Timbre and Technology:
an analytical partnership

The development of an analytical technique and its
application to music by Lutoslawski and Ligeti

Stephen N. Malloch

Text and Examples

PhD
The University of Edinburgh
1997

EX BIBL. UNIV. EDINB.

Declaration

This thesis has been composed by the candidate, and is entirely the candidate's own work.

Recognition is given to A.Murray Campbell for his work in writing computer programmes used in this thesis.

Stephen N. Malloch

Acknowledgements

I would like to acknowledge the help and encouragement that I have unfailingly received from my two supervisors, Mr Peter Nelson in the Faculty of Music, and Dr. Murray Campbell in the Department of Physics of Edinburgh University. Their constant interest and enthusiasm for this project have helped me enormously at times when the difficulties of the research threatened to outweigh my resourcefulness and perseverance.

I also wish to acknowledge the help that A. Murray Campbell has given to me by writing the computer programmes that I used during this research.

Without the assistance and encouragement of my family, I would never have completed this thesis, and I wish to pay particular tribute to my wife, Megan Thorpe, who has had to cope with the stress of being the wife of a PhD student for longer than she (and I) anticipated.

Dedicated to my wife, Megan Thorpe,
and to the memory of my grandfather,
Noel G. Mather

Contents

Examples are contained in volume 1.

Figures are contained in volume 2.

The pitch convention used in this thesis recognises middle-C as C^4.

Abstract

The discipline of Music Analysis consists of "the resolution of a musical structure into relatively simpler constituent elements, and the investigation of the functions of those elements within that structure." [1] What are the 'constituent elements'? Pitch is the element most often studied when a composition is analysed. Yet pitch is only one element of music. Pitches must be sounded, so an indispensable element of music is its timbre. While analysts are able to discuss pitch structures with a great deal of sophistication, their attempts to discuss timbre are few and often rudimentary.

In this thesis, it is recognised that the discipline of Music Analysis would be enhanced if timbre could be discussed with the same degree of precision as pitch. To this end, a number of acoustic methods for analysing music timbre are proposed. The sound of a music performance is analysed through spectral analysis. This shows us what is physically present in the sound. To begin to understand timbre perception, methods of data weighting and reduction are introduced, based on psychoacoustic models. Six measures of timbre are proposed: timbral width, timbral weight, timbral pitch, roughness, sharpness, and observations based on the results of cepstrum analysis. The measures of timbral width, weight and pitch are obtained from a new technique inspired by the Tristimulus Method of Pollard and Janson [2].

The thesis reviews past attempts, both scientific and musicological, to analyse and structure music timbre. Certain pieces appear to use timbre as their principal means of organisation. The abovementioned measures of timbre are applied to two important pieces of this type: the first movement of Lutoslawski's *Jeux vénitiens* (1961) and Ligeti's *Atmosphères* (1961). These analyses, which are the focus of the thesis, incorporate an exploration of timbre with a thorough investigation of the music's other structural elements. It is in the examination of the play between timbral and other structures that we find a level of insight into these pieces which has not, until now, been possible.

[1] Bent, Ian: (1987). *The New Grove Handbooks in Music: Analysis*. London: Macmillan Press, p.1.

[2] Pollard, H., and Janson, E.V.: (1982). "A Tristimulus Method for the Specification of Musical Timbre." *Acustica* 51: 162-71.

CHAPTER 1

Timbre

CHAPTER 1

Timbre

---

1. Definitions

What is music analysis? Ian Bent describes it as "the resolution of a musical structure into relatively simpler constituent elements, and the investigation of the functions of those elements within that structure."[1] What are the 'constituent elements'? Pitch is the element most often studied when a composition is analysed. Schenkerian analysis and pitch-class set analysis are both methods by which pattern and structure are sought in the realm of pitch. But what of the way a composer manipulates *sound*? Pitches must be sounded, whether in a performance or in a person's imagination, and composers usually make decisions about *how* they would like these pitches to be sounded. Yet, while analysts are able to discuss pitch structures with a great deal of

---

[1] Ian Bent (1987): 1.

sophistication, their attempts to search for structure in the way pitches are sounded are few and often rudimentary. This thesis is concerned with this search for structure.

The manner in which pitched or un-pitched notes are sounded, I will call timbre - a term much used, but often ill-defined. The American Standards Association (1960) states that "timbre is that attribute of auditory sensation in terms of which a listener can judge that two sounds similarly presented and having the same loudness and pitch are dissimilar."[2] The definition has the disadvantage of describing what timbre is not, rather than saying what it actually is. In this study, timbre is defined as ***the primary aural information that is used in the perceptual task of assigning an identity to a sound.***

A sound's perceived identity is as distinct as its pitch, loudness or duration. However, unlike the largely independent natures of pitch, loudness and duration (for example, a sound retains its pitch regardless of changes in its length and dynamic [3]), a sound's identity may be influenced by changes to its pitch, loudness or duration.

To illustrate this new definition of timbre, let us consider the famous bassoon solo at the beginning of Stravinsky's *Rite of Spring*. The bassoon plays at the very top of its range, in a flowing rhapsodic style. This is highly atypical when compared with earlier material written for the bassoon - for example, the bass part in a baroque trio sonata. In the *Rite of Spring,* we might describe the sound with such words as 'nasal,' 'hollow,' or 'strained.' In the trio sonata, we might use words such as 'full,' fat,' or 'mellow.' These words attempt to describe the sound's identity, which is determined by its

---

[2] American Standards Association (1960): 45.

[3] This is true for a typical musical context. In experimental situations, however, a large change of dynamic will occasion a change of perceived pitch.

timbre. The sound's identity is perceptually separate from the other characteristics of sound - pitch, loudness and duration. However, changes in the pitch of the bassoon, and changes in the style of playing (which affect duration and loudness) will affect timbre, which in turn affects identity. Thus, timbre is *not* independent of the other characteristics of sound. Indeed, even a sine wave, swept through the audible range, will have a perceived change in timbre (due to the characteristics of the human auditory system).

In this definition, a sound's identity does not equate with our learned ability to classify a sound as being the result of a particular means of production (for example, "that sound is produced by a violin"). In the above case of the sound of a bassoon, the timbre, and thus the identification, is very different between the two different types of music, yet, through experience, we have learned that both timbres are produced by the one instrument. However, the identification of a sound *does* equate with the primary information that we then use in the classification of its means of production. Returning to the American Standards Association definition, if two sounds are presented where pitch, loudness and duration are held constant, it will be the information that we use in identification (the timbre) that allows us firstly to discriminate between them (produce different identifications), and then secondly to judge the probable means of production for each sound (for example, "I notice the timbre of that note; it sounds nasal, hollow and strained; therefore it is a bassoon playing in its topmost register").

From the preceding argument, a three stage process can be postulated that describes the way in which we perceive timbral information - timbral information leads us to the *identification* of a sound, which leads us (usually) to the *classification* of a sound:

timbre → *identification* → *classification*

The step of *classification* need not always happen. In Pierre Schaeffer's *Traité des objets musicaux* (1968), the *classification* of sound is positively discouraged. Rooted in the *musique concrète* tradition, Schaeffer attempts to create a classification system that will encompass not just typically 'musical' sounds, but all sounds, so that they may all be used 'musically.' For this to happen, the listener *must* divorce the sound from identification with its source - it must become an *objet sonore* (a 'sound object'). An *objet sonore* may then be classified as an objectified sound (we stop at the step of *identification*), rather than as a product of a sound source (which is the step of *classification*).

We have seen that timbral information allows us to identify and classify a sound. But what is this timbral information? It must consist of particular acoustic properties that are subsequently processed by the auditory system.

We know that a tone produced by a musical instrument or voice consists essentially of three parts - starting transient (also called onset or attack transient), steady state, and decay. The starting transient consists of the very beginning of the tone, where there is rapid change during a very short period of time (for example, the length of the starting transient for an oboe note is around 15 milliseconds); the steady state portion is the more or less steady region of the tone, typically lasting between 100ms and some seconds (though sounds produced by a single moment of excitation, such as those heard from the harp, piano and percussion, do not strictly contain steady state regions); and the decay is that portion of the tone that occurs after the player has ceased applying energy to create the sound.

Herman von Helmholtz, in the first major publication on music acoustics (1863), believed that timbre related only to continuous, steady state tones, with constant pitch and loudness. This definition of timbre has held sway for many years, and is still found today. Since the work of Helmholtz, it has been shown that the starting transient is often a vital source of information for a listener's classification of an instrument, its brief duration belying its importance in this respect.[4] If this starting transient is cut from a recording of an instrumental sound, it can often make correct classification very difficult. While certainly not ignoring the acoustic information contained within the starting transient, this study does not attempt to deal with the issue of how the starting transient can be analysed separately from the steady state in order to understand its particular contribution. And although the starting transient is often vital for the *classification* of a sound, it is supposed here that it is less important for the *identification* of a sound. For example, the adjectives used above to identify the sound of a bassoon (full, fat, mellow, or nasal, hollow, strained) appear to apply largely to the steady state, and would still be valid if the starting transients were omitted. This thesis attempts to find ways in which the timbral trajectories of compositions can be discussed, rather than to investigate how the information contained in the starting transients specifically influences our perception. It is supposed that our impression of a piece's timbral trajectory, especially within the context of an ensemble sound, consists largely of the information contained in the steady-state portion of sounds, which, in an ensemble sound, combine to create an ensemble timbre.[5]

The timbre of a sound can be very difficult to define. The words used above to describe different aspects of the timbre of the bassoon are highly subjective, and not at all rigorous. In both the fields of music research and acoustics, much work has been

---

[4] See especially Grey (1975).

[5] For support of this view, see Grey (1978), and the review of his work in Section 3.1 of this Chapter.

done over the past twenty-five years or so in trying to determine what the elements of timbre are, and how timbre may usefully be discussed. This study continues in this tradition. In particular, it attempts to answer the question: "Can we find a plane of concepts and a vocabulary that allow timbre to be discussed and analysed within the discipline of music analysis with the same degree of rigour that has been brought to the sphere of pitch relations ?"

2. Hearing and the Physics of Sound

In order to lay the foundations for the new approach to timbre analysis to be presented here, and in order to aid the reader's appreciation of some of the recent literature concerned with research on timbre, let us now review some basics of the physics of sound, and the functioning of the human auditory system.

Virtually all scientific work on timbre has as its basis the work of Helmholtz, who in turn made use of an idea first put forward in 1822 by the French mathematician J.B.J. Fourier. Fourier found that a complex periodic wave form, that is, a complex vibration that repeats itself exactly after a certain period, may be represented as the sum of a set of sine curves, which are the simplest possible type of periodic motion. These sine waves are related to each other such that the frequency of the first is equal to the fundamental frequency, f, of the complex wave form (that is, it has the same repetition frequency), the second has a frequency of $2f$, the third $3f$, and so on, until all the details of the complex wave form are accounted for. Although a pitched musical sound is generally not exactly periodic, and un-pitched sounds are un-pitched precisely because they are non-periodic, it is still possible to analyse such sounds into their sinusoidal, or *Fourier* components, which can subsequently be displayed in a *spectrum analysis*. The quality of a musical sound depends almost entirely on our perception of its spectrum, which comprises its Fourier components - their frequencies, the way they are spaced (their differences in frequency), the way they change through time (though this issue is not directly addressed in this study), their amplitudes, and the way in which the components react with each other. The ways in which the sinusoids 'fit together' - their relative phases - are of secondary importance in a musical context, and in particular are of secondary importance during the steady-state.[6]

Representing the spectrum of a musical sound is only the very beginning of understanding its timbre, for the sound is processed by our hearing mechanism before we perceive it. So, let us now consider what happens to a sound after it enters our ears, as well as discuss some of the influences that the properties of the hearing mechanism have on this sound.

The human ear may be divided into three sections - outer, middle and inner ears. For an understanding of auditory perception, the function of the inner ear is especially important. Within the cochlea is the basilar membrane. A sound wave sets the basilar membrane in motion. The extent of the basilar membrane's movement changes with frequency - a pure tone (sine wave) of low frequency causes much of the membrane to vibrate, with maximum displacement occurring near one end of the membrane; a pure tone of high frequency causes only a small section to vibrate, with maximum displacement near the opposite end of the membrane. The oscillations of the basilar membrane stimulate hair cells which in turn stimulate neurones to fire in a way which codes the information from the basilar membrane. It is probable that the pitch

---

[6] See Plomp (1976): 88-93

perception of pure tones is due to the brain perceiving both the point of maximum displacement, and the pattern of vibration of the basilar membrane.[7]

When more than one pure tone reaches the basilar membrane at the same time, as is the case with a complex tone, every pure tone will contribute to the membrane's pattern of vibration. According to how close two pure tones are to each other, their vibration patterns will overlap to a greater or lesser degree, and their points of maximum displacement will be distinct or will interfere with each other. If these vibration patterns and points of maximum displacement are less than a certain distance from each other, the brain has difficulty sensing each independently, and it is said that the tones which generate these close maxima lie within the same *critical band.*

The critical bands may be understood as auditory filters. Harvey Fletcher (1940), who first discovered the critical bands, explains how he measured the changing threshold of a sine wave tone (the threshold of a sound is its minimum audibly detectable level) in the presence of a band of noise, centred at the frequency of the sine wave. Up to a point, the threshold of the sine tone increases as the noise band increases in bandwidth, which in turn increases the total noise power. However, there comes a point where increasing the noise bandwidth no longer has a significant effect on the tone's threshold. It is this bandwidth that determines the critical band. Each location on the basilar membrane represents a critical band, or filter, with a different centre frequency, so that each location responds to a limited range of frequencies (example 1.1).

---

[7] For a detailed discussion of the place and temporal theories of pitch perception see Moore (1989): chapter 5.

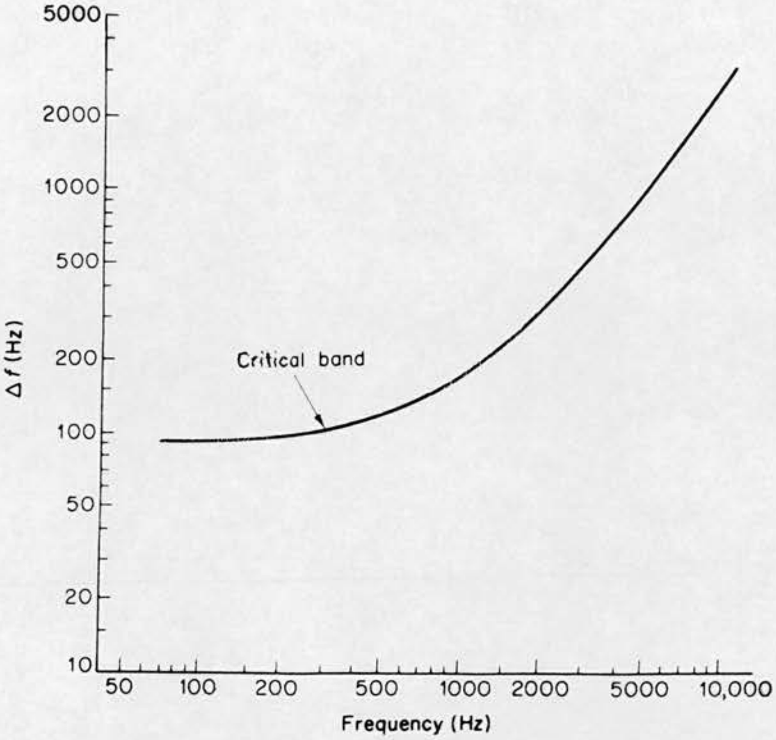

Example 1.1
The width, Δ f, of the critical band. Adapted from Moore (1989): 88; in turn based on Zwicker et al (1957): "Critical Bandwidth in Loudness Summation," *JASA* 29: 548-57.

The concept of the critical band is central to the phenomenon of *masking*, which is "the process by which the threshold of audibility for one sound is raised by the presence of another (masking) sound,"[8] or, put another way, the perceived loudness of a sound may be reduced by the presence of other (simultaneous) sounds. Fletcher's experiment described above is an example of masking - a tone is masked if its frequency and the frequencies of the masking sound fall within the same critical bandwidth. Different degrees of masking (complete masking or a degree of partial masking) take place depending on the extent of envelope overlap along the basilar

---

[8] American Standards Association (1960), quoted by Moore (1989): 84. The masking described here is when the tone and the masker are sounding simultaneously. There is also the phenomenon of non-simultaneous masking, both forwards and backwards, though the process is not well understood. See Moore (1989): chap. 3, part 8.

membrane. Because of the asymmetrical shape of the sound wave along the basilar membrane, lower tones are more likely to mask higher ones, than vice versa.

How is a complex tone perceived? The basilar membrane, upon being excited by a complex waveform, registers the presence of sinusoidal components; in other words, a partial Fourier analysis occurs. The lower components each lie in separate critical bands, and thus are registered as separate pitches. The higher components, however, are separated by less than a critical band, so their resolution is blurred. Why, then, do we hear a single pitch for a complex tone, rather than many different pitches ?

It is, in fact, possible to 'hear out' some of the harmonics in a complex, periodic tone, but it requires practice and concentration. In general, people cannot distinguish harmonics above the eighth, as this and higher harmonics are separated by less than a critical bandwidth. But this still does not explain why, under most circumstances, a complex periodic tone is perceived to have a single definite pitch.

It has been shown that, when presented with a harmonic series in the form of a complex periodic tone, the brain assigns it a single pitch equal to the highest fundamental that fits the harmonic pattern. Thus, if the brain registers the presence of an harmonic series f, $2f...7f$, it recognises that a pitch equal to the fundamental, f, provides the 'best fit'. Pitch recognition occurs when some, or even many, harmonics are missing, and, importantly, still occurs when the fundamental itself is missing. For example, a bassoon playing the note E^3 generates the frequencies 329.6 Hz, 494.4Hz and 659.2Hz. The brain recognises these frequencies as $2f$, $3f$ and $4f$ of a physically non-existent fundamental of 164.8Hz, and thus we hear a pitch equivalent to this frequency - E^3. The brain assigns the *highest* possible fundamental, so the bassoon harmonics are not recognised as $4f$, $6f$ and $8f$ of a non-existent fundamental of 82.4Hz

- E^2. This close octave equivalence may account for a listener's octave confusion in some musical contexts.[9]

The perception of the overall loudness of a complex tone and the relative loudness of its components are particularly influenced by two aspects of auditory perception - masking, and the way loudness perception changes with frequency.

Through perception tests, graphs have been drawn that map the different sound pressure levels (SPL) necessary for tones of various frequencies to be heard at the same loudness. These graphs consist of *equal loudness contours*, and the unit of measurement is the 'phon' (though loudness may also be measured in units called 'sones').[10] As may be seen from the graph below (example 1.2), for two pure tones of 55Hz (A^1) and 3520Hz (A^7) both to have a loudness of 60 phons (roughly equal to the musical dynamic *mp*), there would be a SPL difference between them of about 20dB.

---

[9] The explanation for pitch perception presented here is a greatly simplified version of the pattern recognition model (see E. Terhardt (1974): "Pitch, Consonance and Harmony". *JASA* 55, 1061-1069.

[10] Equal Loudness Contours are produced by asking subjects to vary the loudness of pure tones of various frequencies to match that of a fixed tone of 1000Hz (a sharp B^5).

Example 1.2

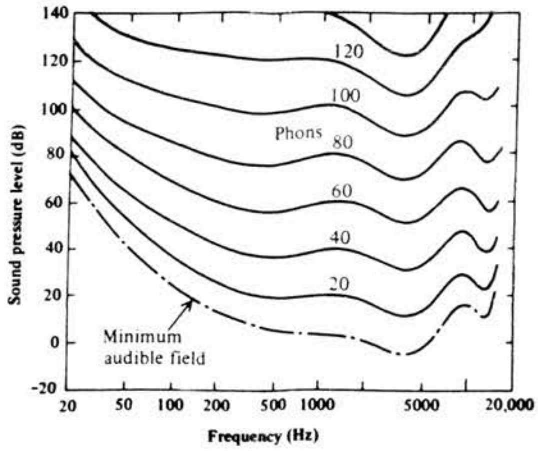

Equal-loudness contours. From Moore (1989): 53, in turn redrawn from Robinson & Dadson (1956) "A Re-determination of the Equal-Loudness Relations for Pure Tones." *British J. of Applied Physics* 7: 166-81.

In addition to considering the equal loudness contours when attempting to represent the perception of a complex tone, the role of the critical bands must be considered. Varying degrees of masking occur between the components of a single complex tone, and also between the components of any other complex tones that sound at the same time. Thus, the loudness of a group of pure tones is greater if they are spread over several critical bands rather than contained within the one critical band. Complex mathematical models have been created by which the loudness of any sound may be calculated.[11] These models, one of which will be described in the next Chapter, take into account the affects of the critical bands, masking, and the equal loudness contours.

In addition to masking, there are the phenomena of *auditory adaptation* and *fatigue*, which, like masking, act to reduce the potential loudness of a sound. According to Scharf (1981), "people differ widely with respect to the degree of adaptation they experience. While most people hear the loudness of a high-frequency, low-level tone

---

[11] See Moore (1989): chapter 4, part 4, and Campbell and Greated (1987): 109-130.

decline by at least half within one minute, some others report no changes in loudness and still others report that the tone disappears."[12] The extent of auditory fatigue, measured by the change in threshold of a tone after exposure to a fatiguing stimulus, increases with the increasing intensity, frequency, and duration of the stimulus, and decreases as the time interval lengthens after the cessation of the stimulus. Neither of these effects is very well understood, and a greater understanding of both phenomena requires further research.[13]

As well as the auditory system influencing the loudness of components of a sound, 'additional' factors are sometimes introduced when two pure tones are perceived simultaneously. These additional ingredients, which will not be observed in any graph of a sound's spectrum, take the form of 'roughness' and added tones.

When two pure tones approach each other and cross into the one critical band a degree of roughness will be perceived, which increases to a maximum at a distance of about a quarter of a critical bandwidth.[14] As the tones become closer, the roughness clarifies itself into separate beats, and two distinct pitches can no longer be perceived - rather a single intermediate pitch is heard. These effects are due to the interaction of the two waves along the basilar membrane. Roughness is an attribute of single complex tones (their upper partials falling within a single critical bandwidth) and complex tones sounding simultaneously (where roughness is closely associated with the musical concepts of consonance and dissonance).[15] Roughness will be discussed

---

[12] Quoted by Moore (1989): 74.

[13] See Moore (1989): chapter 2, part 7.

[14] For a graph of roughness graphed against the distance of two pure tones, see Plomp (1976): 69, as well as this thesis, example 2.2. Also see Campbell & Greated (1987): 59-61.

[15] See Terhardt (1974), Plomp (1976) (esp. p.57-74), and Plomp and Levelt (1965).

further in the following Chapter, and a mathematical model for its calculation will be introduced.

Other added tones arise due to distortion within the ear. These extra tones are called *combination tones*, and are usually heard when two pure tones are separated by more than a critical bandwidth - thus no beating occurs. A combination tone is "a perceived tone not present in the sensations produced by the constituent components of the stimulus when they are presented singly."[16] There are two types of musically important combination tones - the simple difference tone, heard for relatively loud sounds (its frequency equal to the subtraction of the frequency of the lower tone, f_1, from that of the higher tone, f_2: that is, $f_2 - f_1$), and the cubic difference tone, heard for tones of moderate loudness level ($2f_1 - f_2$).[17]

Vibrato and *chorus effect* are another two phenomena which are not directly apparent in a graph of a sound's spectrum. *Vibrato* will be seen as a frequency fluctuation, which indeed it is, though it is usually experienced as a 'warming' of the sound rather than as a succession of pitch changes. Research has shown that our detection of small pitch changes is most sensitive for abrupt changes in a loud tone, and least sensitive for steady fluctuations in a quiet tone. Small pitch fluctuations, occurring faster than around 10 times a second, are usually averaged by the brain to a single pitch with added 'warmth.'

While vibrato typically consists of pitch changes occurring at a rate of around 6 times a second, this still gives a similar perceptual effect when it occurs at small amplitude.[18] The *chorus effect* occurs when a number of instruments or voices

---

[16] Moore (1989): 176.

[17] See Campbell & Greated (1987): 62-67; and Plomp (1976): Chapter 2.

[18] See Campbell & Greated (1987): 95, 162-63.

produce the same perceived pitch, but whose individual frequencies do not exactly match. Perceptually, the different beating rates between the slightly different frequencies produce a characteristic 'silvery shimmer.'

In section 2 of this Chapter, some of the physical components of a musical sound have been discussed, as well as some of the affects that the human auditory system have on these components. The physical components can be represented in a spectrum analysis, while some of the perceptual affects can be estimated through mathematical models based upon extensive perceptual experiments. It is very important to remember that analysis of musical sound consists of these two, distinct tasks - first, measurement of the physical properties of the sound (the results of which can be presented in a spectrum analysis); second, manipulation of this data in order that it might accord more closely with our perception.

Auditory perception theory is a relatively young science, and its findings are often controversial. In discussing the task of the analysis of musical sound, the physicist and acoustician Harold Pollard divides sound analysis into physical analysis, psycho-physical analysis, and feature analysis. Pollard's table of the different types of acoustic analysis is reproduced below (example 1.3).

While there is information in this table that requires further explanation (see Section 3, below), for now it is important to observe that the table stresses the role of data manipulation at the higher levels of analysis. A spectrum analysis of a musical sound is a measure only of some of the physical properties of that sound - as Pollard points out "[the results] can be misleading if attempts are made to assess the response of a listener."[19] If our understanding of the response of the human ear is applied to this

---

[19] Pollard (1988): 234.

data, we begin to move towards comprehension of our perception of this sound - a psycho-physical analysis. In order to move further towards perceptual understanding, we must attempt to discover the level of importance the brain assigns to different sorts of information, and how it groups this information - feature analysis. Feature analysis, arrived at through the weighting and reduction of the data from a physical analysis of a sound (spectral analysis), is vital if we are to begin to close the gap between measurement and perception.

Relationships between the level of analysis, the acoustic mechanisms involved and the type of information that can be extracted

Physical Analysis	Psycho-physical analysis	Feature analysis
Level of analysis:		
Reflects properties of the instrument and its radiation; does not involve assessment by the ear.	Incorporates spectral weighting, masking, characteristic times related to the properties of the ear; data reduction.	Involves computation, comparison, correlation, integration of measures, memory, judgement, etc., leading to the formation of an acoustic 'image'; data reduction.
Mechanisms:		
Musical Instrument →	Inner ear →	Brain
↓	↓	↓
Acoustic output related to the design or testing of a musical instrument.	Physiological processing dependent on the mechanisms in the cochlea region.	Perception and assessment of individual sounds, sequences of notes, musical composition.
Type of information:		
* mean SPL of complete sound * spectrum levels of instrumental resonances or radiated sound at specific times during starting transient, steady state and decay (landscape plot). * fluctuations in spectrum levels and frequencies during steady state.	* mean loudness of complete sound * band loudness spectra at specific times either as 1/3 octave set of growth curves or as a set of Zwicker critical band diagrams (showing masking effects). * loudness derivatives as a function of time (start-times, rise-rates of partial tones, duration of starting transient).	* assessment of starting transient (early sound, synchronism, dominant tones, duration). * sharpness (loudness centroid) as a function of time. * timbre (synthesis of factors such as sharpness, synchronism, early sound or equivalent measures as in the tristimulus method). * pitch * loudness * fluctuations of loudness, pitch, timbre. * roughness (beats between partials in same critical band.) * other cues (e.g.., compactness, missing partials).

Example 1.3. From Pollard (1988): 233.

3. Review of Timbre Research

We have seen that in order to investigate timbre acoustically, it is necessary - indeed, vital - that the 'raw' data attained through spectrum analysis is 'cooked' through the application of the findings of research into auditory perception. The role of the level of analysis that Pollard calls feature analysis is integral to this endeavour. The work of moving towards successful feature analysis is reviewed in the following summary of the significant work in timbre research over the past 25 years in the disciplines of physics and music. A common feature of this research is the desire to arrive at a set of concepts and a vocabulary that will allow timbre to be discussed in a succinct and detailed way.

3.1 Timbre and Physics

The study in music timbre perception by G. von Bismarck (1974a, b) is an early attempt to match perception (as identified through verbal scales) and physical measurement in a highly methodical way.[20] Bismarck takes a spread of different synthesised timbres, equalised for pitch, loudness and duration, and presents these to subjects for verbal categorisation. Thirty pairs, or scales, of opposite timbral attributes were available to the subjects to describe the type of timbre they heard, these thirty selected by the subjects themselves from an even lengthier list of verbal scales.[21] After statistical analysis of the results, Bismarck concludes that the 35

---

[20] For further reading see both Moore (1989) and Campbell & Greated (1987), both of which contain excellent bibliographies. Stephen McAdams (1987) contains a very good survey of the history of music perception, as well as a detailed bibliography.

[21] Some examples of the 30 timbral pairs are: weak-strong; gentle-violent; fine-coarse; reserved-

different timbres used in his experiment can be almost completely described by using only four of the scales - dull-sharp, compact-scattered, full-empty, and colourful-colourless. Bismarck finds that, in order of importance, the factors that carry most of the variance are the scales dull-sharp (sharpness), and compact-scattered (compactness). Compactness he concludes is a measure of a sound on a scale between complex tone and noise - that is, the difference between discrete and continuous spectra - though he proposes no actual scale. Sharpness he shows to be determined by the frequency location of the overall energy concentration of the spectrum as determined by loudness measurements - the higher in frequency the loudness centroid, the greater the sharpness. Bismarck concludes that sharpness is a perceptual attribute of steady-state sounds separate from pitch and loudness. Sharpness remains today an important measurable timbral attribute, and its use in the approach to timbre analysis developed in this study will be presented in Chapter 2.

Work that marks a significant advancement in the study of timbre is the dissertation (1975) and subsequent published articles by John Grey.[22] He was one of the first to attempt to understand the role of the starting transient in the perception of musical sounds. At the beginning of his dissertation he observes that:

> timbre research has made few advances beyond the classical theory of Helmholtz, and the most widespread definition of timbre has it that it essentially consists of the steady state spectrum of a tone. This ignores the temporal information which has been so strongly implicated in the identification of timbre.[23]

---

obtrusive; low-high; dull-sharp; dark-bright; relaxed-tense; calm-restless; smooth-rough; rounded-angular; clean-dirty; simple-complex.

[22] Grey (1975). Subsequent articles are: Grey (1977), (1978); with Moorer (1977); with Gordon (1978).

[23] Grey (1975): 24-25.

Through the use of computer based analysis/synthesis techniques, Grey explores the extent to which the complexity of an instrumental sound may be reduced before our perception of it begins to change.

Grey also searches for "a theory for the salient [timbral] features of classes of sounds."[24] To do this, he makes use of multidimensional scaling techniques (MDS) for processing the data from experiments in similarity perception. The method does not describe "how a subject perceives a stimulus as an individual entity; rather it is used to illustrate the *relationships* among a group of stimuli... These are psychological relationships, and may or may not correspond to physical relationships."[25] In representing the results of MDS, the aim is to construct a configuration so that the distances between the representations of the stimuli are proportional to "how different the stimuli appear psychologically. The number of dimensions in this configuration is ideally the same as the number of aspects of similarity."[26] An attempt may then be made to interpret the dimensions as particular physical attributes of the stimuli.

Grey's three dimensional MDS solution for 16 instrumental sounds, equalised for pitch, loudness and duration, is shown below (example 1.4).

---

[24] Ibid.: 65.

[25] Gordon and Grey (1978): 24

[26] Ibid.

Example 1.4

Adapted from Grey & Gordon (1978): 26.

Two dimensional projections of the configuration appear on the wall and floor. Abbreviations: O1, O2: oboes; C1, C2: clarinets; TP: trumpet; TM: trombone; S1, S2, S3: strings; FL: flute; X1, X2, X3: saxophones; EH: English horn; FH: French horn; BN: bassoon.

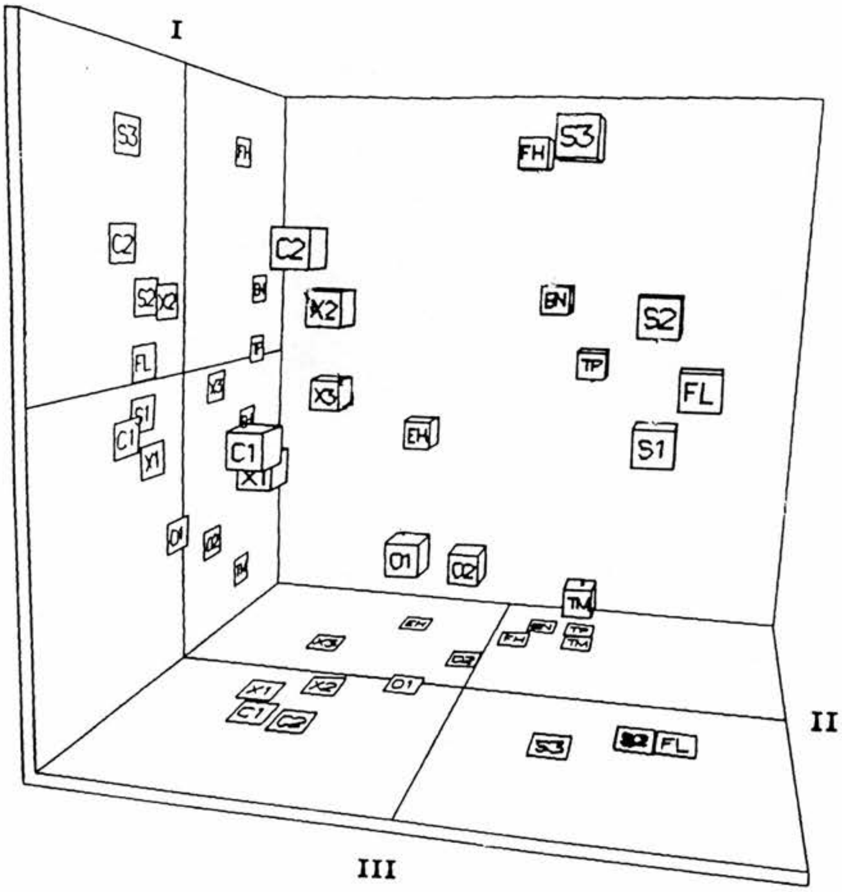

After further experiments, Grey concludes that the three dimensions represented on the graph relate to: I) the overall spectral energy distribution in the steady state section, which is closely related to the attribute sharpness suggested by Bismarck (see above); II) whether the higher partials rise and fall in amplitude together (or possibly a higher level distinction based on recognition of instrumental families); and III) whether low-amplitude high-frequency energy is present at the beginning of the starting transient.[27] Grey's work attempts to move towards the creation of a three-dimensional timbre space, where different timbres can be identified by their position in that space, thereby creating the beginnings of a graphical timbral vocabulary.

David Wessel, who has undertaken similar research to Grey, uses only two dimensions in his MDS configuration - they relate to the overall spectral energy distribution, and the quality of the 'bite' of the starting transient.[28] Wessel postulates that it may be possible to hear progressions across this 'timbre space,' such that the quality of a relationship between two points could be reproduced in another area of the timbre space, just as a pitch-pattern may be transposed (see example 1.5). To test this hypothesis, Wessel presents subjects with a sequential instrumental timbre pair A/B. He then presents a third timbre C, and asks which of several other presented timbres form a relationship with C which is most like the relationship A to B. The listeners chose most often the solution which is closest to the predicted ideal solution, such that the vector C-D is parallel to the vector A-B. Wessel concludes that "though more research needs to be done, the notion of transposing a sequence of timbres by

---

[27] Plomp & Steeneken (1973) also find three (unidentified) dimensions in their attempt to account for their MDS solution on timbre judgements.

[28] Wessel (1979).

forming another sequence geometrically parallel to it in timbre space appears to be a reasonable and musically viable idea."[29]

Example 1.5
From McAdams (1987): 31; based on Wessel (1979).

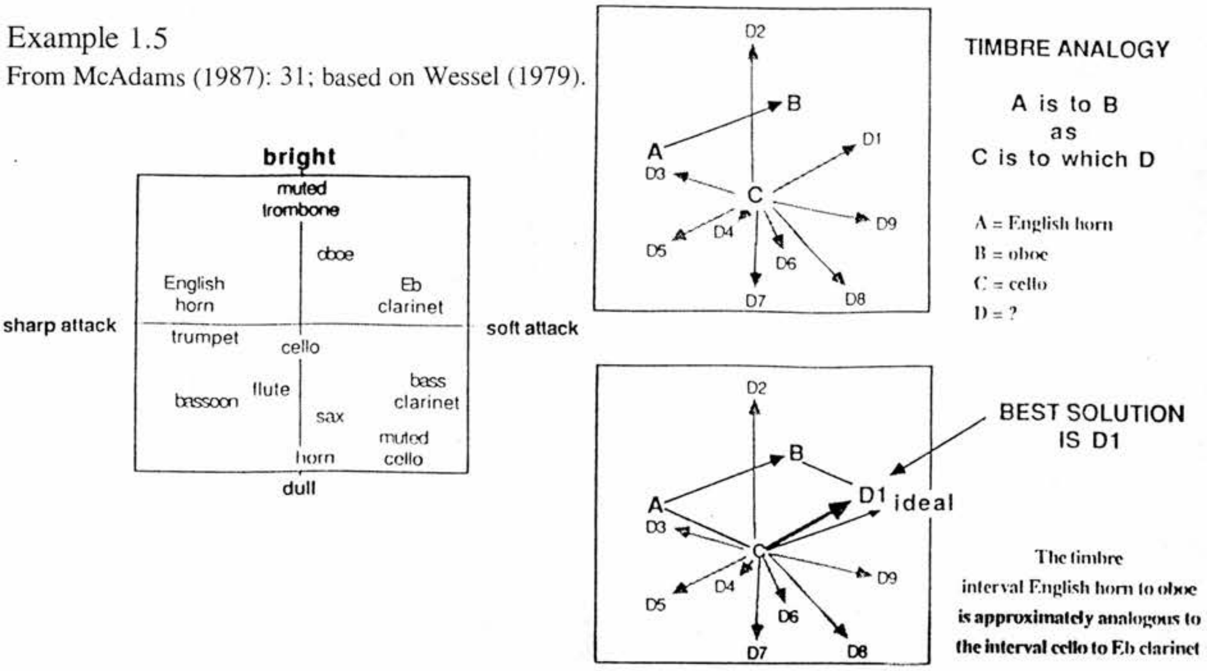

In the work of Bismarck, Grey and Wessel we have found three different suggestions for the optimum number of parameters needed to define one particular timbre from a limited set of timbres. Bismarck suggested that four scales were needed to define his group of steady-state sounds; Grey found that three dimensions were needed to satisfactorily locate his 16 instrumental sounds; Wessel used just two dimensions in his study of timbral 'transposition.' [30]

---

[29] Ibid.: 50. Grey (1975) also found that his three-dimensional timbre space allowed for interpolated sounds, derived by forming timbral sequences of modified sounds between a pair of original sounds.

[30] Three other studies that attempt to identify the parameters of timbre are: Preis (1984), Bruijn (1978), and Padgham (1986). Preis attempts to order his selected complex tones by assigning an 'ideal' spectrum envelope to them, while Bruijn, similarly to Wessel, uses MDS techniques on perception tests of steady tones, and finds that two factors are necessary for their specification. For Padgham, see note 33.

