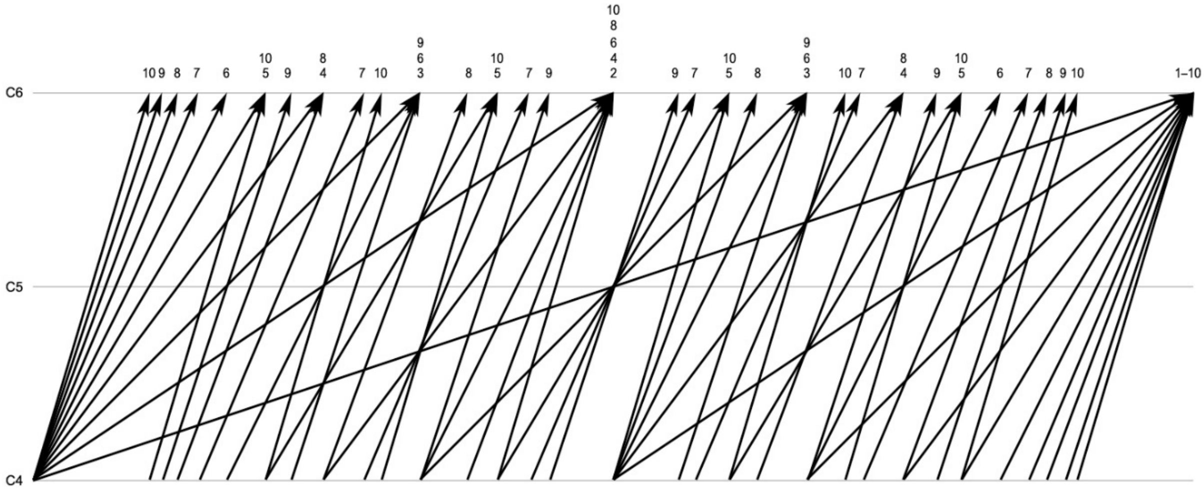

Fig. 2.6. *prime-temps*: pitch structure.[57]

It is worth noting that *prime-temps*, along with the third movement *fall*, are rare exceptions in my work of algorithmic processes that do not involve stochastic calculations. In both cases, their respective algorithms will produce exactly the same result every time.

II. samā

The Sanskrit word for summer, 'samā' (समा) literally means "half-year" and has connotations of the words 'same' and 'equal'. Five flutes are used in this movement, which features a succession of asynchronous, overlapping ascents. The players enter at intervals of roughly 48 seconds (dividing the four-minute duration into fifths), but play independently; timings throughout are therefore approximate. Once begun, each flute performs three ascents, with respect to tendencies beginning on F♯4, F♯5 and F♯6, always finishing at C7; the flutes can deviate away from these tendencies by up to four semitones above or below. The durations of the ascents reduces by a ratio of 3:2:1; since each successive flute participates for a shorter amount of the movement as a whole, this has the effect of gradually accelerating the speed of the ascents. This is reinforced by the players having increasingly intricate quantisation lattices:

Table 2.6. *samā*: distribution of quantisation grids.

player 1	**player 2**	**player 3**	**player 4**	**player 5**
grace, simple	grace, triple	grace, quint	grace, sext	grace, sept

Each ascent follows an increasing dynamic tendency, the first rising from ***pp* – *mp*** at the start to ***mf* – *f*** at its zenith, the second from ***p* – *mf*** to ***f* – *ff*** and the third from ***mp* – *f*** to ***ff* – *fff***.

[57] For clarity, the diagram in Fig. 2.6 is drawn using straight lines; however, as the ascents are accelerating the lines should actually be curved.

Harmony is determined *ex machina*, episodically alternating between primaries of F♯ and C; the duration of these episodes decreases in pairs, in a ratio of 5:4:3:2:1, with a reducing chromatic index, from nine to five degrees.[58] Rhythm, rest probability, articulation probability and dynamics are also determined *ex machina*, chosen at random with respect to limits/ranges that change (at a linear rate) throughout the movement's duration; however, in the case of rest probability, this is fixed at 0% for the final harmonic episode. Articulation, when applied to a note, is either in the form of a trill (longer notes) or a mordent (shorter).

Table 2.7. *samā*: *ex machina* parameter limits/ranges.

parameter	start	end
rhythm	𝅘𝅥𝅯 – 𝅗𝅥.	𝅘𝅥𝅯 – 𝅘𝅥𝅮.
rest probability	40%	20% (0% for final harm. episode) max. 3 successive rests
articulation probability	20%	40%

The overall pitch structure of *samā* is shown in Fig. 2.7.

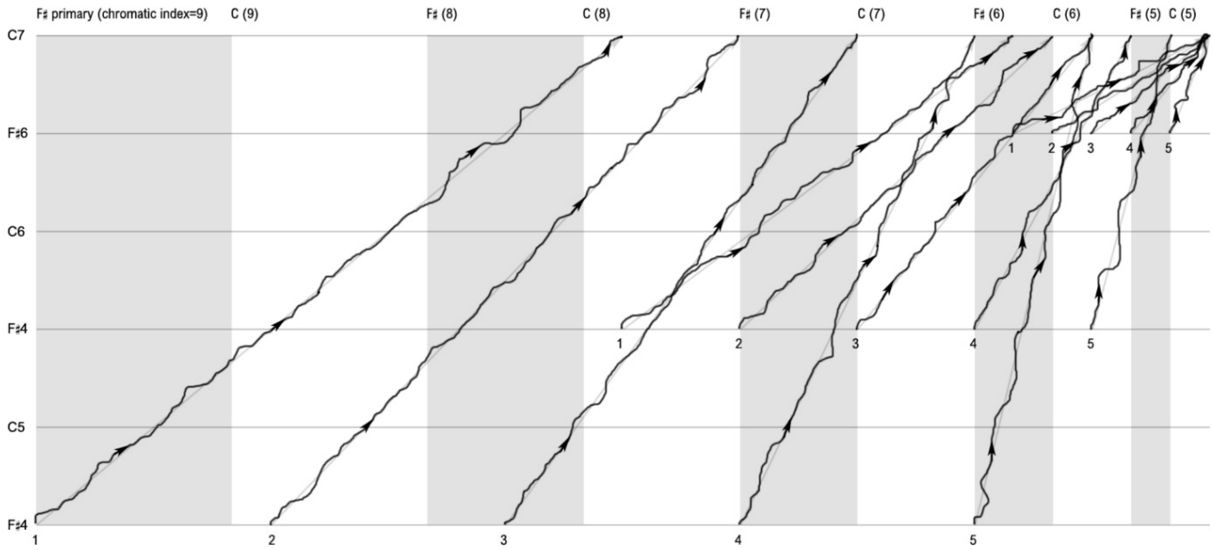

Fig. 2.7. *samā*: pitch structure.

It is worth mentioning that there was a quite specific *deus ex machina* criteria here to choose which material generated by the program would be used. Usually my criteria is intuitive and somewhat hard to define, but on this occasion I specifically wanted:

[58] The original intention was that the pitch deviation from the tendencies should gradually reduce from four to two semitones, but this proved incompatible with the reduction in chromatic index. Trial and error showed that a deviation smaller than four semitones was almost always impossible beyond harmonic episode six, due mainly to an additional rule preventing the same pitch from occurring twice in succession. This is another example of the creative dialogue between sketching and realisation, one of the most stimulating aspects of computer-aided composition.

- instances of pitches lower than F♯ shortly after the start of the first ascent;
- rests to occur in the penultimate harmonic episode;
- few instances of ***fff*** in the penultimate harmonic episode; and
- at least three different pitch-classes among the final few notes.[59]

The criteria were used in order to create the most interesting and dramatic overall effect. They were applied strictly, and material generated by the program that did not satisfy these conditions was rejected.

Although the flutes are not strictly synchronised, I wanted the work to end with an exuberant conclusion with all players finishing together; this is achieved by having the final bar in each part repeated ad lib., excitedly reiterating the culminations of their unique ascents, whereupon (coordinated visually) they all stop.

III. fall

The title is an alternate word for autumn.[60] This movement is for ten flutes, lasts ten minutes, and its process is a linear inversion of that used in *prime-temps*, involving simultaneous, synchronised descents. Each descent is chromatic, falling two octaves from C6 at a linear rate (all durations, for each player, are the same); quantisation is again maximal throughout. Each descent is divided into four dynamic zones, reducing from ***f*** to ***p***. Vibrato and legato are used throughout.

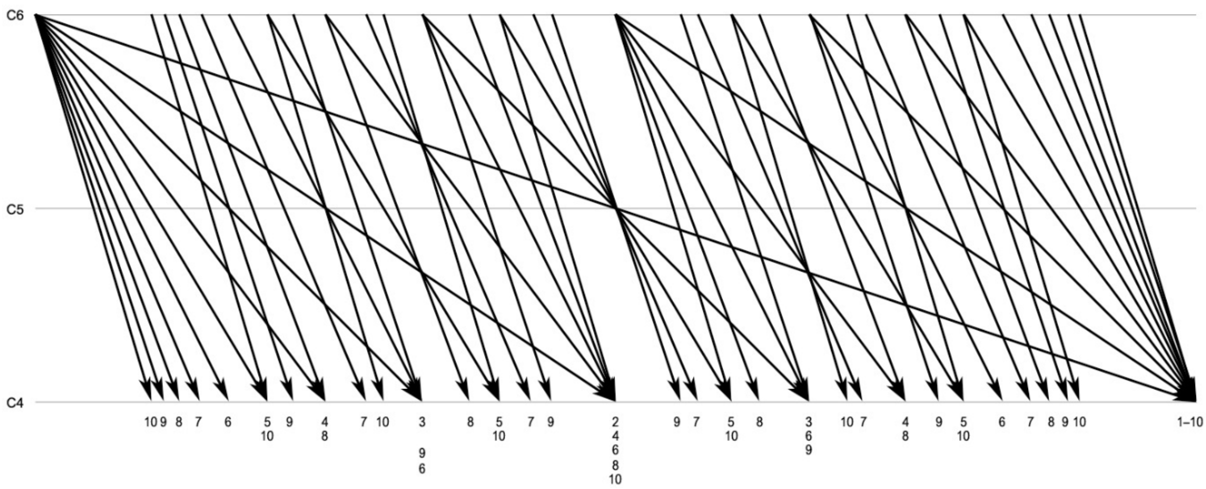

Fig. 2.8. *fall*: pitch structure.

[59] Of course, it would be possible to add these criteria into the program to ensure they occur, but my preference was not actively to impose them onto the ascents but to await them happening naturally.
[60] This use of the word 'fall' originated in 1660s England, as a shortening of the 16th century term 'fall of the leaf'.

IV. vindo

'Vindo' is the Gaulish word for 'white', and a possible early source for the word 'winter'. Lasting six minutes, this movement is for five asynchronous players, exploring gradual descents according to a pitch tendency that begins at G6 and falls three octaves. Each flute plays just a single descent, but the duration it takes is a proportion in fifths of the overall six-minute duration, e.g. flute 1 = one-fifth = 72 seconds, flute 2 = two-fifths = 144 seconds. The players begin on the piccolo and switch to alto flute partway through.

Throughout the descent, the maximum deviation away from this tendency reduces from four to two semitones, and an additional probability that pitches are only white notes is applied, descending from 100% to 0%. Durations are random between limits of a semiquaver and a dotted minim, and the rest probability is 20%. Each descent is structured in five quantisation/dynamic regions, reducing in subdivisional intricacy and from *circa* ***mp*** to *circa* ***pp***.

Table 2.8. *vindo*: the quantisational/dynamic structure of each descent.

simple, sept	simple, sext	simple, quint	simple, triple	simple, grace
circa ***mp***	***p*** – ***mp***	*circa* ***p***	***pp*** – ***p***	*circa* ***pp***

There is a 50% probability that notes will be articulated either with fluttertongue (for shorter durations) or a trill (for longer). The five players start together, but thereafter do not need to synchronise further; approximate points where each player ends are indicated in the remaining parts. The resultant pitch structure of *vindo* is shown in Fig. 2.9.

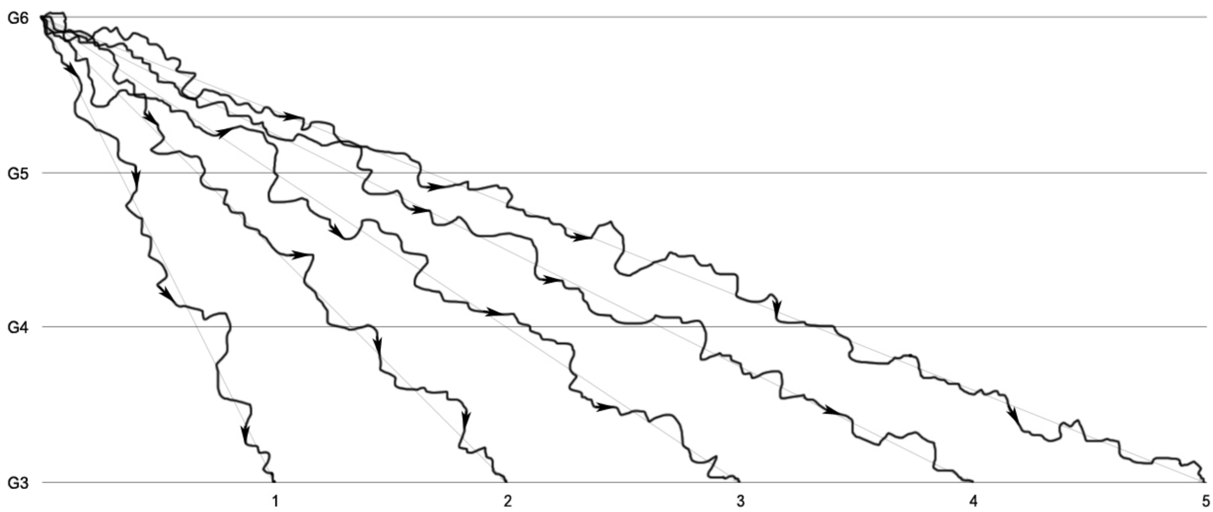

Fig. 2.9. *vindo*: pitch structure.