In Section 2, Howard Pollard's notion of feature analysis was introduced in order to emphasise the importance of the use of models from auditory perception theory when we wish to move from raw spectral data to perceptual understanding. Howard Pollard and Erik Jansson have devised their own method of data weighting and reduction called the *Tristimulus Method* for extracting information on the timbre of a single complex tone. Three numbers (of which two are independent variables - see the discussion below), which indicate loudness distribution across the spectrum, are plotted within a triangular space.[31] Pollard and Jansson demonstrate that this graph corresponds with our timbre perception, but do not claim that three numbers *completely* define the perceived timbre at a particular moment. This approach constitutes a highly successful method for plotting the changing timbre of a single pitched complex tone: by using three parameters to define a particular timbre, it uses a number of parameters that falls within the range seen in the previous reviewed articles; it allows for great flexibility when dealing with *developing* sounds (Pollard and Jansson use it particularly to chart the development of starting transients); the parameters are easily measured; and the method has been successfully used by the authors to graph the differences in the timbre of various organ stops (see Pollard 1990a, 1990b; Pollard & Jansson 1982b). It appears that this method presents a representation of some of the information with which we *identify* sounds. The *Tristimulus Method* will form the theoretical basis for a new form of timbral analysis that will be described in Chapter 2. Because of its importance for this thesis, the method will now be explained in more detail.

---

[31] Pollard and Jansson (1982). The theoretical starting point is that there are only three colour receptors in the eye, which combine to produce any colour. This same theoretical starting point is used by H. Yilmaz (1967/68) to graph changing vowels in a three dimensional vowel space.

Pollard and Jansson's Tristimulus Method divides the total loudness of a single note, as calculated using Stevens' Mark VII Procedure,[32] into three constituents: the loudness of the fundamental (d), the loudness of partials 2 to 4 (p_1), and the loudness of partials 5 and upwards (p_2). It is known that partial numbers 3 and 5 are very important for our perception of an harmonic sound, so, through experimentation, it was decided that these two partials should be incorporated into separate measurements. The total loudness of a note (t) may now be expressed as the sum of three constituents:

$$t = d + p_1 + p_2$$

Three values, x, y, and z, are calculated that represent the loudness of the three loudness constituents as a fraction of the total loudness. Only two of the values are independent variables, and thus only these two need to be plotted, as all three must add-up to 1 (because all three are fractions of the total). Importantly, by the nature of the calculations, the values have been normalised for pitch and total loudness.

$$x = p_2 / t$$

$$y = p_1 / t$$

$$z = d / t$$

As may be seen from example 1.6, Pollard and Jansson plot their co-ordinates within a triangular space in order to represent the changing balance between the three constituent loudnesses. Example 1.6a shows how the triangular space is 'mapped'; example 1.6b shows the tristimulus contours for three different starting transients.[33]

---

[32] Stevens' Mark VII Calculation Procedure is an internationally recognised procedure for loudness calculation. It is described in detail in Chapter 2. See Stevens (1972).

[33] Another attempt at graphical representation of timbre is that by Padagham (1986). The steady-state sections of pipe organ tones were assessed by subjects who were asked to estimate a position for

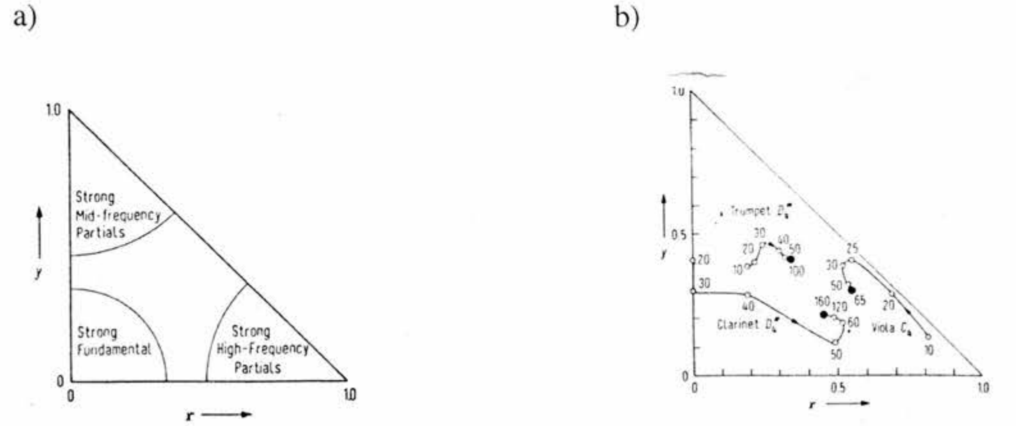

Map of triangular space

Acoustic tristimulus diagram for viola C4, clarinet D#4, and trumpet D#4. Numbers indicate time in milliseconds after the onset of sound.

From Pollard & Jansson (1982): 106, 108.

Example 1.6

Let us refer once more to the Table devised by Pollard that appeared at the end of Section 2 (example 1.3). Now that the terms 'sharpness' and 'tristimulus diagram' have been explained, the reader may assess more accurately what is involved to obtain a feature analysis of a sound. In addition to the physical properties of the sound that can be measured and then weighted, our perception of the timbre of a musical work can be influenced by the composer's ability to alter our listening focus within the overall effect of an ensemble by manipulating, for example, the amount of change in

---

the timbre on a clock-face type chart, the two dimensions being 'tone' (reference points being *flute*, *diapason*, *string*, *trumpet*) and 'complexity' (progressing from 'simple' to 'very complex'). Padagham suggests that 'tone' is related to the strengths of harmonics 1 to 5, and 'complexity' to the strengths of harmonics 6 to *n*.

timbre, dynamic and pitch content.[34] Processes such as short and long-term memory also play a part in the ways we perceive music.

The studies described above all deal with just single instrumental sounds. However, as was implied in the previous paragraph, the situation of listening to combinations of sounds can be very different. What happens when we hear many instrumental or vocal sounds within a musical context ? Stephen McAdams and Albert Bregman state that: "The perceptual effects of a sound are dependent upon the musical context in which that sound is embedded. That is, a given sound's perceived pitch, timbre and loudness are influenced by the sounds that precede it, coincide with it, and even follow it in time."[35] In work subsequent to his dissertation, Grey investigates the way we perceive timbre in musical patterns. At the opening of his article, he observes that:

> there are clearly many facets of timbre presented in typical musical contexts that do not exist for isolated tones. The listener is allowed to compare the different spectra of notes from the same instrument taken at different pitches, and the composite spectral map of the instrument may be a factor in normal timbre recognition and evaluation... Also, perception of a musical context may be different from the perception of isolated phenomena because of various temporal performance constraints. For example, the ability to discriminate two timbres may be affected by the complexity of their context. Likewise, the set of criteria used to form a similarity judgement between two isolated tones may be quite different from those used to rate the similarity of two timbres in musical contexts.[36]

---

[34] Alan Belkin (1988): 49, lists three types of sounds, their information, and their perceptual effects:

Type of sound	*Information*	*Perceptual Effect*
Short, percussive.	Much information in a short time.	Attracts attention.
Sustained, stabilising.	New information quickly diminishes.	Settles into background.
Sustained, evolving.	Continues presenting new information.	Keeps listeners' attention.

[35] McAdams & Bregman (1979): 26.

[36] Grey (1978): 467.

Grey concludes that an important difference in our perception of notes in isolation and notes in context is that "musical contexts seem to extend and amplify spectral differences between notes, while isolated contexts may make slight temporal differences more apparent."[37] This view supports the assertion made near the beginning of the thesis that, when attempting to discuss the timbral trajectory of a piece, which is the primary concern of this study, the importance of the starting transients tends to be subsumed by the perceptual importance of the steady-states.

Grey's observation may be related to the phenomenon known as *stream segregation.* The effect may be heard when a melodic line, consisting of alternating notes of sufficiently different characteristics, is heard to consist of more than one segregated line.[38] Wessel (1979) finds that a melodic line will 'split' when a competing pattern of alternating timbres, based on differences in sharpness, is applied to the melodic pitches. This supports Grey's belief in the importance of *spectral* differences in a musical context.

The question of how we perceive timbre in a musical context is a largely untouched area, but it is, of course, a question addressed by composers with every composition they write. Particularly in orchestral compositions of the 1960s, and in pieces that use electronics, the question of timbral control becomes much more important. Risset and Wessel believe that "with the control of timbre now made possible through analysis and synthesis, composers can ... articulate musical compositions on the basis of

---

[37] Ibid., p.471. Strong & Clark (1967), in a study of single synthetic wind-instrument tones which last longer than 0.5 second, find that the spectral envelope is predominant in aural significance to the temporal envelope for instruments with a unique spectral envelope. Where the spectral envelope is not unique, the temporal envelope is of equal or greater significance.

[38] For example, Bach's solo works for violin and cello often make use of this phenomenon to give the illusion of polyphony. See McAdams & Bregman (1979).

timbral rather than pitch variations." But does composing primarily with timbres offer the same richness of possibilities as composing primarily with pitches? "It has been argued," Risset and Wessel continue, "that timbre perception is too vague to form the basis of elaborate musical communication; however... there already exists an instance of a sophisticated communication system based on timbral differentiation; namely human speech. Hence, it is conceivable that proper timbral control might lead to quite new musical architectures."[39]

3.2 Timbre and Music

In order to assist a movement towards the possibility of new musical architectures based on timbral structures, and in order to understand how timbre has been used in pieces already written, composers and theorists have attempted to formulate a plane of concepts and a vocabulary for timbre. One of the first to address this question was the composer and theorist Arnold Schoenberg in his now much-quoted statement found at the conclusion of his *Harmonielehre*:

> The evaluation of tone colour, the second dimension of tone, is in a much less cultivated, much less organised state than is the aesthetic evaluation of [pitch]. Nevertheless, we go right on boldly connecting sounds with one another, contrasting them with one another, simply by feeling; and it has never yet occurred to anyone to require of a theory that it should determine laws by which one may do that sort of thing... Now, if it is possible to create patterns out [of] pitch[es], patterns we call "melodies," progressions, whose coherence evokes an effect analogous to thought processes, then it must also be possible to make progressions out of... "tone colour," progressions whose relations with one another work with a kind of logic entirely equivalent to that logic which satisfies us in the melody of pitches.[40]

---

[39] Risset & Wessel (1982): 50.

[40] Schoenberg (1978): 421.

Nowhere else does Schoenberg discuss at length what he envisages for *Klangfarbenmelodie* ("tone colour melody"), and it appears he wrote only one composition which obviously used the idea - "Farben," from *Five Pieces for Orchestra*, Op.16 (1909).

Wayne Slawson is a composer who has attempted to provide a systematic theory where Schoenberg did not. In his book *Sound Color*, he presents a theory for composing with a sub-group of timbre, which he calls *sound colour*:

> Sound color is a property or attribute of auditory sensation; it is not an acoustic property... Sound color, like visual color, is abstract; no specific source of energy is implied by either term. Again, like visual color, sound color has no temporal aspect... a sound may be heard to be changing from one color to another, but the change itself is not a sound color. It follows that sound color pertains to the steady-state portions of sounds but not, in general, to their beginnings or endings, where we can sometimes hear rapid changes in the character of the sound. Sound color and visual color are multidimensional.[41]

We can see here that Slawson considers the steady-state to be primary in our perception in a musical context; and this same idea was seen in the quote from Grey (1978) reproduced above.

Slawson provides supporting evidence for his theory from the disciplines of auditory physiology, psychoacoustics, speech and cognitive science, along with examples from the electronic music repertory. Slawson proposes a two-dimensional sound-colour space, which he maps by the placement of different vowel sounds: the x axis represents the frequency (in kHz) of the first resonance, or formant, and the y axis represents the frequency of the second.[42] Slawson then defines certain types of

---

[41] Slawson (1985): 20.

[42] Formant regions are represented by peaks in the spectrum envelope. Even though the pitch of a particular vowel changes, its formant peaks above the fundamental will always be located at the same frequencies. Slawson believes that a vowel sound may be adequately defined by its first two resonances.

movement within this sound-colour space during which a particular property remains invariant. These four properties, which he names *openness*, *acuteness*, *laxness*, and *smallness*, are illustrated in example 1.7. He proposes that movement within the sound-colour space is open to the operations of transposition and inversion.

Example 1.7

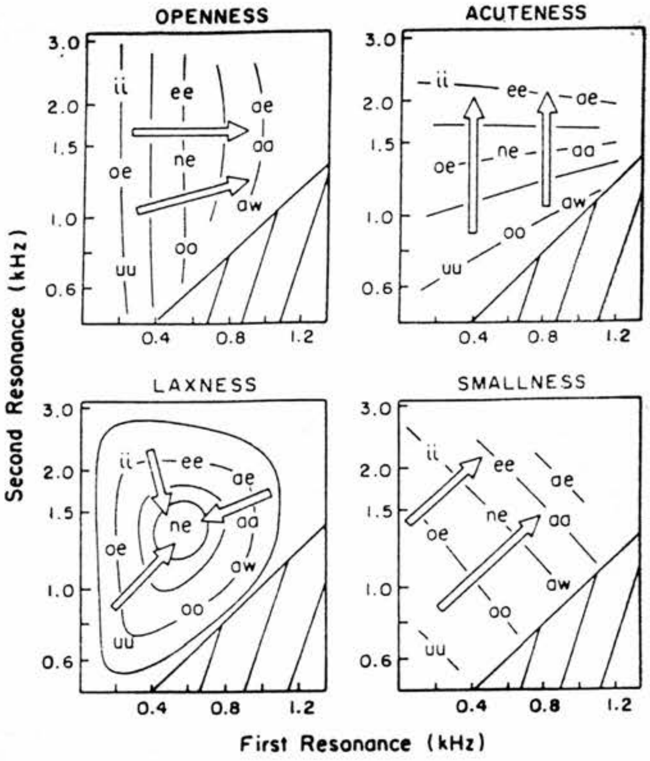

Equal-value contours plotted as functions of the frequencies of the first two resonances. Arrows indicate direction of increasing value. From Slawson (1985): 55.

Slawson's work is rigorous in proofs, and is shown to be musically workable through the inclusion of a discussion and recording of parts of his own composition *Colors*, which he composed using the theory. His attempt to map timbre within a space with a small number of dimensions allies his work with that of Grey (1975) and Wessel (1979). However, while Grey and Wessel were attempting to understand how we perceive timbre, and to provide a basis for timbral analysis, Slawson's theory is designed as a tool for composition, and because of its limited number of timbral possibilities, is not suitable for analysis.

Timbre

The composer and theorist Fred Lerdahl (1987) is concerned with both timbral composition and analysis. Instead of beginning with acoustic theory like Slawson, Lerdahl begins intuitively. He proposes that timbre can be organised hierarchically, and that it is possible to talk of timbral consonance and dissonance.
Lerdahl begins by defining timbral stability conditions, which he discusses in terms of relative consonance and dissonance. Intuitively, he finds a 'brighter' sound more tense or dissonant than a 'dull' one - thus, a violin is more dissonant than a viola; no vibrato, or a wide vibrato, are both more dissonant than the culturally determined optimum vibrato; a sharp attack or release is more dissonant than a smooth attack or gradual release; a pitch with a harmonic spectrum is more consonant than a pitch with an inharmonic spectrum. Thus, the many timbral dimensions are unified by the concept of timbral consonance and dissonance. "Along any particular dimension," Lerdahl writes, "one can *feel* timbres becoming more dissonant or consonant."[43]

With the idea of timbral dissonance and consonance, tensing and relaxing, it is possible to propose timbral 'scales'. A timbral scale "is a linear ordering of timbres at fixed intervals along a given dimension or combination of dimensions" that progress away from or towards the prototypes (the most consonant timbres).[44] The interval size is decided intuitively. For example, for the dimensions of vibrato and harmonicity, their respective scales are shown in example 1.8a. By combining timbral scales into "timbral arrays," the perceptually distinguishable aspects of timbre can be plotted together. The dimensions vibrato and harmonicity are shown combined in example 1.8b. Theoretically, an infinite number of scales, or axes, can be combined.

---

[43] Lerdahl (1987): 142.

[44] Ibid.: 145.

Example 1.8a

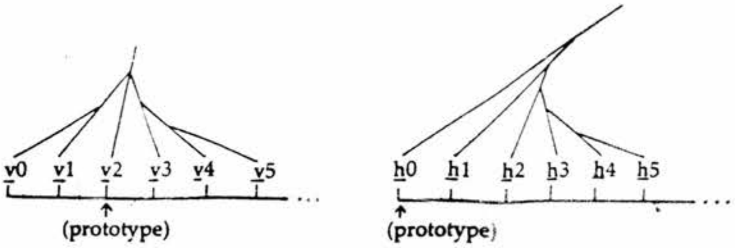

Scales for vibrato and harmonicity, showing tree-branching hierarchies

Example 1.8b

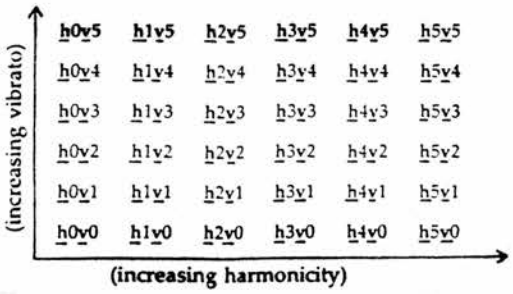

h0v5	h1v5	h2v5	h3v5	h4v5	h5v5
h0v4	h1v4	h2v4	h3v4	h4v4	h5v4
h0v3	h1v3	h2v3	h3v3	h4v3	h5v3
h0v2	h1v2	h2v2	h3v2	h4v2	h5v2
h0v1	h1v1	h2v1	h3v1	h4v1	h5v1
h0v0	h1v0	h2v0	h3v0	h4v0	h5v0

Vibrato and harmonicity scales combined.
Both examples from Lerdahl (1987): 146, 147.

While Lerdahl's theory accommodates the multidimensional aspect of timbre, and does provide the basis for a timbre vocabulary, it is not a vocabulary that could be confidently used by an analyst. Consonance and dissonance are relative ideas - what might be felt as timbrally consonant in one composition, might be felt to be timbrally dissonant in another; and how are we to agree upon the definition of the size of a step along a timbral axis? The theory might perhaps work best as a personal compositional tool, rather than as an analytical method.

Richard Swift, in a review of Robert Erickson's *Sound Structure in Music* (1975), [45] states his requirements for a method of timbre analysis:

---

[45] Swift's aim is to gain "some understanding of the role of timbre in music and [to locate] principles upon which a theory of timbre organisation might be constructed." (p. *ix*). Instead of this, the reader is lead through acoustic theory and musical examples with a joyful abandon that gives rise to an interesting collection of timbral examples, rather than to the beginnings of a timbral theory.

> What is wanted is an account of tone quality that is not antagonistic toward other components of music or to other musical processes, but is congruent and coherent with them, an account which will strive to match, at the least, the level of recent intensive analyses of pitch, contour, harmony and structure.[46]

Both Slawson's and Lerdahl's theories fail as tools for timbre analysis - Slawson's methodology is too limited and specific, Lerdahl's is too subjective and indefinite, and neither matches the rigour of much pitch analysis.

An approach that does contain the potential to live-up to Swift's analytical ideal is that taken by Robert Cogan. Cogan approaches the topic acoustically by examining the sound spectrum of an imaginary or real performance of a composition.[47] His work is closely tied to the technological effectiveness of the means of measuring and representing a sound spectrum. In his article of 1969, Cogan's spectral data for his music analysis consists of a set of spectrographs of vowel sounds produced in 1947.[48] In his work of 1975 and 1976, Cogan presents an analysis of Schoenberg's *Farben* using the spectral data produced by Carl Seashore in 1938. As Seashore only presents data for certain pitches of certain instruments, at the dynamics *forte* and *piano,* Cogan must use spectral estimates when representing the changing timbres.[49] As Cogan himself says:

---

[46] Swift (1975): 158. Four studies which do consider the relationship of timbre to other parameters of a piece (apart from spectrum analysis based studies) are Hatmaker (1985), who finds norms of progression in instrumental combinations, in association with phrase structure, in the Classical symphony; Pellman (1979), who finds vowel-like progressions in Berio's *Sequenza V*; and Chou Wen-Chung (1979) and Francois (1991), who both examine the structural role of timbre in Varèse's *Ionisation.*

[47] Cogan's principal work on this subject is Cogan (1984). His other pertinent writings are: Cogan (1969), (1975), (1980), (1987), (1988), (1991); Cogan & Escot (1976): chapter 4.

[48] His source is Potter, Kopp, and Kopp (1947).

[49] Burkhart (1974) analyses the orchestration of *Farben* in connection with the principal method of pitch organisation, and finds evidence of serial ordering. Unfortunately, Cogan makes no reference to this.

> In preparing the graphs it has been necessary to choose the closest pitch and dynamic for which spectral information is available. Other technical factors (dynamic differences between instruments; masking) have not been accounted for.[50]

An obvious way of improving on the spectrum analysis method would be to have direct measurements of the piece to be analysed. This is what Cogan presents in his principal study involving spectrum analysis (Cogan [1984])[51]. This work marks a huge leap forward in technique, for Cogan generates his own spectral measurements of entire pieces or sections of pieces. To create his spectrum photos, or spectrographs, he uses a "33-millisecond fast Fourier transform instrument, capable of analysing sounds in five contiguous octave registers simultaneously."[52] This is used to drive a cathode ray tube display. The display can retain a section of the changing spectra, which are photographed and the photographs subsequently joined together. On these photographs, the vertical axis represents time, and the horizontal axis represents frequency. To a limited degree, the thickness of the trace indicates amplitude, though, when necessary to his argument, Cogan uses 'detail photos' of the spectral elements to show amplitude at a specific moment (graphs of amplitude versus frequency). An example of Cogan's spectrographs is shown below.

---

[50] Cogan (1975): 63, footnote 9. Examples from Seashore (1938) may be found in Cogan & Escot (1976): 351-54.

[51] Cogan first uses spectrum analysis in Cogan (1980) where he also hints for the first time at the method of analysis using sonic oppositions.

[52] Cogan (1984): 155.

Example 1.9

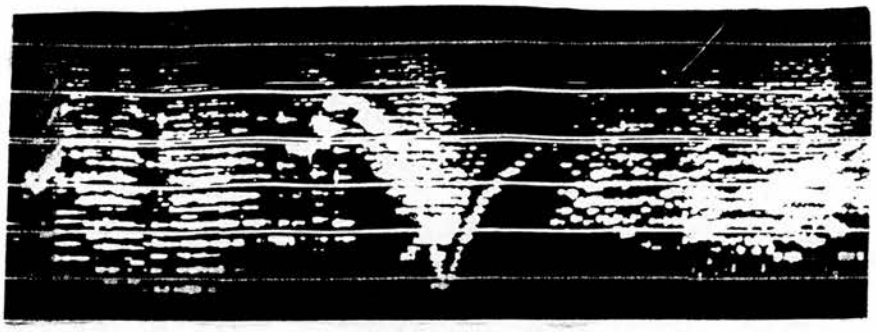

Spectrograph - part of Beethoven's
Op.109/i
From Cogan (1984): 31.

At the end of the book, Cogan states his hopes for the future:

> Ideally, every complete-context spectrum analysis should cover the entire audible space range, respond to the full audible dynamic range, and precisely display the relative loudnesses of all the spectral elements more or less as they are perceived by the human auditory system.[53]

Part II of Cogan's book consists of the presentation of a theory for the analysis of the sonic design of a piece, using the results from the previously analysed spectrographs. The theory, based on linguistic phonology, analyses sonic oppositions. By considering thirteen sonic oppositions (for example, grave/acute, narrow/wide, no-attack/attack, beatless/beating), Cogan assesses to what extent, for a defined segment of a piece, these attributes are present. Thus, for each oppositional pair, a segment of music may be classified as positive (+; a high energy state - e.g., acute), negative (-; a low energy state - e.g., grave), a mixture of positive and negative (±), or neutral (0). These classifications are not absolute, but are judged within the musical context. For example, the decision as to whether a section is grave or acute is based on a division into three of the spectral range of the entire musical segment under consideration. If

---

[53] Ibid.: 156.

there is spectral activity in the upper third of the range, that section of the segment is designated acute (+); spectral activity in the lower third of the range is designated grave (-). If there is activity in both the upper and lower thirds, it is designated grave-acute (±), and activity only in the middle third is designated neutral (0). By considering all the sonic oppositions in a section, as defined by +, -, ±, and 0, Cogan can specify a total number for that section, and thus specify the total activity of a section (as defined by his sonic oppositions) as well as specify ratios of change from section to section.

Cogan's method of ordering and categorising various aspects of his spectrographs is a useful and insightful one. By recognising a limited number of features, and by considering them within the context of a particular segment of music, the spectrographs reveal aspects that add to our understanding of the way the music 'works,' and allow us to compare the workings of different sections of music.[54] Richard Swift, whose review of *Sound Color* was quoted above, believes that "Cogan has formulated a brilliantly viable theory of sound: a theory that... [promises] to provide a basis for the analysis of sound." However, Swift also points out that "it might be objected that... the adjectival characteristics of the terms in which the analysis is couched constitute banal binary oppositions, or that they are narrowly generalized... One views with some misgivings an analytical system that is so restrictive in its alternatives."[55] Cogan also presumes that we identify the sounds of language in a similar manner to the way we listen to music - a point tenaciously criticised by John Strawn (1986 a; b) in his review of Cogan's book.

---

[54] Two brief analyses, by analysts other than Cogan, which unite pitch and formal considerations with spectrum analysis are Gebura (1980) and Fry (1984), who both analyse songs from Stockhausen's *Indianlieder* (1972). The analyses elucidate an important musical structure which could only be discussed concretely with the aid of spectrum analysis.

[55] Swift (1986): 283.

A feature of Cogan's work is that he is interested in the entire 'sonic design' of a piece; so while he often talks of 'tone color,' he does not attempt to define what tone colour is, but rather treats it as an integral aspect of this sonic design. In addition, he treats the spectrum photos as direct representations of our perception, not attempting to interpret the spectral data in the light of perceptual models. In both these cases, Cogan treats the spectra as objects in themselves, rather than as representations of data which can be further manipulated.

In the light of the work reviewed in this Chapter, possibilities suggest themselves for improving upon Cogan's pioneering approach. A spectrograph needs to be recognised as simply a representation of a set of data - data that can be transformed in various ways. This transformation of the data must take into account two interlocking requirements for successful timbre analysis - the data needs to be weighted in order to accord more closely with models of auditory perception, and the mass of data that is represented by a spectrograph needs to be reduced so that we are left with information specifically to do with 'timbre', as defined at the beginning of this thesis. Once the information in a spectrograph has been weighted and reduced, we need methods for displaying and discussing this data that allows the richness and subtlety of timbre to be appreciated by allowing continuous movement between timbral categories. Lastly, the analysis of timbre needs to be integrated into a rigorous analysis of pitch and formal structures so that interaction between structural parameters of a piece can be clearly observed.

In order to arrive at these requirements for successful timbral analysis, we need specific tools for the analysis of timbre. The following Chapter defines these tools.

***

CHAPTER 2

Tools for Exploring Timbre

CHAPTER 2

Tools for Exploring Timbre

---

1. Introduction

In order to arrive at a plane of concepts and a vocabulary with which timbre can be rigorously discussed and analysed, a set of tools needs to be assembled that allows the elements of timbre to be extracted from the mass of data that is represented by a spectrograph. As we saw in Chapter 1, these elements consist of *the primary aural information that is used in the perceptual task of assigning an identity to a sound.* The analytical tools need to bring to light this aural information.

In example 2.1 below, a part of example 1.3 is reproduced - it shows the type of information that Pollard believes is necessary for Feature Analysis:

Type of Information for Feature Analysis:

* assessment of starting transient (early sound, synchronism, dominant tones, duration).
* sharpness (loudness centroid) as a function of time.
* timbre (synthesis of factors such as sharpness, synchronism, early sound or equivalent measures as in the tristimulus method).
* pitch
* loudness
* fluctuations of loudness, pitch, timbre.
* roughness (beats between partials in same critical band.)
* other cues (eg., compactness, missing partials).

Example 2.1

We can extract those elements from Pollard's list that are to do with sound *identification* - that is, those elements that are not pitch, loudness or duration. Because of the focus in this thesis on the steady state element of sound, we can exclude those elements that are to do with the way we perceive temporal change. This procedure produces the following list of elements that are necessary for sound *identification*:

* timbre (as assessed by the tristimulus method);
* sharpness;
* roughness;
* other cues (e.g., compactness, missing partials).

(In this thesis, the tristimulus method is understood as a means of measuring an important aspect of our timbre perception, complementary to other measures; it is not a measure of the entirety of our timbre perception as is suggested by Pollard's wording in example 2.1).

By assembling a collection of tools that allows the above elements to be measured, we will gain access to much of the acoustic information that informs our perception of music timbre.

These tools are summarised below, and described in detail in section 3 of this chapter.

* *Timbre*:

In Chapter 1, the Tristimulus Method was judged to be a successful method for representing some of the information with which we *identify* sounds. The Tristimulus Method was designed specifically to analyse the timbre of a single tone - not the timbre of an ensemble sound, which is the more usual situation in Western music, and thus the more usual situation to be encountered in music analysis. In order to make it possible to move towards the analysis of the changing timbre of ensemble music, a new and original method has been devised in the spirit of the Tristimulus Method. This method is called *Loudness Distribution Analysis*. It produces three timbral measures - *timbral width*, *timbral weight*, and *timbral pitch*.

* *Sharpness*:

There have been a number of studies that have shown sharpness to be an important perceptual attribute[1], but the notion of sharpness has not before entered the discipline

---

[1] See von Bismarck (1974a) & (1974b); Grey (1975); and J.W. Beauchamp (1982): "Synthesis by spectral amplitude and 'brightness' matching of analyzed musical instrument tones." *J. of the Audio Engineering Soc.* 30/6: 396-406; W. von Aures (1985): "Berechnungsverfahren für den sensorischen

of music analysis. In this thesis, sharpness is defined through the frequency position of the loudness centroid.

* *Roughness*:

Roughness is a perceptual term related to the musical term 'dissonance.'[2] Roughness is caused through beating between partials. A method for calculating roughness has been proposed by Hutchinson and Knopoff (1978), and roughness measurements of entire movements of ensemble music are presented here for the first time.

* Harmonicity and *other cues*:

A method for assessing the degree of *harmonicity* is also employed in this study. This is related to the measure of roughness, but is not perceptually based. Called *cepstrum analysis*, it assesses to what extent a spectrum diverges from an 'ideal' harmonic spectrum. It will be remembered from Chapter 1 that Bismarck defined compactness as a measure of a sound on a scale between complex tone and noise - that is, the difference between discrete and continuous spectra - though Bismarck proposed no actual scale.[3] Cepstrum analysis, being a measure of the physical make-up of the spectrum, is a method by which we can measure compactness, as well as, to a limited degree, see the influence of missing partials (if partials are missing, then this will contribute to the degree of divergence from an ideal harmonic spectrum).

---

wohlklang beliebiger schallsignale" [Procedure for calculating the sensory euphony of arbitrary sound signals]. *Acustica* 59: 130-141; Goad (1992).

Both Aures and Goad offer more complex models for the calculation of sharpness than that offered by Bismarck (loudness centroid). However, Bismarck's model is used in this study in order to reduce the complexity of calculating the sharpness of ensemble sounds.

[2] Roughness is a measurable psychoacoustic phenonemon. How dissonance is perceived is dependent on its historical musical context.

[3] See Chapter 1, Section 3.1.

Thus, four complementary tools are used in this thesis to analyse timbre - Loudness Distribution Analysis ('LD-analysis'), Sharpness Analysis, Roughness Analysis, and Harmonicity (Cepstrum) Analysis. These tools provide us with the means to gain access to the information with which we *identify* sounds. With these analytical tools, a plane of concepts and a vocabulary for timbre analysis can be firmly established. Instead of music timbre analysis relying solely on vague descriptive terms, as it has so often in the past, these measures enable timbre analysis to be executed with the same degree of rigour as we find in methods of pitch analysis. Because of this, the interaction between the structures of timbre, pitch and form can also be rigorously investigated. The methods are also repeatable, and thus can be verified, and possibly refined, by other analysts. It is hoped that through this type of work, timbre, that aspect that is so vital to our experience of music, will no longer be ignored in future music analytical research.

2. Technical Background

The analysis procedure begins with a sound source - in this study, a commercially available CD recording (manufactured at a sampling rate of 44.1 kHz). The piece to be analysed is then transferred to the hard disk of a *DEC*pc 425i computer (CPU i486DX2 - 50MHz) using software and sound cards designed by *Digital Audio Labs* (*The EdDitor* vs. 2.31, Waveform Editor and Catalogue Manager; *The CardD* sound card, sampling at 16 bits, at 44.1 kHz)[4]. The sound file is acquired by the Digital Signal Processing package *Hypersignal*,[5] where a Fast Fourier Transform (FFT) is executed on the waveform file, and the resultant data on frequency can be displayed as a spectrograph. The phase information is discarded.

The FFT parameters are as follows:

Transform size	: 4096
Overlap	: 2048
Window	: Hamming

These parameters were chosen to provide optimum time and frequency resolution within the constraints of the software.

3. Methods

3.1 Spectrum Analysis

The immediate predecessor of this study is that by Cogan (1984) in which spectrum analysis plays such a prominent role. The timbre analytic tools described above all use as their input data derived from spectrum analysis. So, let us begin the detailed examination of the methods to be used in this study with a look at spectrum analysis.

---

[4] Digital Audio Labs: 6311 Wayzata Boulevard, Suite 330, St. Louis Park, Minnesota 55416, U.S.A.

[5] *Hypersignal* is produced by Hyperception, 9550 Skillman LB125, Dallas, TX 75243, U.S.A.

We saw an example in Chapter 1 (example 1.9) of a spectrum analysis reproduced from Cogan (1984). Since Cogan's work, computer software and hardware development has been swift, so that spectral displays are now greatly enhanced from the type of display seen in example 1.9. Figure 2.1 shows the spectra of two contrasting sounds.[6] The first is the spectrum of a Hemony carillon bell, with a strike note pitch around 500 Hz. (hum note 251 Hz).[7] The second is a guitar tone with a fundamental frequency of 251 Hz. (a sharp B^3).[8] Frequency, on a logarithmic scale, is represented on the vertical axis of the graph, and the pitch C^4 is shown as a point of reference (the marks on this axis above and below C^4 are pitch-class C's at various octaves). Time is represented along the horizontal axis. Amplitude is shown by a grey scale, the calibration of which is given at the top-left of the graph. In this instance, there is a difference between each of the nine shades of grey of 6.5 dB. The greatest amplitude is assigned the colour white. An examination of the two spectra reveals that the second (guitar tone) is very largely composed of harmonic components, while the first (bell tone) contains many non-harmonic partials causing it to have a less definite pitch. Note the presence of noise at the very beginning of the bell tone, and the presence of a component in the guitar tone, especially prominent during the starting transient, which lies an octave below the fundamental.

It was emphasised in Chapter 1 that a spectrograph is *not* a direct representation of the way we hear a sound: it is possible that a low frequency that appears to have a high amplitude on a spectrograph may, after processing by the inner ear, virtually

---

[6] All spectrographs in this thesis are produced by the software *Hypersignal*.

[7] A description of the components of this spectrum will be found in Section 3.5 of this Chapter.

[8] Both sounds are recorded from *Phillips* Auditory Demonstrations CD (1126-061) (1987): demonstration 28. [Available from the *Acoustical Society of America*.]

'disappear' as far as our perception is concerned. It will be remembered that perceptual tests have been carried out that have allowed equal-loudness contours to be mapped (see Chapter 1, Section 2). A relatively simple way of moving the spectrograph closer to a representation of our perception is by applying what is known as 'A-weighting.' A-weighting weights all spectral components in a way roughly equivalent to the 40 phon equal loudness contour, but has been found to correlate well with the perceived loudness of many different types of sounds and sound levels.[9] Figure 2.2 shows the same two spectra after A-weighting has been applied by a program written specifically for this research.[10] Generally, the effect is that low and very high frequencies are attenuated, while the mid-range is boosted. Comparing the two figures, we can see that the noise associated with the beginning of the bell tone and the lower octave associated with the starting transient of the guitar tone are both attenuated, while the frequencies between C^5 and C^7 are boosted (especially evident in the bell tone).

Spectrum analysis can be very useful for representing the 'acoustic outline' of a sound, or the 'sonic design' of a piece of music (as Cogan has shown). By weighting the spectrograph, we move a little closer to a visual representation of perception (subsequent spectrographs in this study are all A-weighted). Some timbral information can also be deduced from a spectrograph - the first sound of figure 2.2 is visually less consonant than the second. However, the information in a spectrograph can be reduced and displayed in such a way that we come still closer to a representation of music timbre. As was noted at the start of this Chapter, for this task we need certain tools, and these tools are described in detail below.

---

[9] See Parkin, P.H. (1965). "On the accuracy of simple weighting networks for loudness estimates of some urban noises." *Journal of Sound and Vibration* 2: 86-88.

[10] All computer programmes created for this thesis were written by A. Murray Campbell.

3.2 Loudness Distribution Analysis

The Tristimulus Method, which was described in Chapter 1, produces three values derived from the loudness distribution in the spectrum of a single note. It will be remembered that these values are derived as follows:

$$x = p_2 / t$$
$$y = p_1 / t$$
$$z = d / t$$

where p_2 represents the relative loudness of partials 5 and upwards; p_1 the relative loudness of partials 2 to 4; d the relative loudness of the fundamental; and t the total loudness.

In the tristimulus method, loudness is grouped and measured from the fundamental upwards. In an ensemble sound, there is no single fundamental from which measurements can be based. However, in an ensemble sound our attention is likely to be drawn to the loudest element. Thus, in the Loudness Distribution Method, the loudest element is taken as the basis for the measurements. The value of the loudness of the fundamental in the Tristimulus Method is replaced with the value m - the loudness of the loudest 1/3 octave frequency band.[11] Two other variables can then be derived - n - the loudness of all frequencies above the loudest band, and o - the loudness of all frequencies below the loudest band. Thus, just as the total loudness (t) in the Tristimulus Method is the sum of the loudness of three sections of the spectrum

---

[11] 1/3 octave frequency bands are chosen because this is the unit used by Stevens in his Mark VII procedure, which in turn is the basis of the Tristimulus Method. The Stevens' Procedure for loudness calculation is especially appropriate for the Loudness Distribution Method because in the procedure the step prior to the calculation of the total loudness involves calculating the loudness of the loudest 1/3 octave band.

(d, p_2, p_1), the total loudness (T) in the Loudness Distribution Method is the sum of m, n and o:

$$T = m + n + o.$$

Values for l, a and b may then be derived. l is the value of the relative loudness of the loudest band; a is the value of the relative loudness of all frequencies above the loudest band; b is the value of the relative loudness of all frequencies below the loudest band.

$$l = m / T$$
$$a = n / T$$
$$b = o / T$$

The graph produced from this data is called a Loudness Distribution Graph (LD-graph).

As in the Tristimulus Method, loudness is calculated using Stevens' Mark VII Procedure. This consists of calculating the loudness of all 1/3 octave bands in sones, taking into account frequency-weighting functions consisting of equal loudness contours. A summation rule is then applied:

$$S_t = S_m + F(\Sigma S - S_m)$$

where S_t is the total loudness, S_m is the loudness of the loudest band, ΣS is the sum of the loudnesses of all the bands, and F is a factor that accounts for masking - F varies with S_m in a way defined by Stevens. Steven's Mark VII Procedure was incorporated into a series of specially written programs for this thesis to calculate Loudness Distribution Graphs.[12]

Figure 2.3 shows an LD-graph of the bell and guitar tones discussed above. The triangular form of the original Tristimulus Diagram has been abandoned because of the need for the inclusion of a time axis (time is represented on the horizontal axis). Each sound is graphed using two separate plots The upper line shows the change over time for the sum of a and b (see the equations above). The lower line shows the change over time for the subtraction of b from a. Thus, the upper line shows the total fraction of loudness that lies outside of the loudest 1/3 octave band (the spread of the loudness), and the lower line shows the degree to which this fraction of loudness lies above or below this loudest band - on the graph, above or below 0.0 (the skew of the loudness). To use words suggestive of the quality of these quantities - $a + b$ is timbral *width* (focused/diffuse); $a-b$ is timbral *weight* (heavy/light). Because l, a and b must add up to 1, only two of the values need to be plotted. However, the frequency of the loudest band needs to be considered, as even a sine wave, swept through the audible frequency range, has a perceptible change of timbre. Timbral *pitch* is the loudest 1/3 octave band number (whose loudness is l) graphed against time. This is shown in figure 2.4. A description of the movement of timbral width, weight and pitch will form the beginnings of a timbral vocabulary.

Let us look first at the bell sound of figure 2.3 (the first sound represented on the graph). Timbral width ($a + b$) lies in the mid-range; it shows a steady decline during

---

[12] The amplitudes that result from an FFT are valid only in relationship to each other - not in any absolute sense. For the Stevens' loudness calculations to relate to perceived loudness, a figure is entered into the calculations that is an estimate of the value in sones of the loudest moment in the music that is being analysed. This figure appears at the bottom right-corner of the LD-graphs. On the bottom left corner is a figure that indicates the degree of averaging that has been used and the resulting time period between each data-point on the graph.

the first half of the sound (up to letter y), and minimum change during the second half. In other words, during the first half of the sound the timbre becomes more focused - a greater fraction of the total loudness lies within the loudest band than at the beginning. This conclusion is supported by a comparison with figure 2.2. (It is important to remember that the frequencies in the spectrum have been A-weighted, while the data that is used in calculating the LD-graphs has been weighted using Stevens' procedure of 1/3 octave banding. Therefore, close, but not exact agreement is what we will look for.) The moment shown by letter y on figure 2.3 is also shown on figure 2.2. At the beginning of the bell tone there is a lot of inharmonic spectral activity, especially in the upper part of the spectrum. As we move towards letter y, this activity lessens in amplitude. At letter y, we are basically left with only seven partials evident in the spectrographic display. From this point until the end of the sound, all partials lessen in amplitude together, with some falling below the dynamic range shown here. However, the measure of timbral width remains relatively static over this time. Why is this? Timbral width is a measure of the fraction of loudness that lies outside of the loudest band, relative to the total loudness. Thus, although the absolute level of loudness decreases in bands other than the loudest, the relative level of loudness in the loudest band compared with that in other bands remains static.

Timbral weight (*a* - *b*) lies below 0.0 during the starting transient and for the first 0.4 seconds; it rises sharply, then, except for a couple of peaks, shows a smooth descent, till towards the end of the sound it rises in a leap once again. In other words, the timbre is relatively heavy at the beginning, and becomes lighter after the first 0.4 seconds (there is now a greater fraction of loudness above the loudest band than below). The timbre then becomes progressively heavier (the upper partials becoming weaker), until a few seconds from the end a lighter timbre returns.

Timbral pitch (figure 2.4) drops by 4 bands after the first 0.4 seconds, and then drops again by 1 band just prior to the 3rd second. Again, let us compare these findings, along with those on timbral weight, with the spectrograph, in order to confirm these measurements.

After briefly touching on band number 33, timbral pitch lies at band number 32 for the first 0.4 seconds, and then drops to band number 28. Band number 32 centres on 1600 Hz (a sharp G^6), and has a band width stretching approximately from 1415 Hz to 1785 Hz. Band number 28 centres on 630 Hz (a sharp $D\#^5$) and has a band width stretching approximately from 557.5 Hz to 702.5 Hz.[13] On the spectrographic analysis shown in figure 2.2 it is difficult to pick out particular frequencies. Figure 2.5a shows a spectra analysis of the A-weighted bell tone, taken at frame number 2 (92.88 m.s.), not long after the loudest band has changed from band number 33 to 32. There is a very strong peak at 1507 Hz, which falls into band number 32. The peak at 603 Hz falls into band number 28. The peaks either side of these peaks do not fall into the same bands. Below the graph is the value that Stevens' procedure gives for the sone value for each of these bands.[14] Figure 2.5b is taken at frame number 7, very near to where the loudest band moves to band number 28. The peak at 1507 Hz is now greatly reduced, while that at 603 Hz is only slightly reduced in amplitude when compared with frame 2. The Stevens' method shows band 32 reduced in sone value, while band 28 has the same loudness in sones. Figure 2.5c is taken at the change to band number 28, at frame number 10. Here we can see that the peak at 1507 Hz has continued to decrease in amplitude so that it is now below that at 603 Hz. The results

---

[13] See Appendix for a complete list of 1/3 octave band numbers.

[14] These sone values are calculated from the FFT's values before the subjective sone value is incorporated into the calculations (see note 12) - the sone values are only valid relative to one another.

from Stevens' method reflect this.[15] Relating these findings to the spectrograph in figure 2.2, the peak at 1507 Hz (band 32) is the fifth partial (which at its beginning has the colour white), and the drop to band 28 (which includes the peak at 603 Hz - the third partial) occurs just before the line representing partial 5 changes to an obviously darker grey. Relating timbral weight to the spectral analyses, it will be remembered that timbral weight was shown to be relatively heavy during the first 0.4 seconds. This is supported by figure 2.5a which shows more spectral energy lying below the frequency spike with greatest amplitude than above it. Just after 0.4 seconds, the timbre was shown to become lighter. Once more, this is supported by figure 2.5c, which shows more spectral energy above the spike of greatest amplitude than below it.

The difference between a heavy and a light timbre is often one of a sense of definition or 'crispness' in the lower frequencies. In a situation of a light timbre (when there is very little energy below the loudest band) the lower frequencies can sound muddy and ill-defined. This is the case in the bell tone where, apart from the first two and last two seconds of the sound, the graph of timbral weight shows a steady decline from lighter to heavier. This follows the characteristics of the sound as it departs form its initial 'clang' towards a sound with a more definite sense of pitch and a more definite sense of 'bass.'

Sudden changes in the plot of timbral weight, *a* - *b* (figure 2.3), are due to shifts in the frequency position of the loudest band (figure 2.4). For example, the sudden rise after

---

[15] It is important to remember that, while it has been shown that the results from A-weighting and Stevens' Mark VII Procedure are similar, they do not produce the same results. Stevens' method shows that in frame 9, band 32 still has a greater sone value than band 28, whereas the A-weighted spectra analysis shows the peak at 1507 Hz (band 32) to be slightly below the peak at 603 Hz (band 28).

0.4 seconds corresponds to a drop in the position of the loudest band from band-number 32 to 28. Similarly, the two upward peaks that interrupt the steady descent correspond to the two closely-spaced downward peaks in the graph of the loudest-band position. No satisfactory damping method has been found so that audibly insignificant shifts can be filtered from the results. Comparing LD-graphs with the auditory experience, however, shows that momentary shifts in timbral weight (such as occur just after the 2nd second) can be successfully interpreted within the overall trend of the graph. It is important to note that although shifts in timbral pitch and weight often reflect each other, they also function independently.

Similarly to the bell, the timbral width of the guitar tone shows a steady decline, the timbre gradually becoming more focused. The exception lies in the starting transient, where the timbre is initially more diffuse. Note the oscillation of timbral weight during the second half of the sound. This corresponds to an audible 'wow' as the fundamental oscillates in intensity, and as timbral pitch moves back-and-forth between band-numbers 24 and 28. The oscillation of the fundamental's amplitude can also be seen in the spectrograph (figure 2.2).

If readers care to imagine the sound of a typical church bell and the sound of a plucked guitar string, they will find that a significant part of their timbral experience has been described in the terms set-out above. Thus, we are establishing a timbral vocabulary.

Let us take as our next example an LD-graph of a section of ensemble music that has a very distinct change of timbre (figure 2.6). The central point of György Ligeti's *Atmosphères* (written in 1961) consists of a very high-pitched note-cluster played by four piccolos (pitch-class set 4-1, G^7 to Bb^7) followed by a very low-pitched cluster played by eight double-basses (pitch-class set 8-1, C#^{1} to G#^{1}).[16]

This graph consists of two sets of lines. The upper set represents timbral width, and the lower set represents timbral weight, which is similar to the graph in figure 2.3. The thicker line in both sets joins data points 0.464 seconds apart (an average of 10). The thinner line joins data points 0.0464 seconds apart (an average of 1). Thus, we can observe both detail structure, and larger-scale movement. Wherever averaging is found in an LD-graph, the averaging process occurs at the very beginning of the calculation procedure, at the point loudness levels in 1/3 octave bands are calculated over a time interval that is a multiple of 0.0464 seconds (the smallest time-interval available within the limits of the software). In this instance, the averaging multiples are 1 and 10. From a perceptual viewpoint, it was considered better to average over the sound that we hear, rather than average over the results at the end of a complex series of calculations Letters on the graph enclosed in square brackets indicate the placement of the rehearsal letters in the score. The relevant section from the score is reproduced in figure 2.8.

Up to the point on the graph before the 260th second, where the timbre suddenly becomes more diffuse and heavier, the piccolos, oboes, clarinets and trumpets have been playing, and becoming increasingly louder and higher in pitch. This has caused the timbre to become increasingly narrower and heavier. At the 248th second, all wind instruments, but the piccolos, stop playing, while the piccolos continue to rise in pitch and increase in dynamic. At the 250th second, the value for timbral width

---

[16] The recording used for these measurements is the C.D. *Wien Modern*, Deutsche Grammophon 429 260-2 (1990): Wiener Philharmoniker, conducted by Claudio Abbado. The same recording is used in Chapter 4 of this thesis.

reaches its lowest point in this extract, showing that most of the loudness energy is now concentrated in the loudest 1/3 octave band (variable *l*), and the value for timbral weight is very close to zero, showing that what energy there is outside of the loudest band is distributed almost equally above and below the loudest band. We hear a sound similar to a cluster of sine-waves.

A moment before letter [G], there is a sudden drop in timbral weight, plus a rise in timbral width. This is the point of the very beginning of the double bass entry. A bit over a second later, timbral weight rises once again. What is happening here?

Figure 2.7 shows timbral pitch during this section (the movement of the loudest 1/3 octave band). Here we can see that timbral pitch has been rising throughout this extract, and for virtually all of the piccolo section (from the 248th second) the position of the loudest band has been static - band number 35, which is centred on 3150 Hz (a sharp G^7), the region of the piccolo fundamentals. Then at 260.35 seconds it suddenly drops to band number 21, centred on 125 Hz. (a flat B^2), which is in the frequency region an octave-and-a-fifth above the fundamentals that the double-basses are playing - the region of the third harmonics. This is at a point around half a second after the double-basses have entered. This evidence suggests that for about half a second after the double basses enter, the piccolos are still heard as the loudest sound. After this, the piccolos are no longer the loudest sound - in fact, they are no longer heard at all.

From this information, we can say that for about half a second after the double basses have entered, the loudest band still lies in the region of the piccolos' note-cluster, and the amount of loudness energy below this band has increased greatly due to the entry of the double-basses - thus, timbral weight suddenly becomes heavier, and timbral width becomes more diffuse. When the piccolo fundamentals cease to be the loudest sound and are replaced by the third harmonics of the double-bass fundamentals, timbral pitch drops, and timbral weight becomes suddenly lighter, as there is now a lot of loudness energy above the loudest band. Timbral width remains unaffected as the *amount* of energy outside of the loudest band has not changed.

What has just been discussed describes, in an overall way, the way we hear the timbre change at this point. At first we hear the *focused* sound of the piccolos; then the timbre becomes suddenly much *heavier* (more 'solid') and *more diffuse* as the double basses overlap with the end of the piccolos' part ; then we settle in to the richness of the double bass sound, with its correspondingly *lighter* timbre (less 'solid'). Timbral *heaviness* is associated with a solid, well-defined sound. What the graph cannot show is how we perceive the piccolo timbre as a separate entity from the double-bass timbre. The graph can only show a representation of our 'overall' timbre perception.

3.2.1 A Note on Timbral Weight

Because timbral weight ($a - b$) is so very strongly influenced by movement of the loudest band (timbral pitch), it often shows a high degree of instability. Any movement in timbral pitch is very likely to alter a-b, but it can leave timbral width ($a + b$) unchanged. Because of this, timbral weight needs to be interpreted with caution. It is most reliable when timbral pitch is static - any change in timbral weight is then due to changes in the skew of the loudness energy around a fixed point. It is most difficult to interpret when timbral pitch shows a high degree of fluctuation - the often rapid changes in timbral weight are simply mirroring the fluctuations in timbral pitch. When interpreting the graph of timbral weight, long term tendencies need to be looked for, rather than small-scale changes, and results need to be compared with movements in timbral pitch.

3.3 Sharpness Analysis

In the above section, we saw how LD-analysis is based on locating the loudest 1/3 octave band, and that this position is interpreted as timbral pitch. In this study, sharpness measurement consists in locating the 1/3 octave band that corresponds with the loudness centroid. Thus, we have two similar processes. One locates the loudest band through straightforward comparison with all other bands; the other locates a band that lies 'centrally' in comparison to the overall spread of loudness. Why was the loudest band chosen as the basis of LD-analysis, rather than the band that represents the loudness centroid?

Figure 2.9 shows an LD-graph of the same two bell and guitar tones that were analysed above, but this time measurements are not based on the position of the loudest band, but on the position of the loudness centroid band. Figure 2.10 shows the position of the loudness centroid band. It was demonstrated that the LD-graph (based on the loudest band: figure 2.3) was a good visual representation of the way we hear the changing timbre of these two tones. The LD-graph of figure 2.9 (based on the loudness centroid band) loses virtually all of the elements that were timbrally descriptive in figure 2.3. The line representing timbral width ($a + b$) now exhibits large jumps - this is especially evident in the guitar tone. These jumps in no way correspond with perception. Timbral weight ($a - b$), which has demonstrated some problems for interpretation, now becomes more problematic by losing those aspects that gave us insight into the sound - especially note that the 'wow' effect that was reflected in the timbral weight of the guitar in figure 2.3, is now lost in figure 2.9. Why do these changes occur? Firstly, the position of the loudness centroid band will, for the most part, be different from that of the loudest band; thus, the two graphs must be different. Note that in places where the loudness centroid band and the loudest band share the same band number (for example, the section in the bell tone from 1.4s to 2s) the two graphs are exactly the same. Secondly, if the centroid band is not the loudest band, then timbral width will be greater when the centroid is used as reference. This is demonstrated if we compare the graphs of the guitar tone. In figure 2.9 we see timbral width for the guitar tone oscillating a number of times. The bottom of the first large oscillation measures ca. 0.45 and lasts from 5.1s to 5.2s. On figure 2.10 we find that during this same period, the loudness centroid band lies at band number 27. Turning to figures 2.3 and 2.4 (loudest band measurements), we find that over this same period the measurements are exactly the same. At 5.3s, the loudness centroid band has risen to lie at band number 28, and timbral width in figure 2.9 has leapt to 1. On figure 2.4, the loudest band still lies at band number 27, and on figure 2.3, timbral width has continued its steady fall. From this we can see that when the loudness centroid band leaves the position of the loudest band, this allows the loudness energy present in the loudest band to greatly influence the measurement of timbral width (which is the loudness energy outside of the loudest / loudness centroid band).

Figure 2.9 has been shown to be less useful than figure 2.3. However, figure 2.10 is a very useful measure of changing sharpness. The sharpness of the bell slowly falls, while that of the guitar rises towards the end of the sound. The graph of sharpness is much less angular than that of timbral pitch (figure 2.4). Perceptually, timbral pitch and sharpness measure different aspects of timbre. Timbral pitch seems to measure the placement of the centre of gravity of a sound, whereas sharpness measures the changing overall 'focus' of the sound. The usefulness of each measure will become clearer during the course of examples in later chapters.

Figure 2.11 shows the position of the loudness centroid for the central section of *Atmosphères* that was discussed above. If we compare this graph with that shown in figure 2.7, we can see the characteristic 'smoothing' that occurs when centroid calculations are used compared with loudest band measurements - the graph is subject to fewer fluctuations, and the overall range is more restricted. Sharpness is a very useful perceptual measure in music timbre analysis.

3.4 Roughness Analysis

Roughness, or acoustic dissonance, has been defined as the beating that occurs between simultaneously sounding partials.[17] Roughness is not synonymous with musical dissonance - what is musically dissonant or consonant is defined through the practice of the era. Roughness is a perceptual measure inhabiting a world larger than that of musical practice.

A measure for roughness has been proposed by Plomp and Levelt.[18] Through psychoacoustic testing, they determined a curve that shows the degree of perceived dissonance of two simultaneously sounding sinusoids of equal amplitude as a function of critical bandwidth (see example 2.2, below). When two sinusoids have a pitch difference equal to or greater than 1.2 times the critical bandwidth (a similar measure to the one-third octave frequency band), the dissonance factor is 0. A difference of roughly 0.25 times the critical bandwidth produces maximum dissonance (arbitrarily given the value 1), and as the sinusoids become even closer together, the dissonance factor rapidly decreases.[19]

---

[17] See Helmholtz (trans., 1954), and Plomp and Levelt (1965).

[18] Plomp and Levelt (1965).

[19] Terhardt (1974) proposes a method whereby roughness of amplitude and frequency modulated tones can be calculated as a function of the degree of modulation.

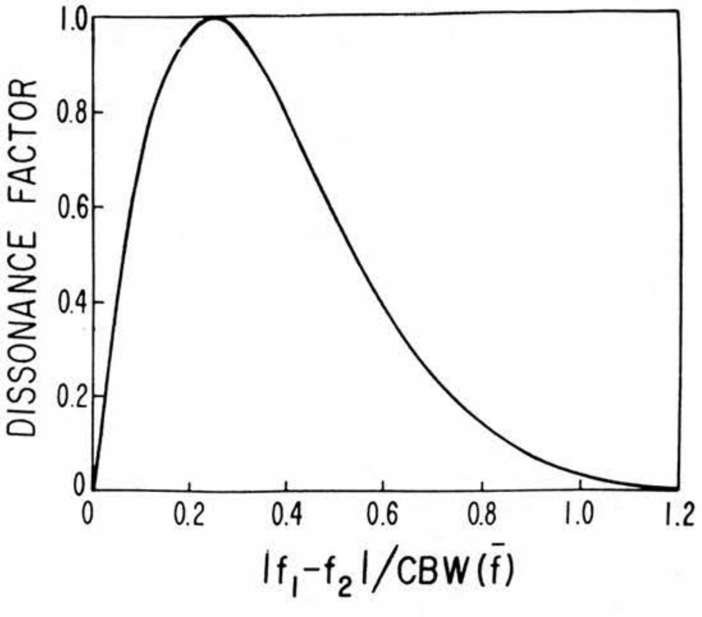

Degree of perceived dissonance of two simultaneously sounding sinusoids of equal amplitude as a function of critical bandwidth (after Plomp & Levelt, 1965).

Example 2.2

Sethares (1993) believes that "the consonance theory of Plomp and Levelt is probably the most important current consonance theory, but it is not uncontroversial."[20] Sethares uses the Plomp and Levelt dissonance curve as the basis for exploring synthetic scale forms that provide optimal consonance for a given timbre.[21]

---

[20] Sethares (1993): 1226.

[21] Sethares (1994): 10.

Kameoka and Kuriyagawa (1969a and b) propose a highly complex formula for the calculation of 'dissonance intensity,' based on their own extensive psychoacoustic testing that produces results comparable with those presented by Plomp and Levelt. As Plomp and Levelt's measurements are the basis for more recent work on dissonance measurement (cited below), and because Kameoka and Kuriyagawa's work has not been utilised by more recent researchers, and because their method is so highly complex, it was decided not to pursue their dissonance theory in this thesis.

Hutchinson and Knopoff (1978) have proposed a formula based on Plomp and Levelt's work for calculating the degree of beating between any number of sinusoids (and thus, between any number of complex tones). The roughness (or dissonance, as they call it) between *two* sinusoids may be calculated by:

$$D = \frac{A_1 A_2 g(f_1 - f_2)}{N}$$

where D is the dissonance factor, A_1, A_2 are the acoustic amplitudes of the two sinusoids, and $g(f_1 - f_2)$ is a weighting factor (derived from the work of Plomp and Levelt) which is a function of the difference in the two sinusoidal frequencies f_1 and f_2. N is a normalising factor proportional to the total intensity:

$$N = A_1^2 + A_2^2 .$$

The weighting factor, g, is calculated as:

$$g = |f_1 - f_2| / \mathrm{CBW}(\bar{f})$$

where the critical bandwidth (CBW) is calculated as a function of the mean frequency of two simple tones:

$$\mathrm{CBW} = 1.72(\bar{f})^{0.65}\ .$$

Example 2.3 shows the plot of this CBW curve.

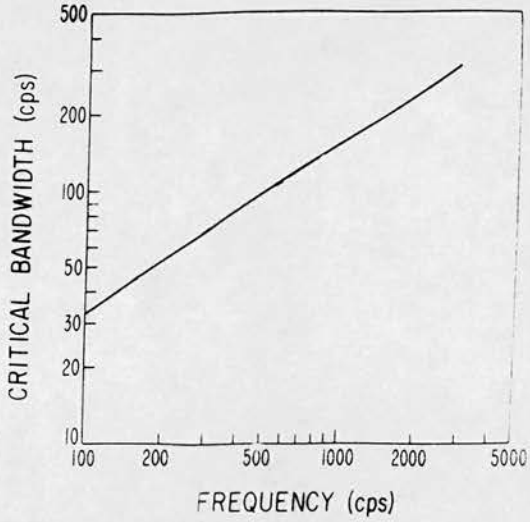

Critical bandwidth for two simple tones as a function of their mean frequency (from Hutchinson & Knopoff, 1978).

Example 2.3

The formula for calculating *D* can then be modified for *complex tones* so that the numerator becomes the sum of all sinusoid-pair interactions within that complex tone, and the denominator the total power of the sound:

$$D = \frac{1/2 \sum_{i=1}^{N} \sum_{j=1}^{N} A_i A_j g_{ij}}{\sum_{i=1}^{N} A_i^2}$$

As Hutchinson and Knopoff themselves say, this formula is not ideal, for it does not deal with the problem of phase shifts, nor does it address the problems associated with simply adding the roughness interactions of a sound that has very many components. However, the results reported by Hutchinson and Knopoff, as well as the results presented in this thesis, suggest that the formula is reliable when different sounds are compared.

Figure 2.12 shows a plot of roughness against time for the bell and guitar tones discussed earlier, calculated using a program that incorporates the formula described above.[22]

The graph confirms expectations. We see that the bell tone shows about twice the roughness of the guitar tone if the beginnings of the sounds are compared. Fluctuations in the graphs are due to the varying relative strengths of the sounds' components over the length of the decay.

Let us now return to the musical example taken from the centre of *Atmosphères*. Figure 2.13 shows a graph of Roughness against time for the same section of music represented in the LD-graph shown in figure 2.6 (the score for this section of the music was shown in figure 2.8).

As the wind instruments increase in dynamic and tessitura, the amount of roughness increases, moving from a roughness measure of around 4 at letter [F], to around 8 half way through bar 36 (marked as *x* on the graph), where, overall, the instruments are playing at their loudest and highest. At the point aurally where the piccolos are left to play on their own (marked as *y*), there is a sudden drop in the roughness graph, and the roughness measure continues to become lower as the piccolos continue to climb in pitch. The group of piccolo notes sounds increasingly like a group of sine

---

[22] The program was written specifically for this thesis.

waves. This is due to fewer partials of the piccolo notes lying in such a way as to produce a significant level of roughness. Note the slight increase in roughness during bar 39 (*z*), where the piccolos all crescendo together to *fff*. This corresponds to rising timbral width and weight that starts from the beginning of bar 38 (see figure 2.6). At the entry of the double-basses [G], there is a marked rise in roughness to a level equal to that heard at letter [F].

A Roughness measure is a very different but also complementary acoustic attribute from that shown by an LD-graph or by Sharpness analysis. Thus, the attribute roughness is a very useful addition to the vocabulary of timbre analysis.

3.5 Cepstrum Analysis

A measure of roughness involves a calculation of the degree of beating we perceive when listening to a sound. However, this quality of sound can be approached from another direction - we can analyse a spectrum through cepstrum analysis to ascertain how strongly the spectrum consists of harmonically related components. This type of analysis is not based on a perceptual model, but on a physical examination of the sound's spectrum.

Cepstrum analysis consists in calculating the FFT of the power-spectrum of a sound. Thus, the second FFT will show to what extent periodicity exists in the power spectrum. Figure 2.14 shows cepstrum analyses of white noise, the previously discussed bell sound, and the previously discussed guitar sound. The power-spectrum was calculated using a 4096 transform size on a waveform sampled at 44.1 kHz. The second FFT consisted of a transform size of 2048, this being the number of points that exist in each frame of the power spectrum.

Cepstrum analysis shows evidence of periodicity (harmonicity) through the existence of spikes. The higher the spike the stronger the periodicity associated with that spike. A series of spikes usually exists, lying equidistant from each other, which confirm the existence of the periodicity at 2x, 3x, etc., where x is the initial periodicity. The members of this family of spikes lie nearer to each other for high pitched tones than for low pitched tones. When a number of pitched tones sound simultaneously, each will have its own family of cepstrum peaks (*rahmonics*). If these pitched tones lie in a harmonic relationship (e.g., a major triad) then their harmonics will tend to share frequencies - and the rahmonics associated with the tones will share *quefrencies*. If these pitched tones form a pitch cluster with very little harmonic identity (e.g., pitch-class set 3-1), then the harmonics of each note will lie at slightly different frequencies, as will the families of rahmonics calculated from these tones. Thus, if in a cepstrum analysis of a sound we see a pronounced peak, followed by a family of equally spaced peaks, probably decreasing in amplitude, then it can be deduced that the sound has a high degree of harmonicity. If, on the other hand, we see a number of peaks, not equally spaced, or a number of competing families of peaks, then we can deduce that the overall sound has a low degree of harmonicity, even though it is comprised of individual tones that are harmonic.

Returning to figure 2.14, we see three contrasting cepstra. A spectrum of white noise shows no evidence of harmonicity - thus, its cepstrum (topmost graph) shows no evidence of a periodic structure. A spectrum of the carillon bell shows evidence of an harmonic structure, but with a strong presence of other partials. Thus, its cepstrum (middle graph) shows clusters of peaks, indicating that there is more than one set of harmonically related partials. A spectrum of the guitar tone shows strong evidence of a clear harmonic structure. Thus, its cepstrum (lowest graph) shows a clear set of peaks, with no other significant families of peaks present. The degree of harmonicity present in a sound is measured by the signal-to-noise ratio in its cepstrum, however no quantifiable measure is offered here beyond a visual examination of a sound's cepstrum.

The cepstrum examples discussed above are of discrete moments in time. In order to analyse a section of music, the type of display seen in the previous spectrum analyses will be used: amplitude is shown by a grey scale, thus freeing the horizontal axis for time. In cepstrum graphs, we are looking for the presence of sets of parallel lines (which represent a series of rahmonic peaks). If we see a set of lines (rahmonics) lying equally distant apart, we can deduce that the sound is perceived as relatively harmonic. If the set of lines show competing periodicities, we can deduce that the sound is perceived as relatively non-harmonic.[23]

Figure 2.15 is a spectrum analysis of eight stages in a re-synthesis of the carillon bell tone that has already been analysed above.[24] Original components of the bell tone are used. The stages of re-synthesis are as follows:

1. Hum note (251 Hz)
2. Prime or fundamental (501 Hz)
3. Minor third and fifth (603 & 750 Hz)
4. Nominal or octave (1005 Hz)

---

[23] This display for a cepstrum analysis is very sensitive to the setting for the dB colour scale in the software *Hypersignal*. This setting, whose value is given in the top left corner of the graph, determines the breadth and resolution of the amplitude scale. Too high a setting (showing a broad total dynamic range, with a relatively coarse level of amplitude resolution) produces a 'mess' of detail; too low a setting (showing a narrow total dynamic range, with a relatively fine level of amplitude resolution) will produce large 'blank' areas of cepstrum display. To find a dynamic level for the grey scale that shows clearly whether significant cepstrum peaks are present or absent can be difficult.

[24] Recorded from *Phillips* Auditory Demonstrations CD: demonstration 28.

5. Duodecime or twelfth (1506 Hz)
6. Upper octave (2083 Hz)
7. Next two partials (2421 & 2721 Hz)
8. Remainder of partials.

If 251 Hz were the first partial of a harmonic series, the first eleven elements of that series would be: 251, 502, 753, 1004, 1255, 1506, 1757, 2008, 2259, 2510, 2761 Hz. Let us compare this harmonic series with the tones that are present, and study the cepstrum analysis of the re-synthesis of the bell (figure 2.16).

When the hum note alone is present (1. 251 Hz), there is no evidence of a rahmonic family in figure 2.16. When the prime is added (2. 501 Hz), only 1 Hz lower than an exact octave, a strong family of rahmonics appears. On the addition of the minor third and fifth (3. 603 and 750 Hz), one non-harmonic tone, and one closely harmonic, the rahmonics disappear, as the strongly harmonic relationships have been disrupted. The Nominal (4. 1005 Hz) and the duodecime (5. 1506 Hz) both reinforce harmonicity, and the rahmonic family begins to appear once again (but with competing rahmonics also becoming apparent as the added tones also interact with the non-harmonic elements that are present). The upper octave (6. 2083 Hz) is a sharp 8th harmonic to 251 Hz (sharp by 75 Hz), so this addition both strengthens the members of the family of rahmonics, and adds to the background 'noise.' With the addition of the next two partials (7. 2421 and 2721 Hz), the first is inharmonic, the second harmonic, so, again, the rahmonic family is strengthened as is the 'noise.' With the addition of the remainder of the partials (step 8), which presumably consist of both harmonic and enharmonic elements, once again both the rahmonic family and the 'noise' are strengthened. Note that over the duration of the completely re-synthesised sound (step 8), it increases in harmonicity - the inharmonic elements die away more quickly than the harmonic elements.

Let us now see what cepstrum analysis can tell us about that same section of *Atmosphères* that has been analysed through LD-analysis, Sharpness analysis and Roughness measurement. Figure 2.17 shows a cepstrum analysis of that section from *Atmosphères*.

In the section a little after [F], we can see evidence of harmonicity. The dominant rahmonic family is marked by an asterisk (*). These rahmonics lie at an interval of 0.95 of a division. Knowing that the vertical axis represents 4.656 ms per division, we can calculate the fundamental that this rahmonic family represents:

Separation = 0.95 units;
0.95 units = 4.42 ms;
Frequency of fundamental = 1 / 4.42 = 226 Hz. = a slightly sharp A^3.

There are also other rahmonics present at this point on the graph - but they are not as strong as the family already discussed. As we move towards letter 'a', which is marked on the figure, a greater number of unrelated rahmonics emerge. This indicates that there are a greater number of non-harmonically related tones present, thus the overall level of harmonicity is now lower.

A little further along from 'a', there is a very strong family of rahmonics identified on the graph by 'x'. At this point a very strong single trumpet note is heard to emerge clearly from the mat-like texture, and then dissolve once more back into this texture. Note that the rahmonics are much closer together here than in the family nearer to [F]. This indicates that the trumpet's fundamental responsible for these rahmonics is higher than the fundamental indicated earlier. To calculate the trumpet's fundamental we do the following:

Separation = 0.3 units;
0.3 units = 1.397 ms;
Frequency of fundamental = 1 / 1.397 = 716 Hz. = a pitch between F^5 and $F\#^5$.

The actual pitch that is heard is $F\#^5$ (but see note 11 in Chapter 4). Given the method of calculation, this falls within an acceptable degree of error. In this study, cepstrum analysis is used to identify harmonicity, not to calculate the pitches that are present.

Shortly after this trumpet note, the piccolos are left to play on their own (except for a number of violins - inaudible on this recording). The piccolos' cepstrum pattern consists of a large number of very closely spaced rahmonics, indicating a very high, relatively harmonic sound. It will be remembered that the notes played here consist of pitch-class set 4-1 (G^7 to Bb^7). Because these notes lie very close together, there is only a small amount of roughness (see figure 2.13), and this measurement is supported by our observation of the degree of harmonicity. The sound heard is of a group of fundamentals that *almost* resolve themselves into a single pitch. What principally works against this tendency is the ringing combination tones (see Section 2, Chapter 1) that are audible, but not physically observable.

At letter [G] the piccolos cease, and the double-basses enter (pitch-class set 8-1, $C\#^1$ to $G\#^1$). The double-basses create a much higher level of roughness than the piccolos (figure 2.13). On the cepstrum graph, there is a small 'clump' of rahmonics quite high-up. That there is a 'clump' rather than a series of discrete rahmonic lines, suggests that this sound is non-harmonic. That these rahmonics are so high, with nothing significant below them, suggests that this rahmonic 'clump' is the first of a family of 'clumps', but that the other family members disappear off the top of the scale (it will be remembered that the further the individual components of a rahmonic family lie from each other, the lower the fundamental responsible for them).

It has been demonstrated, in comparison to Roughness analysis, that Cepstrum analysis offers a different and complementary method for the measurement of acoustic consonance and dissonance. While Roughness analysis is based on perceptual data, Cepstrum analysis is based solely on a physical examination of a sound's spectrum. Both tell us about an aspect of timbre complementary to the information obtained from LD-analysis and Sharpness analysis.

4. Conclusions

Four complementary methods have been presented for analysing music timbre - Loudness Distribution analysis, Sharpness analysis, Roughness analysis, and Cepstrum analysis. These methods present us with six aspects of timbre - timbral width, timbral weight, timbral pitch (all derived from LD-analysis), sharpness, roughness and harmonicity. The description of the movement of these measurable aspects of timbre constitutes the timbre vocabulary of this study. This timbre vocabulary allows for the first time a detailed description of a particular musical timbre to be made, and the *development* of this timbre to be followed closely. With the existence of a precise timbre vocabulary, it becomes possible to undertake rigorous timbre analysis - we have already seen small-scale examples of this in the analysis of the bell and guitar tones, and in the analysis of a section from *Atmosphères*.

The aspects of timbre that have been chosen for analysis have been determined by the findings of previous acoustic and psychoacoustic studies (Sharpness, Roughness, Harmonicity) and by the extrapolation from a pre-existing timbre analysis method (the Tristimulus Method leading to Loudness-Distribution analysis). This list of timbre dimensions will not cover the entire scope of our timbre perception, but, from the evidence of the examples given in this chapter, it appears that these aspects of timbre do account for much that we perceive as timbre.

While changes in timbre are important in practically all musical compositions, it is in music of this century that timbre at times becomes elevated to a structural importance equal to or superior to that of pitch structures. It is in this genre of music that timbre analysis techniques are especially needed - an analysis of pitch tells us only a part of the story (and often the less significant part) of how this type of music works.

It was, in fact, the dilemma of gaining analytical access to the timbre structure of works where timbre appears to play such an important role that originally motivated this research. Now that a plane of concepts for timbre has been established, along with a precise vocabulary so that movement in timbre can be described, it now becomes possible to analyse timbre in a musical composition with as much depth and rigour as is possible in the pitch domain. The next two chapters of this thesis comprise detailed analyses of two pieces where timbre appears to play a highly significant role in structuring the music: Chapter 3 consists of an analysis of the first movement of Lutoslawski's *Jeux vénitiens* and Chapter 4 consists of an analysis of Ligeti's *Atmosphères.*

***

CHAPTER 3

Analysis of *Jeux vénitiens*: movement 1

CHAPTER 3

Analysis of *Jeux vénitiens*: movement 1

(1961)

Witold Lutoslawski

---

1. Introduction

Steven Stucky describes *Jeux vénitiens* as the piece that begins Lutoslawski's third major style period.[1] Composed between 1960 and 1961, Lutoslawski uses for the first time the compositional technique of *limited aleatorism* - pitch and rhythm are defined for individual parts, but the exact vertical relationship between parts is not defined. As the composer says in the introduction to the score, "Each musician should play his part as if he were playing it alone." Thus, Lutoslawski retains control of the pitch, timbral and textural 'outline' of a section, but he allows freedom in the detailed structure that is created through the interaction of the parts.

---

[1] Stephen Stucky (1981): *Lutoslawski and his Music*. Cambridge: Cambridge University Press.

2. Pitch and Sectional Description

The first movement of *Jeux vénitiens* uses as its means of organisation the juxtaposition of two different sound-ideas in conjunction with two different methods for organising pitch .

The first sound-idea, which utilises the technique of limited aleatorism, consists of a lively, bubbling texture, at first played only by woodwinds. On each appearance of this texture (from now on, referred to as 'A-texture') an additional instrumental group is added (figure 3.1). The instrumentation of this texture thus proceeds as follows: woodwinds; woodwinds and timpani; woodwinds, timpani and brass; woodwinds, timpani, brass and piano. Thus, the A-textures show a textural and timbral progression.

The pitches of the A-textures are derived from three symmetrical pitch collections, each pitch collection associated with an instrumental group (figure 3.2). The woodwinds play a symmetrical 12-note collection, which is dominated by interval class (i.c.) 2, and the brass play a 4-note collection constructed from i.c. 1. This 4-note collection forms a symmetrical 16-note collection with the woodwinds, as well as combining with the 8-note collection played by the piano, constructed from i.c. 3 and 4, to form another symmetrical 12-note collection. When all three pitch collections come together, a very diffuse sounding 24-note collection results.

Between occurrences of the A-textures, there are quieter and much more static, through-composed sections for strings ('B-textures'), that are punctuated by brief textural and timbral 'exclamation points,' produced, for example, by short *tremolo* interjections, or short, sudden increases in dynamic. These exclamation points have a similar effect within the B-textures to the larger-scale effect of the alternation of the A- and B-textures - an effect of sudden, and sometimes surprising arrival. The pitches of the B-textures are organised in a very general way into expanding and contracting collections centred on the interval of a perfect fifth. The method of pitch organisation is far less rigorous than in the A-textures. The beginning of every section is marked by a single stroke from a group of percussion instruments (drums, glockenspiel and xylophone).

In this movement, two different types of texture are associated with two different types of pitch organisation. However, a description of these methods of pitch organisation does not tell us very much about how the music works - it is in the realm of timbre and texture that we find the strongest structural determinants. Therefore, to discover more about this piece, we need to turn to timbre analysis.

3. Timbral Analysis

Figure 3.3 shows an A-weighted spectrograph of the movement with the individual sections of the piece labelled along the time axis.[2] The spectrograph is produced from a C.D. recording of Witold Rowicki conducting the National Philharmonic Orchestra in Warsaw, recorded in 1962.[3] The 'bubbling,' energetic nature of the A-textures, and their additive instrumental construction, make a sharp visual contrast

---

[2] Sections that have an additive name in the score, for example, A+C+E (as shown in figure 3.1), are labelled simply E on the spectrograph. In the text, the sections are also referred to in this abbreviated way

[3] *W. Lutoslawski* (vol.2): *Venetian Games, Trois Poèmes d'Henri Michaux, Symphony No.2.* Witold Rowicki / Witold Lutoslawski conducting the National Philharmonic Orchestra of Warsaw; Jan Crenz conducting the Polish Radio National Symphony Orchestra in Katowice. *Polskie Nagrania* - PNCD 041 (1989).

with the thinner, discontinuous spectral lines of the much more static B-textures. The large dynamic difference between the sections is clear in the spectrograph, and the four percussion strikes at the very end stand-out against the silence that surrounds them. The percussion strikes that come between sections, though distinctly audible, are largely absorbed into the visual 'activity' of the A-textures as represented on the spectrographs. This applies especially in the later, louder sections.

Figure 3.4 shows an LD-graph of the whole of the movement - the music's sections are marked along the time axis. A surprising feature of this graph is that the different sections are not markedly different. Timbral width hovers around 0.65 until section [H], where it gradually becomes more focused. Except in sections [B] and [H], timbral weight appears to show no pronounced pattern, and, again, there is not the strong delineation between the A-textures and the B-textures that might be expected. Looking at figure 3.5, a graph of timbral pitch, we again do not see the strong differences between the A-textures and the B-textures that we might expect to see. Why is this?