Four Seasons was valuable in clarifying various aspects of my method. The use of self-contained behaviours, used previously in *HELP/ME*, although interesting were to my mind less engaging than those that undergo transitions to contrasting behaviours, where one perceives a more dramatic sense of transformation. Similarly, non-stochastic algorithmic processes (which produce the same results every time), hitherto untested in my work, proved to be less compositionally fulfilling than those involving stochastic variety. More positive were the experiments in self-similarity and asynchronous ensembles in *samā* and *vindo*. The former acts as a kind of 'behavioural canon' redolent of Ligeti's micropolyphony, while the latter, through relatively simple means, clearly offers considerable scope with regard both to harnessing more directly the individuality exhibited in such 'people processes' (see p. 16), as well as the potential for considerable contrapuntal complexity.

Chapter 3. *Cloud Triptych* (2016) for orchestra

Following reflection on the compositions discussed in the previous chapter, various motivations led to the development of the main work in the portfolio, *Cloud Triptych*. The most straightforward aim was simply to compose my first work for full orchestra, as all of my previous work had been for varying sizes of chamber ensemble, composed for specific groups and individuals.[61] The range and ambition of further motivations, though, made it immediately clear that a much larger program would need to be created in order to develop and accommodate them all, essentially written from scratch but once again incorporating code from earlier programs where appropriate. An important goal from the outset was that it should as far as possible obviate the need to write further programs in future, having the capacity and scope to create numerous (perhaps limitless) compositions. Prior to the ensuing discussion, it is worth remarking that, being such a considerable expansion of my existing methods and practices, these developments were by no means as straightforward as they may appear described below; indeed, many aspects were extremely difficult and time-consuming to realise.

An important motivation when planning *Cloud Triptych* was to turn away from composing music loaded with an emotional or otherwise psychological subtext. As discussed previously, several of the preliminary works, although I felt they succeeded on a technical level, had proved problematic in terms of the extra-musical considerations I had hoped to embed within them. I therefore decided more fully to embrace abstraction, focussing on an exploration of musical processes devoid of an imposed extra-musical meaning. An integral part of this involved a parallel urge to reduce and/or undermine the extent to which degrees of control and certainty are imposed on the composition process. This would not simply manifest itself in complex systems with less predictable outcomes, but more significantly in terms of an indifference to – or, perhaps more accurately, a fundamental questioning of – conventional notions of musical focus, structure, narrative and coherence. After all, to quote Robert Hasegawa writing of Xenakis, "[w]hat does 'coherence' mean in the work of a composer whose music is often written with stochastic processes and random walks?"[62] Put simply, I wanted to explore ways to thwart my own inclinations to control, throwing out the extra-musical bathwater while retaining the absolute musical baby. I regard this not as a negation of composition as an act of self-expression, but rather a means toward a more unpredictable, spontaneous and intuitive manifestation of it, of the kind advocated by Roger Alsop:

[61] Strictly speaking, I have composed for orchestral-scale forces twice before, but *Chorus angelorum te suscipiat* (2009) is an orchestral 'explosion' of an earlier piano work, whereas *L'Ensemble Mystique* (2010) is a transcription of various organ works by Charles Tournemire. *Cloud Triptych* is therefore my first wholly original work for orchestra.

[62] Kanach 2012, p. 233.

> By creating or subjecting oneself to a set of rules, or a set of processes such as an algorithm, one hopes to confine one's actions to an area of truer self-expression. [...] By using [algorithms], the composer can make worlds of complex interrelations, generating cascading actions that trigger other foreseen and unforeseen reactions. It is possible [...] to easily preview the results of these interactions. This ease allows me to examine my motives and actions and their results while working, without the preciousness of prolonged, painstaking and single-minded efforts required by more traditional compositional methods. As each action of a composition algorithm has foreseen and unforeseen reactions, one tends to want to be accountable for these reactions. However, this sense of accountability is tempered as the process becomes more practiced and intuitive, much like the sense of accountability for each sound a novice violinist makes is reduced as the playing becomes more practiced and intuitive. At this point of expertise, the building of a computer algorithm becomes equivalent to making an instrument, learning how to play it and creating a composition all at the same time.[63]

Inspiration and Architecture of *Cloud Triptych*

Here I present the elements fundamental to *Cloud Triptych*, chiefly the collection of cloud types from which its musical behaviours were inspired and subsequently derived. This is followed by a discussion of how I developed them into the architecture of the CloudCube program in order to generate the instrumental material.

The specific inspiration that provided the context for a more abstract compositional outlook came from my interest in cloud formations. The starting point for the work's musical behaviours was Richard Hamblyn's *The Cloud Book*, which contains a detailed description and photographs of each type of cloud formation. These are broadly classified into twenty-seven discrete varieties, grouped in three tiers of nine, according to low, medium and high altitudes (referred to via the abbreviations 'CL', 'CM' and 'CH' respectively). I decided to use this arrangement as the compositional basis for a large-scale work for orchestra, to be of around thirty minutes' duration.

First, I made brief written summaries of all cloud types, encapsulating their main behavioural characteristics. These summaries then served as the basis for sketches, visualising the behaviour. Rather than attempting to create a literal aural equivalent of each cloud type, I used the written summaries simply as inspirational launch pads for imaginative flights of fancy, resulting in quickly-drawn sketches that illustrate an overview of the orchestral behaviour. Many of the cloud types bear similarities to or are altered versions of other types, and I reflected this in the way each behaviour was designed. An example of one of these summaries with its accompanying sketch is shown in Fig. 3.1.

[63] Alsop 1999, p. 90.

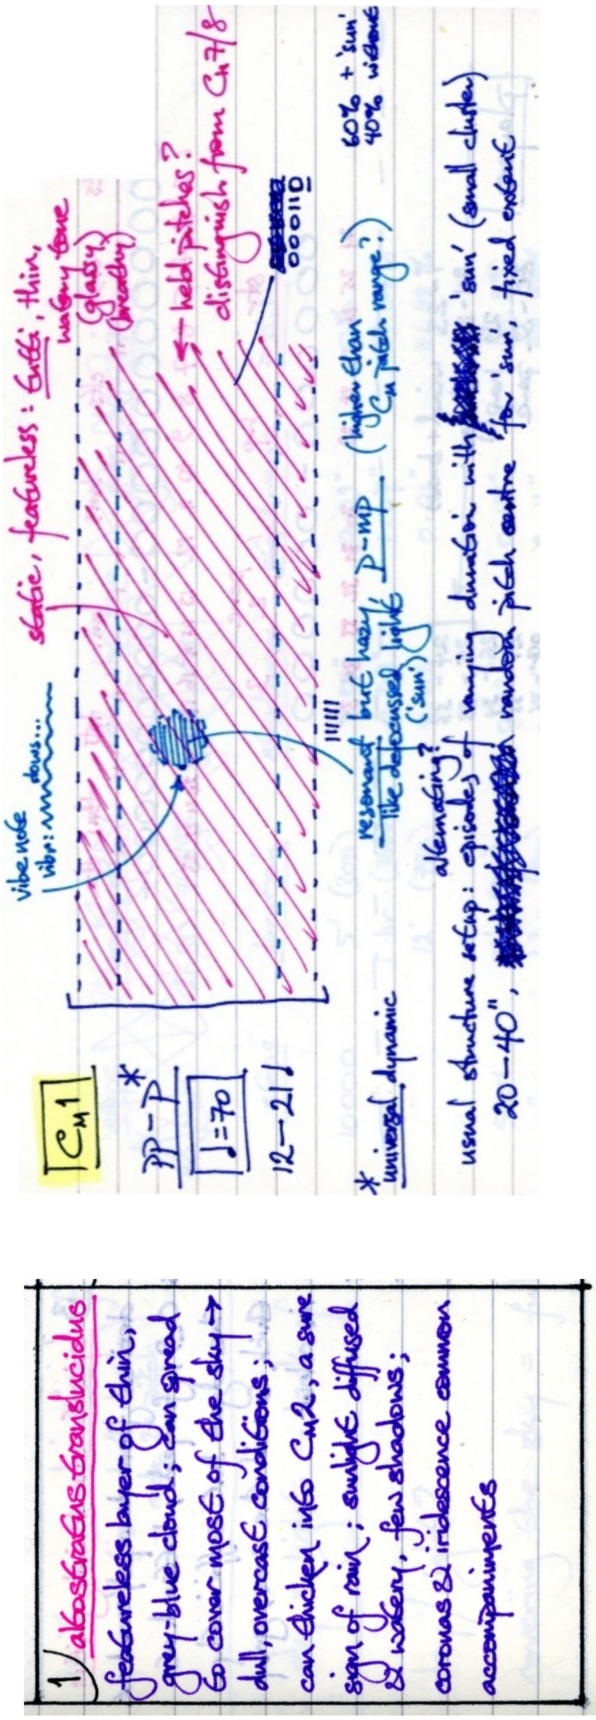

Fig. 3.1. Summary and visualisation sketch for cloud type CM1.

The sketch visualises the behaviour of cloud type CM1, characterised by a large background textural 'sheet' of string glissandi with a defocussed 'sun' comprising upper wind and brass melodies clustered around a vibraphone pitch centre. Among other things, the sketch also indicates that the 'sun' is present 60% of the time.

Through this process of sketching, the instrumental forces necessary for the work became clarified, resulting in a moderately large orchestra:

woodwind	3 flutes (1/2=picc, 2/3=alto) 3 oboes (2/3=cor) 3 B♭ clarinets (1=E♭, 2/3=bass) 3 bassoons (2/3=contra)	**percussion (2 players)**	glockenspiel vibraphone 3 triangles (small–large) 3 Chinese cymbals (small–large) bass drum (+ rute) 12 wine glasses (high–low) 6 rototoms (small–large) 8 tam-tams (small–very large)
brass	4 horns 3 trumpets (1=picc) 3 trombones bass tuba		
strings	16.14.12.10.8		

In all of the preliminary works described in the previous chapter, behaviours were unvarying in both their essential character and the way they operate, since they consisted of an unchanging set of rules. I wanted here to make the behaviours mutable, with potentially various modes of operation, changing over time (reflecting the real-world character of clouds). The following sections provide a brief description of all twenty-seven behaviours. For the sake of clarity and brevity, not all aspects of each behaviour are discussed. I decided at the outset that each altitude should have a distinct focus, instrumentation, register and range of tempi.

Low altitude (CL)

The low altitude uses the full orchestra, with a modest emphasis on the brass. It has a relatively low pitch domain and a slow tempo range, ♩ = 40–50. The focus is on harmony (in conjunction with melody and texture), and to that end the low altitude behaviours share an underlying, repeating series of primaries – acting here as deep fundamentals – that determines their weighted tonality. The series uses all twelve chromatic notes:[64]

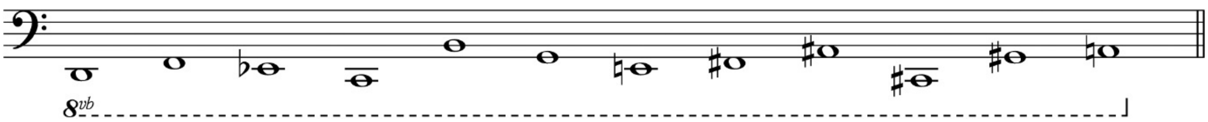

Fig. 3.2. Low altitude primary pitch series.

[64] The series is taken from an abandoned undergraduate composition titled *Serial Demons*.

Since CL9 is *cumulonimbus capillatus*, the most dominant cloud type, the duration of each primary is determined solely by CL9's behavioural rules.[65] The extent to which these primaries are emphasised varies from behaviour to behaviour.

Table 3.1. CloudCube: summary of low altitude (CL) behaviours.

CL7 *fractus* • ♩ = 50 • Thin, insubstantial, ragged viola/ cello melodies, sul pont. • Staccato muted brass notes, reinforced by bass drum with rute. • Episodes: 30% violas only, 30% cellos only, 40% all.	CL8 *cumulus & stratocumulus at different heights* • ♩ = 50 • Top half: two narrow bands of upper wind melodies. • Bottom half: lower wind melodies. • Plus either homophonic upper strings (episodes divided into two to six equal durations) or muted brass/lower string chords (chromatic index=3).	CL9 *cumulonimbus capillatus* • ♩ = 40 • 40%: vast tutti melody in octaves. • 40%: melody in octaves plus low held chords. • 30%: 'drama': 'poised' violin 1 moment – tutti outburst (random mayhem) – brief silence.
CL4 *stratocumulus cumulogenitus* • ♩ = 40 • Individual lower wind/brass melodies and lower string chords. • Primaries articulated by tam-tams. • Episodes: melodies 20% silent, 80% present (chords always present).	CL5 *stratocumulus stratiformis/ castellanus/lenticularis* • ♩ = 40 • Three independent bands of wind/brass melodies, with 25% probability of articulations (accent, heavy accent, timbral trill). • String chords in adjacent pitch bands (each ~1 octave wide). • Melodies occur in odd-numbered episodes only.	CL6 *stratus nebulosus* • ♩ = 40 • Background mixture of wind/brass air noise and low string clusters. • Foregrounded 'sun' comprising violin cluster (variable pitch centre) and three triangles. • Episodes: sun 60% present, articulated by Chinese cymbals.
CL1 *cumulus fractus/humilis* • ♩ = 50 • Individual brass melodies. • Primaries articulated by handbells.	CL2 *cumulus congestus/mediocris* • ♩ = 50 • Individual lower wind melodies and brass chords. • Primaries as CL1 plus tam-tams. • Episodes: 20% melodies only, 20% chords only, 60% both.	CL3 *cumulonimbus calvus* • ♩ = 50 • Individual upper wind/trumpet melodies and bassoon/brass chords. • Primaries as CL2, with additional contrabassoon/tuba pedals. • Episodes as CL2, but higher pitch domain and pedals always present.

Medium altitude (CM)

The medium altitude uses a reduced size orchestra, rarely involving all sections at once. In particular, there is much less use of the strings. It has an intermediate pitch domain and its tempo

[65] Although it has a very low base (hence classed as low altitude), *cumulonimbus capillatus* ascends higher than any other cloud type; its designation as CL9 – the ninth of the low altitude clouds – is the origin of the expression "cloud nine".

range is moderate, ♩ = 60–90. The general focus of this altitude is on melody, though textural emphases also feature in some of the behaviours, with and without melodic content. Harmony – in the sense of weighted tonality – is not explored here.