At the beginning of this chapter, the A-textures were described as consisting of a 'lively, bubbling texture,' and the B-textures as 'quieter and much more static, through-composed sections for strings.' A very important difference in these contrasting types of music lies, as the labels suggest, in texture. Texture can be defined as consisting of the interactions of discrete musical events. The timbre of each discrete event is part of the interaction. Thus, timbre is a subset of texture. The A-textures consist of fast moving individual lines that create a complex web of counterpoint. The B-textures contain far fewer individual 'events' - sustained notes are much in evidence. So, because an important difference between these sections from *Jeux vénitiens* is their contrasting texture, as well as their contrasting dynamic, these facets will not be measured by a method designed to quantify differences in timbre. Indeed, the differences in texture and dynamic are best borne out by the spectrograph shown in figure 3.3.

So what of the differences in *timbre* between these sections? After all, the A-textures are played loudly by a variety of wind instruments, and the B-textures are played quietly by strings.

In Chapter 2, we saw that LD-analysis provides a good measure of timbre perception of a bell and a guitar tone. However, timbre is a multidimensional attribute, and in the first movement of *Jeux vénitiens* it appears that the timbral differences that we hear between the A- and B-textures are not readily evident in LD-analysis. This suggests that the differences may be more evident if other timbre analysis tools are employed.

3.1 Roughness Analysis

Figure 3.6 shows a roughness graph of *Jeux vénitiens*. Here we can see very clear differences between the A- and B-textures. Sections [A], [C], [E] and [G] (the A-textures) have a much higher degree of roughness than sections [B], [D], [F] and [H] (the B-textures). Moreover, the degree of roughness in the A-textures increases from section to section - as would be expected, as more instruments are added. This increase in roughness over the A-textures is seen more clearly in figure 3.7. Here, the A-textures have been spliced together (omitting the percussion strikes that mark the beginning of sections), and a spline fit and a straight-line average applied to the data points so that a better idea of the overall direction of roughness movement can be discerned. (In the following discussion of the A-textures we are interested only in overall effect so an average of 10 is used in these graphs - a data point every 0.464 seconds.) An interesting pattern emerges. We can see that during sections [A] and

[E] roughness increases. Measuring from the straight-line average, during section [A], roughness increases from 5.5 to 6.9 (an increase of 1.4); during section [E], roughness increases from 6.7 to 7.8 (an increase of 1.1). Thus, not only do these sections both show an increase of roughness, but they also show a similar *degree* of increase in roughness. In contrast to this increase in roughness, sections [C] and [G] show over their length a small decrease in roughness (measuring from the straight-line average, section [C] shows a decrease of 0.2: from 6.4 to 6.2; section [G] shows a decrease of 0.1: from 7.7 to 7.6). Thus, a pattern emerges. Section [A] increases roughness from the very beginning of the piece; section [C] maintains a much more steady measure of roughness, keeping roughness at a level near to where section [A] finished (measuring from the straight line averages, section [A] finishes at 6.9, section [C] begins at 6.4); during section [E] roughness increases again; and section [G], like section [C], maintains a steady level of roughness at the level (this time matching exactly) established at the conclusion of section [E]. Thus, sections [A] and [E] are the 'roughness stairs,' and sections [C] and [G] are the 'roughness floors' to which the stairs lead.

This would appear to be an extraordinary achievement on the part of Lutoslawski. Through using a conventional orchestra, employing conventional playing techniques, he appears to have achieved an extremely high degree of control of sound - in particular, roughness. Let us examine the score to see what orchestration techniques are used.

Section [A] shows an increase in roughness. Yet, looking at the score of section A (figure 3.8; the music for section A is labelled [A][C][E][G]) we search in vain for evidence that shows any indication that Lutoslawski planned this increase in roughness. The parts, considered both individually and collectively, show no progression through tessitura, dynamic or style of attack that would suggest an increase in roughness. And, it is very difficult to say at any particular moment of the recording where in the score the instrumentalists are - "Each musician should play his part with the same freedom as if he were playing it alone" Lutoslawski instructs on this opening page. The same situation applies to section [E] (which consists of the sections in the score labelled [A][C][E][G], [C][E][G], and [E][G]). Here it is even more difficult to say where the instrumentalists are in the score as the instruction Lutoslawski gives for playing sections [C], [E] and [G] is "the individual parts ought not to be played from the beginning but from any other phrase between two caesuras." Given this condition, it would be impossible for Lutoslawski to plan any systematic increase in roughness over the length of a section. Indeed, the way in which Lutoslawski uses the technique of limited aleatorism shows that his aim is to create a section of music that remains 'static' in overall pitch, dynamic and timbral range, yet which has movement within set boundaries.

However, in this recording of the piece, there is an increase in overall roughness during sections [A] and [E] - an increase that appears to match very well the level of roughness shown in sections [C] and [G]. Why?

A particular recording of a piece is, obviously, only a single instance of many possible interpretations. Thus, the above observations on the movement in roughness in this recording of *Jeux vénitiens* might, or might not, be particular to this recording. It will be instructive - indeed analytically vital - to compare the findings on this recording with the findings on another recording. Thus, before attempting to answer the question posited at the end of the previous paragraph, let us look at another recording of *Jeux vénitiens* - a C.D. recording of Lutoslawski himself conducting the Polish Radio National Symphony Orchestra, recorded in 1977.[4] Hereafter, the Rowicki

---

[4] *Matrix 13 - Lutoslawski: Concerto for Orchestra, Jeux vénitiens, Livre pour Orchestre, Mi-parti.*

recording will be referred to as recording 1, and the Lutoslawski recording referred to as recording 2.

Figure 3.9 shows the roughness graph of the A-textures from recording 2. The differences between this graph and the roughness graph of the A-textures from recording 1 (figure 3.7) raise matters concerning not only differences in interpretation, but also differences in the 'sound style' of the recordings.

The A-textures in the second recording all show higher measures for roughness than those in the first recording. Indeed, to accommodate this, the roughness axis is calibrated from 0 - 8 for recording 1, but from 0 - 15 for recording 2. The reason for this is not in the style of playing, but in the 'sound style' of the recording. Recording 1 has a much duller and muddier sound than that heard in recording 2. Recording 1 is both older (1962, as opposed to 1977 for recording 2) and is produced by a different recording company (*Polskie Nagrania*, as opposed to *EMI*). Evidence for this duller sound can be seen when a comparison is made on an A-weighted spectrograph of section [A] from the two recordings. Figure 3.10 shows that there is significantly less energy in the higher frequency range in recording 1 than in recording 2, and this is typical of the situation found throughout the two recordings.

Looking now at the roughness structure of figure 3.9, we find both similarities and differences with figure 3.7. Sections [C] and [G] show a decrease in roughness in the straight-line average in both recordings. In recording 1, section [C] decreases in roughness by a very small amount (0.2). In recording 2, section [C] shows a decrease of 0.5 (8 - 7.5). Section [G] in recording 1 decreases in roughness by only 0.1 - in

---

Lutoslawski conducting the Polish Radio National Symphony Orchestra. *EMI Classics* - CDM 5 65305 2 (1994).

recording 2, section [G] decreases by 0.9 (11.3 - 10.4). Thus, the decrease in roughness is more pronounced in recording 2 than in recording 1. In recording 1, it was seen that the roughness of sections [A] and [E] both increased over their length. In recording 2, this increase is still seen in section [E], (an increase of 1.1 in recording 1; an increase of 1 in recording 2: 9.6 - 10.6) In recording 2 of section [A], there is a small *decrease* in roughness of 0.2 (8.2 - 8), as opposed to the sharp increase (1.4) in section [A] of recording 1.

What can we deduce from these results?

The roughness measurements of figure 3.9 are not as 'neat' as those of figure 3.7. When we first saw figure 3.7, there seemed to be evidence of extraordinary precision in Lutoslawski's manipulation of roughness, especially in the way in which sections [A] and [E] acted as ladders to the floors of sections [C] and [G]. Yet, when we examined the score, we could find no evidence that an increase in roughness in sections [A] and [E] was planned by the composer. When we examine another recording of the work, conducted by the composer, we find that the roughness structure suggested by figure 3.7 is only partly supported. In recording 2, section [G] still has a higher level of roughness than section [C], and section [E] still shows an increase in roughness (though the end of its straight-line average does not match so closely the beginning of section [G]). As has been mentioned, the roughness of section [A] is very different in the two recordings. In the second recording, there is no sense of this section acting as a 'ladder' to section [C].

Both recordings show that roughness generally increases through the piece as more instruments are added on each appearance of the A-textures. In section [C], three timpani are added. However, in both recordings, these instruments appear to add nothing to the roughness measurements in comparison with section [A] (due, perhaps, to the timpani's overall 'duller' sound). In section [E], three trumpets are added. These instruments make a distinct contribution to the degree of roughness (due, perhaps, to their overall 'brighter' sound). In section [G], a piano is added, which adds a small amount to the roughness measurements. Overall, when the two recordings are compared, it appears that sections [A] and [C] lie at one level of roughness, and section [E] lifts the piece to a higher level which is continued in section [G].

This still leaves us with the question of why sections [A] and [E] in recording 1, and section [E] in recording 2, all show increasing levels of roughness over time, when there is no evidence in the score as to why this should be.

A possible theory is that in recording 1, at the start of the piece (section [A]), the musicians carefully observe the dynamic and other markings in the music, possibly even playing a little tentatively. Because of the nature of the individual lines, however, which contain wide leaps and involve complex rhythms, as the music progresses the musicians begin to play more loudly and with stronger articulations. This tendency is encouraged because there is no obvious balance in the music (for example, melody and accompaniment) for which the individual musicians can aim. Possibly with the composer as conductor, this tendency is checked more carefully. This same effect appears to happen with the brass players in section [E] (this time seen in both recordings). Why do we not see this effect in sections [C] and [G]? The music that is added in these sections does not lend itself to the type of 'performance development' described above. The addition in section [C] is for one player playing three timpani. The rhythms are relatively simple, and the music is uniformly marked *piano*, except for one brief phrase marked *mf pp cresc. mp* - very straightforward dynamics when compared to what is required from individual instrumentalists in section [A]. The addition in section [G] is of a piano (two players), again playing relatively straightforward rhythms, and in a dynamic constantly *mezzo forte*, except for two instances of *sforzando* at the very beginning of the section. But what of the sections that are repeated, and that have already shown a tendency for increasing roughness? Dynamic and articulatory instructions are not precise, but cover a generally agreed range of possible interpretation. It is possible that when it comes time for the musicians to play a section of music for a second time, they have already pushed the dynamic and the forcefulness of the articulation as far as their musicianship (and the conductor) will allow them. All this is, of course, merely supposition, and it is not possible to confirm it by 'listening in' to individual performers in the recording, because of the muddiness of the recording. However, my own experience as an orchestral musician and conductor supports this explanation.

Examining the roughness structure of the A-textures, separate from the B-textures, has made it easier to uncover a roughness structure in the A-textures. Let us now examine the roughness graphs of the B-textures.

The B-textures are sections for strings. Changes in roughness are due to changes in playing technique, changes in tessitura, and changes in instrument combination. Figure 3.11 shows roughness measurements for the B-textures of recording 1, with straight-line averages. Unlike the analysis of the A-textures, for the B-textures we are interested in locating individual events (remember that the B-textures are through-composed), so the data-averaging here is 5 - there is a data point every 0.232 seconds. It was found that an averaging greater than this caused the individual musical events to be 'ironed out' in the graph, and an averaging less than this introduced too much 'noise' (or jitter) in the graph for the musical events to stand out clearly from the background.

Looking at figure 3.11, a definite roughness 'shape' becomes apparent from the straight-line averages. (A straight-line average is not calculated for section [F] as there are too few data points for a meaningful result to be obtained.) Roughness increases in section [B]; it continues to increase over section [D]; in the brief section [F], roughness appears to lie at a point roughly equal to that reached at the end of section [D]; and during section [H] roughness gradually decreases. What are the reasons for these changes in roughness level? The changes within sections [B] and [D] are largely due to changes in tessitura. Figure 3.12 shows a plot of the pitch structure of section [B]. The coloured ellipses represent the pitches that are sounding. If an ellipse is off-set from the centre of a bar, it indicates that the particular pitch to which it refers only appears during the first or second half of the bar. It was said at the opening of this chapter that the pitch structure of the B-textures is loosely based on the idea of expanding and then contracting collections beginning and ending with the interval of a perfect fifth. Here, we see that the overall movement is upwards. With instruments playing in a higher tessitura at the end of the section than at the beginning, there is an increase in the degree of roughness. The other element that increases overall roughness as the section progresses is the greater density of occurrences of staccato notes, played with *ricochet* bowing (see the score of this section, reproduced in figure 3.15).

The pitch plot of section [D] shows a similar rise in overall tessitura during the course of the section (figure 3.13). Instead of *ricochet* staccato notes, in this section the timbral 'exclamation marks' consist of short *mf* bursts of *sul ponticello tremolo* bowing. Like section [B], occurrences of these 'exclamations' come closer together towards the end of the section (see figure 3.15).

Section [F] consists of a brief sounding of a symmetrical eight-note collection (shown in figure 3.13). There are no timbral 'exclamation marks' in this section.

Section [H] is the longest of the B-textures, and over its length the level of roughness gradually decreases. The pitch-plot of this section (figure 3.14) shows an initial drop in tessitura (that coincides with the trough in roughness level around 70 seconds). As the pitches gradually step higher once more, instruments begin to drop out, so that by the end of the section (which is also the end of the first movement) there are only two violins left playing. The 'exclamation marks' in this section, which consist of brief *sul ponticello* notes, played *tremolo*, and harmonics played pizzicato (all played by the cello), and single *mf* semiquavers within the context of sustained notes played *pp*, remain evenly spaced through the section rather than increasing in density. All these factors contribute to the overall decrease in roughness.

In recording 1 of the B-textures, we have observed an increase in roughness over the first two sections (see figure 3.11), the maintenance of that roughness level in the brief third section, and a decrease in roughness in the last section. This may indicate a very structured approach to the changing level of roughness over these sections.

Figure 3.16 shows the roughness levels of these same sections as measured from recording 2. Like recording 1, the straight-line average in section [B] shows an overall increase in roughness. Sections [H] and [F], too, shows similar tendencies to that seen in recording 1 - the overall degree of roughness in section [H] decreases during its length, and the level of roughness in section [F] appears to be very close to that at the end of section [D]. Thus, roughness tendencies for sections [B], [F] and [H] agree well between the two recordings. The obvious difference between the recordings is the straight-line average for the roughness measurement of section [D]. Instead of it increasing at roughly the same rate as in section [B] (which we saw in recording 1), it shows a large increase in roughness at the start of the section, followed by a rapid fall towards its end. Why is this?

It will be remembered that an important difference between the two recordings is the presence of a greater amount of energy in the higher frequencies in recording 2 (refer back to figure 3.10). Another important difference is the 'closeness' with which instrumental solos are recorded, and the related phenomenon of a much greater dynamic range in recording 2. Recording 2 brings all solos (which includes the 'exclamation marks') much more to the foreground, especially the violin solo of the repeated E^6's at the end of section [B], and in the first half of section [D] (see figure 3.15). The effect of these closely-recorded repeated E^6's on the roughness measurements in recording 2 can be clearly seen in figure 3.16. The entry of the solo violin at the end of section [B] causes the degree of roughness to increase dramatically at around the 22nd second (marked 'a'). In comparison, the same moment in recording 1 (figure 3.11) (at around the 24th second; also marked 'a') shows a rise that is comparatively far smaller. At the beginning of section [D] in recording 2, the entry of the repeated E^6's coupled with the 'exclamation marks' causes the roughness measurements here to be relatively high (marked 'b'). The corresponding points in the first recording show increases, but on a much more modest scale. Thus, because of the differences in recording technique, section [D] in recording 2 shows a decrease in roughness over its length, while the same section in recording 1 shows a small increase.

Is it possible to say which recording is a better representation of the piece? Which is more realistic? Both recordings have their strengths and faults. The sound of recording 1 is dull, and somewhat 'blurred,' yet the balance between instruments is realistic; the sound of recording 2 is much brighter, and distinction between instruments is good, yet at times the sound is too 'sharp,' and individual instruments are too far forward. This state of affairs gives us very different roughness measurements in section [D] - but it does not appear to influence the roughness shape of the other sections. In section [D] I feel that recording 1 gives a more reliable roughness graph than recording 2. The sound of the solo violin playing E^6 is more realistically balanced, and thus its influence is more realistically represented in the roughness measurements of figure 3.11.

Having examined the roughness structure of the A- and B-textures in two different recordings, the following conclusions can be drawn:

1) The A-textures increase in roughness during the course of the movement. Sections [A] and [C] have similar roughness measurements; section [E] shows increased roughness (plus increasing roughness over its length in both recordings, though this is *not* supported by an examination of the score); and section [G] shows, on average, a higher level of roughness than any previous section. These findings are summarised below:

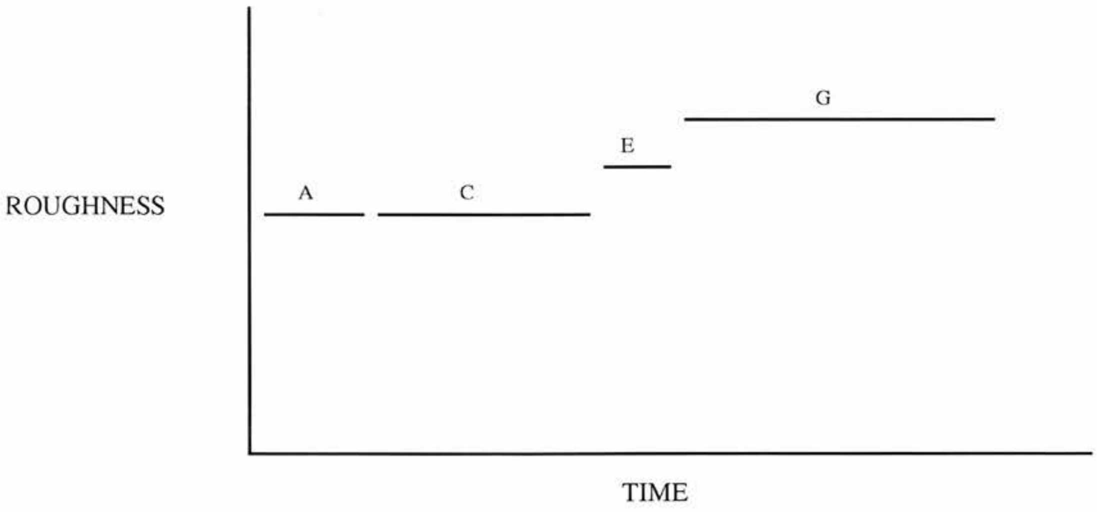

2) The B-textures show an increase then a decrease in roughness during the course of the movement. Sections [A] and [D] both increase in roughness over time

(despite the evidence of recording 2); section [F] appears to have a level of roughness approximately equal to that of section [D]; and the roughness of section [H] decreases over time. This is summarised below:

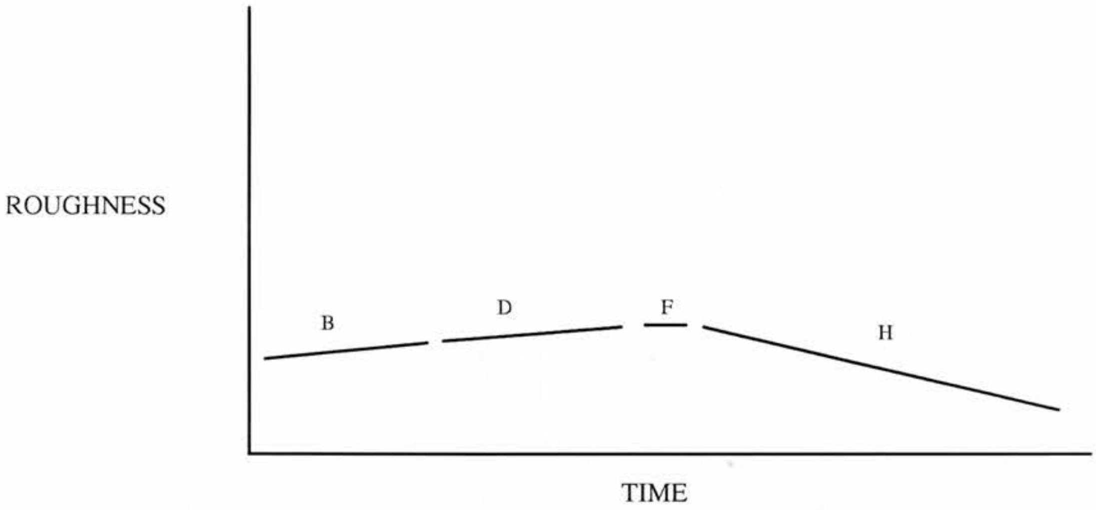

Roughness is an attribute of this piece that appears to have been controlled by the composer. The above roughness contours show how roughness contributes to our feeling of shape in the first movement. As the A-textures progress, we hear the texture become more dense, the level of activity increase, and the timbre become rougher; the A-textures are about the process of addition and growth. As the B-textures progress, we hear the character and the density of the 'exclamations' change, an important affect of which is to shape our perceived degree of roughness; the B-textures are about the process of controlled shaping, with a sense of beginning, middle and ending.

Figures 3.17a and 3.17b show the same roughness graphs as we saw in figures 3.11 and 3.16 respectively. Precise indications are now given on the graphs of where the 'textural exclamation points' occur, in the form of a bar number and below it the name of the instrument that plays. For example, '3 vl I' in section [B] indicates the demisemiquavers played *ricochet* in bar 3 by viola I; '103 vn II' in section [H] indicates the crotchet played *tremolo* in bar 103 by violin II. Thus it is possible to follow these graphs of roughness in detail through with the score (figure 3.15). It is clear that there are many increases and decreases in the degree of roughness that are not associated with any notated event in the score. For example, at the beginning of figure 3.17a (recording 1) the increase in roughness associated with '3 vl I' has a smaller roughness value than that of much of the music that comes before it, which consists of a sustained low-pitched cluster, but which is not performed completely smoothly; '10 vn II' in section [H] causes a rise in roughness, but the degree of roughness keeps climbing because the background roughness level increases at the same time.

For the most part these two graphs are self explanatory, however there are a couple of interesting anomalies that are worthy of comment. Near the end of section [B] the pitch E^6 appears as a sustained note in the violin. In both recordings this causes a pronounced increase in roughness. At the beginning of section [D], this same pitch appears, played in the same way, and marked with the same dynamic in the score. However, the sudden increase in roughness that appears in both recordings just after the start of this note is caused by '102 vc II,' not E^6. This can be seen if the portion of the roughness graph between the start of the first E^6 of section [D] and the entry of '102 vc II' is examined. During this part of the graph the degree of roughness is much lower than that associated with the appearance of E^6 near the end of section [B]. A reason for this appears to be that, in both recordings, the E^6 at the start of section [D] is played slightly more quietly and with less intensity than the E^6 in section [B]. However, the degree of reduction in roughness seems to be much greater than the perceived degree in reduction in volume and intensity suggests. An additional reason for the reduced roughness level could be that the harmonics of the E^6 in section [D] do not set up as much beating with the harmonics of its chordal background as happens with the E^6 in section [B] - but this is only speculation. In both recordings the E^6 at the start of section [D] sounds more integrated (less harsh) than the E^6 in section [B], and this perception is borne out in the graphs. Another curiosity is the way in which the *very* beginning of the pizzicato, which is '11 vc I' (shown on figure 3.17a as lasting for the duration indicated by the two dashed lines), decreases the degree of roughness, before roughness increases during the sounding of the *pitch* of the plucked note. This can indeed be heard - the initial 'dull thud' of the very beginning of the pizzicato sounds less rough than the immediately preceding texture, and the 'ring' of the pizzicato increases the degree of roughness. An important point that these annotated graphs of roughness bring out is the role of the interplay between the particular moments of changes in roughness as shown on the graphs, and the extent to which we perceive these as particular moments, or as events that become subsumed into our perception of an overall degree of roughness. The roughness graphs cannot show this, and an answer to this problem is beyond the scope of this study. However, the play between the particular perception, which may be perceptually weighted by dramatic effect, and the overall perception, which will be influenced by memory, is a constant factor in the interpretation of these graphs of timbre perception.

3.2 Sharpness Analysis

(and some further discussion of Roughness)

Figure 3.18 shows the position of the loudness centroid by 1/3 octave band number for the whole of the first movement as heard in recording 1. In other words, this graph shows a measure of sharpness. When interpreting the graph, it is important to realise that an oscillation in the graph (that is, when the graph moves back and forth between two adjacent band numbers) is most likely to signify that the true loudness centroid lies *between* two adjacent bands, and that very small frequency and/or amplitude changes are causing the position of the loudness to oscillate back and forth. In the sharpness graphs that follow figure 3.18, two different degrees of data averaging are displayed - the thinner line represents an average of 5, and the thicker line represents an average of 10 - thus, both overall trends and smaller scale details can be seen. Unless otherwise stated, in the following discussion the position of the loudness centroid is taken from the thicker line.

Let us begin our examination of sharpness by briefly looking at each section in turn in figure 3.18. In Section [A], the centroid oscillates between two band numbers. In section [B] the centroid climbs from a lower to a higher band. Section [C] sees the centroid oscillating once more between two bands. In section [D] the overall movement of the centroid is from a higher to a lower band. In section [E], the centroid remains static. In section [F], the centroid also remains static (this is evident if this section is viewed in greater detail - the apparent rise in the position of the loudness centroid in this section is caused by the influence of the data either side of this section when averaging takes place). In Section [G], the centroid once again oscillates between two bands. And in section [H], the overall movement of the centroid is first to fall, and then to rise. The large centroid movements at the end of section [H] represent the final four percussion strikes that end the movement.

The movement of the loudness centroid in recording 1 can be summarised as follows:

Position of loudness centroid

(recording 1)

A-textures	B-textures
Oscillates between two band-numbers, or centroid remains static.	[B]: rises [D]: falls [F]: remains static [H]: falls - rises

When we examined roughness, it was found that the A-textures formed a roughness structure, and the B-textures formed a different roughness structure. From the above table, it appears that sharpness also forms separate structures in the A- and B-textures. Because of this, the A- and B-textures will, once again, be examined in separate groups.

3.2.1 The A-textures

Figure 3.19 shows the position of the loudness centroid for the A-textures (sharpness), minus the percussion strikes between each section. The graph confirms the initial statement that the A-textures in this recording generally show an oscillating or static position for the loudness centroid. The information in this graph can be summarised as follows:

Position of loudness centroid
(recording 1):

A-textures

[A]	29/28
[C]	28/27
[E]	28
[G]	29/28/27

The single peaks represented by the thinner line (average 5) in section [A] (reaching to band number 30) and in section [E] (reaching to band number 29) are not considered significant. The movement to band number 29 in section [G] is considered important because of the number of times it occurs, and because the thicker line also moves there.

There appear to be no significant directional movements of the loudness centroid within the A-textures. There is, however, a trend over the duration of all the A-textures for the loudness centroid to move downwards - the beginning of section [A] has the loudness centroid largely in band 29, while the end of section [G] sees the loudness centroid lying largely in band number 27. This tendency will be discussed below.

Figure 3.20 shows the position of the loudness centroid (sharpness) in the A-textures as presented in recording 2. The information in this graph is summarised below:

Position of loudness centroid
(recording 2):

A-textures

[A]	31/30/29
[C]	30/29/28
[E]	32/31/30
[G]	31/30/29

Comparing the position of the loudness centroid in the A-textures in the two recordings, we find that, on average, the loudness centroid lies one to three bands higher in recording 2. This is not surprising since recording 2 sounds, overall, much brighter than recording 1, and we saw in figure 3.10 that recording 2 has more energy in the high frequencies than recording 1, and this will influence the position of the loudness centroid. Because recording 2 brings the instrumental solos much more to the foreground, and because it has a larger dynamic range than recording 1, it is also not surprising that the position of the loudness centroid should, on average, cover a wider range in each section than that seen in recording 1.

In recording 1 we saw that there was a tendency for the loudness centroid to become lower as we progress from one A-texture to the next. This pattern is not reproduced here. For the most part, the loudness centroid of section [A] lies in band number 30; that of section [C] lies in band number 29; and that of sections [E] and [G] lies in band number 30 once again. Why is this pattern not reproduced in recording 2?

In order to answer this question, let us consider more carefully what happens from one A-texture to the next. Moving from section [A] to section [C] sees the addition of timpani. If we compare figure 3.19 (recording 1) with figure 3.20 (recording 2) we can see that, as we would expect, the addition of timpani in both recordings causes the position of the loudness centroid to move downwards - from band numbers 29/28 to 28/27 in recording 1, and from band numbers 31-29 to 30-28 in recording 2. In section [E], brass instruments are added. Brass, especially trumpets which are present here, have a bright sound, so we might expect the position of the loudness centroid to rise. And we do indeed see a small rise in both recordings - from band numbers 28/27 to 28 in recording 1; and from band numbers 30-28 to 32-30 in recording 2. In section [G] a piano is added. The difference in the way the piano is recorded determines whether the loudness centroid becomes lower in this section (recording 1), or remains at the same level as section [E] (recording 2). The piano part is written so that there is an eveness of distribution of pitches between the upper and lower registers of the piano. In recording 1 this balance is represented well. The ensemble timbre already contains the brightness of the brass, so that the upper register of the piano seems not to have much more to add here in terms of sharpness. The lower register of the piano, however, seems to add significantly to the ensemble timbre, and lowers the position of the loudness centroid. Figure 3.2, which sets-out the pitch structure of the A-textures, demonstrates how, in terms of pitch-space, the piano's upper register is partly covered by existing pitches in the woodwind (and thus will be partly masked) , while the piano's lower register takes-up previously unoccupied pitch-space.

In recording 2, which tends to be quite bright, the upper register of the piano is heard very distinctly, while the lower register is scarcely audible. This means that the lower register's contribution to sharpness in this recording is negligible. Thus, the position of the loudness centroid remains largely the same as in section [E]. As was the case when discussing the significance of the solo E^6's on roughness in the B-textures, it seems that although recording 1 tends to sound somewhat muddy and dull, its better sense of instrumental balance means that figure 3.19 (recording 1), which shows an overall drop in sharpness over the course of the A-textures, is a better representation of the composer's intention than figure 3.20 (recording 2), where this drop in sharpness is not evident.

3.2.2 The B-textures

Figure 3.21 shows the position of the loudness centroid for the B-textures as calculated from recording 1, and figure 3.22 shows the same information for recording 2. As with the graphs of the A-textures, two different degrees of averaging are shown.

Let us begin the examination of these graphs by reminding ourselves of the overall movement of the loudness centroid in each section (which was described in Section 3.2 of this Chapter). Importantly, the same overall description applies to the sections from *both* recordings.

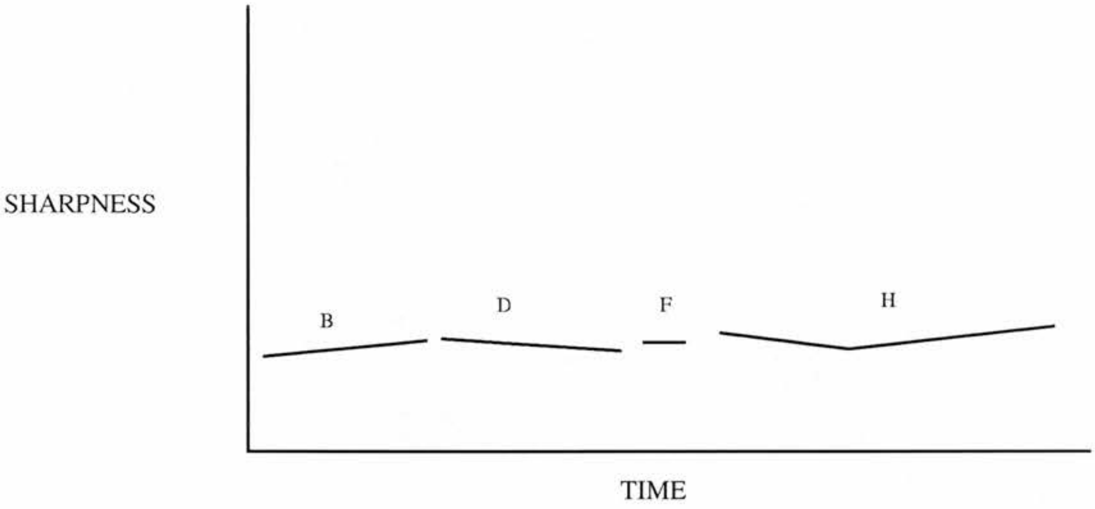

As would be expected from the previous observations about the different natures of the two recordings, individual movements of the loudness centroid in recording 2 often cover a wider range, and the number of the loudest band for any given moment of the piece is always higher in recording 2. Despite this, as observed above, the shape of the movement of the loudness centroid is the same in both recordings.

Let us now examine the individual sections in detail, comparing and contrasting the differences and similarities that exist between the two recordings.

In recording 1 (figure 3.21), section [B] can be seen to be in three parts (as marked on the figure). Part 1 lies at band number 26/27; part 2 lies at band number 27/28; and part 3 falls from band number 29 to number 28 (the movement to band number 27 belongs to the beginning of section [D]). In recording 2, section [B] once again can be seen to consist of three parts. Part 1 lies at band number 27/28; part 2 lies at band number 28/29; and part 3 falls from band number 32 to number 29. This three-part division in each graph corresponds to the same moments in the music. The three parts are shown on the score in figure 3.15 and on the pitch-plot of the section (figure 3.12). From an examination both of the score and of the pitch-plot, the relationship between the three-part structure and the pitch-structure can be established, and the musical causes of the changes in loudness centroid, and thus of the three-part division, can be found.

Let us look first at the relationship to the pitch structure. The beginning of part 2 coincides with the initial entry of the violin (only one violin plays in this section). The violin plays the pitch E^4 with *ricochet* bowing, and this style of playing E^4 is repeated by the violin in the following bar. The beginning of part 3 coincides with the violin playing the first appearance of the only other note it plays in this section, E^6, played *legato*, the complete aural opposite of *staccato ricochet*. The pitch-class E appears in no other instruments or octaves in this section. Thus, the boundaries of the three-part structure coincide with the appearance of the pitch-class E, its two appearances presented in contrasting styles and octaves.

The musical reasons for the changes in the placement of the loudness centroid are also related to the pitch structure of the section. The boundary to part 2 is marked in both loudness centroid graphs (but especially in the graph of recording 1) with a leap upwards. This coincides with, and is initially caused by, the *staccato ricochet* playing of E^4. The reason that the overall level of the loudness centroid remains raised after this moment is because of the pitch change of cello II from $A\#^2$ up to $C\#^4$ (from bar 3 to 4). Nothing else changes that would account for the rise in loudness centroid. The boundary to part 3 is also marked in both centroid graphs by a leap upwards - but a larger leap than that at the boundary to part 2. This is due to the entry of the pitch E^6.

That part 3 continues to show a higher level than part 2 for its loudness centroid after this entry is again due to the general rise in pitch that occurs from bar 7 to 8 (see the pitch-plot, figure 3.12).

Section [B] divides into three parts, primarily defined through the position of the loudness centroid (sharpness), and this division is supported by the pitch structure. How does this three-part structure compare with the roughness measurements for this section?

When the roughness measurements were discussed in Section 3.1, the focus was on the overall trend in each section as defined by the straight-line averages. If we now look in more detail at the movement of the graph we find that roughness measurements support the three-part structure. It has already been pointed out that the beginning of part 2 is signalled by the pitch E^4 played *staccato ricochet*. This event will have an impact on the level of roughness. What is interesting is the extent of its impact. If we look at the roughness at this moment in recording 2 (figure 3.17b; the event is labelled '4 vl I' at 8.5 seconds - remember that these sections have been spliced together, so that timings do not relate to the piece as a whole) we find that the level of roughness here is the third highest for this section. The only higher level of roughness for a note played *staccato ricochet* is '6 vn I' (at 14.5 seconds), which is also the pitch E^4 played by the violin. The moment of highest roughness in this section is that associated with the appearance of the (in this recording, overly prominent) high E^6 (at 22 seconds), which marks the beginning of part 3. Thus, in recording 2, the three part-division is supported by the highest and third highest points of roughness of section [B]. In recording 1 (figure 3.17a), the *staccato ricochet* E^4 is also the third highest point of roughness. The highest, which is the violin's E^6 (at 25 seconds), has a roughness measurement just above that of the bass's *staccato ricochet* G3 ('7 cb I' at 20 seconds). Thus, here too, the three part structure is supported by prominent moments of roughness measurements.

This discussion of the roughness measurements raises the question of why the *staccato ricochet* E4 has such a high roughness measure in both recordings in comparison to other pitches played in the same way? The pitches played as *staccato ricochet* in section [B] are as follows: F#3, E4, A#3, Ab3, G3 and F3. Because of the sympathetic resonance with the open string, two of these notes will have a tendency to sound with extra force - E4 (played by the violin) and G3 (played by the double bass). In recording 2, it is E4 that measures a greater degree of roughness; in recording 1, it is G3. Chance, and the recording engineers, probably determine which note came out with the greater roughness.

Let us now look at the sharpness measurements of section [D] in figures 3.21 and 3.22. This section can be divided into two parts from the evidence of both recordings. This two-part division is shown on the graphs along with indications of the musical events that directly relate to the movements in the position of the loudness centroid. Just prior to the entry of the violin's E6 (see also the score, figure 3.15), the loudness centroid lies at band number 27 in recording 1, and at band number 29 in recording 2. Upon the violin's entry, the loudness centroid immediately rises, and continues to oscillate around band number 29 (recording 1) or 31 (recording 2) till the cessation of the pitch E6, now played *sul ponticello tremolo*. Note that the movement around the band is more often to move above it than to move below it. In recording 1, the cello II's B2 ('102 vc II' at the 34th second on the graph) also makes an impact on the position of the loudness centroid. Part 2 is placed at the end of the *sul ponticello tremolo* E6. From here till the end of this section, the tendency is for the loudness centroid to move below the band number established during the *tremolo* E6. The loudness centroid moves decisively lower at the entry of the double bass which plays E^3 *sul ponticello tremolo* ('107 cb'). The loudness centroid then returns to a higher position right at the end of the section, when the double bass has ceased to play.

There have been two other occurrences of a low note played *sul ponticello tremolo* which have not produced such a dramatic lowering of the loudness centroid - '102 vc II' (which lowers the loudness centroid in recording 1 but not recording 2) and '106 cb I' (which lowers the loudness centroid in both recordings). Why? In both instances there has been another event that has caused the loudness centroid not to be as completely influenced by the low note. In the case of '102 vc II' the effect of the B^2 is reduced by the violin's strong E^6. In the case of '106 cb I' the effect of the E^3 is reduced by the virtually simultaneous sounding of C^5 at '106 vl I.'

Figure 3.13 shows the two-part division on the pitch-plot of section [D]. Just as in the pitch-plot for section [B] (figure 3.12), the overall movement of the pitches over the duration of section [D] is upwards. So, why do we not see an overall rise in the loudness centroid as we did in section [B]? The reason is primarily because of the influence of the violin's E^6, and, to a lesser degree, the influence of the double-bass's E^3. Just as in section [B], it is the pitch-class E that has a pivotal role in defining the sections by the position of the loudness centroid, and by reinforcing these divisions by its role in the pitch organisation of the section - the boundary of the two-part structure coincides with the cessation of the sounding of E^6 in the violin, and, just after this boundary, the pitch E^3 is played by the double-bass (the p-c E appears in no other instrument or octave).

We can consider what the loudness centroid profile would be like if the influence of the p-c E were removed. This can be done if we look at the opening of the section (either figure 3.21 or 3.22), the end of part 1, and the end of the section. These are moments when neither E^6 nor E^4 are sounding. The loudness centroid is initially low, it rises by two band numbers, and, after the double bass's note is no longer sounding, it returns to this higher band at the end of the section. This is the movement that would be expected from the overall pitch movement of the section, without the influence of p-c E. Indeed, if we compare the pitch-plots of both sections [B] and [D], we see that the pitch movement of section [D] begins from a position very close to where the pitch movement of section [B] ended - the pitch movement of the two sections can be seen as one span of overall upwards movement. This overall upwards movement is reflected in the profile of the position of the loudness centroid over the two sections.