Table 3.2. CloudCube: summary of medium altitude (CM) behaviours.

CM7 *altocumulus stratiformis duplicatus* • ♩ = 90 • Two to seven overlapping bands of wind/brass melodies, variable pitch centres (as CM2/CM5 but more rhythmically intricate). • Thick veiled string texture. • Episodes: each band 50% present (strings always present).	CM8 *altocumulus castellanus/floccus* • ♩ = 60 • One to three narrow bands of repeated brass notes, variable pitch centres. • Occasional descending divisi violin 'streaks', sul tasto, 10% brief tremolandi, plus wine glasses. • Episodes: 'streaks' 30% present (bands always present).	CM9 *altocumulus of a chaotic sky* • No unique definition: randomly combines the other eight CM behaviours (each player has complete behaviour) • Combined tutti effect: chaotic, unpredictable simultaneity of behaviours.
CM4 *altocumulus lenticularis* • ♩ = 60 • Isolated tapering clusters of barely-moving brass, becoming air noise (timbral rather than dynamic fade).	CM5 *altocumulus stratiformis* • ♩ = 80 • Four to seven overlapping bands of wind/brass melodies, variable pitch centres, 10% accents. • Episodes: each band 50% present.	CM6 *altocumulus cumulogenitus* • ♩ = 70 • Tutti swelling harmonic series (primary not included; lowest note heard is first harmonic). • Independent vibraphone melody, 20% accent, tremolando or damped notes.
CM1 *altostratus translucidus* • ♩ = 70 • Thin background 'sheet' of divisi string glissandi, sul tasto. • Defocussed 'sun': upper wind/brass melodies around variable pitch centre, reinforced with vibraphone and three triangles. • Episodes: 'sun' 60% present.	CM2 *altostratus opacus/nimbostratus* • ♩ = 70 • As CM1 but 'sun' obscured more. • Thick 'veil' of divisi string glissandi, sul pont./tremolando. • Wind/brass as CM1 but quieter, triangles muted with rubber beater.	CM3 *altocumulus translucidus/stratiformis* • ♩ = 80 • One to four overlapping bands of wind melodies, variable pitch centres, 10% trills. • Episodes: each band 60% present (all four can be absent, resulting in silence).

High altitude (CH)

The high altitude employs a minimal orchestra, omitting the brass. It has a high pitch domain and a fast tempo range, ♩ = 100–140. Melody and harmony are essentially absent from this altitude, which focuses almost entirely on texture.

Table 3.3. CloudCube: summary of high altitude (CH) behaviours.

CH7 *cirrostratus fibratus/nebulosus* • ♩ = 110 • Static divisi string 'veil' with tam-tam tremolandi. • Sporadic wind clusters (60% present), variable pitch centres, 50% timbral trill, with Chinese cymbals.	CH8 *cirrostratus not progressively invading the sky* • ♩ = 110 • As CH7 in narrower pitch domain with reduced string divisi.	CH9 *cirrocumulus stratiformis/ floccus/lenticularis* • ♩ = 140 • Upper wind/violas: lozenge-shaped pitch formations (60% present), variable pitch centres.
CH4 *cirrus uncinus/fibratus* • ♩ = 120 • As CH1, with increasing pitch deviation from tendencies and thicker texture.	CH5 *cirrostratus* (<45°) • ♩ = 100 • Wind/strings: three registral/ durational bands: 50%: high/short, 33%: middle/medium, 17%: low/long. • Plus lowest three rototoms.	CH6 *cirrostratus* (>45°) • ♩ = 100 • Wind/strings: as CH5 in higher pitch domain with thicker texture. • Plus highest three rototoms.
CH1 *cirrus uncinus/fibratus* • ♩ = 120 • Delicate string motes of pitch, downward glissandi, one or three descending tendencies. • Episodes: 50% one/three tendencies.	CH2 *cirrus spissatus/castellanus/ floccus* • ♩ = 130 • 50%: as CH1, one tendency. • 50%: upper wind clusters (60% present) and lower motes of pitch, non glissando.	CH3 *cirrus spissatus cumulonimbogenitus* • ♩ = 140 • Upper wind/string phrases, roaming pitch centres, plus bass drum (50% present).

Following the process of sketching, I turned my attention to the architecture of the software environment in which the behaviours would be codified and used to generate musical material.

CloudCube

As already described, clouds are classified into twenty-seven types, arranged in three altitudinal tiers of nine. I decided to use this arrangement as the model for an ambitious expansion of my existing method. In the preliminary works, a maximum of two or three behaviours were used, moving between them in a single dimension, along a straight line. For *Cloud Triptych* I designed a three-dimensional framework, comprising twenty-seven behaviours occupying a three-dimensional 'metacube'. Due to the nicely incongruous combination of the morphological fluidity of cloud formations within a fixed geometric environment, I opted to name the new program written to generate these pieces CloudCube.

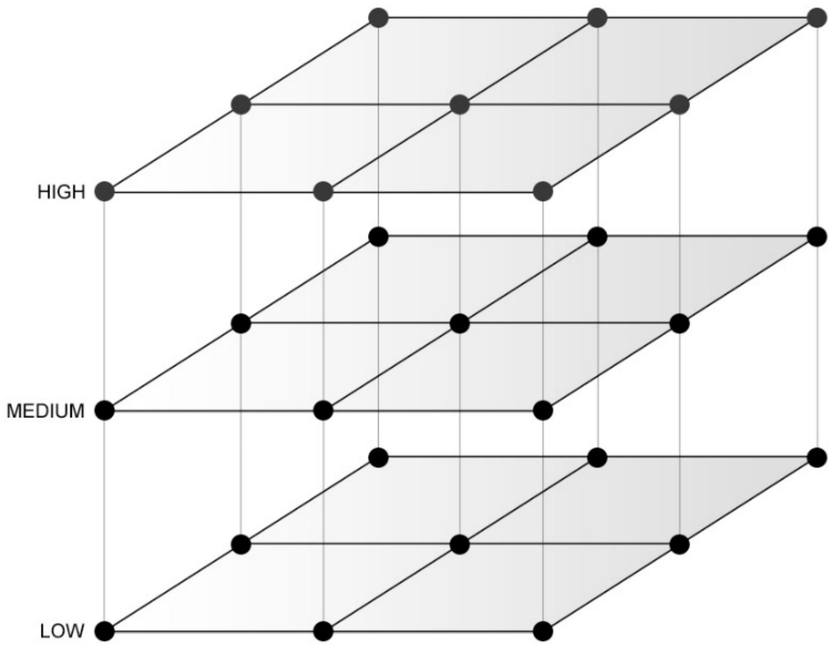

Fig. 3.3. The three-dimensional twenty-seven-behaviour framework used in CloudCube.

In addition to the motivations outlined at the start of this chapter, another primary aim was to improve the nature of my musical transitions. In all the preliminary works, their probability calculations led to the choice of just a single behaviour, the rules of which were then used to generate musical material. This 'jump-cutting' between behaviours results in what may be described as a 'coarse' transition, since at any given moment (within each part) only one behaviour is being heard.[66] The intention now was to upgrade this decision-making process to facilitate the combining of aspects from multiple behaviours at once. The intended result of this combinatoriality would be more smooth, subtle and complex musical transformations.

In order to introduce behavioural combinatoriality, I decided that the use of probability, instead of selecting an entire behaviour, should instead select individual modules, in order to produce a composite made up of parameters from different behaviours. The number of modules to be used was determined by the probability distribution within the metacube. I fixed each behaviour's range of influence to extend as far as its adjacent neighbours, resulting in a spherical range henceforth referred to as its *orbit*. At the centre of this orbit I established a *core*, a proximity threshold (equal to one-tenth of the orbit's radius) that acts as a kind of 'event horizon', within which the probability instantly becomes 100%, ensuring the behaviour is unavoidable.

[66] Of course, within the context of works for multiple players, there is the possibility of more than one behaviour being heard at any particular instant, due to the players' independent movement between behaviours.

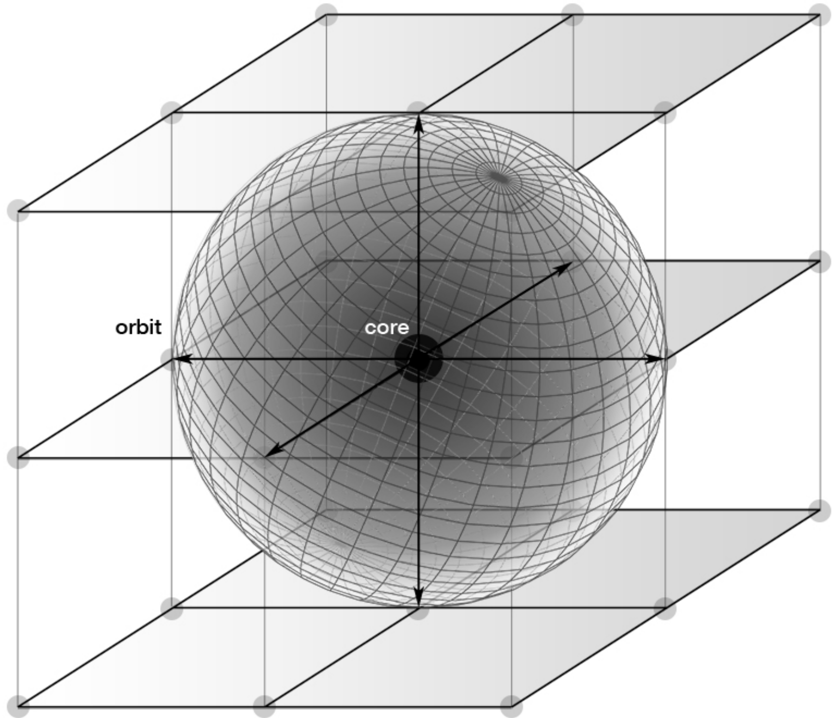

Fig. 3.4. Orbit and core of the central behaviour.

The upshot of this probability distribution is that wherever the cursor moves inside the 3×3×3 metacube, the only behaviours that can probabilistically be in range are those whose nodes are at the eight corners of the smaller 2×2×2 cube the cursor is presently within, as shown in Fig. 3.5. This region is henceforth referred to as a *metarange*.

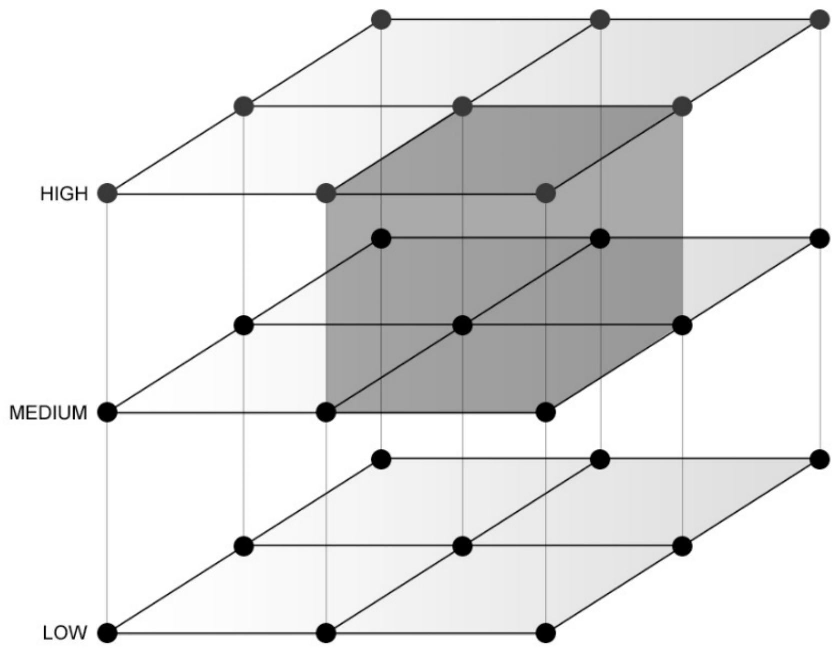

Fig. 3.5. One of the 2×2×2 cubes circumscribing a metarange.

If the cursor were to find itself at the centre of a metarange, equidistant from all eight behavioural nodes, each would have an equal probability (12.5%). For this reason, every behavioural definition comprises eight modules, enabling combinatoriality of up to eight different behaviours at once. The function of the eight modules is summarised in Table 3.4.

Table 3.4. CloudCube: behavioural modules.

module	details
1. Instrumentation	which players/instruments are used
2. Structural duration	how much material is created
3. Quantisation	how time is subdivided
4. Rest probability	likelihood of notes being rests
5. Rhythm	durational limits/rhythmic language
6. Pitch	pitch limits/harmonic language
7. Articulation	how material is played
8. Dynamic	dynamic limits

I now explain in a little more detail the main aspects of each of the modules.

1. Instrumentation

This specifies which players/instruments are used. Each behaviour has a specific instrumentation; if, when generating material, the selected player is absent from this instrumentation, the remaining modules generate silence. For the sake of simplicity, within each behaviour players only use a single instrument; changes in instrument (e.g. from flute to piccolo) occur as the music moves between behaviours.

2. Structural duration

This determines the quantity of material to be created, specified as a duration in seconds.[67] Each altitude has the following upper/lower structural durational limits:

[67] This module was changed drastically quite late in the composition process. Initially, as described in Chapter 1 (p. 14), structural durational limits were specified in beats, and each behaviour had unique limits, but this turned out to be over-complicated. More importantly, the low altitude nodes, which were given a generous and lengthy range, proved problematic to the overall length of compositions, dramatically increasing them beyond the durations specified in the trajectories. For example, in a test using a three-link trajectory chain (CL1 → transition → CL2), where each link lasted thirty seconds, instead of producing a piece lasting approximately ninety seconds, the result had a duration of over two minutes; this was due to CloudCube only being able to choose long structural durations, stretching the overall duration.

Table 3.5. CloudCube: structural durational limits.

altitude	structural durational limits
high	4–10"
medium	8–14"
low	12–18"

CloudCube randomly chooses a duration from within these limits, and the result is converted into a number of beats, taking the prevailing tempo(s) into account (see 'Trajectories, Time, Tempo' below).

3. Quantisation

The usual quantisation grids are available, as described on page 18. However, I decided to omit the appoggiatura option from the grace grid, as I had tended to avoid using it in previous works since it often felt somewhat clumsy, and in the context of *Cloud Triptych*'s polyphony seemed redundant. The grace grid therefore consists only of the acciaccatura (beat position 0.95).

4. Rest probability

Many behaviours have a simple, fixed value for the probability that its rhythmic values will be rests. In some more complex instances a behaviour has various values depending on the instrument. For example, due to the way it behaves I decided that CL8 should have these values:

Table 3.6. CloudCube: CL8 rest probabilities.

players	rest probability
flutes, oboes, clarinet 1, bassoon 1	14%
clarinets 2/3, bassoons 2/3	25%
violins, violas	20%
trumpets, trombones, lower strings, percussion	0%

5. Rhythm

The rhythm and pitch modules together comprise the most significant parts of each behavioural definition. In theory, the rhythm module simply specifies the durational limits and then makes selections accordingly, but in practice it contains by necessity a great deal of information about how the behaviour operates over time. Almost all of the behaviours have different sets of rules for the various sections of the orchestra.

Although the content of this module is varied and complex, a recurring procedure is to fill the current structural duration with rhythms, the durations of which are determined in the usual way as random values within the behaviour's upper/lower limits. The way this is done is as follows:

1. Durations are chosen until the length of the structural duration has been reached. If the final duration is outside the behaviour's durational limits, the process is repeated as necessary. Each duration is either a note or a rest; this is determined by the current rest probability.
2. The order of the durations is then randomised. This is essential as the process in step 1 often results in the shortest durations being at the end.
3. Finally the durations are aligned from floating-point to beat positions according to the quantisation lattice in the usual way.

It should be noted that the rhythm module does not make decisions about phrasing; this is determined intuitively, usually by considering where rests fall in conjunction with changes in dynamic.

6. Pitch

Having established all the rhythmic information, the pitch module assigns pitches and registers to each note. The overall pitch domain spans eight octaves, from C0 to C8, which are divided between the three altitudes according to overlapping pitch limits:

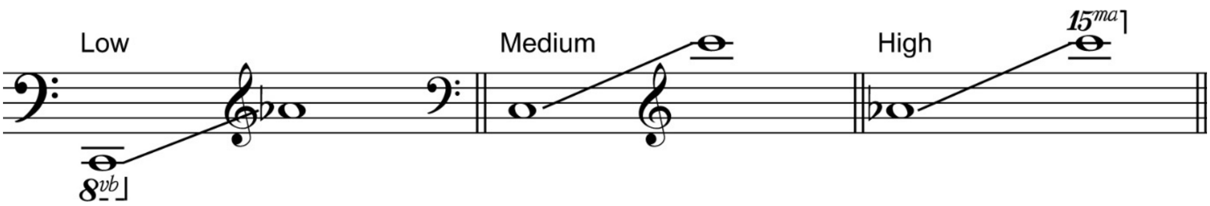

Fig. 3.6. CloudCube: altitudinal pitch limits.

Beyond this, like its rhythmic counterpart the pitch module contains intricate, detailed behavioural information, as well as providing alternate rules for certain players or sections of the orchestra. For the most part, each rhythmic note is assigned to a single pitch (in the form of a number between 0 and 96), but chords of up to four notes can also be generated for multiple-stopped string notes and degrees of relative pitch for unpitched instruments. The pitch module is unique in that it can, in response to its calculations, make retroactive amendments to decisions made at earlier stages in the modular chain; if, for example, numerous attempts fail to find an appropriate pitch (i.e. within the range of the present instrument), then the note is converted to a rest.

7. Articulation

Articulations can be of two types, structural and individual. Structural articulation is applied to a complete structural duration, affecting every note. Instructions such as 'legato', 'con sord.' and so on are examples of this. Individual articulation is dealt with on a note-by-note basis, encompassing a wide variety of possibilities, including staccato, accents, harmonics, trills and tremolandi, fluttertongue, indeterminate pitches, ricochet, cuivré and types of vibrato.

8. Dynamic

The dynamic module is one of the simplest, comprising upper/lower limits according to the following scale (again, the limits can be the same):

Table 3.7. CloudCube: range of dynamics.

1	2	3	4	5	6	7	8	9
n	*ppp*	*pp*	*p*	*mp*	*mf*	*f*	*ff*	*fff*

The module can also indicate dynamics that are part of an ongoing crescendo/diminuendo. Usually, though, gradations of dynamic are composed intuitively. In *Cloud Triptych*, the general approach taken has been to emphasise gradual changes within a phrase and sudden changes from phrase to phrase.

Due to the complex interactions now taking place between the behaviours, musical outcomes of trajectories through the metacube are difficult to predict at the micro-level, and situations can inevitably arise where conflicts may be unresolvable. Simple examples already noted are where the present instrument is not included in the selected instrumentation module, in which case no material can be created and its structural duration will therefore be silent. Likewise, the pitch module may not provide anything usable within the present instrument's range; in this instance, if 500 attempts to calculate a usable pitch result in failure, the note is converted to a rest and the program moves on. There are also more complex 'coping mechanisms' that, in certain circumstances, permit parametric limits to be systematically widened as a preliminary attempt at a solution. In this respect, although CloudCube includes aspects of computation that bear resemblance to constraint programming – searching for a solution that satisfies a set of initial conditions, the approach taken in software such as OpenMusic[68] and Strasheela[69] – CloudCube's primary purpose in generating the material for an entire composition means that strategies to resolve irreconcilable conditions must be integrated into the program in order that it will continue

[68] cf. Bresson, Carlos and Assayag 2011.
[69] cf. Anders 2011.

in any circumstance. (This is easy to say, but in practice is rather more complex and difficult to achieve!)

Having presented the mechanics by which CloudCube determines the musical parameters of each behaviour, I now turn to the means by which movement takes place through the metacube, thereby generating the composition.

Trajectories, Time, Tempo

Within CloudCube the relationship between behaviour and time is fundamentally different from that in the preliminary works. Hitherto, the timeline was established first and the behaviours then assigned nodal positions along it; in CloudCube this situation is reversed. The timeline is now, in essence, a timecurve, and its trajectory through the metacube must be defined. This is done via a graphical representation of the metacube, displaying a large horizontal view from above – showing the behavioural nodes and their cores and orbits – beside a narrow sideways vertical view (Fig. 3.7).

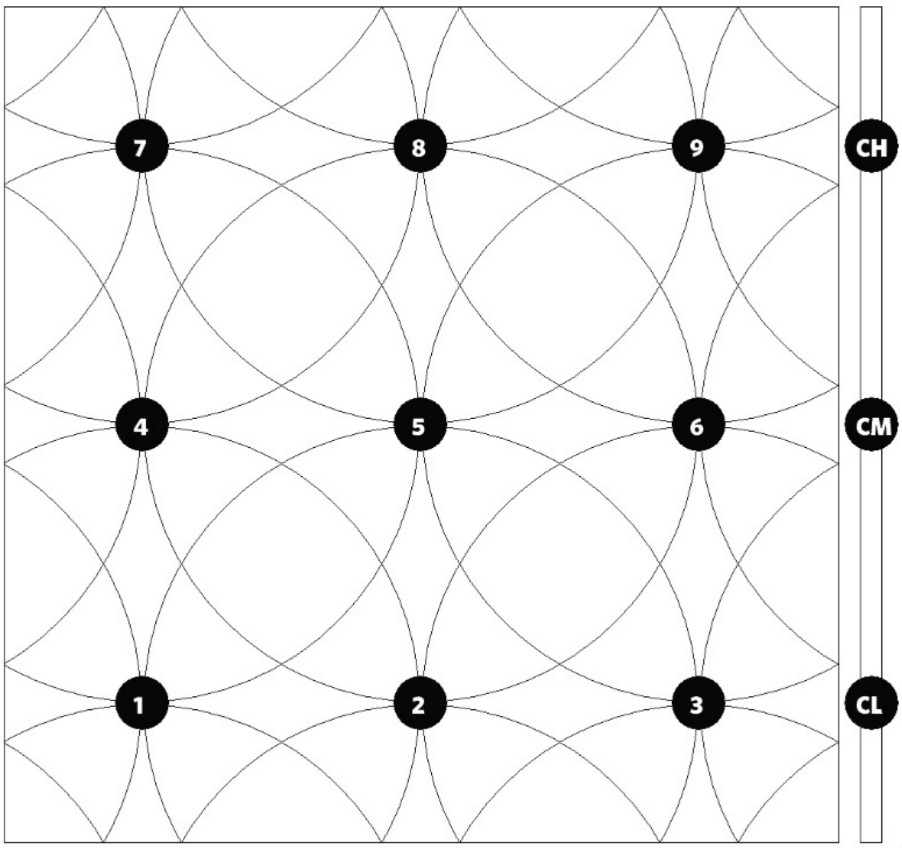

Fig. 3.7. The CloudCube graphical interface.

Two or more 'control points' are placed onto this – the *x*/*y* coordinates on the left followed by the *z* coordinate on the right – which are then used to construct a trajectory through the metacube. This can be done either with straight lines or using the control points as the basis for a Bezier curve.[70] When completed, the duration of this trajectory in seconds is specified; CloudCube then creates a timeline by dividing up the trajectory into a large number of small time-steps, such that no individual step is longer than one second. The necessity of these time-steps is to ensure ensemble synchronisation. As in the preliminary works, each instrument's material is created individually, and the subdivision into time-steps ensures that all instruments move along the timeline at an equivalent overall rate. When the material for a particular structural duration has been generated, CloudCube checks at what time position the material finishes and rounds this to the nearest time-step, guaranteeing overall temporal consistency. As the timeline is being created from the control points, CloudCube simultaneously calculates the probabilities of all the behaviours at every time-step and, in the usual way, comparing to a random number, makes the corresponding decisions for each of the eight modules.[71] Although different in methodology and function, this approach bears striking similarities to the two-dimensional 'metasurface' implemented in the live performance software Audiomulch. Sets of parameters are assigned to 'snapshots' positioned freely on the metasurface; Audiomulch then interpolates parametrical values in real-time as the mouse is moved with respect to these snapshots.[72]

Tempo is dealt with separately as it must be determined *ex machina*, applied to all instruments the same, otherwise material could be created at different tempi simultaneously. Each behaviour includes a rule about tempo (in all instances a fixed value); CloudCube determines the overall tempo according to whichever node is closest to the position of the current time-step. In the (unlikely but not impossible) case of more than one node being equally close, the lowest node number is chosen. Changes of tempo must be at least three seconds apart. Use of tempo changes marks a departure from my previous work which, as already discussed, has tended to employ a fixed tempo throughout. However, I deemed flexibility of tempo not just desirable but unavoidable; clouds, after all, are ever in flux, subject to extreme variations of wind speed. In general, I have mirrored the way wind speed operates in nature, with the fastest speeds found at the highest altitudes. Three tempo ranges are used, shown in Table 3.8. My approach has been to decide intuitively whether tempo changes should be sudden or gradual, according to context.

---

[70] A Bezier curve is a curve defined by a number of 'control points' (typically three or four). A recursive algorithm (invented by French mathematician Paul de Casteljau) uses these control points to generate individual points on the Bezier curve; the resolution of the curve can be scaled indefinitely.

[71] These probability calculations are based on the principles of trilinear interpolation. This uses the cursor's present position as an axial point for dividing the metarange into eight smaller cuboids. The volumes of these cuboids as a proportion of the total volume of the metarange determines their probability (with each cuboid corresponding to the node diagonally opposite).

[72] cf. Bencina 2005. Unlike the linear probabilistic orbits used in CloudCube, Audiomulch's metasurface uses a Voronoi diagram created using the snapshots as points.

Table 3.8. CloudCube: altitudinal tempo ranges.

altitude	tempo range
high	♩ = 40–50
medium	♩ = 60–90
low	♩ = 100–140

The results of all these calculations are saved in data files, since CloudCube needs to move along the same trajectory multiple times, in order to create the material for each instrument.[73] If desired, many of these trajectories can be linked together to create a chain of complex movements through the metacube at different relative rates.

Behavioural Mutability

As already discussed, behaviours include a variety of parameters the specific values of which are determined randomly from within predefined upper/lower limits. I wanted to extend this by making the behaviours mutable, giving them multiple modes of action, exhibited episodically, with the lengths of these episodes also determined randomly. This presented an interesting problem, since the mutability must be precisely the same for all instruments. In other words, aspects of this endemic randomness, above all the episode durations, need to be exactly repeatable. For example, if the randomly-determined durations of a behaviour's first three episodes were twelve, ten and nineteen seconds respectively, these same durations must apply to every player, otherwise the behaviour cannot establish a coherent identity. CloudCube therefore needed to be able to make decisions that were both random yet capable of being duplicated. This was achieved by using the decimal places of pi as source for the random numbers.[74] Each behaviour is assigned a cache of 5,000 digits, which are then used for all random calculations that require exact duplication (these groups are successive, so there are no overlaps or duplications among the digits). The decimal places are used in pairs to form a number between 0 and 99; I refer to these resultant numbers as *pirns* (*pi r*andom *n*umbers).

[73] This is akin to the use of 'motion control' in film-making, where camera movements are captured as data in order to be precisely repeated numerous times.

[74] The question of how to generate truly random numbers, as well as the extent to which the digits of pi are truly random, is one that continues to be regularly argued about in mathematics. Charles Ames has addressed their use in an artistic context on more than one occasion, concluding that "[t]he subject of algorithmic independence is a philosophically tricky one, since in one sense any sequence of numbers calculated by a computer must be wholly dependent, while in a second sense a computational procedure that responds to ill-defined criteria does not strictly qualify as an algorithm. [...] [f]rom a human perspective, a sequence is independent or random over a given range when it is wholly unpredictable from a perceptual standpoint." (Ames 1992, p. 56) The arguments show no sign of being resolved – Tu and Fischbach 2004 and Marsaglia 2006 present a more recent example of the ongoing cut-and-thrust with regard to the veracity of pi's randomness – but as far as anyone has been able to discern, the digits of pi *are* random, and within the context of my own work are certainly "wholly unpredictable".

The pirns can be used within each behaviour in any number of ways, but the principle is very simple: limits are established and the pirn is then converted to an equivalent value between these limits. Fig. 3.8 illustrates an example with regard to episode duration, using pirns derived from the first eight decimal places of pi to create durations of between five and ten seconds.