In section [B], we found that the graph of roughness supported the divisions implied by the graph of the loudness centroid. Do we find the same support between loudness centroid and roughness in section [D]?

The two part division of this section is marked by the cessation of the violin's *tremolo*, rather than the entry of an event, so it is unlikely that the division will be marked by a point of high roughness. The moments that do have a high degree of roughness, in both recordings, are the 'exclamation marks' (the *mf sul ponticello tremolo* notes) and sometimes an occurrence of the note E^6,which is always played by violin I. For example, figure 3.17a shows a high roughness measure for the first *mf* quaver of the violin I line (bar 103, marked with a small quaver on the figure), while figure 3.17b shows a high roughness measure for the second *mf* crotchet of the violin I line (bar 103, marked by a small crotchet). The moment in both recordings that marks the division between the two parts is a point of low roughness, marked on the graphs by an unbroken line intersecting the graph trace. Thus, whereas the division into three of section [B] indicated by the loudness centroid graph was supported by points of *high* roughness, the division of section [D] into two is supported by a point of *low* roughness.

Section [F] is very brief (the timing indicated in the score is only two seconds, which includes the percussion strike; because these percussion strikes have been omitted from the graphs that show the B-textures spliced together, section [F] is even shorter in these graphs). It consists of a single (symmetrical) pitch collection played as a single chord (see figure 3.13). There is no segmentation in this section - it functions as a surprisingly brief interlude between two A-textures.

In the sharpness graphs of both recordings, there appears to be considerable movement in the position of the loudness centroid during section [F]. This is due to the data points at either end of the section being calculated as an average that includes data from the preceding or the following section. After considering graphs with finer resolution (with an average of 1 - a data point every 0.0464 seconds), section [F] in recording 1 lies at band number 29/30, in recording 2 it lies at band number 30.

The graphs of the position of the loudness centroid for Section [H] for both recordings agree remarkably well. On the basis of the position of the loudness centroid, section [H] divides into two sections - the dividing line is the appearance of the first note played as an *harmonic* - cello I playing E^5 in bar 6 ('6 vc I').

Why is this moment chosen as the event that marks the beginning of part 2 of this section? The 'exclamation marks' of section [D] (the *tremolo* notes, sudden *mf*'s, and so on) are performed against a background of sound with a high loudness centroid. When one of these 'exclamation marks' occurs, the loudness centroid moves temporarily downwards. For example, the dip in loudness centroid associated with '5 vc I' is caused by cello I's *mf* semiquaver (E^4) in bar 5. All the 'exclamation marks' have a tendency to cause the loudness centroid to move downwards, with the exception of the bowed *harmonic* E^5's of cello I ('6 vc I' and '7 vc I'). These notes tend to cause the loudness centroid to rise. In figure 3.21, the loudness centroid rises on both occasions that these notes are sounded; in figure 3.22, only the second note played as an harmonic causes the loudness centroid to rise; during the first, the loudness centroid remains static. Overall, the structure of section [H] consists first of a downwards movement of the loudness centroid from the section's opening to '5 vc I'. From the entry of cello I's *harmonic* E^5 ('6 vc I'), the loudness centroid's overall movement is upwards, that movement continuing until the end of the section. Thus, the first part of the section moves downwards towards the lowest moment of the section, and the beginning of the second part is marked by the first upwards moving 'exclamation mark', and this overall upwards movement continues until the end of the section. This movement of the loudness centroid is more apparent in the graph of recording 2 than in that of recording 1.

Figure 3.14 is a pitch plot of section [H]. During section [H], the pitch space gradually contracts and moves upwards. As with the preceding sections, p-c E is pivotal to the structure already outlined - the downwards movement of the loudness centroid in the first part completes its descent with the sounding of E^4, and the next section begins with the appearance for the first time in this section of E^5. Indeed, p-c E is played prominently throughout this section by cello I.[5] The roughness

---

[5] An interesting detail of pitch organisation can be seen in the second half of bar 4, where the following pattern is formed:

A
3
Gb
4
D5
3
B4
3
G#
3
F
3
Db4
4
Bb3
3

This is the same symmetrical interval pattern that is used to order the pitches in section [F].

prominence of this pitch-class can be seen on figures 3.17a and 3.17b, where each occurrence of the cello's p-c E is shown. Unlike the previous sections, the point of division into two of this section, based on the position of the loudness centroid, is not supported by a corresponding moment of relatively high or low roughness.

We have seen how changing levels of sharpness help to structure this movement, both on the small and the large scale. We have also seen how sharpness and roughness can work together to create structural partitions, and that these partitions are supported by the pitch structure of the movement. Let us now turn our attention to harmonicity.

3.3 Cepstrum analysis

Figures 3.23 and 3.24 show cepstrum analyses of the A-textures in recordings 1 and 2 respectively. As would be expected, occasional suggestions of rahmonic families occur throughout the A-textures, as individual notes are occasionally heard through the often dense texture. However, the distinguishing feature of these cepstrum analyses is the cepstral line that begins in section [E] and continues through section [G]. In both graphs, there are suggestions of further cepstral lines lying in a rahmonic family relationship above this line. The beginning of this line coincides with the entrance of the brass, playing a four-note pitch collection G^4-Ab^4-A^4-Bb^4. When the brass enter with this 'narrow-band' pitch collection, the degree of harmonicity increases. This can be summarised as follows:

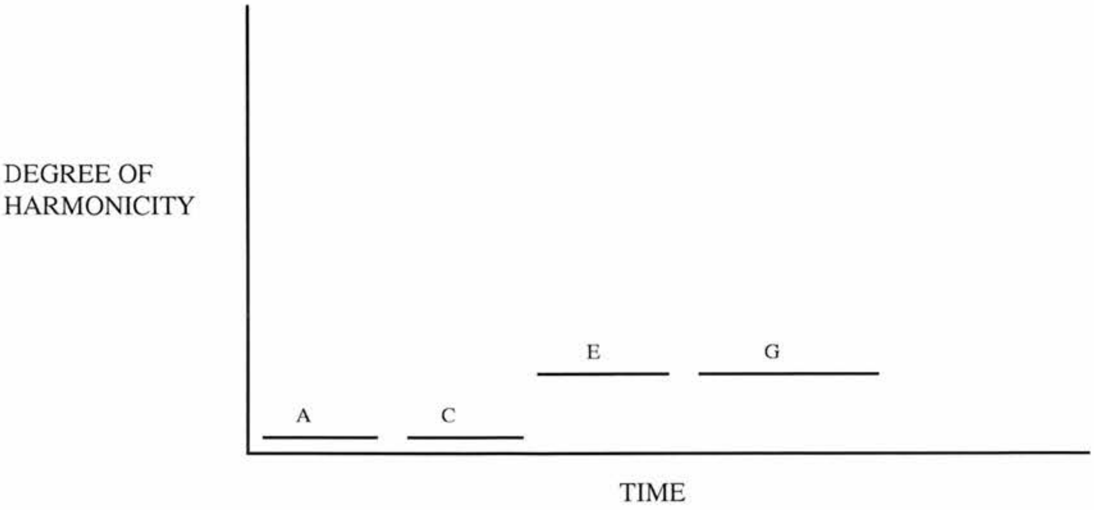

Figures 3.25 and 3.26 show cepstrum analyses of the B-textures in recordings 1 and 2 respectively. The graphs are very similar. We see the occurrence of the violin's E^6 at the end of section [B] as a distinct set of parallel lines. As shown on the graph, this is the point for the beginning of part 3 of this section. The sparse cepstral peaks prior to this occurrence of E^6 are due to some to the brief 'exclamation points' that momentarily increase the degree of harmonicity. One of these, '4 vn I' (particularly prominent in recording 1) marks the beginning of part 2. The repeated violin E^6's in the first half of section [D] show-up clearly in both graphs as a pronounced rahmonic family. The dividing line between parts 1 and 2 is the point where the E^6 *tremolo* ceases. Thus, the dividing line between the two parts is where this rahmonic family ceases. This is clear in the cepstrum of recording 2, but is not what we see in the graph of recording 1 because it is necessary to use a low degree of resolution in order to suppress the 'noisy' background. Towards the end of section [D] the degree of harmonicity becomes less. Section [F] shows no strong evidence of a rahmonic family in either recording. Section [H] shows evidence of rahmonic families coming and going. In recording 1, these moments of greater harmonicity come and leave quite suddenly, whereas in recording 2, the rahmonic families appear to merge one into another. The degree of harmonicity becomes less as we move towards the end of the

section. One prominent rahmonic family (especially in recording 1) belongs to '6 vc I' - the dividing line between parts 1 and 2 of this section.

Although no actual measure of harmonicity is offered here, overall trends in harmonicity can be observed. Looking at the cepstrum structure of the B-textures as a whole, it appears that section [B] moves to a point of higher harmonicity; section [D] begins with a higher level of harmonicity, and ends with a low level; and section [H] has moments of higher harmonicity distributed over its length, until a decrease in harmonicity near its end. This can be summarised in the following way:

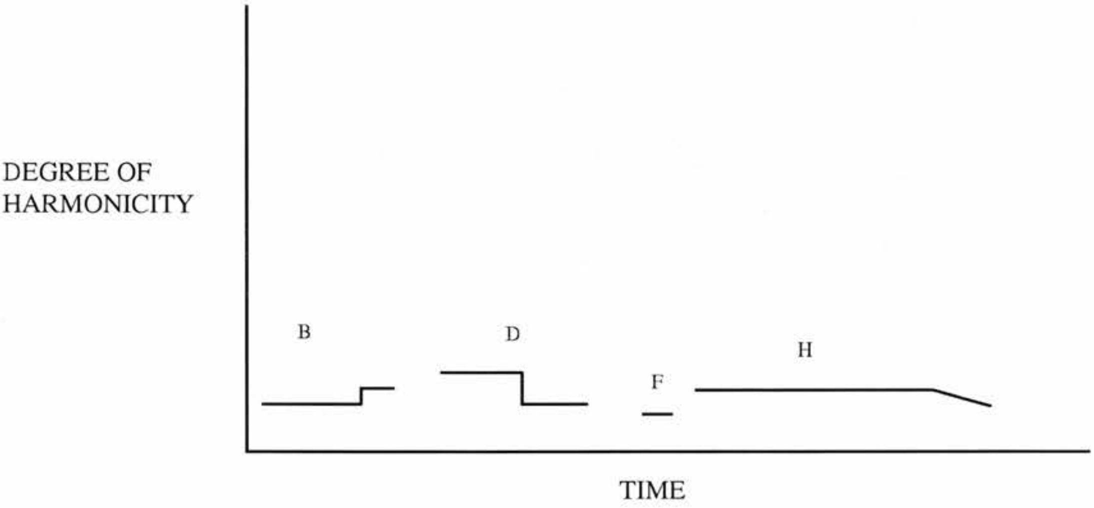

Comparing the cepstral structure of the B-textures with the part-structure of the B-textures proposed in earlier sections of this chapter, we see that this part-structure is supported by the cepstral (harmonicity) structure.

Cepstrum analysis is a very successful way of determining the degree of harmonicity of a sound. This degree of harmonicity is strongly related to the extent to which we perceive a sound as pitched - in other words, Bismarck's parameter of 'compactness' (see Chapter 1, Section 3.1). It is possible for a strongly pitched sound also to have a relatively high degree of roughness. This is what we see during the occurrences of the violin's E^6 - they have a relatively high degree of both harmonicity and roughness. Cepstrum (harmonicity) analysis and roughness analysis are therefore measuring different, but complementary, aspects of what musicians refer to as consonance and dissonance.

3.4 Loudness Distribution Analysis

When the results of Loudness Distribution Analysis of recording 1 were considered at the beginning of Section 3 of this Chapter, a preliminary conclusion was made that the differences between the timbres of the A- and B-textures were not borne out convincingly by this method of analysis. We then turned to other timbre analysis tools that have shown very well some of the timbral differences that exist in this piece. It is now time to return to LD-analysis to see what it has to tell us.

3.4.1 The A-textures

Figures 3.27 and 3.28 are graphs of timbral weight for the A-textures of recordings 1 and 2 respectively. Looking again at figure 3.4, it will be seen that the graph of timbral weight (the lower line) often fluctuates dramatically, so that it is very difficult to see any evidence of systematic movement. In an attempt to overcome this problem, figures 3.27 and 3.28 present graphs of timbral weight with an averaging value of 50 - data points are separated by 2.32 seconds. Thus, we are looking for evidence of overall trends, rather than fine details of movement in timbral weight.

In the analysis of *sharpness* (position of the loudness centroid) it will be remembered that recording 2 of *Jeux vénitiens* consistently had a sharper sound when compared with recording 1. This was inherent in the different 'sound styles' of the two recordings. The measurement *timbral weight* is heavily dependent on the position of the loudest band (which is the measure of timbral pitch). The position of the loudest band in the same averaging as seen in figures 3.27 and 3.28 is shown in the graphs of timbral pitch in figures 3.29 and 3.30. It can be seen quite readily that, on average, the position of the loudest band in recording 2 of the A-textures (figure 3.30) is higher than that of recording 1 (fig. 3.29). Thus, it follows that, on average, the timbre of the A-textures in recording 1 will be lighter in weight than that in recording 2.

In Section 3.3 of Chapter 2, timbral pitch was described as a measure of the placement of the centre of gravity of a sound, as opposed to sharpness which measures the changing overall 'focus' of the sound. In the measurement of timbral weight a heavy timbre has more of a sense of definition or 'crispness' in the lower frequencies when compared with a light timbre - a sound with a heavy timbre has more 'presence' than a sound with a light timbre. Thus, it follows that, overall, recording 1 has a lower 'centre of gravity' and a less-well defined sound in the lower frequencies when compared with recording 2. This corresponds well with my perception of the difference between the 'sound styles' of the two recordings. At the beginning of this Chapter, recording 1 was described as having a much 'duller' and 'muddier' sound than that heard in recording 2. This observation is now confirmed by the graphs of timbral pitch and weight. The overall sound of recording 1 is 'deeper' than that of recording 2 (its overall position of timbral pitch is lower) and it has less definition, with a 'muddier' sound (its timbral weight is lighter).

Comparing figure 3.29 (recording 1) and 3.30 (recording 2), we can see that the graphs of timbral pitch move in different ways. In figure 3.29, timbral pitch begins at band number 29 and ends on band number 25. There is thus a tendency to move downwards in section [G], though overall the graph of timbral pitch oscillates either side of band number 27. In figure 3.30, timbral pitch begins at band number 29 and moves upwards as high as band number 35, before concluding on band number 33. Here, there is a definite tendency for timbral pitch to rise. Why does timbral pitch move in different ways when we are dealing with the same piece of music? The answer lies in the 'sound style' of the recordings. As has been said, recording 1 has a 'muddier', less well defined sound when compared with recording 2. With the addition of more instruments as the piece progresses this 'muddiness' does not improve - if anything it becomes slightly worse. Thus, we see the graph of timbral pitch oscillating either side of band number 27, and in section [G], with the addition of piano, timbral pitch dropping a little. In the graph of timbral weight in recording 1, (figure 3.27) weight oscillates either side of a measurement of 1, tending to become lighter in section [G].

We see a very different tendency in the graphs taken from recording 2. As has been said, figure 3.30 shows timbral pitch rising during the course of the A-textures. With the addition of timpani in section [C], timbral pitch (the centre of gravity of the sound) rises a little; with the addition of brass in section [E], timbral pitch rises still more, and largely remains at this level during section [G], where the piano is added. In figure 3.28, we see the corresponding movement in timbral weight. The graph shows an increase in heaviness in section [C] in comparison with section [A], a further increase in section [E], and a general maintenance of this level in section [G]. It will be remembered that it was shown that recording 2 contains more energy in the upper frequency range than recording 1 (see figure 3.10). In recording 2, as more instruments are added, the loudest band moves steadily higher as the additional instruments add energy to this upper frequency range. Along with this movement, timbral weight becomes heavier, and the 'definition' in the sound remains strong even though the texture becomes increasingly complex. In recording 1, where this upper frequency range is not available because of the 'dull' nature of the recording, the addition of instruments makes little difference to the position of the loudest band: the tendency is for the position of the loudest band to drop slightly due to the energy of the fundamentals. In recording 1, timbral weight remains relatively static, and rises slightly in section [G]. With the increase in textural complexity the 'definition' of the sound becomes slightly weaker.

To end this discussion of timbral pitch and weight in the A-textures, let us briefly compare the movement of timbral pitch with that of sharpness, as measured by the movement of the loudness centroid (the changing 'focus' of the sound). In the discussion of sharpness in the A-textures, recording 1 was shown to have a small movement downwards in the loudness centroid over the course of the A-textures (from band number 29 to band number 27), while in recording 2 the position of the loudness centroid oscillated about band number 30. Thus, in recording 1 we find a small downwards movement in both sharpness (loudness centroid) and timbral pitch (loudest band), whereas in recording 2 we find sharpness to be relatively static, while timbral pitch rises significantly. Therefore, while the overall perceived 'brightness' of the sound remains constant during the A-textures of recording 2, the measure of timbral pitch detects the addition of energy into the upper frequency range as instruments are added to the texture, which causes the upper frequencies to become more insistent on our perception.

To sum-up the information on timbral pitch and weight in the A-textures, the information from recording 2 rather than recording 1 has been chosen as the basis for the summary. Previously, recording 1 has been chosen as the more representative recording because the instruments have been judged to be in better balance with one another. However, with measurements of timbral weight and pitch I feel that the information in the upper frequency range is very important, and recording 2 reproduces this information more faithfully than recording 1.

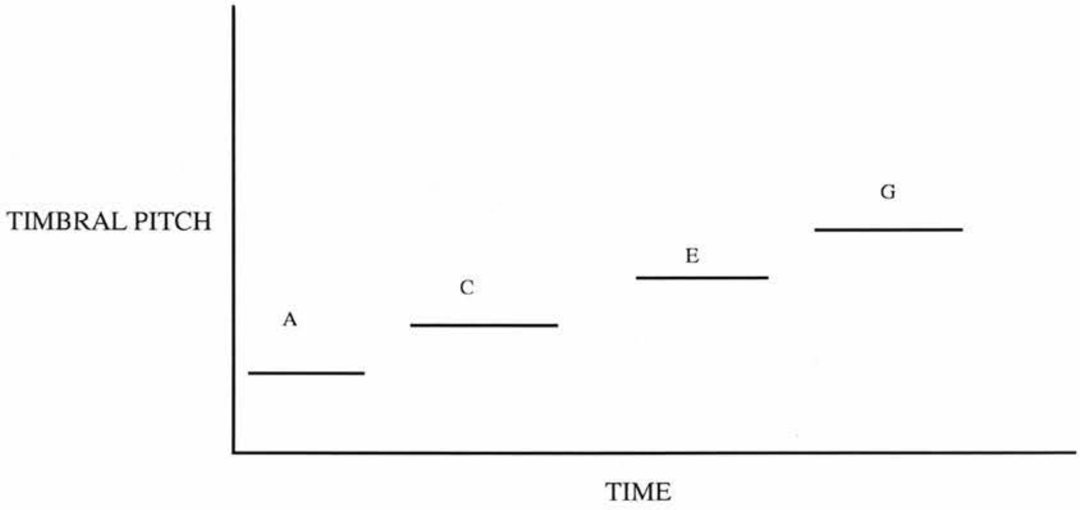

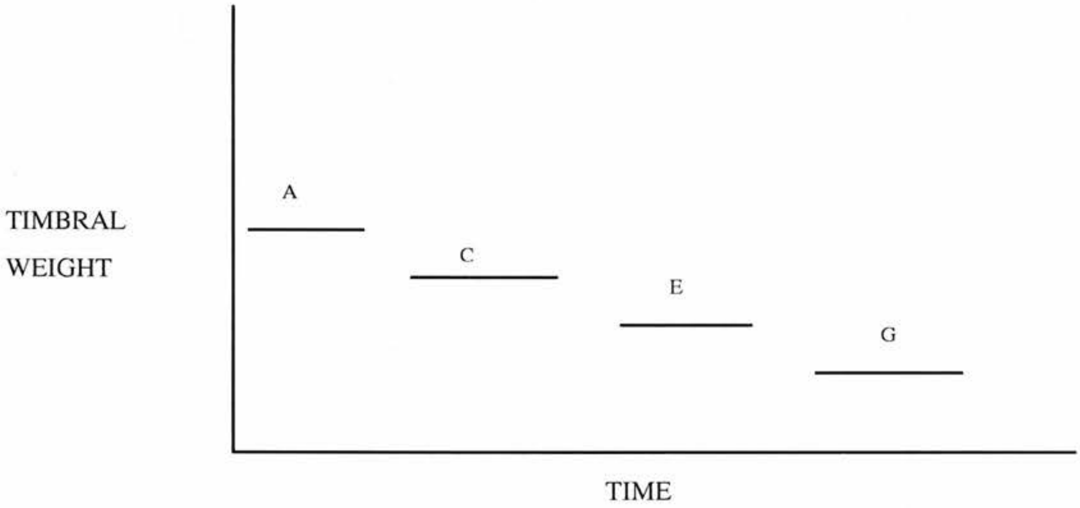

The final measurement obtained from LD-analysis to be examined for the A-textures is the measurement of timbral width. Figures 3.31 and 3.32 are graphs of timbral width in the A-textures for recordings 1 and 2 respectively. The results include straight-line averages through the data points. These are presented so that the overall direction of the data can be seen, while the finer movement of the individual data points can still be taken into consideration.

As might be expected, timbral width increases from section to section in both recordings as more instruments are added. Measured from the straight-line averages, recording 1 (figure 3.31) starts at a measure of timbral width of 0.55 and finishes at the end of section [G] at 0.65; recording 2 (figure 3.32) begins at 0.5 and ends on 0.675. Thus, as also to be expected, recording 2 shows a wider range of timbral width than recording1. Comparing the different recordings more closely, we find that an overall increase in width in recording 1 is not borne out by the position of the individual data points - many points in section [C] are as high, and some higher than points in section [G]. Indeed, the end of the straight-line average of section [A] is as high as the end of the straight-line average in section [G]. In the graph of recording 2, however, we find that an overall increase in width is borne out by the individual data points as well as by the end-points of all of the straight-line averages. The progression of timbral width in the A-textures, as seen in recording 2 (for reasons of its better frequency range), is summarised below:

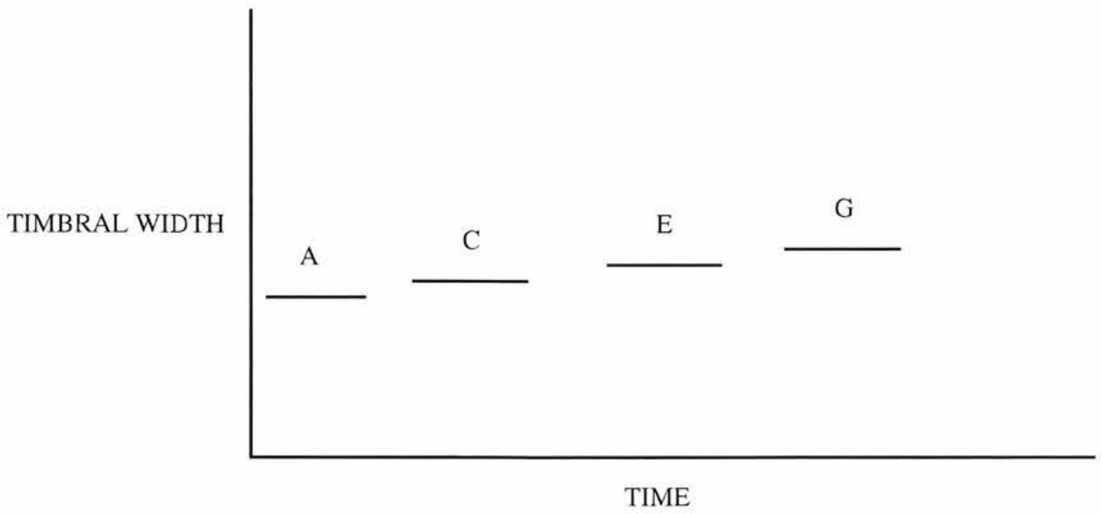

3.4.2 The B-textures

Figures 3.33 and 3.34 are graphs of timbral pitch for the B-textures from recordings 1 and 2 respectively. Figures 3.35 and 3.36 are the corresponding graphs of timbral weight. Just as with the graphs of pitch and weight for the A-textures where we found that the two recordings produced very different tendencies in the graphs, so too here we find very different tendencies in the graphs representing the two recordings.

Timbral pitch in recording 1 (figure 3.33) keeps returning to band number 26, with sudden leaps at the end of section [B] and in the first half of section [D] to band number 34 (these leaps are caused by the playing of the pitch E^6 by the violin). Timbral pitch in recording 2 (figure 3.34) begins very similarly to recording 1 - it moves around band number 25, with a sudden leap near the end of section [B] to band number 31, once again caused by the playing of the note E^6. However, in the following sections in recording 2, timbral pitch lies largely around band numbers 33 to 35, with occasional excursions down to band number 26 - the band number that recording 1 stays near to in this section. Why do these differences exist?

Once again, the answer lies in the different 'sound styles' of the two recordings. In general, the pitch range of each B-texture becomes higher with each successive section (see figures 3.12, 3.13 and 3.14). It will be remembered that recording 2 has much more energy in the higher frequency range than recording 1. In section [B], with the pitch range lying relatively low, there is very little difference between the graphs of timbral pitch for the two recordings. When the pitches move higher in subsequent sections, significant energy from the harmonics of these pitches lie in the frequency area where recording 2 demonstrates greater sensitivity than recording 1. Thus, the loudest band in recording 2 shifts from around 25 in section [B] to around 34 for subsequent sections. In recording 1, however, the loudest band only moves into this frequency area when a strong E^6 is played. The sharp dips in the position of the loudest band in recording 2 are due to low string notes bringing the loudest band down into this lower frequency range. Specifically, the low peak at 47 seconds is due to the bass harmonic in bar 107 of section [D] ('107 cb I'); that at 64 seconds is due to the sudden *mf* cello 1 note in bar 5 of section [H] ('5 vc I'); and that at 87 seconds is due to the pizzicato harmonic cello note in bar 11 of section [H] ('11 vc I') (see the score, figure 3.15).

Let us now turn to the graphs of timbral weight in figures 3.35 and 3.36. Comparing these two graphs, we can see that section [B] in recording 1 is a little heavier than section [B] in recording 2. Indeed, this is the only time when, perceptually, recording 1 has a sound that seems a little more 'focused' than recording 2. From section [D] onwards, recording 2 has a generally heavier timbre than recording 1, the effect of which is heard in the more focused nature of the recorded sound of recording 2. Note the high upward spikes in the graph of timbral weight in recording 2. These correspond to the downward spikes of timbral pitch shown in figure 3.34, which were related to specific events in the score. It is interesting to note that the occurrences that cause timbral weight to become momentarily lighter are occurrences that perceptually lack a strong sense of focus. Also interesting to note is the disparity in size between the large upward spike of timbral pitch in section [B] of recording 2 (figure 3.34 at 21 seconds) and the corresponding but smaller downward spike of timbral weight at the same time in figure 3.36. We have been used to seeing the graphs of timbral pitch and weight generally mirroring one another because as the position of the loudest band shifts the general spread of spectral energy remains much the same. At this point, however, the two graphs do not mirror one another. This indicates that when the violin plays the note E^6 the position of the loudest band rises, but as this happens the proportion of energy above the loudest band increases - this is shown by the downward timbral weight spike being smaller than the upward timbral pitch spike. Why might this be? Again, this is related to the property of recording 2 having much greater sensitivity in the higher frequency area than recording 1. (Indeed, if the same moment in recording 1 is examined (see figures 3.33 and 3.35 at 23 seconds) it will be seen that the spikes are the same size.) Presumably, the violin note E^6 generates a lot of energy in its harmonics (the position of the loudest band at this point - band number 31 - is at the frequency of the fundamental of E^6), thus providing a greater proportion of energy above the loudest band at this moment than just prior to the sounding of E^6.

The tendencies noted in the graphs of timbral pitch and weight for recording 2 are summarised below:

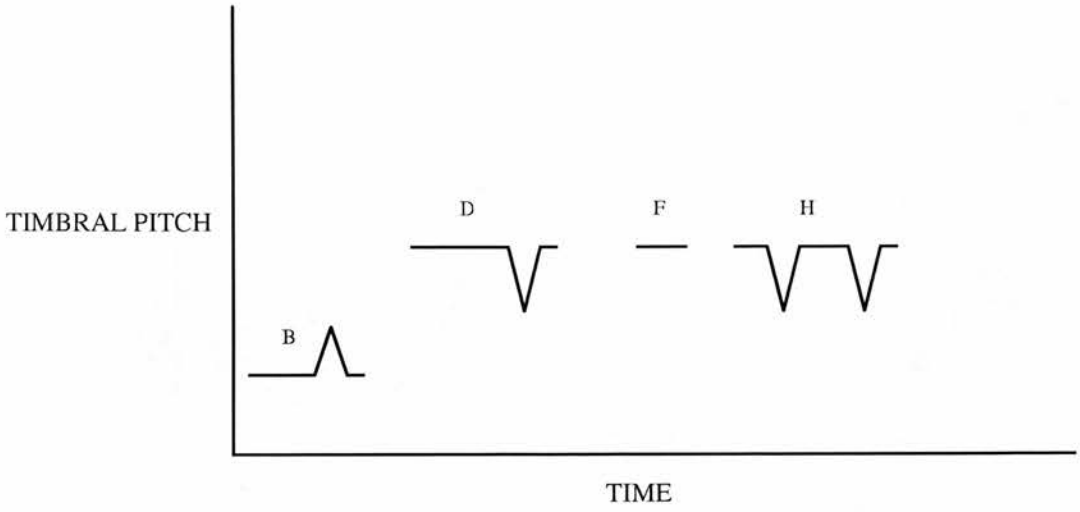

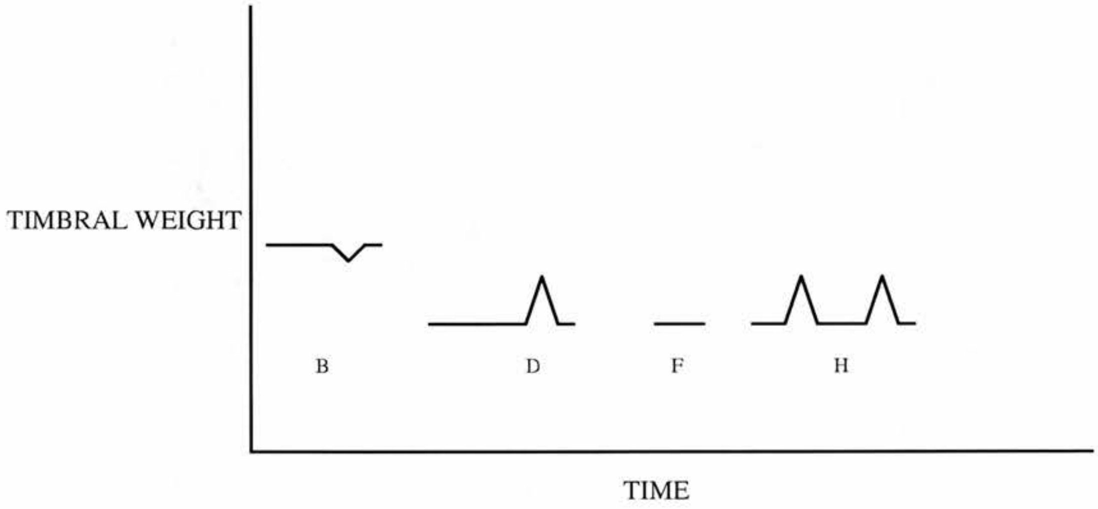

Graphs of timbral width for the B-textures are presented for recordings 1 and 2 in figures 3.37 and 3.38 respectively. The averaging here is 5, so that there is a data point every 0.232 seconds. Similarly to the graphs for roughness, it was found that an averaging greater than this caused the individual musical events to be 'ironed out' in the graph, and an averaging less than this introduced too much 'noise' (or jitter) in the graph for the musical events to stand out clearly from the background. Precise indications are given on the graphs of where the 'textural exclamation points' occur, in the form of a bar number and below it the name of the instrument that plays.

If the graph of timbral width moves upwards, it indicates that the timbre is more diffuse - that there is a greater amount of total energy outside of the loudest band. If the graph moves downwards, it indicates that the timbre is more focused - that there is now less energy outside of the loudest band. Let us go through each section in turn, looking at how the different types of textural exclamations affect timbral width in different ways, and examining how the two recordings differ.

In section [B], the textural exclamations consist of small groups of demisemiquavers played with *ricochet* bowing, usually with the dynamic marking of *mf*, *decrescendo* to *pp*. In bar 8, the pitch E^6 appears, played first *mf*, and then *pp*, as an *harmonic*.

Ricochet bowing entails hitting the string with the bow and then allowing it to bounce as the bow is moved. Thus, the start of a note played *ricochet* often has quite a 'harsh' sound. This harshness manifests itself in a graph of timbral width by showing width becoming more diffuse - the graph moves upwards. In recording 1 this is especially apparent in '6 vn I'. Sometimes the movement is very small, as in '3 vl I', because the *ricochet* note is played with very little force, so that it does not stand out very much from the background. All the 'exclamations' in section [B] of recording 1 cause width to become more diffuse, except '6 vc III' - here, the timbre becomes more focused. The reason is that it is the weakest of all the *ricochet* bowings in this section, so that the note is clearly heard through the texture, but with none of the 'harshness' associated with *ricochet* bowing - there is present only the clarifying effect of hearing a pitched note appear in front of the diffuse background. Near the end of the section, the pitch E^6 appears - shown on the graph by two minims, and dotted lines either side of them to show their duration. As would be expected, both appearances cause the timbre to become more focused, but when the note is played as an *harmonic* there is a smaller movement towards a more focused timbre.

Turning to recording 2, we find that in section [B] the movements in timbral width are sometimes greater than in recording 1 - compare the two recordings, for instance, at the very first 'exclamation,' 3 vl I. Unlike the differences in timbral weight discussed above, which were often largely due to the difference between the recording 'style' of the two recordings, differences in timbral width appear to be very largely due to differences in playing technique, with differences in the recorded sound playing a secondary role (the 'closeness' of musical events in recording 2 influences width to a degree). Thus, the principal reason that recording 2 of '3 vl I' shows a greater movement towards being more diffuse than the same moment in recording 1 is because the violinist in recording 2 attacks the *ricochet* much more robustly than the player in recording 1. Overall, it appears that the interpretation of the [B] section in recording 1 favours lighter *ricochet* attacks, and the interpretation in recording 2 (conducted by the composer) favours heavier *ricochet* attacks. Note especially that the occurrence of '6 vc III', which caused, atypically, a more focused timbre in recording 1, causes a more diffuse timbre in recording 2. Note, too, that the first appearance of E^6 in recording 2 produces a greater movement towards a more focused timbre than is the case in recording 1, and that the second appearance (the *harmonic*) produces a smaller movement. This is because in recording 2, E^6 is heard very loudly in comparison with the background texture - thus the background texture largely disappears. The second E^6 is played very quietly, thus the background texture remains more audible. It will be noticed that the marked points on the graphs do not, by any means, account for all the movement in timbral width. This is because the graphs are also measuring the changes in width of the background sustained string clusters. Individual notes are played with slight variations during the course of a section, and notes sometimes move towards and away from our perception as they move in and out of the diffuse background. Two clear examples of this can be seen in recording 2 at the points in section [B] marked with an asterisk. As can be seen from the graph, the effect of hearing a clear pitch emerge from the diffuse background is to make timbral width more focused. An example of this is also marked in recording 1 in section [H] between '8 vc I' and '9 vc I.'

The 'exclamations' in section [D] consist of crotchets played *sul ponticello*, *tremolo*. According to the graphs of timbral width, sometimes these 'exclamations' cause the timbre to become more focused, and sometimes more diffuse. Let us examine the individual cases.

At the start of section [D], there is a series of repeated E^6's, the notes played alternately *ordinare* and as *harmonics*, the alternation becoming increasingly rapid, until it dissolves into *tremolo*. As is shown on the graphs, the *tremolo* exclamations in the first half of the section occur while the E^6 is being played. Thus, it needs to be remembered that the graphs are showing the combined effect of the presence of the note E^6 along with the *tremolo* exclamations. It can be seen that E^6 played *ordinare* makes the timbre more focused, while, in comparison, the same note played as an *harmonic* makes the timbre more diffuse. This alternation can be seen very clearly in section [D] of recording 2 between the points '103 vn III' and '104 vc II'. Every peak here, both upwards and downwards, corresponds to this alternation in the method of playing E^6.

In recordings 1 and 2, we see that usually the effect of the *tremolo* exclamations is to make the timbre more diffuse. This is because of the 'noisiness' of the *sul ponticello tremolo*, just as the *ricochet* bowing in the previous section caused timbre to become more diffuse because of the 'harshness' of the sound. Exceptions to this in recording 1 are '102 vc III,' '106 vc I,' and the beginning of '107 cb.' These are all instances where there is less of a leaning towards a *sul ponticello* effect. This is especially true of

'102 vc III' where it does not sound as if the note is played *sul ponticello* at all. Exceptions in recording 2 are '106 vc I' and the end of '107 cb.' Why should '106 vc I' and part of '107 cb' appear as exceptions in both recordings ? '106 vc I' consists of a cello playing a high A^4. It is possible that the players in both recordings were reluctant to play this note too close to the bridge in case the combination of the high tessitura coupled with *sul ponticello* were to produce an overly harsh, 'scraping' sound. '107 cb' consists of a double bass playing an harmonic E^4. The fact that this is an harmonic, even though played *sul ponticello*, appears to take some of the harshness off the sound, and during the longer duration that this *tremolo* note is played, variation in the timbral width is very likely.

In section [H] the exclamations are provided by sudden, brief *mf* semiquavers, *tremolo* notes at the start of the section (recalling the previous section [D]) and bowed *harmonic* notes and *harmonic pizzicato* notes near the middle and the end of the section. The sudden *mf* semiquavers have very clean attacks, incorporating not nearly as much noise as the previous exclamations in sections [B] and [D]. From this, we can conclude that the timbre they produce is going to be more focused, rather than more diffuse - and this is exactly what we find.

In our discussion of section [H], let us examine it in smaller parts, and compare these smaller parts between recordings. The section begins with an 'exclamation' of a single *tremolo* note, followed by a single *mf* semiquaver, followed by a brief succession of *sul ponticello tremolo* notes played by the first cello. As expected, the single *tremolo* note causes the timbre to become more diffuse, the single *mf* semiquaver makes the timbre more focused, and the brief succession of four *tremolo* notes causes the timbre to move to a more diffuse position and then back again, due to the breaks between the notes (this is particularly evident in recording 1). These observations hold for both recordings. Indeed, the remainder of the sudden *mf* semiquavers in both recordings all create, to a greater or lesser degree, a more focused timbre. That is, all but one of them - '10 vn III' in recording 1 is the exception. The probable reason for this will be discussed in a moment. The bowed *harmonic* notes played by the first cello ('6 vc I,' and '7 vc I') both show a more focused timbre, and the *pizzicato harmonic* notes all show a more diffuse timbre - the attack of the *pizzicato* making the sound less focused than a bowed note. As I said above, the one point of disagreement between the two graphs is '10 vn III.' Why?

In the previous sections, we have seen that notes played with *ricochet* bowing and notes played *sul ponticello tremolo*, result in the timbre becoming more diffuse. Notes played as harmonics have shown a tendency to have a focused timbre - though this is also dependent on the volume at which the harmonic is played. In the graph of section [H] of recording 1, the beginning and end of '6 vc I' is precisely marked. The dynamic of the note is *pp crescendo mf*. It can be seen from the graph that at the start of the harmonic timbral width is unaltered, it then becomes more focused when the note becomes louder, and then returns to its previous level of timbral width as the note dies away. Thus, the quiet portion of this harmonic has very little effect on the measure of timbral width. Looking at section [D] in recording 2, between '103 vn III' and '104 vc II,' we see the pitch E^6 played alternately *ordinare* and *harmonic*. It was observed that the downward peaks (more focused) coincide with the note played *ordinare*, the upward peaks (more diffuse) coincide with the note played as an *harmonic*. Even though '10 vn III' in section [H] is not marked as such in the score, it sounds to me that in recording 1 this note is played as an *harmonic*, and is also not played particularly loudly. Thus, instead of the timbre becoming more focused, as we would expect from looking at the score, and as happens in recording 2, we see the timbre become more diffuse.

From a detailed examination of the B-textures, let us now stand back a little and look at their overall timbral structure. Straight-line averages are shown on both graphs for sections [B], [D] [F] and [H]. The movement of the straight-line averages is summarised below:

<u>Movement of straight-line averages in B-textures:</u>

	<u>Recording 1</u>	<u>Recording 2</u>
[B]	0.625 - 0.65	0.725 - 0.7
[D]	0.575 - 0.65	0.0.575 - 0.725
[F]	0.65 - 0.65	0.725 - 0.75
[H]	0.65 - 0.475	0.65 - 0.65

In recording 1, section [B] shows a small movement towards a more diffuse timbre, while in recording 2 there is a small movement towards a more focused timbre. Because these movements are so slight, and because recordings 1 and 2 move in opposite directions, it is concluded that, overall, timbral width remains static in section [B]. In both recordings, section [D] shows a substantial overall movement towards a more diffuse timbre. Section [F] in recording 1 shows no movement for timbral width, while in recording 2, there is a very small movement towards a more diffuse timbre. It is concluded that section [F] is best represented by a straight-line average that does not slope. Section [H] in recording 1 shows a substantial movement towards a more focused timbre, while in recording 2 the straight-line average shows no overall movement. In my opinion, recording 1 sounds more realistic than recording 2 in its timbral progression during section [H] as the number of instruments gradually becomes fewer, and the volume slowly dies away. The sound of the instruments in recording 2 seems to remain too 'present,' too near to the listener, and this affects the degree to which the timbre sounds focused. Thus, in the summary of timbral width below, the results from recording 1 are used:

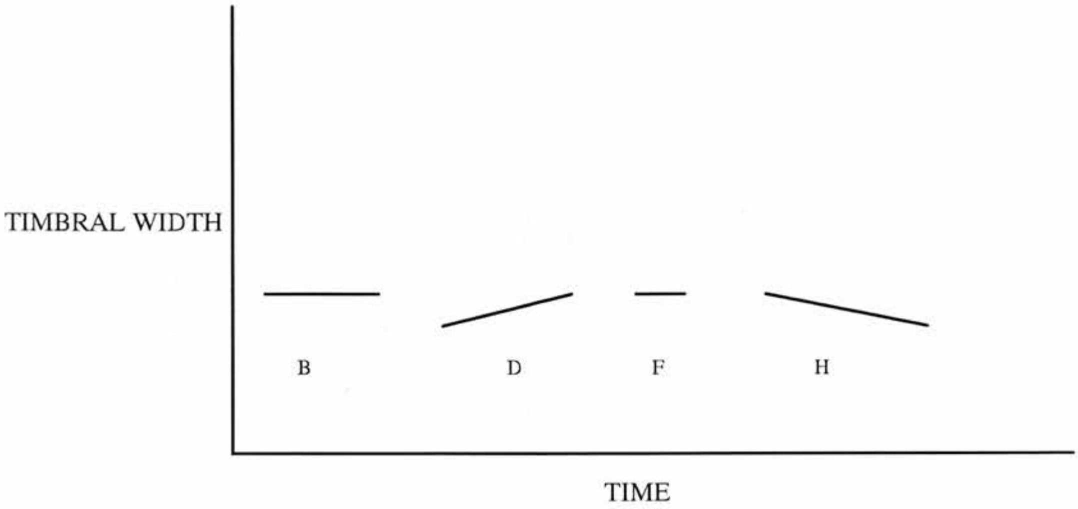

During the course of the detailed discussion of timbral width in the B-textures, it was found that certain types of 'exclamations' cause the timbre to become more diffuse, while other types cause the timbre to become more focused. In section [B] the exclamations consist of *ricochet* demisemiquavers, which make the timbre more diffuse. In section [D] the exclamations consist of crotchets played *sul ponticello tremolo*, which also make the timbre more diffuse. During both these sections, the pitch E^6 appears, played in a way which makes the timbre more focused. Thus, there are two different types of foreground events in sections [B] and [D] - those that are very brief and sudden, which cause the timbre to become more diffuse, and, in contrast, those that are long and sustained, which cause the timbre to become more focused. These long sustained events themselves divide into two contrasting types - E^6 played *ordinare* and E^6 played as an *harmonic* - the former being more focused, the latter being more diffuse, but both forms being more focused than the surrounding background timbre. In section [H] the sustained E^6 pitch does not appear, and now most of the interruptions cause the timbre to become more focused - if the sustained E^6 were to appear here it would not contrast in timbral width with the principal method of exclamation. Section [H] also contains *tremolo* notes and *pizzicato* *harmonics* that provide diffuse timbral contrast to the focused *mf* semiquavers and the focused *harmonic crescendo* crotchets.

All this means that very specific timbral contrasts are at play here - the play of focused against diffuse timbres. Because of this, in the interests of a successful interpretation of the music, it would be beneficial to make clear these contrasts. Both recordings do this. However recording 2, conducted by the composer, more consistently achieves the correct timbral flavour in the exclamations - it is the timbre of these exclamations that help give a sense of structure to the movement.

The partitioning of the B-textures is shown on the graphs of timbral width. It appears that changes in width do not support the points of division that are supported by sharpness, roughness and harmonicity.

4. Combining Timbral Measures

Below, the graphs are reproduced that were considered representative of the timbre of the piece after the differences between the two recordings were considered. We can appreciate how the different timbral aspects work with each other. It is clear that the succession of A-textures forms a single directed trajectory, which is partly created by, and complements the additive instrumental and pitch structure. The graphs of the through-composed B-textures, however, often demonstrate a timbre shape that conveys a sense of beginning, middle and end.

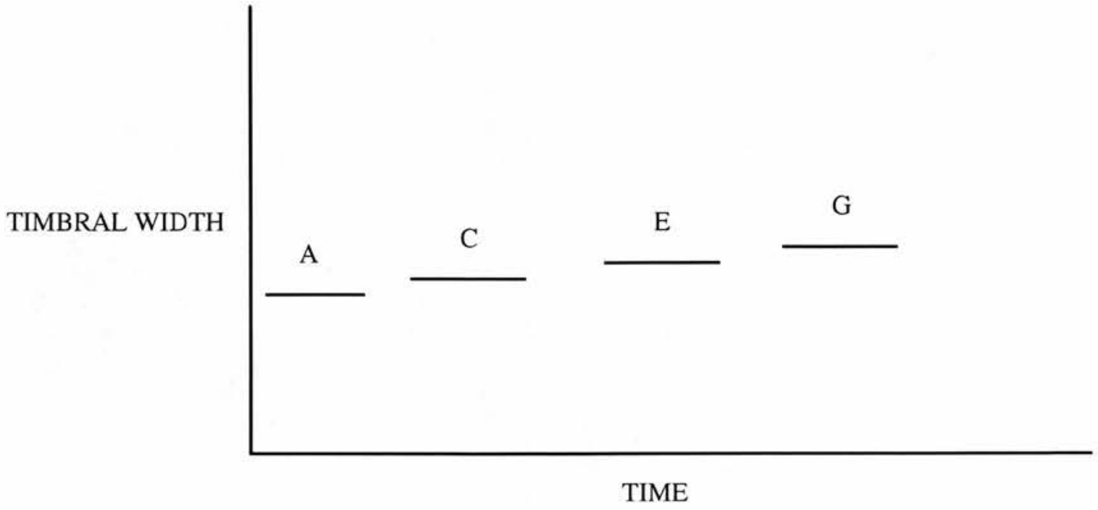
TIMBRAL WIDTH
A
C
E
G
TIME

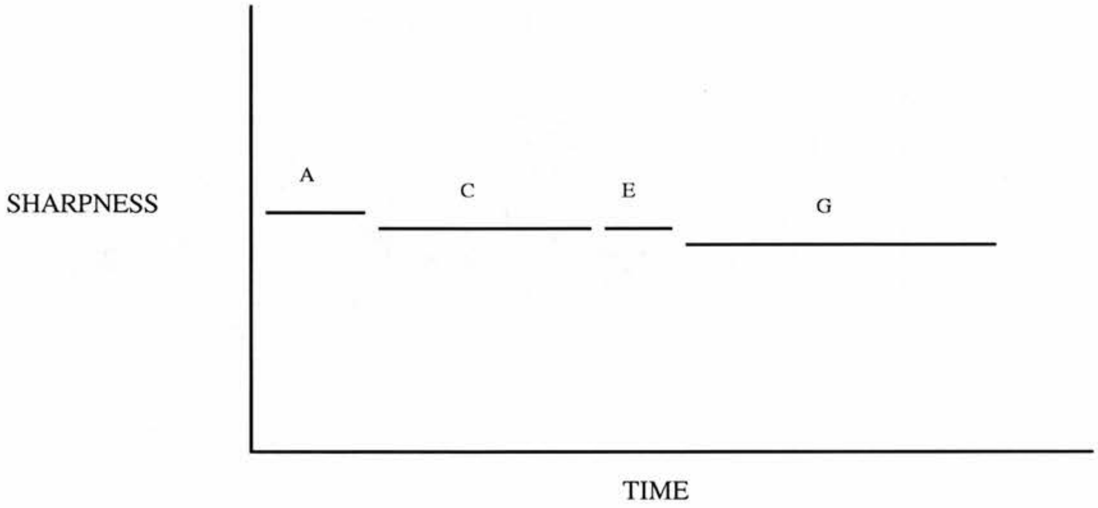
SHARPNESS
A
C
E
G
TIME

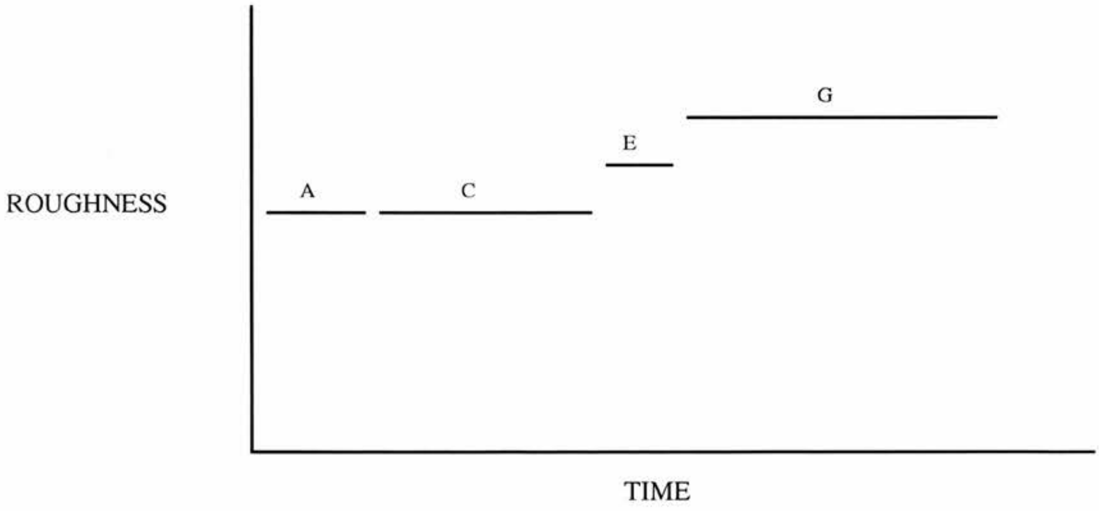
ROUGHNESS
G
E
A
C
TIME

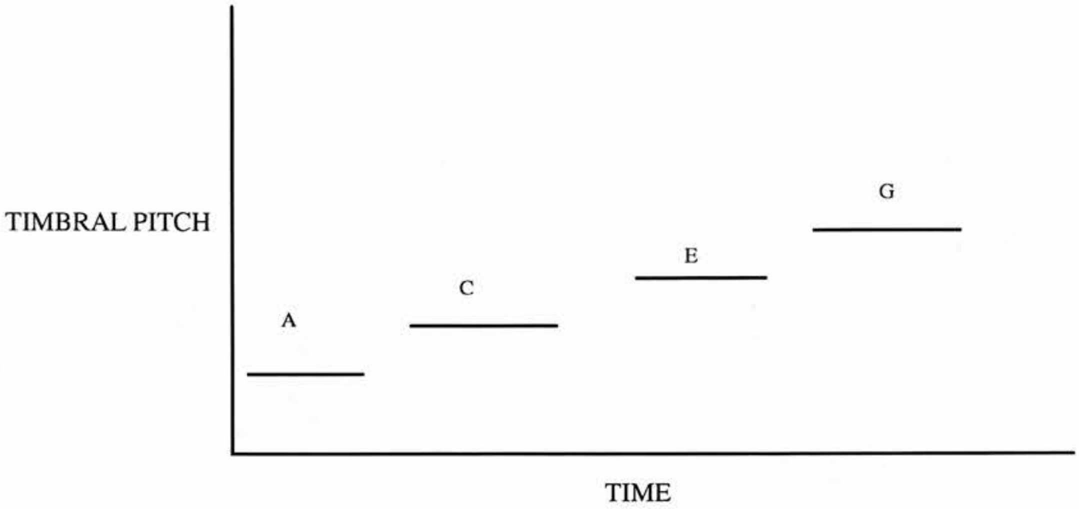
TIMBRAL PITCH
A
C
E
G
TIME

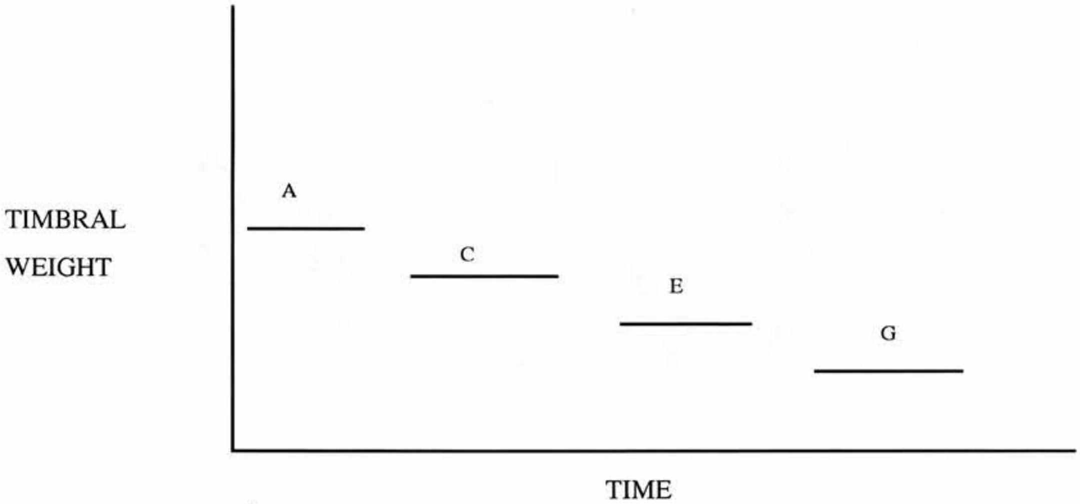
TIMBRAL
WEIGHT
A
C
E
G
TIME

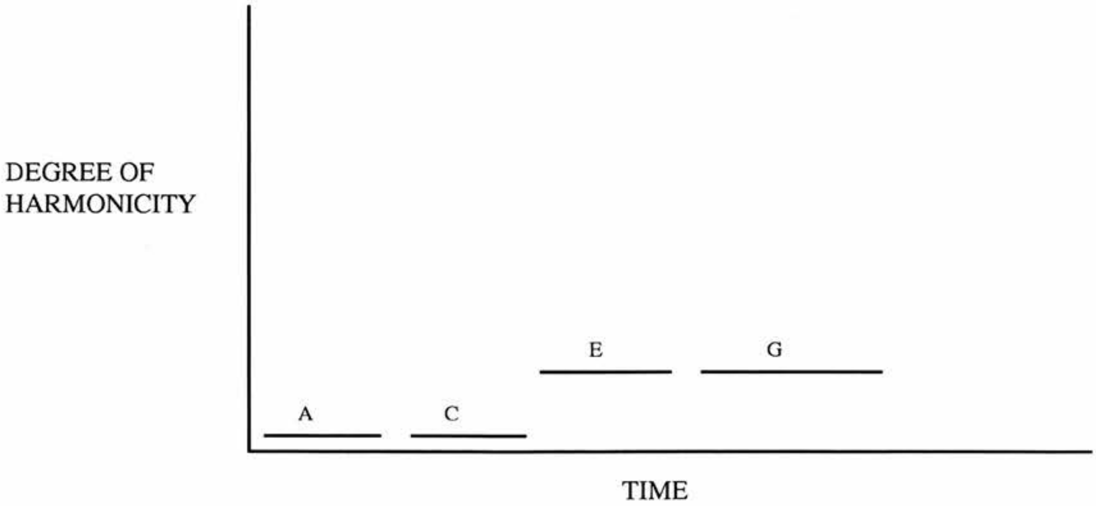
DEGREE OF
HARMONICITY
A
C
E
G
TIME

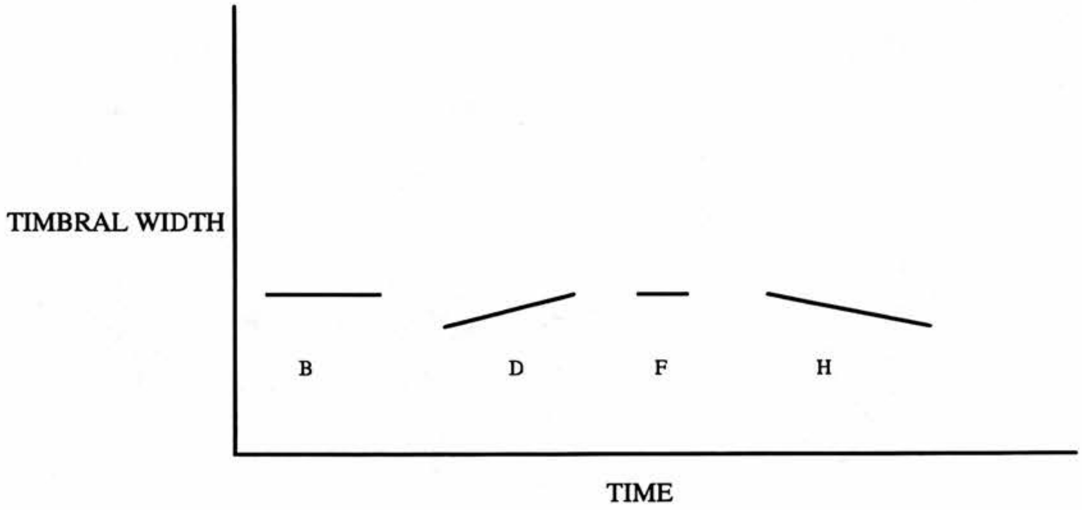
TIMBRAL WIDTH
B
D
F
H
TIME

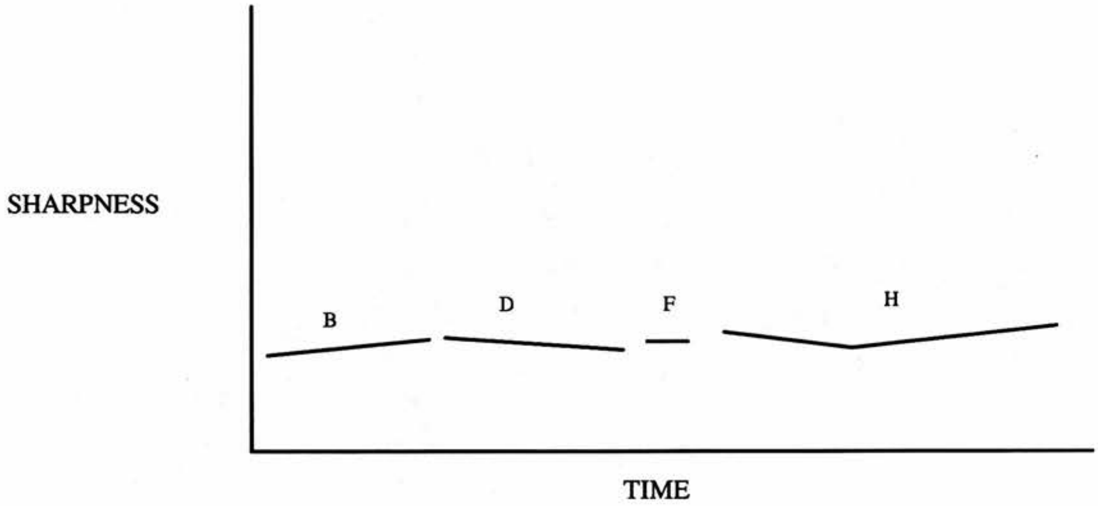
SHARPNESS
B
D
F
H
TIME

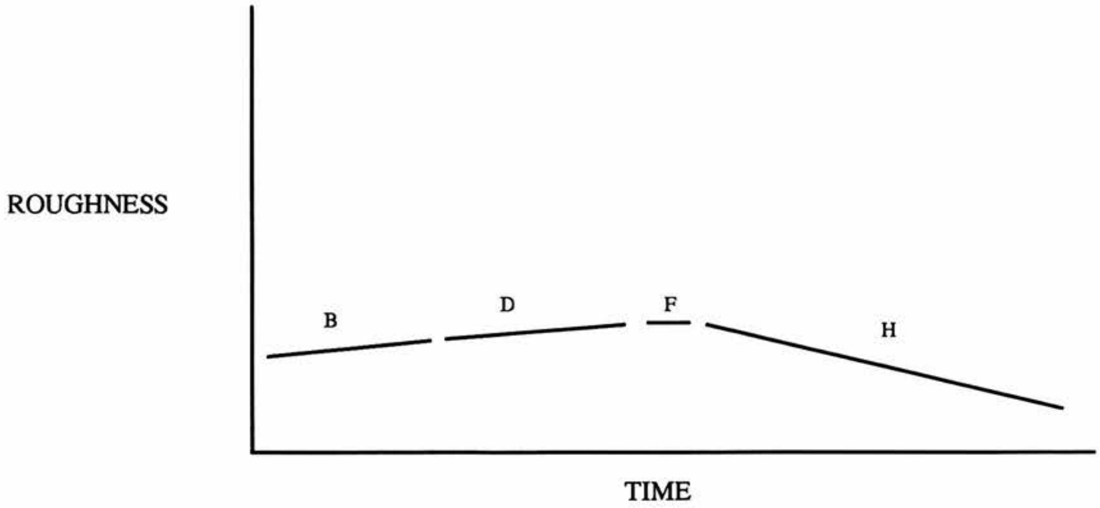
ROUGHNESS
B
D
F
H
TIME

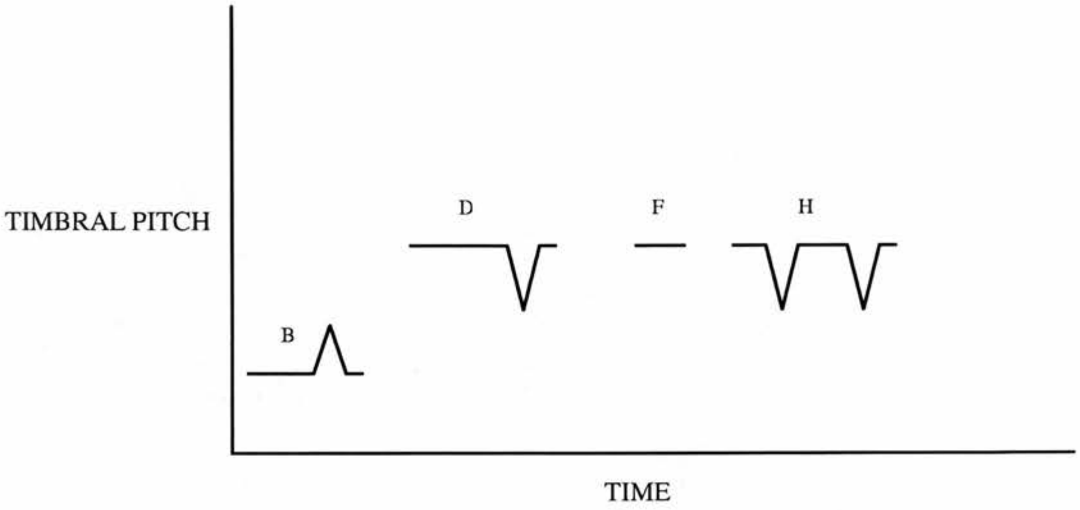
TIMBRAL PITCH
B
D
F
H
TIME

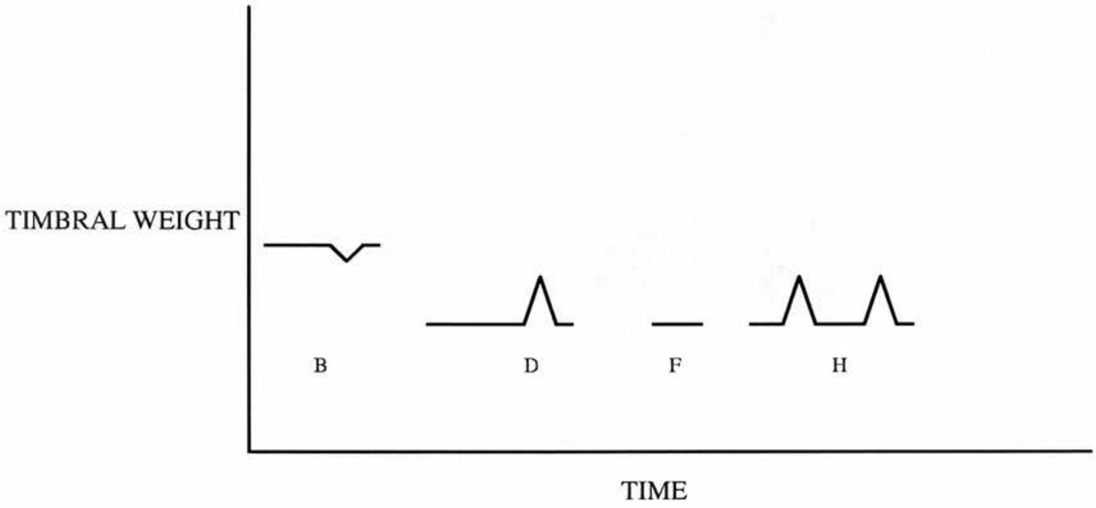
TIMBRAL WEIGHT
B
D
F
H
TIME

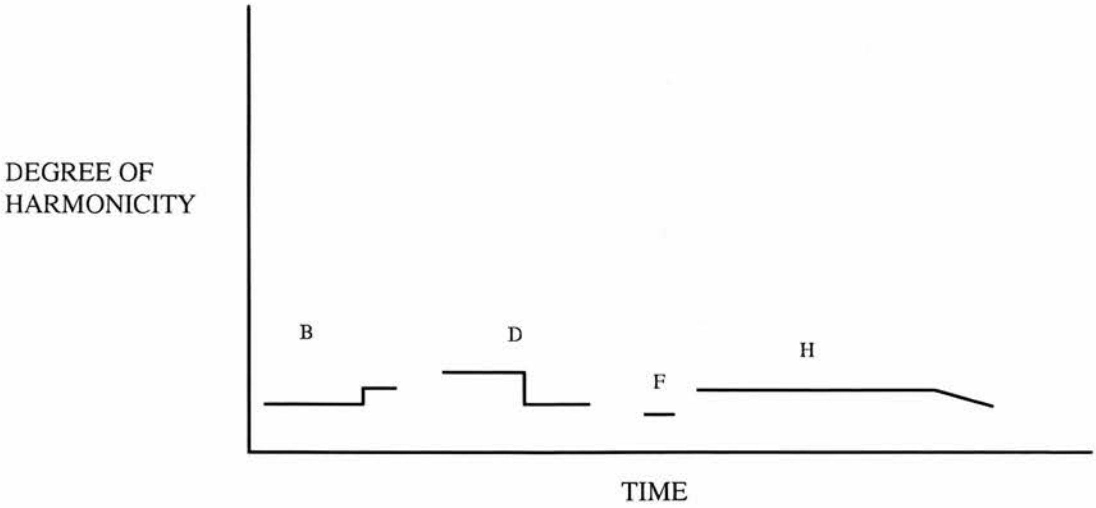
DEGREE OF HARMONICITY
B
D
F
H
TIME

We can also study the combined 'detail' graphs of the different recordings. In figures 3.39 and 3.40, we see, from top to bottom, graphs of timbral width, sharpness and roughness combined along the one time axis for the B-textures of recordings 1 and 2 respectively. These three measures represent three quite different aspects of our timbre perception, and thus complement each other well. The vertical axis has no scale because the graphs all use different scales, so the placement of the graphs relative to the vertical axis is based solely on considerations of presentation. The graphs, however, have the same relative dimensions as when they were shown separately.

When considering just one measure of timbre, there will be times when two events sound largely dissimilar, yet have the same measurement for that particular timbral parameter. Timbre is a multidimensional attribute, which is why a variety of timbre analytical tools are employed in this study. If we consider a number of timbral parameters together, we can begin to move towards a unique timbral identifier for a particular event. For example, '8 vc I' and '9 vc I' in section [H] are both *harmonic pizzicati* on E^5, played by the cello. In both recordings, its timbral profile consists of a rise in width, a fall in sharpness, and a rise in roughness.[6] '5 vc I' in section [H] consists of a *mf* semiquaver E^4 played by the cello - in both recordings its timbral profile consists of a *fall* in width, a fall in sharpness, and a *large* rise in roughness. These two profiles can be represented as follows:

---

[6] Note that '9 vc I' in recording 2 does not coincide with the rise in sharpness just prior to it. Although '9 vc I' does not coincide with a fall in sharpness for data at an average of 10, at an average of 5 the event does cause sharpness to return down after a brief rise.

	'8 vc I' and '9 vc I'	'5 vc I'
width -	↑	↓
sharpness -	↓	↓
roughness -	↑	↑ (+)

Thus, when these three timbral parameters are considered, we can see that most of the timbral difference between these two types of sound lies in timbral width, and the degree of roughness. Sometimes the same sound event, when compared between the recordings, has a different overall timbral profile - which will be due to differences in the way it is played or differences in recording style. For example, '10 vn III' in section [H] has the following timbral profile in recordings 1 and 2:

'10 vn III'

	Recording 1	Recording 2
width -	↑	↓
sharpness -	– (static)	↑
roughness -	↑	↑

Some of the possible reasons for these timbral differences were discussed above. Even when the movement of the timbral measures are in the same direction, important differences in degree of movement must be considered. For example, '3 vl I' in section [B] consists, in both recordings, of the three timbral parameters all moving upwards - but the degree to which they move upwards is very different between the recordings. The upwards movement of roughness in recording 1 associated with '3 vl I' is smaller than the immediately preceding roughness spikes, which are caused by roughness fluctuations in the 'background' string cluster. The same is true for the upwards movement of timbral width in recording 1. This is in marked contrast to the way the upwards movement in timbral width and roughness caused by '3 vl I' in recording 2 stands out markedly against the background fluctuations. Thus, these graphs show that each musical event represented here has a distinctive set of timbral parameters, and we can begin to see how it would be possible to designate degrees of timbral similarity between events based on these measures.

Figures 3.41, 3.42, 3.43 and 3.44 show graphs for timbral width, roughness and sharpness for the A- and B-textures in recordings 1 and 2. Figures 3.43 and 3.44 show the same graphs as figures 3.39 and 3.40, except without annotation, and with the addition of straight-line averages. All the graphs are presented in the same way as we saw in figures 3.39 and 3.40. In the graphs of the A-textures (figures 3.41 and 3.42), the upper-most set of data points and straight-line averages represent timbral width, the middle set of data represents roughness, and the lowest, connected points represent sharpness. In the graphs of the B-textures (figures 3.43 and 3.44) the positions of the separate graphs is changed in order to accommodate the graphs more easily. The upper-most set of data is still timbral width, but the graph of sharpness lies in the middle, and the graph of roughness lies at the bottom.

The detail of these individual graphs has been discussed at some length, and reasons have been put forward as to why the graphs move in the ways they do. The reader is now invited to browse these graphs in the same manner as one might read a musical score. As the lines of timbral movement are followed, I hope the reader will find that it is possible to hear inwardly how the timbre changes through time - timbre divorced from the perceptual considerations of pitch, loudness and texture. This ability may take practice, just as the ability to read a score fluently only comes with practice, but the achievement of this skill consolidates the ability to hear the timbre structures clearly, just as the ability to read a score consolidates the ability to hear pitch structures clearly.

Timbre analysis has allowed the timbral structure of the first movement of *Jeux vénitiens* to be concretely understood. The analysis has not had to resort to vague, colourful adjectives when describing the particular types of timbre that are present in this music, but has made use of finely discriminating analytical tools that enable the change from one type of timbre to another to be precisely measured in a way that relates directly to our perception of that change. Through this concrete understanding of the timbre in this music, we have uncovered a clear timbre structure on both the small and the large scale. In the B-textures, for example, on the small scale we have seen the timbral contributions of individual 'exclamations', and on the large scale, we have seen how both the 'exclamations' and the 'background' timbres go together to create overall timbral projections through the first movement. It is the timbre structure that gives form to the first movement - the pitch structure supports and helps to create the timbre structure - and it is only through this type of analysis that this aspect of the music can be clearly demonstrated.

***

CHAPTER 4

Analysis of *Atmosphères*

CHAPTER 4

Analysis of *Atmosphères*

(1961)

György Ligeti

1. Introduction

Atmosphères, scored for large symphony orchestra, is constructed from shifting clouds of notes which expand and contract, and which exhibit a distinctive internal life created through meticulous rhythmic and melodic counterpoint. This type of counterpoint Ligeti named micropolyphony. Ligeti himself said of this work: "*Atmosphères* is just a floating, fluctuating sound, although it is polyphonic."[1] The clouds of notes consist of two interlocking characteristics - a pitch structure, and a timbral structure. The timbral gestures require a pitch structure so that the gestures

---

[1] *György Ligeti in Conversation*: p.14.

may be realised. The use of subtle timbral shading and dramatic timbral contrast generates the energy and life of the piece. On the role of timbre in this work, Ligeti said: "*Atmosphères* is a composition in tone-colours *par excellence* and is closely connected with Schoenberg's third orchestral piece from his opus 16 [*Farben*]."[2]

Music requires time to unfold, and it creates a sense of its own metaphorical space as it does so. This especially applies to music, like *Atmosphères*, whose form arises out of the evolution and contrast of *plane*s of sound. Concerning this idea, Ligeti quotes Adorno on the

> conversion of temporal relationships into spatial ones [in contemporary serial music]. The course of the form is no longer experienced as a 'process of congestion and relaxation, but as a juxtaposition of colours and surfaces, just as in a picture.'[3]

In *Atmosphères*, form is created not out of melody and functional harmony, but rather, in the words of Ligeti, the "shaping of form... has been handed over in serial music to more complex categories, such as Groups, Structures or Textures, and, because of this, the way these are woven now takes over a very eminent role in the compositional design."[4] This weaving together of the strands of timbral and textural gestures forms the design of *Atmosphères*. Supporting this design is the meticulous pitch structure of the work. Ligeti has said that "the formal characteristics of *Atmosphères* operate on two levels, internal structure and audible form. The internal structure does not come through, you cannot actually hear it."[5] Except for the

---

[2] Ibid., p.86.

[3] Ligeti, *Die Reihe* (1960): p.15-16. Quoted from Adorno (1956): "Die Philosophie der Neuen Musik," from *Dissonanzen*, Göttingen: p.125-126.

[4] Ligeti (1960): p.14.

[5] *György Ligeti in Conversation*, p.43.

obvious shifts in pitch space, the pitch structure forms the 'internal structure,' and the textural and timbral gestures create the 'audible form.' The following analysis will attempt to show how these two structures operate and interlink.

2. Introduction to the Pitch Analysis

At the opening of this chapter, *Atmosphères* was said to consist of shifting clouds of notes. What method can we use to uncover pattern in the pitch movements?[6] The two methods which suggest themselves are pitch-class set theory and analysis based on note clusters. In this analysis, set theory proves occasionally useful for identifying particular collections of notes, but the most useful approach is to see the often very large note collections as clusters which change size, mutate and migrate through pitch space.

2.1 Defining a cluster

To assist the search for pattern in the sea of pitches of *Atmosphères*, I have formulated the following guidelines for cluster analysis.

A cluster can be defined by its central pitch or pitches, by the pitches which define its boundaries, and by its density. The vertical space between the outer pitches of a cluster can either be completely filled (a complete chromatic cluster) or there can be gaps, the notes which are present possibly falling into a symmetrical pattern, or some

---

[6] The pitch analysis will concentrate on the larger scale pitch distribution in the piece, rather than the rhythmic and melodic counterpoint between individual instrumental lines.

other structurally important shape. A cluster can move either by expansion or contraction (symmetrical or otherwise), by simple vertical movement (maintaining the interval between top and bottom), or a combination of these. The density of a cluster can wax and wane through the addition or subtraction of pitches. It is also quite possible for subgroups to exist within a cluster, defined through orchestration, style of playing, or pre-existence.

2.2 Analysing a succession of clusters

The similarities which are possible between clusters are pitch and pitch-class recurrences, recurrences of intervals (or interval-classes) of transposition, or recurrences of the intervallic pattern which forms a non fully-chromatic cluster. The overall movement of a cluster can also be significant - there can be a trend of upwards or downwards movement, of expansion or contraction, or a combination of these, the succession of moves itself possibly forming a pattern.

When looking for links between pitch clusters, it is important to be able to distinguish between important structural relationships, and those relationships which are of a more passing, local nature. To this end, the following guidelines are offered:

1) A similarity between clusters becomes more significant the greater the number of times that similarity occurs. A similarity may consist of the recurrence of pitches or intervals (or pitch- or interval-classes) or the recurrence of a process (for example, an interval of transposition).

2) A pitch recurrence that remains at the same octave is more significant than one that involves octave transposition.

3) An interval of transposition, or intervals of contraction or expansion, share the same level of structural significance as a pitch centre.

4) The structural importance of a subgroup is defined by the number of pitches it contains relative to other subgroups of the same cluster - the greater the number of pitches, the higher the level of significance. The same rule applies to the significance of the pitches forming the boundaries of different sized subgroups.

5) The significant pitches of an entire cluster have a higher level of significance than the significant pitches of a subgroup of that cluster.

6) Pitch centres have a higher level of significance than pitches that define boundaries.

7) Pitches that define boundaries have a higher level of structural significance than pitches emphasised through omission.

8) The hierarchy defined in 2) to 7) will be influenced by 1).

9) A single pitch-class or dyad, while recurring, can change function (for example, a pitch-class dyad may change from forming the centre of one cluster to forming the boundaries of another). Its significance is still defined by the guidelines above.

These analytical guidelines will allow a structural hierarchy to be proposed.

3. Pitch Analysis - the internal structure

An important compositional process in *Atmosphères* consists of the designation of certain pitches and pitch-classes as structurally important, while their structural role is free to change from cluster to cluster. For example, a particular dyad will be the centre of an entire cluster at one point, and then, in a subsequent cluster, one of the pitch-classes will be emphasised through omission (within an otherwise complete chromatic cluster) and the other pitch-class will be the cluster's lowest note (refer to analysis guideline 9).