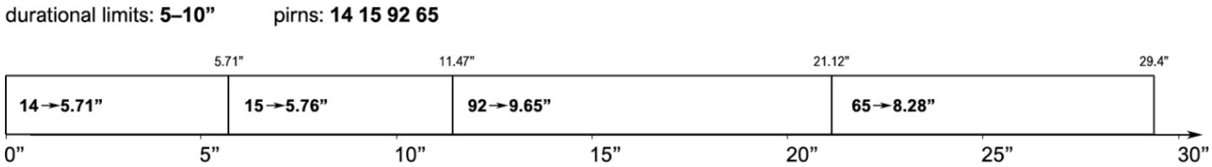

Fig. 3.8. Episode durations determined from pirns converted into limited values.

As CloudCube moves along the timecurve, it can locate its position within these episodes by jumping through the decimal places (determining the duration of each episode as it does so) until it reaches the episode that contains the current time position. In practice, the behaviours utilise many more than just two pirns (CL9 uses forty-one!); the number of pirns used must be declared within each behavioural definition so that CloudCube knows how to skip through them to locate its current position.

Output

In order to understand better the relationship between movement within the metacube and the resultant music, I designed the output from CloudCube to show both the trajectory chain and the material being generated, along with various associated calculations. Illus. 3.1 shows a typical screenshot. The trajectory chain is on the left, above which can be seen various calculations pertaining to the current point in the chain, including the timestep (65/450), the *x*/*y*/*z* coordinates (1621.6, 1023.68, 1625), the behaviour determining the tempo (3.6, i.e. CH6), the breakdown of the eight modules (all CH6) and the behavioural probabilities across all three altitudes. On the right is the musical material, showing the instrument, tempo, beats, time positions and a breakdown of each note and rest (the small digits indicate the rhythmic positions) followed by articulations and dynamics. CloudCube saves a series of screenshots for each instrument which form the basis for notation.

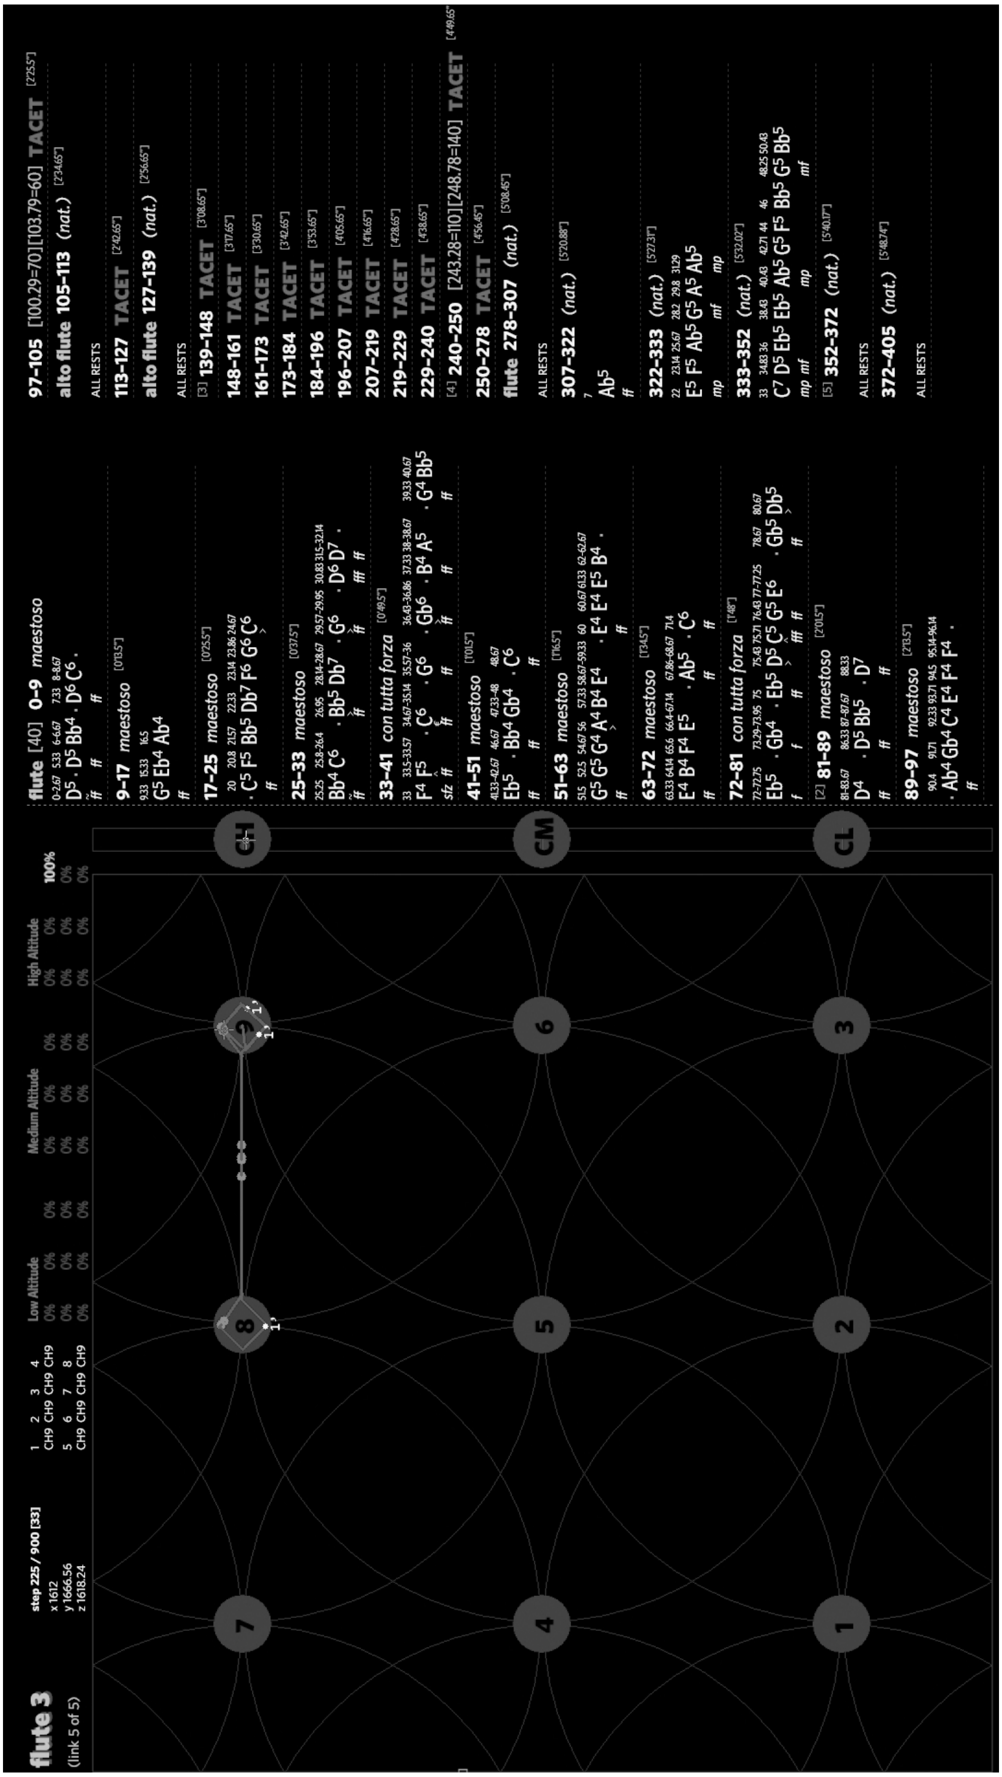

Illus. 3.1. CloudCube: a typical screenshot.[75]

[75] This screenshot shows Flute 3's material for the first 405 beats of the third movement of *Cloud Triptych*.

However, while this visual output goes some way to clarify the relationship between movement and music, the complexity within the CloudCube environment – in terms of the way rules from different behaviours impact upon and affect each other – is such that it is very much more difficult both to predict what will happen in advance and also to understand retrospectively which specific decisions led to the resultant outcomes. This is a deliberate consequence of my decision to design CloudCube in such a way as to limit the composer's ability to control everything, and thereby approach composition less as a definitive act than as a speculative one.[76]

Once the architecture and functionality of CloudCube had been designed and created, the behaviours were coded into the program, one module at a time. Each behaviour was tested in isolation to ensure it was acting correctly (this was done using trajectories that remain within the behaviour core).

Cloud Triptych (2016) for large orchestra

Having discussed the architecture and aspects of the functionality of CloudCube, I now present the end result of its use: the main work in the portfolio, *Cloud Triptych*. My original intention had been to compose a single, large-scale work, but I opted instead for a three-movement scheme, as I felt this would facilitate a more dramatic presentation, exploring different parts of the metacube in each movement.

Due to the complexity of the music at certain points, resulting in a large number of individual staves, I decided – drawing on my experiences as a conductor of contemporary music – to create the score in two volumes:

- a conductor's score, in which the string sections are reduced to a single stave with indications of divisi and a brief mnemonic summary of their behaviour shown in a box followed by a wavy line, indicating continuation of the behaviour (where the music can be shown on a single stave, this is done);
- a strings score, containing the complete string music as a reference for the conductor.

The mnemonic system used in the conductor's score first indicates the approximate extents of duration and pitch. Duration is expressed as **S**, **M** and **L** (short, medium and long), pitch as **L**, **M** and **H** (low, middle and high), the latter according to the following registral divisions:

[76] It is theoretically possible to peer 'under the hood' at some of CloudCube's calculations—the specific modular selections for each timestep, for example—as a certain amount of data is stored in the data files saved when the trajectory chain is created. However, no record at all is kept of the subsequent multitude of stochastic and otherwise random decisions made throughout the material generation process, so it is essentially impossible to reverse engineer the exact circumstances that led to the material.

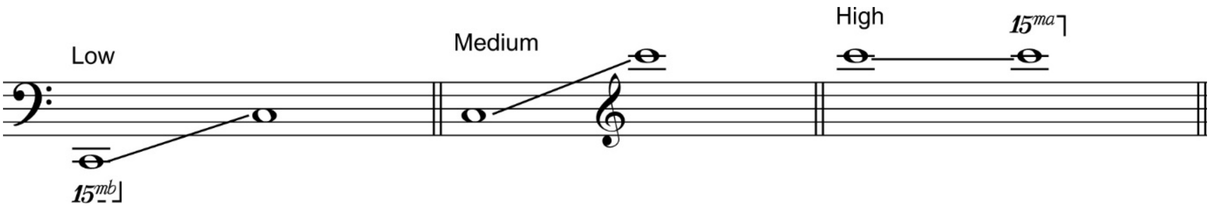

Fig. 3.9. *Cloud Triptych*: conductor's score string registral divisions.

The mnemonics are contained within a text box along with any additional information, such as poco sul pont., gliss., arco nat. and so on; articulations that appear sporadically are shown in parentheses, e.g. (accents). Divisi is declared prominently in a circle above the box; dynamics are shown beneath the box, usually in the form of a range, e.g. ***p* – *mp***. If the mnemonic occupies only part of a bar, this is clarified with an indication of its duration above the box. A complete mnemonic is shown in Fig. 3.10, taken from the second movement of *Cloud Triptych*. It indicates the second violins are divisi a 14, playing long durations in the middle register with glissandi and sporadic tremolandi, all poco sul pont., between ***p*** and ***mp***.

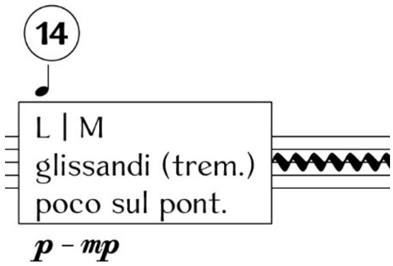

Fig. 3.10. *Cloud Triptych*, second movement, b. 56, violin 2: conductor's score mnemonic.

In the following discussion of the three movements of *Cloud Triptych*, reference is made to positions in the music when certain links in the trajectory chains begin. However, it must be remembered that, since each instrument's music is generated separately, the precise point when each link begins is unique for each player; for clarity, the discussion generalises these points.

First Movement

In the first movement, which lasts around five-and-a-half minutes, I wanted to focus on texture, beginning *Cloud Triptych* in a nebulous, somewhat intangible atmosphere. To this end I decided to restrict myself to the high altitude (which has texture as its main characteristic; cf. p. 50) and explore movement between pairs of behaviours that share strong similarities. To create the material I designed the following eight-link trajectory chain:

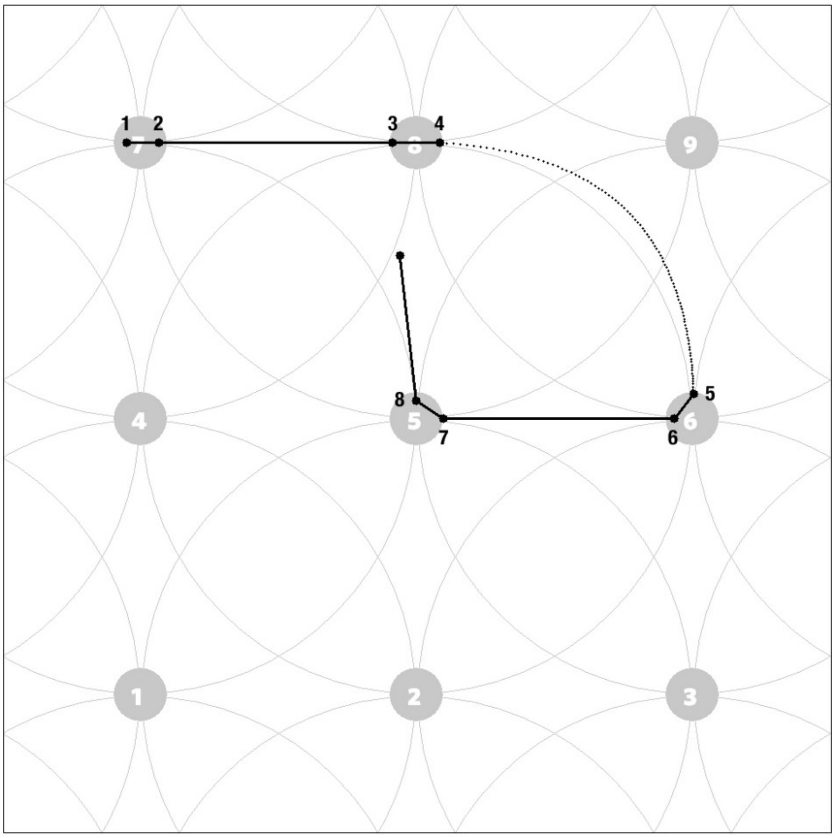

link	behaviour	duration
1	CH7	45”
2	CH7→CH8	45”
3	CH8	30”
4	CH8→CH6	30”
5	CH6	45”
6	CH6→CH5	45”
7	CH5	30”
8	→CH8	30”

Fig. 3.11. *Cloud Triptych*, first movement: trajectory chain.

The movement is broadly structured in two main sequences followed by a coda. The first sequence explores CH7 and CH8, the latter of which is a pitch-constrained version of the former, focussing for a time on each of them with a transition in between. The music therefore begins with a dense, static veil of string pitches – all sections tutti divisi, plus the additional colouration of tam-tam tremolandi – upon which a series of soft, clustered woodwind ‘accumulations’ takes place, gently embellished with Chinese cymbals. These accumulations can be seen beginning in bars 1, 9 and 21. At figure A, the transition towards CH8 begins, immediately evident in the shifts in string divisi (reduced by a half) but manifesting most in a gradual narrowing of the pitch domain. By around bar 43, CH8 has been reached, and the sequence concludes with a section in which wind clusters are absent, featuring just the string and tam-tam ‘veil’ (bb. 47–56).

The second sequence explores CH5 and CH6, which behave the same but in different pitch regions. Figure B marks the point where the transition from the first to the second sequence begins. This transition (link 4, using a curved trajectory) is complex, encroaching closely upon CH9, notably introducing glockenspiel and wine glasses in bars 67–75 (a brief hint at the work’s overall conclusion). From figure C, CH6 is established, the winds become non legato with sudden changes of dynamic, and the three highest-pitched rototoms begin. (This behaviour’s three pitch tiers are somewhat indistinct due to a relatively narrow pitch domain at this point.) The transition from

CH6 to CH5, which takes place from figures D to E, has a similar effect to that in the first sequence, bringing down the music into a lower pitch region, evident in all parts but prominent in the shift to the lowest-pitched rototoms. The arrival of CH5 at figure E is reinforced by the exit of the violins and entrance of the cellos.

The end of the movement was designed to be a vague, allusive kind of recapitulation, transitioning from CH5 to a point loosely in the direction of CH8, thereby introducing hints of the first sequence. This process begins around bar 134, becoming more apparent at figure F when the tam-tams return and the winds revert to accumulations, this time without the Chinese cymbals.

Second Movement

Cloud Triptych was designed to have an extended central panel that would enable me to explore a large number of behaviours in more complex ways. The second movement therefore uses a trajectory chain comprising eighteen links, focussing most on the medium altitude, with a total duration, following tempo adjustments, of around sixteen minutes.

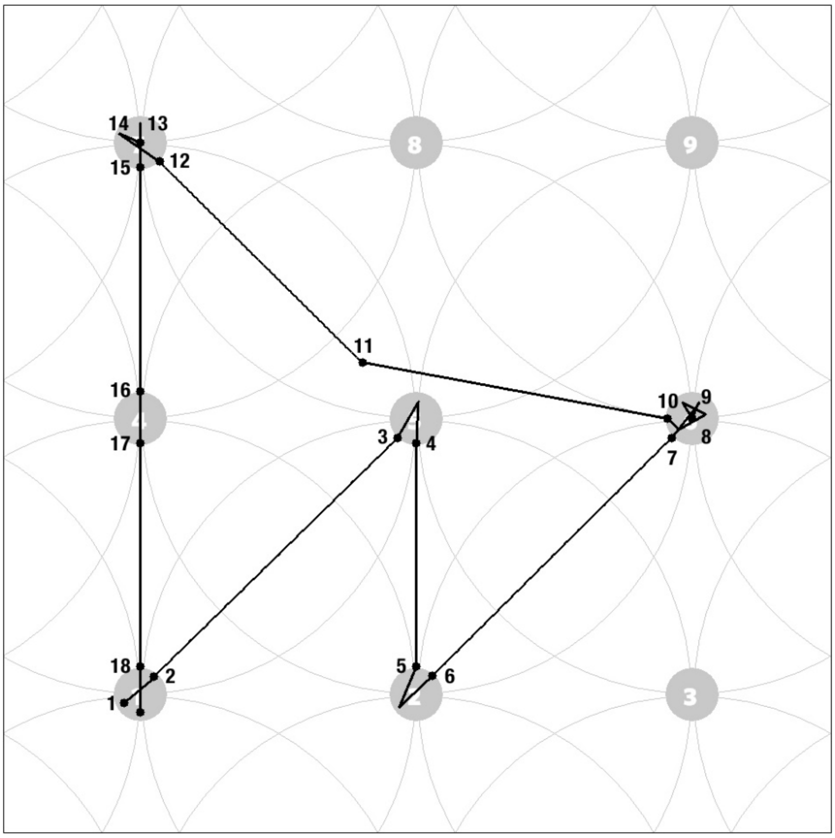

link	behaviour	duration
1	CM1	1’
2	CM1→CM5	30”
3	CM5	45”
4	CM5→CM2	45”
5	CM2	45”
6	CM2→CL6	30”
7	CL6	1’
8	CL6→CM6	30”
9	CM6	2’
10	→	40”
11	→CL7	20”
12	CL7	1’
13	CL7→CM7	20”
14	CM7	45”
15	CM7→CM4	30”
16	CM4	45”
17	CM4→CH1	30”
18	CH1	1’

Fig. 3.12. *Cloud Triptych*, second movement: trajectory chain.

Having presented relatively subtle transitions in the first movement, here I wanted to explore more contrasting behavioural juxtapositions, while at the same time using their inherent similarities to balance the movement as a whole, and also create unifying links with the outer movements. It is structured as two sequences around a lengthy centrepiece.

The first sequence plays on the similarities between CM1, CM2 and CL6, each of which involves the appearance of a 'sun' in the midst of a pitch space filled with nebulous, essentially static textures. Beginning this movement with CM1 establishes an aesthetic connection with the start of the first movement, the wind 'accumulations' and string 'veil' of the latter here becoming wind and brass 'suns' over a dense network of string glissandi. The 'suns' – beginning in bars 3, 9 and 15 – are clarified by vibraphone pitches at their centre and adorned by three triangles. Having established this soundworld, at figure A the music begins to transition to CM5, a highly contrasting behaviour comprising bands of wind and brass melodies, the first shift in *Cloud Triptych* away from focussing on texture. This transition, as a result of moving diagonally through the metacube, is infused en route with elements from both CM2 and CM4, both of which prefigure music later in the movement, such as the instances of air noise and indeterminate pitches from bar 25 onwards, derived from CM4. The bands of melodic counterpoint persist from figures B to C, during which the strings, having been present since the start of the work, finally fall silent. Beginning around figure C, the transition to CM2 is heralded by wind and brass trills associated with its 'suns'. Initially incoherent, they become suddenly clear at figure D, when CM2 is reached. The sequence is concluded with a transition to CL6 which here acts as a kind of inverted, somewhat 'wraith-like' version of what has gone before. Wind and brass now create a dense texture made up of 'melodies' articulated as indeterminate air noise, muddied with low cello and double bass clusters, above all of which occur various muted violin 'suns'.

The movement now enters its centrepiece, transitioning vertically upward from CL6 to CM6, resulting in behaviourally ambiguous material from around bar 97 until figure F. There now follows an extended two-minute tutti burst of CM6, featuring shifting harmonic series behind a flamboyant vibraphone solo in the foreground. I wanted this section to last as long as it does to enjoy the stability of both the material and the movement, which has hitherto been behaviourally fluid.

Since the first half of this movement explored behaviours with broadly similar aesthetic qualities, culminating in a lengthy central period of stability, I wanted the latter half to be altogether more disorienting. To that end, its second sequence transitions between four extremely different behaviours, beginning with one of the more unusual, CL7. This is reached via a lengthy transition (links 10 and 11, together lasting 60 seconds) that encroaches upon numerous behaviours across two altitudes, producing music in an intense state of flux, from figures G to I. Growling, soloistic low wind and brass and ragged string melodies with loud brass and bass drum staccato notes gradually yield to the thick opacity of CM7, deliberately chosen to provide a semi-obscured connection to the first half of the movement, since it employs the melodic bands of CM5. From figures J to L the music transitions to perhaps the most vague of all the behaviours, CM4, populating the brass with tapering clusters dissipating into air noise, in another rare passage that excludes the strings. Having explored a variety of low and medium altitude behaviours, the movement concludes with a final upward transition to CH1. As a result, from figure L the music is at first behaviourally complex and then increasingly delicate, displaying a kind of 'textural lyricism' that, in its use of wine glasses (bars 253ff.) and glockenspiel provides a faint echo of the first movement. The few remaining instruments fall away until only the glockenspiel remains, a fragile end to a movement that has encompassed extremes of behavioural diversity, and which in turn sets the scene for the dramatic opening of the final movement.

Third Movement

From early on in the process of designing the behaviours, I felt sure that CL9 needed to occupy a special place in *Cloud Triptych*. Although I had initially conceived it as having a climactic function, in order to exploit its power in the most impressive way I decided it should begin the final movement. The trajectory of this movement was designed to be the most simple, ascending upwards through just three behaviours, encountering one at each altitude. Lasting a little over seven minutes, it uses a five-link trajectory chain, shown in Fig. 3.13.

The first twenty-six bars, lasting a little over two minutes, clearly demonstrate CL9's unique behavioural modus operandi, based around a vast melodic line. Examples of this melody performed by the whole orchestra in octaves can be seen in bars 1–4, 6 and 12–19; the alternate version of the melody, performed in upper registers only in conjunction with low held chords, is in bars 4–5, 7–8, 11–12, 16–17 and 24–26. The third aspect, consisting of a three-part 'drama' – poised (violin 1 only), outburst (tutti), brief silence – occurs in bars 8–10 and 20–23. Coming straight after the soft coda of the second movement, these opening two minutes are intended to be an intense, awesome experience, tapping into something of the overwhelming, stormy qualities of cumulonimbus. From figure B, the movement abruptly begins an aurally complex transition toward CM8, along the way manifesting aspects of both CL8 (e.g. oboe 3, bb. 27ff. and violin 2, bb. 35–37) and CM9 (which draws randomly on the other eight CM behaviours). The music stabilises around figure C, settling into several bands of repeating muted brass pitches with abrupt changes of dynamic, plus occurrences of downward violin 'streaks' (violin 1, bb. 55–57; violin 2, bb. 57–60).

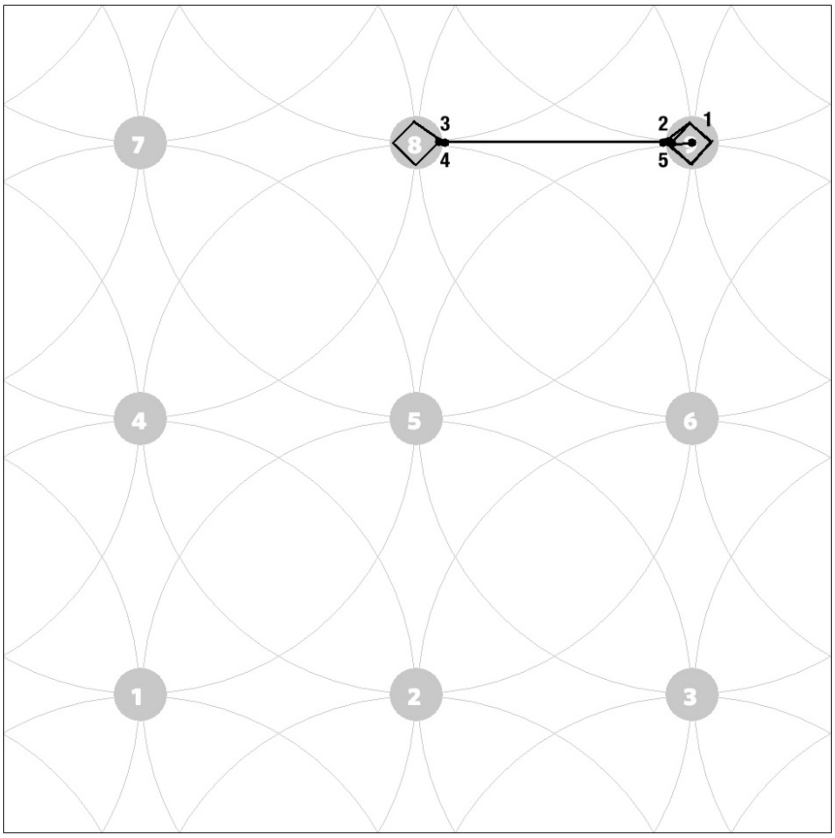

link	behaviour	duration
1	CL9	2'
2	CL9→CM8	1'
3	CM8	1'30"
4	CM8→CH9	45"
5	CH9	1'30"

Fig. 3.13. *Cloud Triptych*, third movement: trajectory chain.

Cloud Triptych concludes with a final transition up to CH9, beginning at figure D. This introduces elements from CH8 (e.g. clarinet 1, bb. 91–96 and tam-tams, bb. 89–92), featured prominently in the first movement, and, again, CM9. The overall effect of the transition is gradually to dissipate the brass material and introduce the glockenspiel (b. 77) followed by the upper winds (b. 82), violas (b. 97) and, finally, the return of the wine glasses (b. 104). The combination of glockenspiel and wine glasses thereby becomes a unifying element in the piece, occurring in all three movements. The counterpoint produced by CH9, maintaining a heightened atmosphere, is allowed to play out for a minute and a half, thinning slightly at the very end, bringing *Cloud Triptych* to a close.

Evaluation

My relationship with the music generated by CloudCube has been interesting and unexpected. I had been concerned beforehand that the sheer complexity of the system would inhibit or even preclude my intuitive involvement. This was, of course, one of the very reasons I had devised it the way I did, as a way of undermining my controlling instincts. Yet while the raw output from CloudCube and the material content of *Cloud Triptych* are almost identical, numerous concerns and considerations, both musical and technical, affected the work's eventual outcome.