Figure 4.1 (level 4 of a four level pitch analysis) shows a representation of the pitches which are present in *Atmosphères*. The graph shows pitch on the vertical axis (C^1, C^2, C^3, etc., and the 6th semitone of each octave is shown on the axis for reference) and time is represented along the horizontal axis. The time axis is to scale. To my knowledge, this is the first detailed representation of the entire pitch structure of *Atmosphères*.[7] The representation of the pitches in *Atmosphères* by this method is invaluable for a subsequent pitch analysis, as it is impossible to tell what pitch clusters are present, and where the chromatic gaps occur in these clusters through a simple examination of the score.

The pitch collections are shown either by vertical lines (often with gaps representing gaps in the chromatic collection) or by enclosed areas where this shows more clearly the movement in the pitch collection. Vertical lines are either drawn in brown (showing a wind collection) or in blue (showing a string collection). Where a brown and a blue line are side-by-side it signifies that the collections sound simultaneously, but the lines are separated for clarity. Although a little simplification of the pitch information is incorporated in this graph, generally a collection represented by a vertical line remains sounding until the next vertical line. Significant horizontal connections are shown by dashed green lines, and significant octave transpositions are shown by dashed red lines (or by dashed red arrows). In part 2 of the 3 part graphic

---

[7] A diagramatic pitch analysis is offered by Erki Salmenhaara (1969), *Das musikalische Material und seine Behandlung in den Werken "Apparitions", "Atmosphères", "Aventures", und "Requien", von György Ligeti*, Forschungsbeiträge zur Musikwissenschaft, Band XIX: Regensburg: p.86, as reproduced in Floros (1996): p.96. The boundaries of the clusters are represented, but there is no indication of the chromatic gaps in the clusters, no indication of orchestration, and no indication of structural relationships between clusters.

analysis, some pitches are shown in different colours in order to demonstrate pitch-class recurrences. Rehearsal letters from the score are placed above the graph.

The text of the following pitch analysis is a commentary on the graph of the pitch relationships shown in Figure 4.1. First a section is described in detail, then structural implications are drawn out from this description. The structural implications of certain pitches gradually become apparent as their influence is observed to stretch across the work.

The opening of the music consists of a pitch cluster spanning D^2 to $C\#^7$. The strings fall into three sections divided by gaps in the cluster. The violins span Bb^4 - $C\#^7$, the violas and cellos C^3 - G^4, and the double basses Eb^2 - Bb^2. The main wind cluster sits inside this string span: Ab^3 - C^5. The pitch D^2, played by the contrabassoon, sits below the string cluster. The dyad forming the centre of the violin subgroup (B^5/C^6) and the dyad forming the centre of the viola/cello subgroup (A/Bb^3) are each 'sandwiched' between a change of instrument type (see examples 4.1a and b respectively). This is an interesting example of Ligeti emphasising a structural point through the text of the score (the 'internal structure') rather than audibly (the 'audible form').

Example 4.1

	:				:		
a)	:			b)	:		
	violin 1	Eb^6			viola	$C\#^4$	
↑	violin 1	D^6			viola	C^4	
	violin 1	$C\#^6$			viola	B^3	
pitch	**violin 2**	C^6	} central		**cello**	Bb^3	} central
	violin 1	B^5	} dyad		**viola**	A^3	} dyad
	violin 2	Bb^5			cello	Ab^3	
↓	violin 2	A^5			cello	G^3	
	violin 2	Ab^5			cello	$F\#^3$	
	:				:		

At letter [A], the cluster contracts from 60 semitones in height to just 20 semitones. This collection is played solely by the violas and cellos. At letter [B], the cluster expands, covering a range from Eb^1 to E^7 (74 semitones). As at the opening, the wind subgroup sits inside the span of the string subgroup, and there are various gaps in the total chromatic of the cluster. At letter [C], the cluster contracts once again (to 33 semitones). In the first of two points to be used in this analysis which have been made previously by Jonathan Bernard,[8] a structurally important symmetry occurs during these expansions and contractions. From the wind subgroup at the opening (D^2 - C^5) to the wind subgroup at [B] (Eb^1 - C#^{6}) there is a cluster expansion downwards of 11 semitones and an expansion upwards of 13 semitones - an [11,13] expansion, as Bernard calls it. From the string subgroup at [B] (Ab^2 - E^7) to the string cluster at [C] (G^3 - Eb^6) there is an [11,13] contraction. This is the only instance that I have found in this piece of a symmetrical arrangement of expansion and contraction.

Let us go back to examine more closely the structure of these pitch collections in order to see the relationships that exist between them.

The boundary of the opening pitch collection consists of D^2 and C#^{7}. These same pitch-classes (p-c's) are found at the centre of the viola/cello collection at [A] - C#/D^4 - sandwiched by the orchestration in the same way as shown in example 4.1b. The topmost note at [A], B^4, has become the topmost note of the wind subgroup just before the entry of the cluster at [A] (this is shown on the graph by a cross above the pitch B^4, showing that C^5 ceases at this point). The lowest note at [A], E^3, does not appear to have a derivation from the opening, unless it be from E^4, the centre of the upper-wind subgroup at the opening (this upper subgroup includes all wind pitches

---

[8] Bernard (1987): p.210

but the lowest - D^2). This relationship consists of a note which is a centre of a subgroup and a note which defines a lower boundary, and nothing happens subsequently to suggest this relationship is part of a distinctive pattern involving the p-c E in this way. The first relationship noted between the opening and [B] was between the boundaries of a complete cluster and the centre of a complete cluster (involving octave transposition), and the second relationship was between the top note of a significant subgroup (the wind subgroup) and the top note of a complete cluster (involving no octave transposition). According to the guidelines proposed earlier (in particular, numbers 2, 5, and 6), the E^3 / E^4 relationship will be of much lesser significance than the other two relationships described above.

Looking now at the pitches at letter [B], we can observe the beginnings of a network of relationships. E^4, the centre of the upper wind subgroup at the opening, now forms the upper note of the central dyad of the entire cluster at [B], Eb/E^4. C^5, the upper note of the wind subgroup at the opening, is now the sole gap in an otherwise complete chromatic string subgroup, and is the centre of this string subgroup. Looking at the structure of the clusters before and after [B], we find that p-c C moves through the work by octave transposition between the opening and [N]. C^5 is the top note of the wind subgroup at the opening; C^6 is represented by a gap in the wind subgroup at [B]; C^7 is defined by a similar gap at [E]; and C^8 is the top of the cluster at [N], and the highest pitch of the entire piece. P-c Ab moves through the piece in a similar way, but balances the movement of p-c C by moving through *downward* octave transposition. Ab^3, the bass of the wind subgroup at the opening, drops to Ab^2, which is the lowest note of the string subgroup at [B]; Ab^1 is then found at the top of the double bass cluster at [G].

Returning to the cluster at [B], B^4, the upper note of the string cluster at [A], is emphasised by a gap at [B] in the wind subgroup. A further gap in the wind subgroup, G^3, is the pitch which lies at the centre of the entire wind subgroup at the opening. Looking at the boundaries of the pitch collection at [B], we find that the central p-c's (Eb/E) reappear here - Eb^1 at the bottom and E^7 at the top. This emphasis on these p-c's is significant, for, as demonstrated by Bernard (1987) in this second point to be quoted from his study of this piece, the pitch structure of the music from [B] to [D] is determined very largely through a process of symmetrical contraction centred on Eb/E^4. This contraction causes the tops of the clusters to move from E^7 to D^6 to Ab^4. The bottom of the clusters travel from Eb^1 to F^2 to B^3. The collection at the beginning of [C], which is not symmetrical around Eb/E^4, is the starting point for a downward sweep of pitches (shown by the dashed line) which brings the music to the symmetrical arrangement that we find prior to [D].

So far, we have seen C^5 emphasised at the opening, B^4 emphasised at [A], and both these pitches emphasised at [B]. This dyad appears again at [C], where C^5 is the sole gap in a chromatic string cluster, and B^4 is this cluster's centre. The dyad symmetrically opposite B^4/C^5, taking the dyad Eb/E^4 as the centre, is G/Ab^3. We find this dyad at the opening - G^3 lying at the centre of the principal wind subgroup and Ab^3 at the bottom of the upper wind subgroup (which excludes D^2). This dyad is found once more at [B] - Ab^3 at the centre of the wind subgroup, and G^3 emphasised through omission. At [C], Ab^3 is no longer emphasised, but G^3 is present as the bass of the cluster. Thus, these two dyads, B^4/C^5 and G/Ab^3, lying symmetrically around Eb/E^4, are emphasised in various ways from the opening to [C]. At a point just prior to [D], the end-point of the symmetrical contraction around Eb/E^4, these two symmetrical dyads move through octave displacement so that their position relative to the central dyad is reversed. B^4/C^5 now lies at B^3/C^4 - the pitches forming the bottom of the double bass and wind subgroups respectively - and G/Ab^3 now lies at G/Ab^4 - the pitches forming the top of the wind and double bass subgroups respectively. (The section from [D] to [E] consists solely of the double bass cluster, B^3 - Ab^4; the wind and double bass clusters sound together towards the end of section [C], within the large string cluster - F^2 - D^6). These dyadic transpositional relationships are shown by dashed red lines joining B^4/C^5 at [C] to B^3/C^4 just prior to [D], and G/Ab^3 at [B] to G/Ab^4 just prior to [D]. Because G^3 also appears at [C], its relationship to G^4 at [D] is shown as well. The dyad G/Ab^3 at [B] not only has a relationship to this same dyad at the opening, but also to the dyad G/Ab at the opening which is found at the fourth octave. This dyad forms the centre of the entire opening cluster.

Looking at centres of subgroups which have not yet been discussed, the dyad that forms the centre of the cello/viola subgroup at the Opening, A/Bb^3, is found an octave higher at [B], A/Bb^4, where it forms the centre of the upper wind subgroup. This octave movement forms a balance to the G/Ab octave relationship discussed above which occurs between the Opening and [B]. The centre of the violin subgroup at the Opening is B^5/C^6. C^6 is emphasised through omission at [B]. C#^{6} at [B] lies an octave below C#^{7} found at the top of the clusters of both the opening and [E], thus forming a bridge between these two occurrences of C#^{7}.

It will be remembered that the symmetrical dyads G/Ab and B/C cross over just prior to [D]. This crossover movement is in fact extended through to letter [H] with one note of each dyad. The note G^3 at [C] climbs the octave to G^4 prior to [D] (the crossover), to G^5 at [F] (the lower limit of the string sub-group, which has Ab^6 as its top-most note - see below), and to G^6 at [H] (the top note of the cluster). The pitch C^5 at [C] undertakes a complementary movement. It moves to C^4 just prior to [D] (the crossover), and then to C^3 at [H] (the bottom note of the cluster) - see example 4.2 .

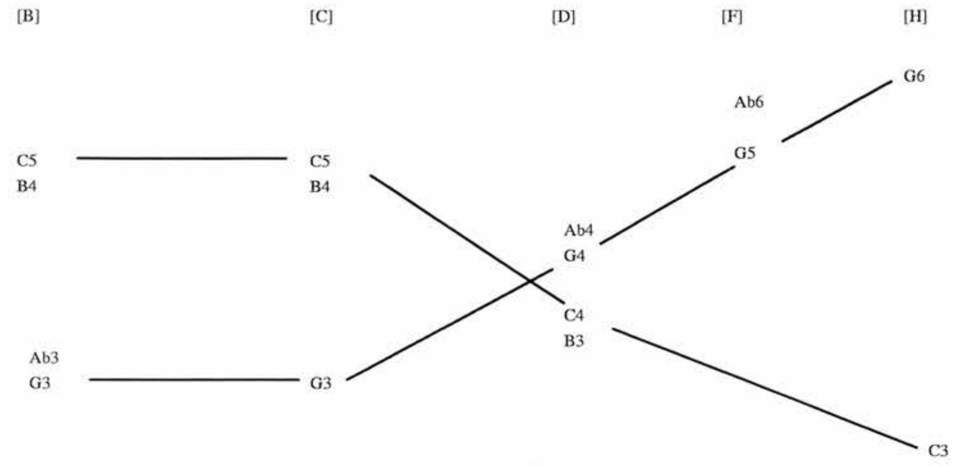

Example 4.2

The dyad Ab/G is also found at [F], split apart so that its pitches are Ab^6 and G^5 (shown in example 4.2 as part of the upward octave transposition of p-c G discussed above). They form the top and bottom boundaries to the string subgroup at this point.

The dyad Ab/G^4, found just prior to [D], plays an important role at letter [G], the centre of the work. From [F] to [G] the piccolos have slowly climbed upwards till they reach the chromatic cluster whose range is G^7 to Bb^7. Immediately the piccolos cease, the double basses enter with their chromatic cluster six octaves lower - C#^{1} to Ab^1. The central dyad between the 'inner edges' of these two clusters - between G^7 and Ab^1 - is G/Ab^4, the same p-c's which form the 'inner edge' boundary around this centre. This organisational technique was previously found at [B], using the p-c's Eb/E. G^4 also plays a role between the appearances of G/Ab^4 just prior to [D] and at [G]: G^4 forms the centre of all the cluster movements between [E] and [F] - in other words, it is the centre between C#^{2} and C#^{7}.

From [E], the dyad Eb/E^4 ceases to be the chief influence, and the dyad C/C# becomes more structurally important. The p-c C# is found at the bottom of two cluster movements - the first is a sudden chromatic 'stab' downwards by the cellos to C#^{2}, the second a more gradual downward movement by the violas to C#^{3}. The p-c C# is also found at the bottom of the static cluster of double basses at [G] - C#^{1}- as well as at the top of the entire cluster at [E] - C#^{7} - which is a reference back to the top note of the cluster at the opening of the work. This C#^{7} is accompanied by a gap in the cluster at C^{7}, thus forming the dyad C/C#^{7}. This dyad is also found in the fifth octave - C^{5} is the bass of a wind sub-group, and C#^{5} is the centre of the cluster which excludes the aforementioned cello chromatic 'stab' downwards to C#^{2}. The wind subgroup mentioned above covers an octave from bottom to top by the time the movement of this cluster has concluded - C^{5} to C^{6} (a relationship shown by a dashed red line on the graph). This C^{6} is also the centre of the entire cluster at [F]. Just above it, there is the centre of the string subgroup - C#/D^{6}. Thus, here too, we have C associated with C#, even though the C# is associated with another dyad - a dyad previously found in the fourth octave at letters [A] and [B] as the centre of the cluster played by the violas and cellos.

At letter [H], the boundaries of the cluster are C^{3} and G^{6}, and the derivation of these pitches has already been discussed (see example 4.2). One bar before [I], the notes cascade inwards towards each other to form the cluster at [I] - a chromatic cluster spanning G^{3} to B^{4}. B^{4} is the same pitch that we saw with C^{5} at [A], [B] and [C], and G^{3} is the same pitch we saw with Ab^{3} at the opening, at [B], and by itself at [C]. Indeed, Ab^{3} soon enters: it is the lowest note of the cluster that forms two bars later. The highest note of this cluster is F^{4} which is also found as the lower note of the central dyad at [G] - the centre that *includes* the pitches above G^{7} and below Ab^{1}. The centre of the cluster after [I] is C/C#^{4}, which are the sole two pitches which remain after the cluster funnels towards [L]. This dyad at [L], not surprisingly, is the smallest cluster of the work. From [L] to [M], the cluster expands upwards to F^{4} (the same top note found at [I]), till at [M], another cluster enters below it.

Between [H] and [M] a collection of pitch centres has developed, with C/C#^{4} at its centre:

Eb^4
D^4
C#^{4}
C^4
B^3
Bb^3

D^4 is the centre of the cluster at [I], and D/Eb^4 is the centre of the upper wind subgroup at [M]. B^3/C^4 is the centre of the double bass subgroup at [I] and the centre of the entire cluster at [M]. Bb^3 is the centre of the entire cluster at [H], and occupies a gap (along with B^3) at [M].

With the entry of the lower cluster at [M], the dyad E/Eb makes a brief reappearance (in the third octave) as the centre of the expanded cluster, just after [M]. The dyad Ab/G^3 reappears as the centre of the lower cluster at [M] before its expansion. There is also an interesting 'cross' formation between this cluster and the next at [N]: A^3 is found at the top of the lower cluster at [M], A^2 at the bottom of the cluster at [N]; Bb^2 is found at the bottom of the cluster at [M] after its expansion, and Bb^3 is represented by a gap in the cluster at [N] (this cross is shown on the graph).

The large cluster at [N] slowly mutates into a 'white note' collection that arrives fully at [O]. The top note of this cluster has already been shown to be part of a large scale progression of the p-c C. The lowest note of this cluster, A^2, is the same as the lowest note of the previous 'white note' cluster that occurred between [B] and [C] (as part of a section that mutates from a 'white note' to a 'black note' cluster). The centre of the cluster at [N] (and at [O]) is the dyad E/F^5. This p-c dyad is also found at the centre of the double bass cluster at [G] (E/F^1) and the p-c F is also emphasised through its omission (F^3) from the otherwise complete 'white note' collection at [O]. Just after [O], the 'white note' collection narrows to just four 'white notes' ranging from F^4 to B^4. B^4 was the pitch at the top of the cluster at [I], and was also part of the important dyad C^5/B^4 found in the first half of the piece. B^4 is also the centre of the cluster at [Q]. The p-c F has played important roles at [G] (F^1), after [I] (F^4), and at [M] (F^4), [N] (F^5), [O] (F^3), and just after [O] (F^4)(see above). This means that from [G] to [O] p-c F has been emphasised in octaves 1, 3, 4 and 5.

At [R], the flutes enter above a cluster in the strings. The bottom note of the flutes is the same as the top note of the string cluster at [Q] - Ab^5; the centre of the entire cluster at [R] (which includes the strings) is A^3 - the same as the top note of the lower cluster at [M]; and the top note of the strings is A^2 - the same note as the bottom note of the cluster between [N] and [O], and of the 'white note' cluster at [B]. After shifting up, down, and up again, the flutes join a string cluster at [T] made up of clouds of rapid note movement. At [T] there is a single very brief appearance of the note A^3 - a discussion of the previous appearances of this note is found above. Apart from this single appearance of A^3, the lowest two notes of the string cluster at [T] are C^4 and $C\#^4$ - the dyad that played an important role in the middle section of the work. The string cluster itself omits the pitches Ab^4, F^4 and Eb^4 from the complete chromatic. P-c Ab has featured with G to form a significant dyad through the first half of the piece, and more recently Ab^4 was found as the top note of the central dyad at [G]. F^4 was the bottom pitch of the cluster to which the large cluster at [O] contracts just prior to [P], and also featured at [G], in section [I], and at [M]. And Eb^4, with its neighbour E^4, has formed the important dyad around which the pitch structure of the first half of the piece has been built. On the entrance of the trombone and tuba cluster near the end of the work (whose top note, G^1, is the same as the bottom note of the string cluster at [R]) the entire cluster (excluding the pitches of the swept piano strings - the pitches of which are not precisely specified) has the same boundaries and central dyad as were found at letter [B]. In both cases the lowest note is played by the tuba. The top note of the flute subgroup, C^6, is the same pitch which at [B] was among those pitches emphasised through omission, and at [F] where C^6 was the centre of the entire cluster.

Despite the occurrence of these pitch references, from [N] until the final cluster of the work, and to a lesser degree from [G] onwards, the pitch structure appears much freer than at the beginning of the work. Local pitch references are certainly made, but the large-scale organising principles discussed in relation to the first half of the piece up to [G] are not in evidence again until the very end of the work where the same central dyad and boundary pitches return as were previously found at [B]. As we will see in the discussion of timbre that follows, this freer pitch structure mirrors the freer timbral structure in the second half of the piece.

The analysis so far has demonstrated the connections that can be drawn between particular pitch-classes, and has shown the structural importance of some of these pitch-classes. I shall now move to a discussion of a reductionist analysis - roughly Schenkerian in approach. In the process of this discussion, the level of significance of pitch-class relationships will be ascertained using the previously proposed guidelines for cluster analysis.

Figure 4.2 shows the third level of my analysis - the next higher level of analysis from that shown in Figure 1. This level of analysis is greatly simplified when compared with the previous level. All pitches now fall into a range of a little under one-and-a-half octaves, there are far fewer pitches represented, and the time axis is no longer to scale.

How has this level been derived from level 4 ?

Looking at the guidelines for cluster analysis proposed earlier, let us examine the impact that the application of guidelines 1 and 2 have on the process of deriving level 3.

Guideline 1 states that 'A similarity between clusters becomes more significant the greater the number of times that similarity occurs. A similarity may be the recurrence of pitches or intervals (or pitch- or interval-classes) or the recurrence of a process (for example, an interval of transposition).'

Guideline 2 states that 'A pitch recurrence that remains at the same octave is more significant than one that involves octave transposition.'

The dyad Eb/E^4 is a significant factor in the ordering of pitches between [B] and [E] - it is the centre of a large-scale symmetrical pitch contraction. Further, the pitch E^4 is referred to at the opening as the centre of the upper wind subgroup. Thus, apart from section [A] and the cluster heard initially in section [C], the dyad Eb/E^4 lies at the centre of an entire cluster, or the upper note of the dyad lies at the centre of a significant subgroup (also see guidelines 5 and 6).

During the period in which Eb/E^4 is functioning as a significant organisational factor, the pitches that comprise the dyads which are placed symmetrically around it - G/Ab^3 and B^4/C^5 - are functioning as centres of clusters or cluster subgroups, boundary limits of clusters, or as gaps in otherwise complete or near-complete chromatic clusters. These dyads remain without significant octave transposition until just prior to [D] where they cross over at the end-point of the symmetrical cluster contraction.

Apart from the central dyad and its two symmetrical satellites, there are no other structurally important pitch connections up to [E] which remain at the one octave for as long. Thus, these three dyads are chosen to represent the music at level 3 up to [E], and the crossover of the satellite dyads is also included because of the significance of its placement (coming at the conclusion of a large-scale symmetrical contraction) and because of later significant statements of these dyads at these octave placements.

At [E], the dyad C/C#^{5} becomes an important structural consideration. There are many statements of p-c's C and C# between [E] and the point up to and including [G]. The fifth octave is initially chosen for this dyad at level 3 because of the important role C^{5} has played previously (guideline 2; at [E], C^{5} is the lowest note of the wind subgroup) and because of the placement of C#^{5} at the centre of most of the cluster movement between [E] and [F]. At [G], G/Ab^{4}, the position of this dyad after the crossover, plays an important structural role as centre of the 'inner edges' of the two far-flung clusters. Thus, this dyad is restated at this point at level 3.

Between [I] and a point a little after [M], the dyad G/Ab^{3} (its position before the crossover) reappears first as the successive pitches at the bottom of two adjoining clusters and second as the centre of a subgroup. The dyad C/C#^{4} appears as the centre of a cluster, and as the point of focus of cluster contraction. Both these dyads are shown at Level 3. Other important pitches in this passage either do not play a structural role for as long, or do not refer back to previous statements at the same octave.

As mentioned previously, from [N] until the reappearance of the central dyad Eb/E^{4} during the closing stages of the work, the pitch structure becomes much harder to characterise. Structural pitches change with the shifting clusters, and only local pitch references appear to be at work. This structural discontinuity is shown at level 3 by the placement of A^2 at [N] and B^4 and F^4 at [O]. A^2 is the lowest note of the cluster at [N] through to [O], and is the top note of the string subgroup at [R]. This pitch has no previous references at level 3. B^4 is the top note of the cluster to which the large 'white note' cluster at [O] contracts, and is the centre of the cluster at [Q]. At level 3 it has been associated with the dyad B^4/C^5. F^4 is the bottom note of the cluster to which the cluster at [O] contracts, and has previously been found at [G], after [I], and at [M], but has not previously appeared at level 3. A^2 and B^4 have an influence on the pitch structure from [N] to [T] for the greatest period of time when compared to other structural pitches. F^4 is included because of its importance from [G] to [O], and because it represents the other significant occurrences of p-c F which have been discussed above.

From [T] to the conclusion of the piece, the dyad Eb/E^4 has the most significant influence on the pitch structure. The dyad actually becomes the centre of the cluster at the entrance of the trombone and tuba subgroup which has G^1 as its top note. However, its influence is felt at [T], as it is here that the shape of the upper subgroup of the final cluster is established. Thus, the dyad Eb/E^4 is placed at [T] at level 3.

The structure of level 3 is thus based around the dyad Eb/E^4, and its two satellite dyads B^4/C^5 and G/Ab^3. The dyad C/C# moves up from B/C at both the fourth and fifth octaves, C acting as a 'pivot' to the new dyad. A degree of pitch discontinuity occurs at [N], and to a lesser extent at [O]. At [T] we return to the pitch structure associated particularly with [B].

This notion of pivoting to a new dyad around the p-c C leads to the idea of 'dyadic modulation.'

Up to letter [E], the following dyads were found to be significant at levels 3 and 4:

B/C C#/D Eb/E G/Ab A/Bb

Out of the 12 p-c's, this list only omits the dyad F/F#. This dyad appears only once as a structural unit - at the centre of the cluster at [G] (the centre which includes the pitches above G^{7} and below Ab^{1}). Thus, all the significant dyads up till [E] fall into a regular dyadic partitioning of the 12 chromatic pitch classes. At [E], the dyad C/C# becomes important - a dyad which does not fall into the partitioning described above. At [G], the 12 note collection of dyads becomes complete with the inclusion of F/F#. After this completion, from [H] to [L] the clusters funnel towards the 'new' dyad C/C#^{4}. At the conclusion of the work, we return to the pitch structure found at [B]. Along with the appearance of the dyad C/C#, other dyads appear in the work, at around the same time as the dyad C/C#, which are not part of the partitioning described above, and therefore they do not relate back to the important structural dyads in the first half of the work.

Thus, a loose 'modulation' is proposed that takes place in the music at the point the dyad C/C# becomes structurally important. From a dyadic partitioning of:

B/C C#/D Eb/E F/F# G/Ab A/Bb

the pitch-set transposes, or 'modulates' to

C/C# D/Eb E/F F#/G Ab/A Bb/B

the dyad C/C# acting as the 'pivot'. This is not a strict modulation, in the sense that *only* the above partitioning is found from [E] onwards. Indeed, out of these six new dyads, only four are found to be significant - C/C#, D/Eb (not previously mentioned - centre of the upper-wind cluster at [M]), E/F, F#/G (not previously mentioned - centre of the cluster at [T]), and these dyads are of a lesser significance than those designated significant in the first half of the piece (prior to [G]) because the time over which they exert an influence is shorter. However, their presence is noteworthy due to their absence from the first half, and because of their relationship to the important dyad C/C#. At level 3 of my reductionist analysis, only the dyad C/C#, the 'pivot', is shown.

At level 2 (figure 4.3) we see that the satellite dyads to Eb/E^4 now only appear once, and the crossover is omitted (see guideline 2). The movement to the dyad C/C# still occurs at two octaves, however the upper octave is placed in square brackets as I consider the move to C/C# at the lower octave to be more structurally important. This is because C/C#^{4} is structurally significant for a longer period of time than C/C#^{5}, and because the music emphasises C/C#^{4} very clearly as the pitches funnel towards this dyad at [L]. All that remains at level 2 is for the music to move back to the central dyad Eb/E^4 at the end.

Level 1 (figure 4.4) shows the highest level of pitch analysis. The satellite dyads are omitted because they are dependent on Eb/E^4, and the movement to C/C#^{4} is now only shown in the structurally more important octave.

If we were to go to an even higher level of abstraction, it would be to a statement of the p-c dyad Eb/E, without reference to octave placement, or to the movement to C/C#. This dyad can be seen as the generator of the pitch structure - a sort of *ursatz* at the root of the cluster movement.

In this pitch analysis, various traces have been uncovered of how the composer may have ordered the placement of pitches. Sometimes it appears that the composer's methods are relatively clear, and sometimes methods of pitch organisation are not apparent. A quotation relating to this point is found in "Metamorphoses of Musical Form", written by Ligeti in 1960 (*Atmosphères* was written in 1961):

> ...the relation between the pre-formed plan of control [of contemporary serial compositions] and the form that proceeds from it is no longer fixed and unequivocal; the musical realisation seems rather to have a constant feed-back effect on the control plan itself... A method of composition - and with it, a form - will only free itself from automatism and dependence on home-made material, when at every moment the composer has the possibility of taking a decision that will alter the future course of the piece entirely. Such heterogeneous happenings [if the form of the work is to be preserved] should have the faculty of influencing each other back and forth, thus making it possible for gradual transformations to take place as well as sudden mutations.[9]

This may partly explain why conjectured methods of pitch organisation at times appear easily out of the above analysis, and why at other times they remain submerged beneath the surface of the composition. The constant tugging back and forth between Ligeti's presumed method of pitch organisation, arrived at through analysis, and the given realisation of that organisation (the score), brings up uncertainties as to how to explain the structure of *Atmosphères*. These uncertainties are partly confronted in the sectional analysis of the piece which appears below, based on the piece's textural gestures - its 'audible form.' This is where a conventional music analytical endeavour would cease. In this thesis, however, we are then able to explore specifically the realm of timbre - the aspect of the piece which Ligeti himself stated was realised in *Atmosphères "par excellence."*

---

[9] Ligeti (1960): p.12-13.

4. Sectional Analysis - the audible form

In the 1960 article quoted above, Ligeti addresses the new emerging freedom of form evident in contemporary serial compositions. He talks of the concepts of *permeability* and of *insensibility to intervals*:

> structures of different textures can run concurrently, penetrate each other and even merge into one another completely [the concept of *permeability*], whereby the horizontal and vertical density-relationships are altered, it is true, but it is a matter of indifference which intervals coincide in the thick of the fray [the concept of *insensibility to intervals*].[10]

Atmosphères is not concerned with different textures running concurrently, and merging with each other, but rather with the next evolutionary step: the use of a single mutating texture. *Atmosphères* is a textural monody consisting of the play of changing timbre against the background of changing pitch. The idea of insensibility to intervals is developed to the extent that Ligeti achieves an effect where there is no foreground or background - at any moment all intervals are as important as each other - and the predominant interval, the semitone, is virtually never heard within the piece's texture. The play of pitch, rhythm and timbre produces a particular shape, or gesture, and from an examination of these gestures, it is possible to propose a way that the work may be heard to fall into a number of sections. The role of timbre in the following sectional analysis is considered from a subjective point-of-view as a part of the gestures of the work. An examination of timbre in its own right follows this sectional analysis.

---

[10] Ibid.,p.8.

I propose the following as a formal plan of the piece:

Example 4.3

Opening	;	[A] [B] [C]	;	[D]	;	[E] [F] [G]	;
Introduction	;	Section 1	;	Interlude	;	Section 2 (bridge)	;

[H] [I]	;	[J] [K] [L] [M]	;	[N] [O] [P]	;	[Q] - end
Episode 1	;	Episode 2	;	Episode 3	;	Coda

Sections consist of three parts in the relationship of beginning, middle and end; *Episodes* consist of a single span directed towards an end-point.

The Opening consists of a diffuse texture without any articulatory features - the aural equivalent of a static fog. The volume dies away as we approach letter [A].

At [A], the pitches undergo a pronounced contraction which has the effect of creating a much more concentrated texture. Between [A] and [B] there is a series of crescendos and diminuendos, which fan out symmetrically from the centre of the pitch cluster (C#/D^{4}). See example 4.4.

Example 4.4

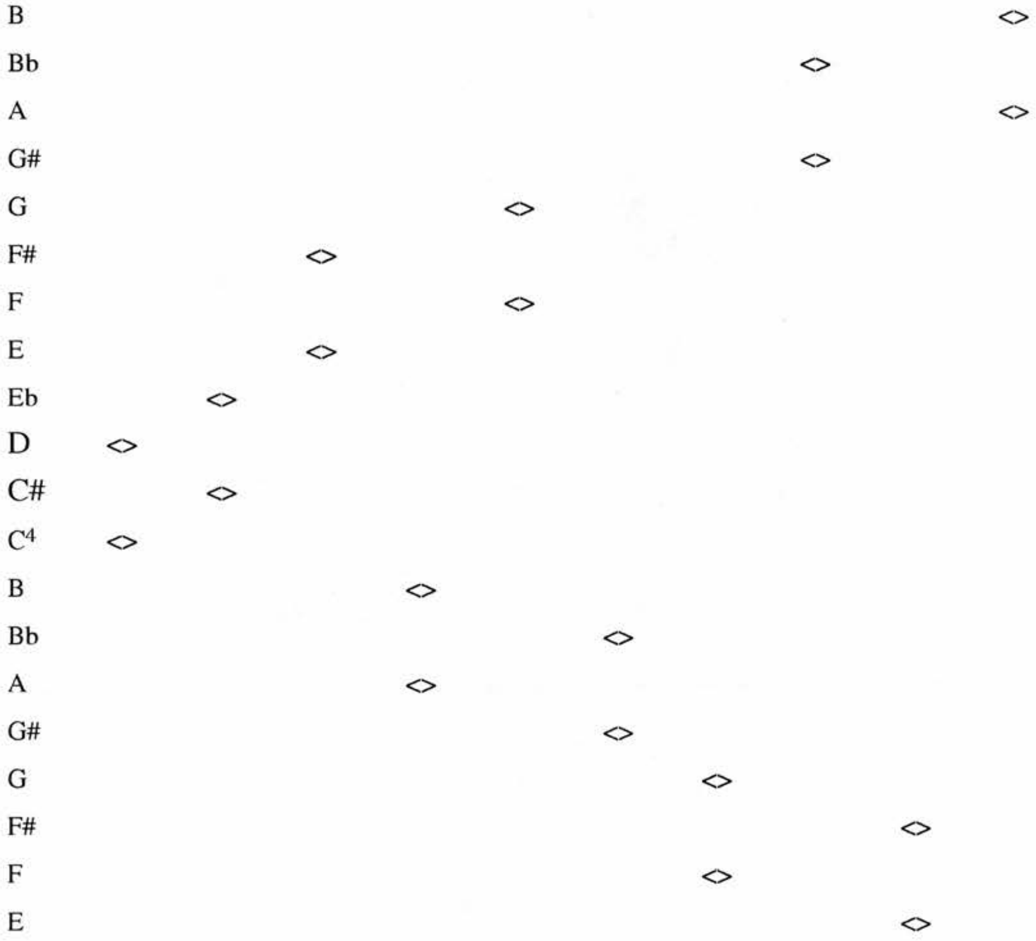

These crescendos and diminuendos propel the music into the much expanded cluster at [B]. During [B] there is a series of re-articulations of pitches, however they are all marked to be played "as imperceptibly as possible," so the effect is more of a hushed ripple than of obvious beginnings and endings. There are 33 re-articulations in this section. Their order relative to the pitches that are re-articulated is shown in figure 4.5. A loose patterning of the re-articulations is clear from this figure.

It is during section [B] that the loudest moments of Section 1 occur, with the alternation of the white and black note collections through the use of crescendos and diminuendos - notes don't drop out, but are dominated by the much greater loudness of the particular chosen collection. Section [C] concludes Section 1 with a progressively busier murmur of strings and winds which coincides with a progressive lowering and spreading of the cluster and the entry of the bass and wind subgroup.

The Interlude of section [D] consists of the static double bass cluster, without any notated articulatory features.

Section 2 shows a much greater degree of movement in the clusters in comparison with Section 1. Movement of pitch has, up until now, either been a matter of changing the cluster at a definite moment to a very different cluster, and then maintaining that cluster for a time, or has consisted of gradual shifts. From letter [E], there is much more direction and urgency in the manner of pitch changes, as may be seen from the level 4 analysis. This movement, which is principally upwards, is achieved through a web of rhythmic counterpoint, so that the music slowly 'crawls' upwards, until we are left with the high piccolo cluster, which becomes continuously louder as we approach [G]. [G] is the emotional climax of the work - the now quadruple *f* piccolos suddenly give way to a chromatic double bass cluster over six octaves lower, also played quadruple *f,* and marked *tutta la forza, tenuto*. In the performance instructions written by Ligeti at the front of the score, it says of this moment that "The contrabass entrance follows the piccolos immediately, without a caesura," which supports my belief that letter [G] does not mark the beginning of a new section, but rather the climax of a process begun at [E]. The double bass cluster after [G] is held for 14 seconds (at crotchet = 60) without any notated articulations, though the randomly occurring bow changes are heard. I hear the section between [G] and [H] as initially a dramatic climax, and then as an interlude, or bridge, until the appearance of the new ideas presented at [H].

To recapitulate, Sections 1 and 2 both have a structure suggestive of an introduction, a climax, and an ending. In Section 1 section [A] acts as the beginning, section [B] the climax of section 1, and section [C] the end, where the music quietens to a murmur. The opening of Section 2 launches the piece into a much greater degree of activity, pushing the pitches higher and higher, until the climax at letter [G]. We then have a period of relative stasis and calm - the ending of Section 2. We will find that this structure is supported by the timbral analysis that follows.

The remainder of the piece is partitioned into three Episodes and a Coda. Episode 1 comprises sections [H] and [I]. Section [H] consists of descending and ascending strands forming a contrapuntal web. In the violins, every instrumental part descends from the top of the violin cluster to its bottom, and then begins its descent once again from the top of the cluster - each part is off-set from all other parts, and there is close, but not exact, contrapuntal imitation in rhythm and pitch. Although the individual parts descend, the overall pitch range of the violins remains the same until one bar and one crotchet before [I]. At this point, the violins cascade downwards towards the bar of [I]. The violas and cellos, on the other hand, play a rising contrapuntal figure that comes up to and crosses the activity of the violins. From letter [I], the violin, viola and cello parts increasingly overlap. From letters [I] to [J], all parts converge towards a cluster ranging from Bb^{3} to C#^{4} (see figure 4.1). Episode 1 begins quietly (all parts marked quadruple *p*), and then becomes increasingly loud as we approach letter [J] (eventually reaching quadruple *f*). Thus, the direction of this episode consists of a contraction from a cluster 44 semitones in height to a cluster just 4 semitones high, and a gradual increase in loudness. During this process, there is constant movement within each of the instrumental parts which comprise the clusters.

Episode 2 consists of sections [J], [K], [L] and [M]. Although the contraction in the height of the cluster continues from Episode 1 through to the end of section [K], so as to suggest that this is a continuation of the previous episode, the instrumentation and melodic figures suggest otherwise. Episode 2 emphasises the sound of wind instruments, whereas Episode 1 consists solely of strings. Section [J] introduces long, sustained notes in the flutes, unlike anything heard in Episode 1, below which is a murmur of clarinets, playing:

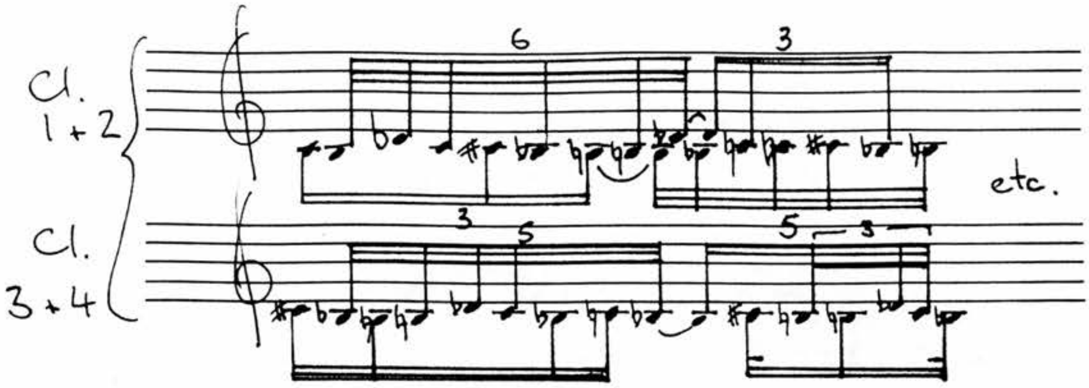

Example 4.5

The clarinet figure is a quiet reminiscence of the loud, frenetic movement that formed the end of the previous Episode. In Section [K], the sustained flute notes become shorter and rests are introduced, as well as quiet, brief interjections in other wind, brass and string instruments:

Bar 55.