An important problem, not apparent until after much of the material for the opening movement (which was composed first) had been notated, was that CloudCube's time-keeping calculations were faulty. As a result, the material throughout the first movement is not strictly accurate in terms of the underlying processes within each behaviour.[77] However, rather than generating fresh material and beginning again I decided it would be more interesting to work with what CloudCube had given me, and allow the errors to become an integral aspect of the movement. Broadly speaking, it does not appear to have had a significant impact on the behavioural trajectories of the music. In the second movement, I had concerns over the material generated by CM6, which forms the lengthy centrepiece. A quick check revealed an error in the program sufficiently significant that, even though by this stage this movement had also been notated, I generated additional material which was then used to replace part of the faulty original. Even after this, the way I had set up CloudCube to report about dynamics meant it was difficult to interpret whether there were errors present or whether I was simply misinterpreting certain results. I therefore ascertained where the correct points of crescendo and diminuendo should take place and imposed these to obviate any ambiguity. CL7 seems to have experienced hiccups too, as its staccato brass notes (third movement, bb. 155ff.) should in theory be vertically aligned, but I allowed these notes to stand as it was impossible to establish what was causing the misalignment, and in any case I very much liked the 'wrong' result. These situations have been invaluable in assessing both my compositional method as it presently stands and its realisation within CloudCube.

Most of the amendments made, though, were to do with musicality and practicality. Early on it became clear that there was no need for the third percussion part which had originally been provided for. It was easily assimilated into the other two parts, which in turn had their material moved around somewhat to accommodate better movement between instruments and improve balance between the players. In the second movement, from figure M, the violins and cellos were divided such that half the players were tacet, but I realised it would be altogether more effective to redistribute the material between the players at each desk. In the final movement some woodwind material was swapped to better suit the instruments.

Having designed twenty-seven unique behaviours, I felt it was important that *Cloud Triptych* encounter as many of them as possible. To this end, I wrote a short program that could evaluate all the behavioural decisions within trajectory data files, and used this to help design the three movements. Unfortunately, the minimal use of some of the low altitude nodes (only appearing

[77] The accuracy decreases following each tempo change, as CloudCube was not at this point correctly taking account of these changes.

incidentally in transitions) meant that handbells, used in their instrumentations, were never used.[78] However, as it stands, all of the behaviours are in some capacity represented at some point within *Cloud Triptych* – with one unfortunate and accidental exception: CL1, which clearly deserves to have a composition all to itself in the future! Despite these concerns and the various minor nips and tucks made to CloudCube's generated material, my attitude towards both confirmed and suspected computational errors was to regard them as consonant with the behaviour of cloud formations generally, which are constantly impinging upon and corrupting each other. Having satisfied myself as far as possible with the correct functionality of the system, I felt it was important to have confidence in CloudCube's results rather than questioning their veracity.

Overall, taken as a compositional, computational and ultimately musical whole, I believe *Cloud Triptych* to be a convincing example of the way discrete behaviours can be used as the basis for large-scale musical structures and transformations. It is true that the behaviours' intrinsic nature and also the way transitions take place between them are highly complex and subtle, far more so than in the preliminary works. But this delicate blurring of where stability ends and states of flux begin makes for a fascinating engagement with the music at both superficial (micro) and deeper (macro) levels of listening. It also mirrors the work's inspirational origins, resembling the way in which cloud formations move seamlessly between concrete and amorphous behaviours and shapes.

[78] It is clear from CloudCube's results that handbells were selected to be used for part of the second movement, but the way the behaviours were interacting at this point (particularly the dominance of CL5) meant that no material was generated for them.

Chapter 4. Conclusions

Achievements

The primary objective in this research has been the composition of perceptible, large-scale musical transformations. Outlined in Chapter 1, my composition method has sought to extend and explore the possibilities, consequences and implications of a variety of algorithmic techniques, using stochastic processes of movement between discrete musical behaviours. Embodying collections of rules determining all aspects of the music, these behaviours are used as probabilistic epicentres of behavioural stability, between which transitional episodes take place. The behaviours and the algorithmic processes governing movement between them are codified in bespoke computer programs.

In the preliminary works, discussed in Chapter 2, I explored this method in a variety of ways. In *'unredeemed' self-)portrait (in the form of an eagle*, changes in behaviour are abrupt and immediate, akin to phase transitions. *Blesi (Five Transitions of the Soul)* introduced transformations between pairs of polarised behaviours, while *HELP/ME: the soul-machine of the cosmology of grief* examined internalised processes that either intensify or cycle potentially endlessly. All three of these works used computationally generated material in conjunction with intuitively-composed music. The string sextet *Intense quick dream of sentimental groups with people of all possible characters amidst all possible appearances* crystallised the model for my subsequent research, establishing transitions between multiple behaviours as the basis for both the micro- and macro-structural levels of the composition. Much more importantly, this piece proved that a computational model could create an entire composition requiring minimal editing. *Four Seasons* consolidated this use of computational models, incorporating new elements such as non-stochastic, self-similar processes and asynchronous ensembles.

The culmination of this project was the large orchestral piece *Cloud Triptych*, discussed in Chapter 3 which, from both a compositional perspective and computational perspective forms an ambitious synthesis of the preceding research. Extra-musical elements were conclusively jettisoned in favour of a more abstract outlook, in tandem with a questioning of conventional notions of structure, narrative, coherence and above all, compositional control. Drawing on the organisational arrangement and morphological fluidity of cloud formations as a paradigm for my processes, I designed a collection of twenty-seven unique musical behaviours which formed the basis for a new computation system called CloudCube, fixing the behaviours' epicentres as the vertices of a three-dimensional metacube. I created a timeline by designing trajectories moving through this space, and the resultant transitions were upgraded so that, for the first time, aspects from multiple behaviours were combined, producing more complex musical transformations. The CloudCube system was used to create the three movements of *Cloud Triptych*, a work of twenty-eight minutes' duration.

The main achievements of this project chiefly reside in two key areas, pertaining to method and realisation. Compositionally, using stochastic movement between discrete musical behaviours as the basis for both material generation and structural design is a compelling, imaginative and persuasive method. It affords me an unusually rich blend of freedom and constraint. Since it does not of itself guarantee effective results – or even interesting music – the onus is entirely on myself both to define distinctive musical behaviours and then to design engaging trajectories between them. At the same time the complexity of the system forces me to accept that not everything can be controlled or even fully understood. By this method, behavioural certainty and uncertainty are elevated to a compositional device comparable to (and working in tandem with) tension and release, consonance and dissonance, activity and repose and the like, becoming one of my fundamental creative considerations.

The realisation of this method within CloudCube provides a valuable and powerful compositional tool. Once its behaviours have been defined, the scope is virtually unlimited; CloudCube can create music of any length, generating the material quickly.[79] Furthermore, it enables me, through an engaging dialogue with the computer, to explore my musical ideas in a multitude of different ways and from alternate perspectives, maximising their creative potential.

Evaluation

I have mixed feelings about the distancing effect brought about by the complexity of the system embodied within CloudCube, due to the way it deliberately renders impossible the ability fully to understand exactly which decisions led to specific outcomes. This has meant that, although the behaviours and the combinatoriality procedures functioned correctly when tested in isolation, it has been difficult to ascertain whether this is always true when actively generating material. Furthermore, in circumstances where it seems perfectly evident that the generated results are not entirely as they should be, it has sometimes been impossible to understand why. These difficulties are in part due to the time constraints of the project. The situation is hampered further by the (relatively few) areas where CloudCube leaves scope for interpretation, such as gradual changes in dynamics. In these instances, the lack of complete specificity in CloudCube's results could – and, no doubt, did – lead to misinterpretations of computational outcomes, potentially undermining to

[79] The time taken for CloudCube to generate the full orchestral material was typically one-tenth the duration of the composition. In tests, CloudCube was able to create compositions of at least an hour's duration without problems.

some extent the integrity of certain musical behaviours. Yet the inability to disentangle the facts in scenarios like this was fundamental not only to the way in which I intended CloudCube to keep my urge to control everything in check but also, just as importantly, to the principle that my compositional method is a means to an end rather than an end in itself. Whatever else may be uncertain, though, the perceptibility of the processes in the resultant music, shifting in and out of behavioural focus, is nonetheless strikingly clear. The concerns I had had that my compositional method ran the risk of over-complexity – and of it becoming a kind of stochastic equivalent of integral serialism – were assuaged by the vivid cogency of its results.

An interesting side effect of this systemic complexity is the way it intensifies the dialogue between composer and computer. In the preliminary works, mainly due to their relatively simple algorithmic systems, the musical outcomes, although stochastically varied, were nonetheless predictable. The dialogue took place almost entirely during the preparatory and design stages, and there was little additional work required to be done to the generated material since it resembled so exactly what had been envisaged. Working with CloudCube, this is no longer the case. Due to the extent of the system's inherent unknowns – above all the way in which different behaviours will impact on each other – once material has been generated I must then spend time becoming acquainted with it, before any further decisions can be made. Thus, the dialogue continues beyond the generation process, making interaction between composer and computer more stimulating and fulfilling, and raising the machine above a mere tool to something akin to a collaborator.

As far as the CloudCube software itself is concerned, use of the BBC BASIC for Windows programming language turned out to be something of a mixed blessing. My decision to adopt BASIC was due to existing familiarity and having already used it to create a number of compositions.[80] While I can hardly fault it in terms of speed or capabilities – it was perfectly adequate when composing the preliminary works, and did not inhibit any aspects of *Cloud Triptych*'s compositional development or implementation – the necessity to create everything from scratch, plus the fact that the language is something of a niche interest (nowadays without a widespread following or support), arguably made the project more time-consuming and challenging than if a more contemporary language had been adopted, such as Processing, Python or Oz. Even something as trivial as introducing the option of Bezier curves for the design of trajectories involved my having to learn De Casteljau's algorithm and code it myself. Some other languages either have such functions built in already or available via online communities. Nonetheless, the necessity to be involved in every aspect of CloudCube's functionality was highly

[80] *gravest one* for solo cello (2002) is the earliest example in my work of using BASIC to encode algorithmic processes.

beneficial when problems arose, and in any case the bespoke nature of the program meant that any use of existing code would have been minimal.

Future Directions

The extent to which my compositional method has been embodied within CloudCube means that I now tend to think of them as synonymous, and I therefore regard future developments as being similarly mutually inclusive. As already stated, CloudCube in its present form is a program with considerable scope. The behaviours can be whatever the composer wishes them to be, and the use of a 3×3×3 metacube, although relevant to this project, is otherwise an arbitrary choice that can easily be generalised to an $x \times y \times z$ metaspace. The relative positions of the behaviours and the extents of their probabilistic orbits and cores can similarly be modified in any way. Furthermore, importantly, CloudCube is not only useful in situations involving transitions. All music, in a sense, 'behaves' a certain way, and could therefore hypothetically be designed within CloudCube. Indeed, it is conceivable that an entire composition could be generated using just a single, mutable behaviour with multifaceted modes of operation (cf. Table 3.1, p. 49: CL9).

There are, nonetheless, numerous areas where further exploration and refinement is desirable and necessary. Most obvious is the way CloudCube generates output, in the form of screen data that must then be manually notated. As my use of computation has grown in scale and complexity, approaching ever closer to the point where the output requires no manual editing, this method has become increasingly inefficient and onerous. It created a major bottleneck in the composition process of *Cloud Triptych*, requiring over a month devoted to notation. A priority for future development is to enable the output to be in a data form that can be imported directly into notation software. Using an intermediary format such as MusicXML[81] or Lilypond[82] to achieve this would seem to be a promising way forward.

The current approach of generating each instrumental part individually, one at a time, may offer less scope than generating all parts in parallel. The main benefit of serial generation (outlined on p. 16) is that it directly utilises the variety of the stochastic process, employing a kind of 'micropolyphony' where players' material is individually insignificant yet vital to collective behavioural cohesion. On the other hand, an obvious benefit of parallel generation (which I have not yet fully explored) is that it would enable deliberate, in addition to merely incidental, forms of counterpoint. Torsten Anders points out, with reference to examples of harmonic- and counterpoint-focused constraint satisfaction problems (CSPs), that "[a] suitable variable [i.e.

[81] cf. Good 2001, Cunningham 2004 and Good 2006.
[82] cf. Nienhuys and Nieuwenhuizen 2003.

modular] ordering is highly problem-dependent. A contrapuntal CSP [... with] two or more simultaneous voices that accompany each other, but which are rhythmically and melodically relatively independent [...] is often solved best by processing from left-to-right in the score and completing all voices more-or-less in parallel."[83] Anders' considerations are here primarily directed toward the implementation of existing paradigms of musical grammar and not to the composition of new idioms, though his argument is still valid. It would be relatively straightforward to implement this within CloudCube if the instruments moved together in terms of their structural durations. CloudCube already generates material for divided string sections in parallel, so this could be extended to every instrument. It would involve an adjustment in the relationship between players and the modular process; Table 4.1 illustrates a comparison of serial and parallel generation in relation to the modules.

Table 4.1. Comparison of serial and parallel generation modular processes.

||: each instrument :||

1 *instrumentation*	2 *structural duration*	3 *quantisation*	4 *rest probability*	5 *rhythm*	6 *pitch*	7 *articulation*	8 *dynamic*

serial (each player makes their own way along the timeline)

||: all instruments :||

1 *instrumentation*	2 *structural duration*	3 *quantisation*	4 *rest probability*	5 *rhythm*	6 *pitch*	7 *articulation*	8 *dynamic*

parallel (all players move together along the timeline)

Generating material in parallel while allowing it the same freedom as now is somewhat anomalous within the present system, insofar as a player's material is not dependent on that of other players but on the behavioural rules applying to everybody. Behavioural definitions could, though, involve specific reference to other parts, and encompass prescribed (and proscribed) forms of interaction. It is also conceivable that CloudCube could be extended to enable *both* structural durational liberty (as now) and unity (the parallel approach shown in Table 4.1), though this would likely involve predetermining 'gathering points' when switching between them, which CloudCube would need to bear in mind when calculating structural durations. This is not dissimilar to certain aspects of how rhythm is presently determined within CloudCube (cf. p. 55), calculating durations to fill a predefined structural duration. This line of enquiry has potentially interesting implications, as it would establish an additional link between micro (rhythms) and macro (structure).

[83] Anders 2011, p. 12.

At both these levels, it is worth considering a subtle but significant issue that arises from the way trajectories through the metacube are used to calculate behavioural decisions. Cursor movement along the trajectory involves jumping between time positions according to the length of each structural duration. The start of each structural duration is rounded to the nearest trajectory time-step in order to ensure temporal cohesion among the players. The behavioural decisions associated with this time-step – calculated and stored when the trajectory was created – are then used throughout the ensuing structural duration. As a result, each structural duration is an extension of the behavioural decisions established at its start. The fact that the material contained within the structural duration, moving through time, effectively therefore continues also to move along the trajectory – changing its position relative to the surrounding behaviours and thereby potentially leading to quite different decisions – is entirely ignored. This is problematic, preventing transitions from having complete smoothness and fidelity to the behavioural framework. Resolving this would likely involve a shift in focus in the application of behavioural decisions, from the present level of structural durations to the level of individual notes.[84] This would likely make structural durations irrelevant, and CloudCube's modular arrangement would therefore require revision, but I suspect it would be a valuable improvement to the current method. Instead of the present synchronic approach, using the decisions of the time-step closest to the start of each structural duration, CloudCube could instead take a diachronic approach, calculating decisions in real-time (as opposed to when the trajectory is created) at each rhythmic position, as shown in Fig. 4.1.

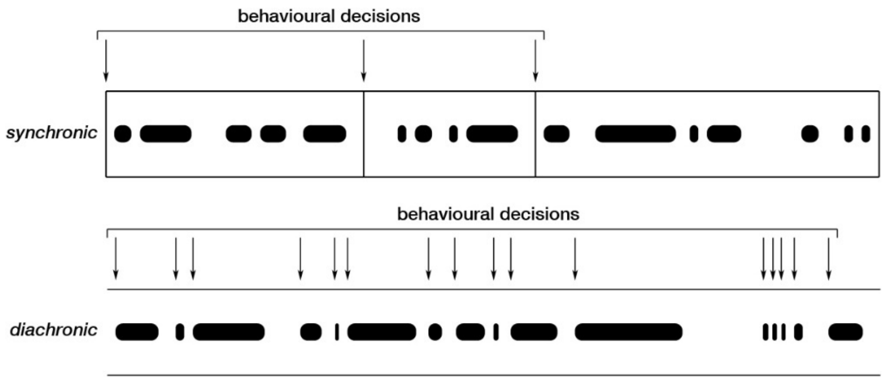

Fig. 4.1. Synchronic and diachronic approaches to behavioural decisions.

These real-time calculations would impact on CloudCube's overall processing speed, but may well prove more efficient than the current approach.[85] This would also make it theoretically possible to

[84] This is essentially a more advanced version of the approach taken in *Intense quick dream*, where most of the material is determined on a note-by-note basis.

[85] This would also necessitate saving the calculated data at each position, otherwise there would be no record of CloudCube's calculations.

introduce a means to reject CloudCube's suggested material also in real-time and generate alternatives, deepening further the level of dialogue with the computer.

Beyond these functional improvements, by far the most valuable future direction for this compositional method would be to develop it, and CloudCube, into a user-friendly computer application that any composer could use, without requiring a detailed understanding of coding. As already indicated, CloudCube has the capacity theoretically to model any kind of musical behaviour. It therefore holds the potential to be developed into something akin to what Torsten Anders has ambitiously called a *most generic* system, which

> would allow its user to implement any music theory model conceivable. [...] such a system is an ideal system [... because] composers usually prefer to make compositional decisions themselves; only with reservations do they accept a composition system which restricts their compositional freedom. A most generic [...] system would not restrict its users...[86]

It is my hope and my intention that the compositional method demonstrated in this research will be developed to this end, in order fully to unleash its creative power and potential.

[86] Anders 2007, p. 4.

Appendix. Selected list of works composed during the Ph.D.

work	year	instrumentation	duration
'unredeemed' self-)portrait (in the form of an eagle	**2008**	**solo flute**	**5'**
… but have not love …	2008	13 percussionists	10'
Extra Abbatia	2009	piano quartet	2'
Chorus angelorum te suscipiat	2009	orchestra	6'
Triptych, May/July 2009	2009	2-channel audio fixed media	24'
Blesi (Five Transitions of the Soul)	**2009**	**6 players**	**6'**
Memories are made of this	2009	2-channel audio fixed media	35'
becauseshewas (veteris vestigia flammæ)	2009	2-channel audio fixed media	19'
Intense quick dream of sentimental groups with people of all possible characters amidst all possible appearances	**2010**	**string sextet**	**6'**
the Ceiling stared at me but i beheld only the Stars	2010	2-channel audio fixed media	35'
Simulated Music	2010	2-channel audio fixed media	58'
HELP/ME: the soul-machine of the cosmology of grief	**2011**	**16 players**	**6'**
Mass in Purple	2011	SATB choir	15'
nolite facere dicunt enim	2012	12 voices	5'
Night Liminal	2012	2-channel audio fixed media	40'
the octave of the grief of the clone that leapt to the remainder of night sky	2012	voice & 5 players	5'
Dither - Pother - Roil	2012	2-channel audio fixed media	49'
Studies 1–22	2014–16	1-/2-/3-/8-channel audio fixed media	115'
Four Seasons	**2015–16**	**flute choir**	**30'**
Cloud Triptych	**2016**	**orchestra**	**28'**
Symptoms	2016	solo cello	10'

The works listed in bold are included in the portfolio.

Bibliography

Books, Journals and Articles

Ames, C. (1987) Automated Composition in Retrospect 1956–1986. *Leonardo*, 20(2), pp. 169–185.

Ames, C. (1990) Statistics and Compositional Balance. *Perspectives of New Music*, 28(1), pp. 80–111.

Ames, C. (1992) A Catalog of Sequence Generators: Accounting for Proximity, Pattern, Exclusion, Balance and/or Randomness. *Leonardo Music Journal*, 2(1), pp. 55–72.

Anders, T. (2007), *Composing Music by Composing Rules: Design and Usage of a Generic Music Constraint System*. Ph.D. Thesis, Queen's University Belfast.

Anders, T. and Miranda, E. (2009) Interfacing Manual and Machine Composition. *Contemporary Music Review*, 28(2), pp. 133–147.

Anders, T. (2011) Variable Orderings for Solving Musical Constraint Satisfaction Problems. In: G. Assayag, ed. *Constraint programming in music*. Columbia: ISTE, pp. 25–54.

Alsop, R. (1999) Exploring the Self Through Algorithmic Composition. *Leonardo Music Journal*, 9, pp. 89–94.

Bencina, R. (2005) The Metasurface – Applying Natural Neighbour Interpolation to Two-to-Many Mapping. *Proceedings of the 2005 International Conference on New Interfaces for Musical Expression*, pp. 101–104.

Berg, P. (2009) Composing Sound Structures with Rules. *Contemporary Music Review*, 28(1), pp. 75–87.

Bischoff, J. (1991) Software as Sculpture: Creating Music from the Ground Up. *Leonardo Music Journal* 1(1), pp. 37–40.

Bresson, J., Agon, C. and Assayag, G. (2011) OpenMusic – Visual Programming Environment for Music Composition, Analysis and Research. *Proceedings of the 19th ACM international conference on Multimedia (MM'11)*, pp. 743–746.

Brooks, S. & Ross, B. (1996) Automated Composition from Computer Models of Biological Behaviour. *Leonardo Music Journal*, 6, pp. 27–31.

Collins, N. and Brown, A. R. (2009) Generative Music Editorial, *Contemporary Music Review*, 28(1), pp. 1–4.

Cummings, E. E. and Firmage, G. J. (1994) *Complete Poems 1904–1962*. New York: Liveright.

Cummings, S. (2014) *Change, chance, character: the use and development of probabilistic 'singularities'*. [pdf] Available at: http://simoncummings.com/writings [Accessed 31 October 2016].

Cummings, S. (2016) *Towards an embodiment of compositional practice: the ongoing evolution of CloudCube*. [pdf] Available at: https://csmc2016.wordpress.com/proceedings/ [Accessed 31 October 2016].

Cunningham, S. (2004) *Suitability of MusicXML as a format for computer music notation and interchange*. [pdf] Available at: http://www.glyndwr.ac.uk/cunninghams/research/SuitabilityMusicXML.pdf [Accessed 1 November 2016].

Edwards, M. (2011) Algorithmic Composition: Computational Thinking in Music. *Communications of the ACM*, 54(7), pp. 58–67.

Good, M. (2001) *MusicXML: An internet-friendly format for sheet music*. [pdf] Available at: http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.118.5431&rep=rep1&type=pdf [Accessed 1 November 2016].

Good, M. (2006) *Lessons from the Adoption of MusicXML as an Interchange Standard*. [pdf] Available at: http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.138.1974&rep=rep1&type=pdf [Accessed 1 November 2016].

Hamblyn, R. (2003) *The Cloud Book*. David & Charles.

Harley, J. (1995) Generative Processes in Algorithmic Composition: Chaos and Music. *Leonardo*, 28(3), pp. 221–224.

Harley, J. (2009) Computational Approaches to Composition of Notated Instrumental Music: Xenakis and Other Pioneers. In: R. Dean, ed. *The Oxford Handbook of Computer Music*. Oxford University Press, pp. 109–132.

Hayden, S. (2008) Computer-assisted composition and aesthetic innovation. *New Notes*, June, pp. 6–7.

Hindemith P. (1945) *The Craft of Musical Composition*. Translated by Arthur Mendel. London: Schott.

Jones, P. W. and Smith, P. (2001) *Stochastic Processes: An Introduction*. London: Arnold.

Kanach, S., ed. (2012), *Xenakis Matters: Contexts, Processes, Applications*. New York: Pendragon Press.

Koenig, G. M. (1967). *The use of computer programmes in creating music*. [pdf] Available at: http://www.koenigproject.nl/texte.htm [Accessed 16 January 2016].

Koenig, G. M. (1968a) *Remarks on Composition Theory*. [pdf] Available at: http://www.koenigproject.nl/texte.htm [Accessed 16 January 2016].

Koenig, G. M. (1968b) *Programmed Music – From the Composer's Viewpoint*. [pdf] Available at: http://www.koenigproject.nl/texte.htm [Accessed 16 January 2016].

Koenig, G. M. (1970) *The use of computer programmes in creating music*. [pdf] Available at: http://www.koenigproject.nl/texte.htm [Accessed 16 January 2016].

Koenig, G. M. (1975) *My Experiences with Programmed Music*. [pdf] Available at: http://www.koenigproject.nl/texte.htm [Accessed 16 January 2016].

Le Mée, K. (2003) *The Benedictine Gift to Music*. Paulist Press.

Locher, J. L. and Veldhuysen, W. F. (2000) *The Magic of M. C. Escher*. London: Thames & Hudson.

Marsaglia, G. (2006) *Refutation of claims such as "Pi is less random than we thought"*. [pdf] Available at: http://interstat.statjournals.net/YEAR/2006/articles/0601001.pdf [Accessed 15 September 2016].

Matossian, N. (1986) *Xenakis*. London: Kahn & Averill.

Nienhuys, H. and Nieuwenhuizen, J. (2003) *Lilypond, a System for Automated Music Engraving*. [pdf] Available at: http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.100.6160&rep=rep1&type=pdf [Accessed 1 November 2016].

Nyman, M. (1999) *Experimental Music: Cage and Beyond*. 2nd edn. Cambridge University Press.

Reich, S. (2004) *Writings on Music 1965–2000*. Oxford University Press.

Sadgrove, M. (1995) *A Picture of Faith*. Bury St Edmunds: Kevin Mayhew.

Siratori, K. (2008) *This Machine Kills*. [pdf] Available at: http://www.noisemachine.info/ [Accessed 9 August 2014].

Tu, S. and Fischbach, E. (2005) A Study on the Randomness of the Digits of π. *International Journal of Modern Physics C*, 16(2), pp. 281–294.

Schwinger, W. (1989) *Krzysztof Penderecki: his life and work*. London: Schott.

Steinitz, R. (2003) *György Ligeti: Music of the Imagination*. London: Faber & Faber.

Sturrock, J. (1993) *Structuralism*. 2nd edn. London: Fontana Press.

Varga, B. A. (1996) *Conversations with Iannis Xenakis*. London: Faber & Faber.

Varèse, E. and Wen-chung, C. (1966) The Liberation of Sound. *Perspectives of New Music*, 5(1), pp. 11–19.

Whittall, A. (1999) *Musical Composition in the Twentieth Century*. Oxford University Press.

Xenakis, I (1966) The Origins of Stochastic Music. *Tempo*, 78, pp. 9–12.

Xenakis, I. (1992) *Formalized Music: Thought and Mathematics in Music*. New York: Pendragon Press.

Scores

The Benedictines of Solesmes, eds. (1962) *The Liber Usualis*. Tournai: Desclée & Co.

Bell, J., ed. (2005) *Church Hymnary*. 4th edn. Norwich: Canterbury Press.

Cummings, S. (2002) *gravest one*. Self-published. Available at: http://simoncummings.com/op11.php [Accessed 8 November 2016].

Ligeti, G. (1969) *Lontano*. London: Schott.

Ligeti, G. (1971) *Atmosphères*. Vienna: Universal Edition.

Penderecki, K. (1961) *Threnody to the Victims of Hiroshima*. London: Ernst Eulenberg Ltd.

Reich, S. (1980) *Piano Phase*. London: Universal Edition.

Strauss, R. (1979) *Don Quixote*. Dover Publications.

Audio-visual

BBC Wildlife Specials: Eagle. [documentary, DVD] Directed by John Downer. BBC/2 Entertain Video, UK, 2004. 50 mins.

Hafler Trio, The (2004) *scissors cut arrow*. [CD] (Phonometrography pgram005).

Kirschner, K. (2010) July 17, 2010. [flac] Available at: http://www.kennethkirschner.com/kirschner071710.flac (Accessed 4 November 2016).

Ligeti, G. (2002) *The Ligeti Project II*. [CD] (Teldec Classics 8573-88261-2).

Nine Inch Nails (1999) *The Fragile*. [CD] (Nothing Records 490 473-2).

Reich, S. (1974) *Drumming*. [CD] (Deutsche Grammophon 427 428-2).

Reich, S. (1987) *Early Works*. [CD] (Nonesuch Records 979 169-2).