[K]

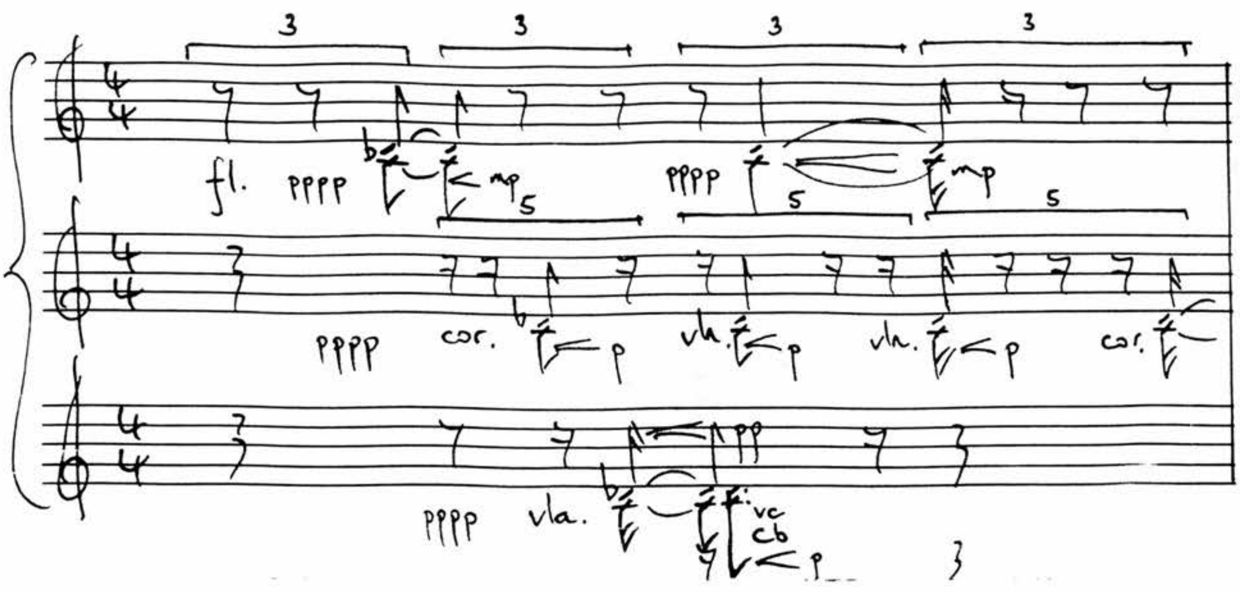

Example 4.6

Section [L] continues the idea of sustained notes in the winds, and the pitch cluster begins to expand again after contracting to the dyad C/C#^{4} at the end of section [K]. Aurally, it is quite apparent when the cluster begins to expand, since there is a very clear D^4 in the flute at letter [L] - the highest note of the three semitone cluster since letter [J] has been C#^{4}. At [M], sustained notes with points of articulation occur in the brass, where each instrument remains on a particular pitch. The horns and trumpets, playing quadruple *f*, and the trombones and tuba, playing double *f*, re-articulate the pitches in such a way that no two instruments re-articulate at the same moment. A constant jabbing at the pitches results.

Episode 3 appears to be concerned with dissolving the energy accumulated in section [M]. This Episode begins with a large, diffuse cluster for strings and winds. The cluster slowly thins, until only a quietly played 'white-note' collection remains. This collection too becomes sparser, until only a four note cluster at the end of [O] remains. This is followed by a small cluster for violas played in such a way that noise is produced rather than a sense of pitch - the instruction in the score is "Do not put the finger of the left hand all the way down. Almost without tone (more of a bowing noise)". This, in turn, is followed by a passage for brass where the players are instructed to blow softly through their instruments without producing a tone. These sustained noises are re-articulated in a way similar to that which we saw in section [M], where moments of re-articulation were placed so as not to coincide between parts. The soft blowing sounds mark the conclusion of Episode 3 - we have travelled from a large diffuse sounding cluster, to articulated noise. Of this point, Ligeti writes at the front of the score - "The 'air passage' in the brass is not to be treated as the beginning of a new section but as the continuation of the viola passage... The impression of a new beginning should not be created until the piano entrance at measure 77 [letter Q]."

From [Q] till the end I hear as the Coda. The music has dissolved into quiet noise at the end of Episode 3, and the entrance of the piano at [Q] sounds very much like the beginning of a new idea - the sound of the music changes from articulated noise to randomly articulated swept piano strings. The Coda doesn't sound like any of the previous sections - it has a static, undirected quality, with constant overlapping of movement in the strings, with staccato articulations, along with the sound of brushed piano strings. These sounds trail away to leave a low, rumbling trombone and tuba cluster at the very end of the work, along with the sound of brushed piano strings in the low register. This sound dies away - *morendo* - until the work finishes with three bars of silence (lasting at least 19 seconds). The static quality of the Coda has allowed the accumulated energy of the music to slowly subside.

The piece can be seen to fall generally into two parts - the first part from the beginning to letter [G], and the second part from letter [H] to the end; the music between letters [G] and [H] - the cluster for double basses - forms the bridge between the two parts. The first part has a feeling of direction and of evolution. The second part consists much more of a series of events which happen to lead one to the other - there is not the strong feeling of direction seen in part one. These different qualities of the two parts are reflected by the division of the first part into Sections, and the division of the second part into Episodes. The difference in the quality of the two parts is seen clearly in the timbral analysis that follows.

I have shown that *Atmosphères* can be divided into sections. The various sections relate to each other through the use of similar gestures of texture. Below is a summary of these different textural gestures.

1) A large, dense, static cluster, played without articulatory features - Opening, and sections [N] and [O].
2) A number of notes are re-articulated, either in an aurally very obvious way or in a subdued way, in combination with dynamic swells, that is <> - Sections [A] and [B] (section [A] only contains the dynamic swells, section [B] contains both the swells and the re-articulations), Episode 2 (sections [J], [K] and [M] contain re-articulations and section [L] contains the dynamic swells and re-articulations), and section [P].
3) A progressively increasing textural density due to increasing rapidity of attack in each part, coupled with an overall drop in pitch - sections [C] and [I].
4) A static, dense cluster, much smaller in height than 1), without notated articulatory features - the first two bars of [A], all of [D] and [G], and the viola passage lying either side of letter [P].
5) A complex contrapuntal web, the contrapuntal lines having a definite sense of directed motion (this is the complete opposite of a static cluster consisting of sustained notes) - sections [E] and [F], and [H].
6) This leaves the Coda, which is texturally very different from what has gone before. In the Coda, the principal feature is that of clouds of pitches washing

back and forth. Also present is a reminiscence of the piccolo texture that occurred at the end of [F], this time played by the flutes, and a gesture played by sweeping the strings of a piano with brushes and cloths.

Below is reproduced my plan of the work, with numbers against the letters indicating the gestural similarities that have been described above. Brackets around numbers indicate that they only apply to a part of the section - the beginning of the section if the bracketed number comes before the letter, the end of the section if the bracketed number comes after the letter.

Example 4.7

Opening^{1} ; $^{(4)}$[A]^{2} [B]^{2} [C]^{3} ; [D]^{4} ; [E]^{5} [F]$^{5\,(6)}$ [G]^{4} ;
Introduction ; Section 1 ; Interlude ; Section 2 ;

[H]^{5} [I]^{3} ; [J]^{2} [K]^{2} [L]^{2} [M]^{2} ; [N]^{1} [O]^{1} $^{(4)}$[P]^{2} ; $^{(6)}$[Q] - end
Episode 1 ; Episode 2 ; Episode 3 ; Coda

This sectional analysis has attempted to show how the various textural gestures relate and interact with one another, forming the sections shown in example 4.7, which in turn make up the larger gestural stages of the entire work. Timbre is a vital element in the structure of *Atmosphères*, but conventional music analytical methods are unable to tackle this aspect of the piece. Here, for the first time, let us examine in detail how the timbres of *Atmosphères* are moulded by the composer with such skill. The following timbral analysis serves both as an analysis of *Atmosphères* in its own right, as well as a further example of the efficacy of the timbral analytical techniques presented in this thesis.

5. Timbral Analysis

We have examined the work from the perspective of a rigorous pitch analysis, and we have explored a structural analysis, based on a consideration of thematic and textural ideas and their accompanying timbres. Let us now turn our attention specifically to an analysis of the work's timbre.

In Chapter 3, we began the discussion of timbre with a look at a spectrograph of *Jeux vénitiens*. Here, we will begin our investigations from an examination of an acoustic representation of the piece at the step prior to a spectrograph - a time waveform.[11]

Just as in Chapter 3, two different recordings of the work will be examined. Figure 4.6a shows two time waveforms of the first half of *Atmosphères*, up to a point a few seconds after letter [G]. Figure 4.6b shows a continuation of the piece to its end. The performance represented by the upper waveform is a live recording by the Vienna Philharmonic, conducted by Claudio Abbado (Deutsche Grammophon, 429 260-2). From now on, this will be referred to as performance 1. The performance represented by the lower waveform is by the Sinfonie-Orchestra des Südwestfunks, Baden-Baden, conducted by Ernst Bour (Virgo). This will be referred to as performance 2.[12]

Performance 1 (upper waveform) is slower, and therefore longer, than performance 2. It can be seen that the events of performance 2 lag behind those of performance 1, and that at the end of the waveform of the first half of performance 2, there is a gap,

---

[11] See Beurmann and Schneider (1991): "Struktur, Klang, Dynamik, Akutische Untersuchungen an Ligetis Atmosphères," *Hamburger Jahrbuch für Musikwissenschaft* 11: 311 - 334, for a further example of the use of a time waveform in analysis.

[12] The different recordings are referred to as performance 1 and 2, as opposed to recording 1 and 2, in order to avoid confusion with the analysis in Chapter 3.

which represents the difference in length of the two performances up to a point a few seconds after letter [G].

A time waveform shows the relative amplitude of the acoustic signal against time. What does this particular acoustic perspective tell us about the piece ? It shows that the piece begins and ends with sections of relatively minimal activity (marked a and a^1); there is a middle section of the piece where amplitude changes quite rapidly, and where amplitude becomes relatively high (c); and both before and after this central section there is a shorter section of increased amplitude (b and b^1). Thus, a loose symmetry is observable in the amplitude of the time waveform, which is heard as a loose symmetry of loudness.

From an examination of the two different time waveforms, we can begin to see how performances 1 and 2 differ from one another. We have already observed that performance 2 is shorter than performance 1. From figures 4.6a and b we can also see that the performances differ in the degree of amplitude difference during a particular section. For example, comparing performances 1 and 2, we see that performance 1 has a greater range of amplitude during section c, while performance 2 has a greater amplitude range in section a^1. To begin to consider timbre, however, we need to move to the next 'level' of acoustic analysis.

Figures 4.7a and b and 4.8a and b are A-weighted spectrographs of the information that we saw in figures 4.6 a and b.

With spectrographic analysis the shape of the piece becomes much more visible. Where with the wave-form display we could only say that different sections showed different degrees of amplitude, with the spectrographic display we can see the way in which the performances move through 'frequency space.'.

Examining the spectrographic analysis, it becomes apparent that within the rough symmetry observed previously (also observable here), there is growth and decay in the way the sound inhabits frequency space. Keeping in mind that two spectrographs of the same sound can visually emphasise very different types of information depending on the parameters of the FFT and of the subsequent spectrographic display - it is informative to investigate the spectral shapes that we see here.

The shape of the spectrograph of the first half of *Atmosphères* as realised in performance 1 (figure 4.7a) is shown in simplified form in the following example. This representation of the spectrograph simplifies both the information on frequency space (by delineating solid spectral shapes) and amplitude and 'density' (only two different shades of grey are used - the darker representing a greater amplitude and density of sound). This type of simplification is similar to stripping away non-harmony notes in a piece of tonal music so that the harmonic progression becomes clearer. In example 4.8 we are left with a representation of the general spectrographic movement.

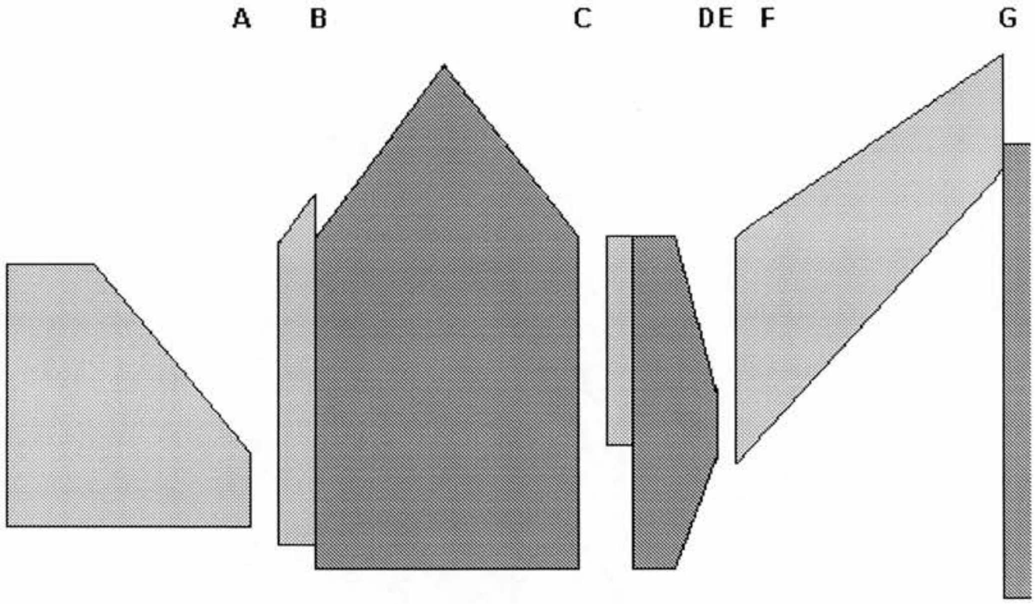

Example 4.8

The above figure shows that there is an initial frequency space that dies away (the first shape). The frequency space then enlarges (the second shape). This enlarged space expands and contracts - this third shape contains the most dynamic movement so far, the high peak of the trumpets forming the apex of this movement. This upwards movement is re-stated by the trajectory of the sound energy from [C] to [G] (the fourth, fifth and sixth shapes). The upwards movement between [E] and [G] is balanced and 'supported' by the strong double-bass entry that forms the seventh and last shape of the first half of the piece.

Example 4.9 below summarises the first half of *Atmosphères* as represented by performance 2 (see figure 4.8a for the spectrograph):

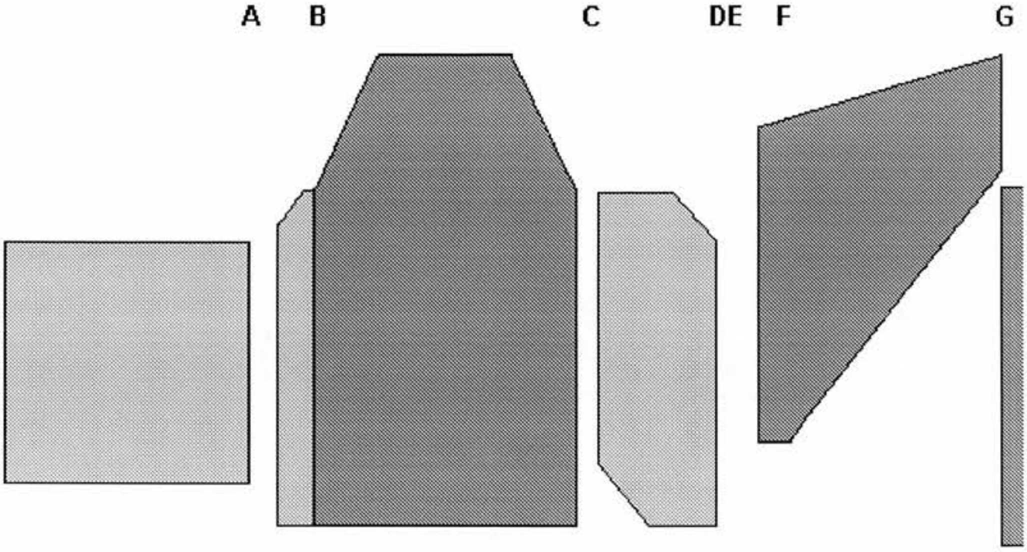

Example 4.9

Here, the spectral shapes and relative amplitudes are very similar to those found in performance 1, but, as can be seen clearly in the above figure, the shapes of sections [B] and [F] are 'blunter' than in performance 1, and the last shape, which acts as the counterweight to the previous shape, is not as 'dense' or as 'high' as in performance 1. Aurally, these differences manifest as performance 1 having sharper, more clearly defined contours in its first half when compared with performance 2. In both performances, the spectral shapes suggest gestures of directed movement, and balance, which are the gestures we discovered in the pitch and sectional analyses.

Example 4.10 summarises the spectral movement in the second half of *Atmosphères* - performance 1 (see figure 4.7b for the spectrograph):

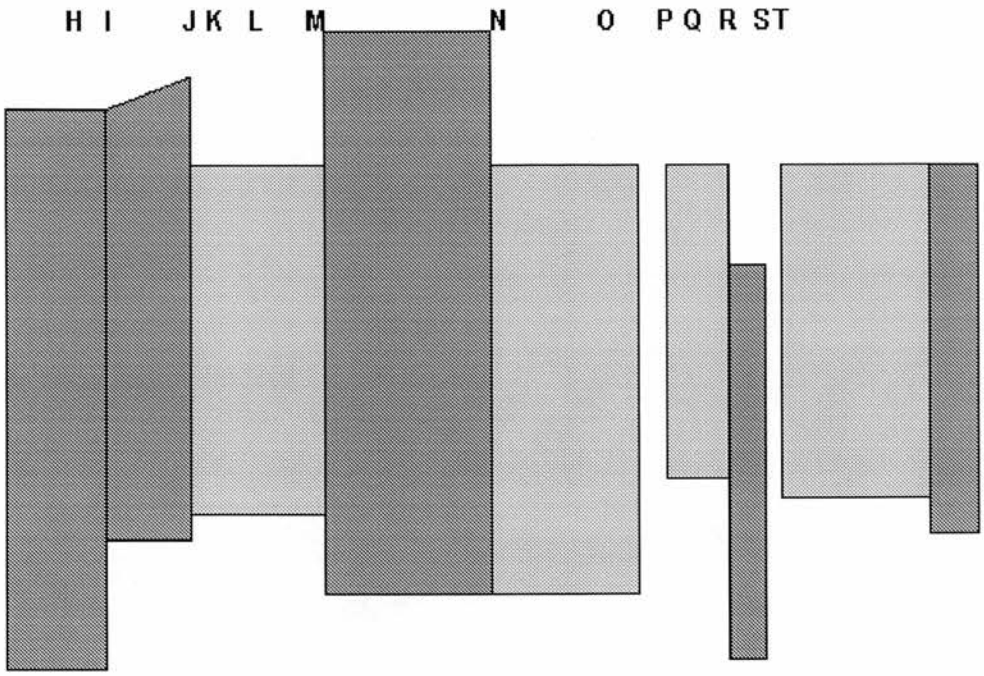

Example 4.10

Here, we see oppositional shapes, and oppositional degrees of density. Unlike the spectrograph of the first half of *Atmosphères*, there is no strong sense of the spectral shapes moving towards a particular goal. Once again, we find that the spectrographs suggest a structure the same as that found in the pitch and sectional analyses.

The spectral movement of the second half of *Atmosphères* (performance 2) is summarised in example 4.11 (see figure 4.8b for the spectrograph):

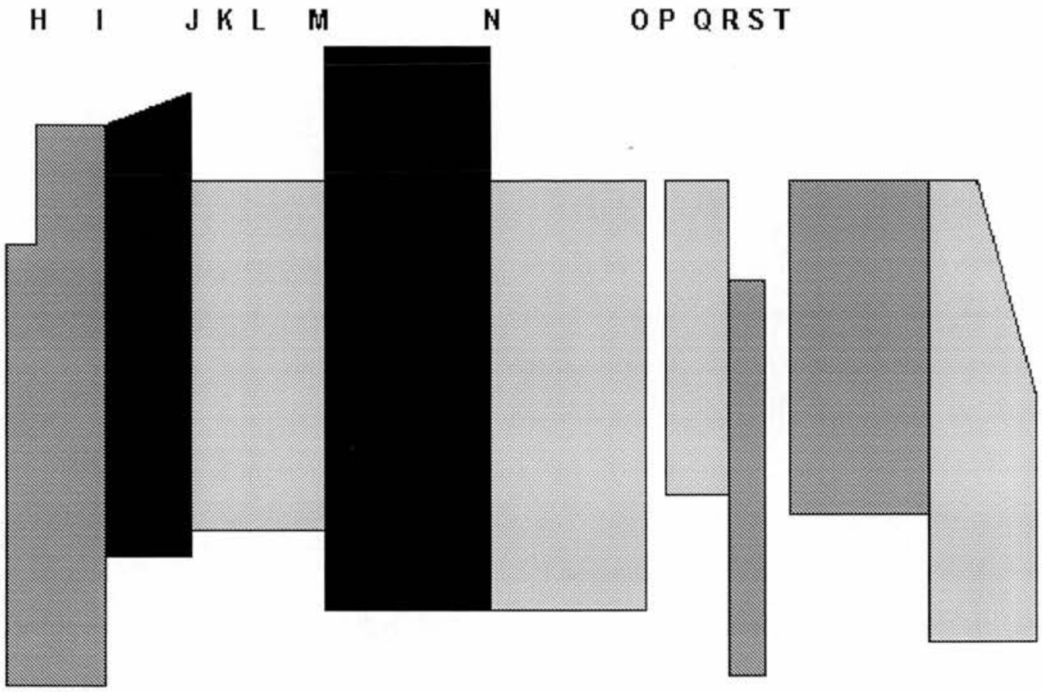

Example 4.11

Although slightly different from the spectrograph of the same section of performance 1 (plus a third level of grey has been introduced in the summary of density), the principal organising feature still appears to be that of the opposition of different sized and placed spectral areas and densities.

A spectrograph is a very good way to observe how a piece occupies frequency space. We can observe the ways in which different frequency spaces and different degrees of density are juxtaposed. These images are also a way in to our experience of hearing the piece. We hear the frequency space to be either broad or narrow, dense or sparse, to be high or low in energy, and to be high or low in frequency. From the above examination of the spectrographs of performances 1 and 2, it appears that the frequency space in the first half of *Atmosphères* is suggestive of directed movement and balance, whereas the frequency space in the second half of *Atmosphères* appears to be based on oppositional relationships.[13]

In order to delve more deeply into the timbral world of *Atmosphères*, let us now move to the next 'level' of acoustic analysis. The following is a detailed examination and explanation of the timbral graphs of *Atmosphères*.

Figures 4.9a and b, and 4.10a and b show timbral measurements for the first and second halves of performances 1 and 2 respectively. The timbral measurements are arranged in the same way as we saw in Chapter 3 in figures 3.38 and 3.39 - three different measures are placed on the one graph so that we may see how these three timbral parameters change in relationship to one another. The vertical axis is dimensionless because the three measures are showing three different quantities - the top trace shows timbral width, the middle trace shows sharpness (the position of the loudness centroid) and the bottom trace shows roughness. Graphs where these traces are presented singly and with measurement on the vertical axis appear after these figures (figures 4.15 - 4.20) - these graphs will not be referenced below, however it is to these graphs that readers are referred if they wish to corroborate the unit measures which are quoted. Figures 4.11 - 4.14 show graphs of timbral weight and pitch, measures that are not included on the 'summary' graphs of figures 4.9 and 4.10.,

As we have seen, the opening of *Atmosphères* consists of a large, dense, static cluster scored for strings and a small group of winds and brass. The cluster slowly decrescendos. The decrescendo is achieved through decreasing the dynamic of individual instrumental groups, and then dropping them from the texture. As shown

---

[13] This type of description of a spectrograph is inspired by the work of Cogan (1984).

on figure 4.9a, in performance 1 this produces a consistently gradual decline in roughness up to letter [A] - performance 1 moves from a roughness measure of 3, down to 0.75 - a difference of 2.25. Performance 2 also shows a decline over the length of the opening (see figure 4.10a) - roughness moves from a measure of 4 down to 2, a difference of 2. Thus, the degree of change in roughness between the two performances is very similar. The slope of the roughness curve, however, is not as consistent in performance 2 as in performance 1. In performance 2, as the winds drop out, the sound of the upper strings increases. This is why the roughness decreases only a little for most of the opening, until the very end of the section, where there is quite a sudden drop at 43 seconds. This decrease in roughness is due to the upper strings suddenly ceasing to be heard in the bar before [A] (bar 8 - this bar is shown on the graphs), even though this bar is marked for the upper strings as *dim...morendo...* - the upper strings should gradually fade over the duration of the bar. It would appear that, at this moment in the piece, performance 1 is more faithful to the score than performance 2, and thus the roughness graph of performance 1 is a better representation of the performance instructions than the roughness graph of performance 2.

Let us now look at the graph of sharpness for the opening section. In performance 1, the graph of sharpness is static for most of the section, lying at band number 27 (whose centre frequency is roughly the pitch B^4). There are a few excursions to either side of this band number, but the fact that these excursions are brief suggests that these movements are not significant - they merely show a momentary shift in balance between instruments. In performance 2 we see a similar situation. The graph lies largely at band number 28 (whose centre frequency is roughly $D^{\#5}$), with some movement on either side of this band number. Although there is more movement in sharpness in performance 2 than in performance1, its momentary nature again suggests that it is not significant, for the same reasons described above. There is, however, a significant movement down one frequency band towards the end of this section. This movement downwards coincides with the dropping out of the upper strings in bar 8, so that only the lower strings are left sounding (see the previous paragraph where this event is also discussed). Although there is only the difference of one band number in overall sharpness between the two recordings in this opening section, the difference is clearly audible. The difference is due both to recording style, and interpretation. Performance 2 is, overall, recorded more 'closely' than performance 1, which, in this instance, increases sharpness. In addition to this, the upper strings dominate the sound in performance 2 in a way that does not happen in performance 1 (the role of the upper strings in performance 2 in roughness measurements was discussed in the previous paragraph).

Comparing the graphs of timbral width in the opening section of the two performances, we find the greatest degree of difference between the timbral measures that have been considered so far. Timbral width is very sensitive to individual notes emerging from an otherwise amorphous texture. We saw this sensitivity in the analysis of *Jeux vénitiens*, where the 'textural exclamations,' which were brief events designed suddenly to emerge from the background texture, showed up very clearly in the measure of timbral width. Ligeti instructs that individual pitches should *not* be heard in *Atmosphères*, and timbral width is a very good indication of moments where this textural and timbral ideal is not realised in performance.

The measure of timbral width in the opening section of performance 1 lies largely around a level of 0.6. In performance 2, it lies around 0.625. So the recordings are very similar in the overall position of timbral width. As has already been mentioned, however, the shape of the graph of timbral width around these general levels is different between the two recordings. In both performances, timbral width climbs for a few seconds from the very beginning of the graph. In both performances, this is due to the instruments taking a moment to settle into a stable balance. It feels as though the instrumentalists start a little tentatively, and then find a comfortable dynamic at which to play. In performance 1, after this initial rise, timbral width drops once more, then rises back to its previous high position, before dropping once more near to the end of the opening section. The first drop in width (at 12 seconds) is due to a small group of pitches becoming more apparent in the texture, so the timbre becomes more focused. The climb back up of timbral width is due to this 'knot' of pitches gradually merging back into the overall texture. The drop in width near to the end of the section is due to the upper strings slowly fading during bar 8 , so that only the lower strings are left playing - thus, the timbre becomes more focused again. In performance 2, there is a steady rise in timbral width for the first 20 seconds. Width then plateaus for 17 seconds, after which point it suddenly drops, before rising back up again just prior to the end of the section. The rise for the first 20 seconds is due to the upper strings slowly increasing in presence over this time. The drop in width at 40 seconds is due to the lower strings being left to play by themselves after the upper strings have swiftly exited during bar 8. This drop is equivalent to the much smoother focusing of timbre that occurs in performance 1 at this same point in the music (due to the upper strings decrescendoing more gradually). That the timbre suddenly becomes more diffuse very near to the end of this section in performance 2 is surprising. It is due to a 'shimmering' sound produced by the lower strings at this moment - the upper harmonics of the strings suddenly become more prominent. This may be due to the slowing of the bow movement in order to 'squeeze' the last moments of the opening into the remaining length of bow (slightly less than half of the string players are asked to play a different note, and to change bow direction, at letter [A]). The concomitant moving of the bow nearer to the bridge that strings players often do instinctively in order to keep the note sounding clearly when the bow moves slowly usually has the effect of emphasising the upper harmonics.

Considering the two performances together, and discounting performance peculiarities, we can conclude that in the opening section the score asks for roughness to fall smoothly for the whole length of the section; sharpness remains fixed at around band number 27 / 28; and width falls near the end of the section, after remaining fairly static at a measurement of 0.6.

Example 4.12 summarises this information:

Example 4.12

Opening section:

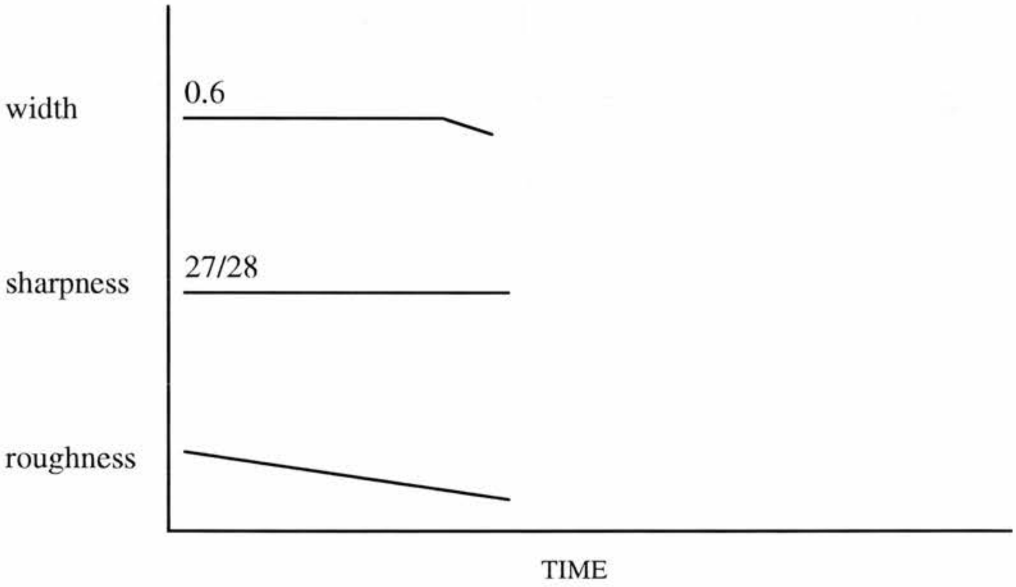

To conclude this discussion of the opening section, let us look at the more unstable measures of timbral weight and timbral pitch (figures 4.11 - 4.14). As in the analysis of *Jeux vénitiens*, the degree of averaging for the graph of timbral weight is very high (an averaging factor of 50: a data point every 2.32 seconds (the black line); along with an averaging factor of 5: a data point every 0.232 seconds (the red line)). Measurements mentioned in the text are taken from the black line, representing an averaging factor of 50. Figure 4.11a shows the graph for timbral weight for the first half of performance 1. For most of the opening section in performance 1, timbral weight becomes heavier. At 50 seconds, weight becomes lighter, and stays lighter until the end of the section. Figure 4.13a shows the graph of timbral pitch for the same section. As would be expected, at the point where timbral weight rises, timbral pitch falls - in this case, it falls by one band number. It falls due to the lower strings remaining after the upper strings have slowly faded out. Thus, as the opening section progresses in performance 1, the timbre becomes heavier, more 'solid' (see comments on timbral weight in Chapter 2). Near the end of the section, we return to the lighter sound we heard at the beginning of the section, though note the fall in weight during bar 8 (marked on figure 4.11a), as the more solid sound of the lower strings becomes more prominent.

Figure 4.12a shows the graph of timbral weight for the first half of performance 2. In the opening section, we see a very different representation of timbral weight from the one we just saw for performance 1 in figure 4.11a. Instead of falling for most of the section, weight generally rises during the opening. Why is there this difference between the performances? It will be remembered that the upper strings in performance 2 appear to maintain, and, indeed, increase their presence during the section. The slight increase in the sound of the upper strings during the section is reflected in the lightening of timbral weight for the first 20 seconds. Timbral pitch remains largely on band number 27 for this time (see figure 4.14a), so the position of the loudest band remains static, while the amount of energy increases above it. At around the 20th second, the lower strings make their presence felt more strongly, and timbral pitch drops by one band. Weight remains relatively static until around the 40th second, which is the beginning of bar 8 (shown on the graph). Here the upper strings drops out, and the sound becomes heavier as the lower strings are left to sound by themselves, albeit with a 'shimmery' sound near the end of the section (not reflected in the measure of timbral weight). To summarise the information on timbral pitch and weight for the opening section, I take performance 1 as the more representative performance of the work. The opening in performance 2 sounds harsh, because it has been recorded quite 'close,' and, in my opinion, the sound of the upper strings is too loud, and variable. Thus, in the opening section, timbral pitch falls by one band, and timbral weight falls until the last moments of the opening, where weight becomes lighter, and then heavier once more as the upper strings continue to die away.

The information on timbral pitch and weight is summarised below in example 4.13:

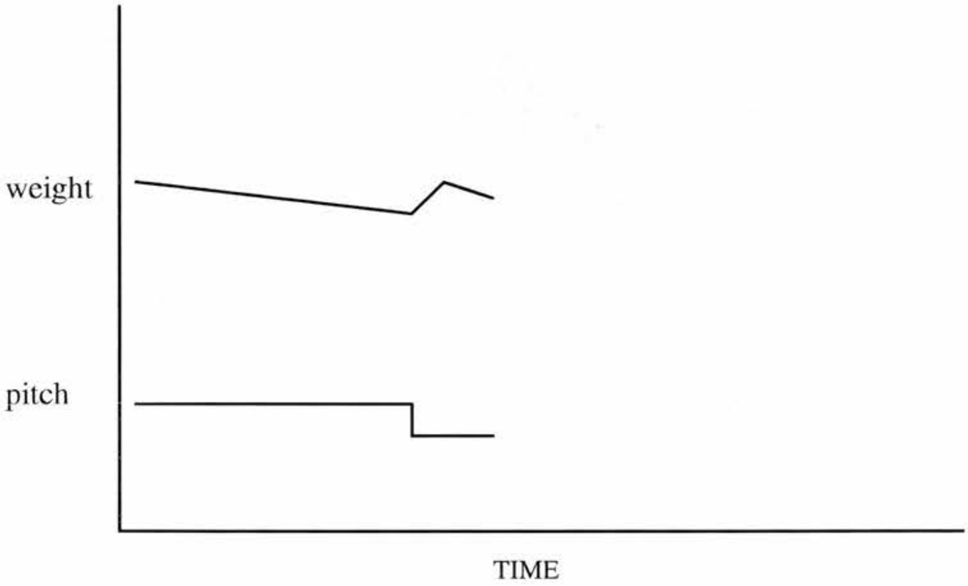

Example 4.13

The section after the opening, Section [A], consists of a static pitch cluster (different from the cluster heard in the opening section) that uses crescendo-decrescendo markings in its second half in order to cause various combinations of pitches to emerge and then re-integrate back into the pervading texture (see example 4.4).

The effect of these dynamic swells can be seen quite readily in the graph of roughness in both recordings (see figures 4.9a and 4.10a). In performance 1, the graph of roughness curves smoothly upwards to a peak at 90 seconds of a roughness measurement of approximately 4.5, and then rapidly falls back to a level of roughness a little higher than the level at which this section began. In performance 2, we see a similar, but 'blunter' roughness shape. The level of roughness at which this section begins is higher in performance 2 - a measurement of 2, as opposed to a measurement of 0.75 in performance 1. It then climbs to a level of roughness of 4, before falling to a level of 2.5. The control of roughness in performance 1 during section [A] both looks and sounds more eloquent than the way in which roughness is shaped in section [A] of performance 2.

Looking at sharpness, we see in both recordings a tendency for the position of the loudness centroid to fluctuate, and to rise as the section progresses - this is due to the effect of the crescendo - decrescendo markings. In example 4.4 it was shown that these dynamic markings fan out symmetrically from the centre of the pitch cluster, the concluding crescendo - decrescendo coming at the highest point of the pitch cluster. This, along with the tendency for the sound of the higher pitches to dominate, is why the sharpness graph of section [A] in performance 1 rises. Looking at the graph of performance 1, the position of the loudness centroid graphed at an average of 10 stays static up until about 73 seconds. This represents the first part of section [A], where no crescendo-decrescendo markings occur. After this point (shown on the graph by '< >'), the movement of the sharpness graph reflects very well the way in which the order of the dynamic markings in the score causes first the higher pitches, then the lower pitches, and then the higher pitches once again to move out from the background texture. The movement of the sharpness graph of section [A] in performance 2 shows that the degree of control of the dynamic markings in this performance is not as precise as in performance 1. The beginning of the section to around 57 seconds represents the first half of the section where no dynamic markings occur. The position of the loudness centroid oscillates between band numbers 27 and 28, which indicates that the loudness centroid probably lies between these two bands. When the dynamic swells begin, sharpness rises, so that the loudness centroid stabilises at band number 28. Sharpness does not fall again, as it does in performance 1, and as the score suggests it should. The rise to band number 30 at around 72 seconds and the fall back to band number 28 in the concluding moments of the section (where the upper notes of the cluster should be emphasised) confirm the aural impression that the balance achieved in the dynamic movement of performance 2 in this section does not follow the indications of the score as well as performance 1.

The graph of timbral width for section [A] of performance 1 does not suggest the presence of the crescendo-decrescendo dynamics as clearly as might be expected. We might expect to see width moving frequently between a more focused timbre and a more diffuse timbre as pitch dyads become louder, and then become quieter again. The reason we do not see this happen frequently (it does happen a few times, and these moments will be discussed shortly) is because of the way the crescendo-decrescendo markings in the individual parts overlap with one another. Thus, there is an increase in volume, but the overlapping crescendos occur on pitch dyads separated in pitch space - width remains diffuse. This effect is also born out aurally - because of the overlapping dynamics, the ear is not drawn towards a particular focused pitch area, but rather the listener is aware of dynamic activity occurring in more than one pitch area. There is an exception to this near to where the dynamic markings first appear at 75 seconds. At this moment one of the first dyadic pairs to crescendo does emerge distinctly from the background, so width briefly becomes more focused. At around 85 seconds width becomes more focused, and remains more focused until near to the end of the section, where it once again moves up to a level of width equivalent to the level at the beginning of the section. The reason this occurs is because at around 85 seconds the dynamic movement focuses on groups of lower pitches (this can be seen in example 4.4 - the section we are discussing begins with the crescendo-decrescendo on the dyad B♭/G$^{\#3}$). These lower pitches sound more strongly than the previous higher pitches, and the aural impression is much more of a focused area of activity than of the diffuse dynamic activity that we heard earlier. As the lower pitches decrescendo, a single, higher dyadic pair crescendos, and width becomes more diffuse once more.

In performance 2, as in performance1, width is relatively static until the crescendo-decrescendo markings begin to make their presence felt at around the 55th second. Although it has been said that the balance achieved in this section in performance 2 is not as precise in following what is indicated in the score as what is heard in performance 1, we see in the graph of timbral width a pattern closer to what we might expect to see - width moving into and out of a more focused position as the crescendo-decrescendo dynamics take effect. This is indeed what we hear in this performance. Thus, performance 1 follows more closely the score in the sequence of the crescendo-decrescendo dynamics, and performance 2 gives a 'closer' aural representation of the detail of the coming and going of the pitch dyads. This is related to the 'closer' style of the recording adopted in performance 2.

The information on timbral width, sharpness and roughness, summarised for the two performances, is shown below in example 4.14:

Example 4.14

Section [A]:

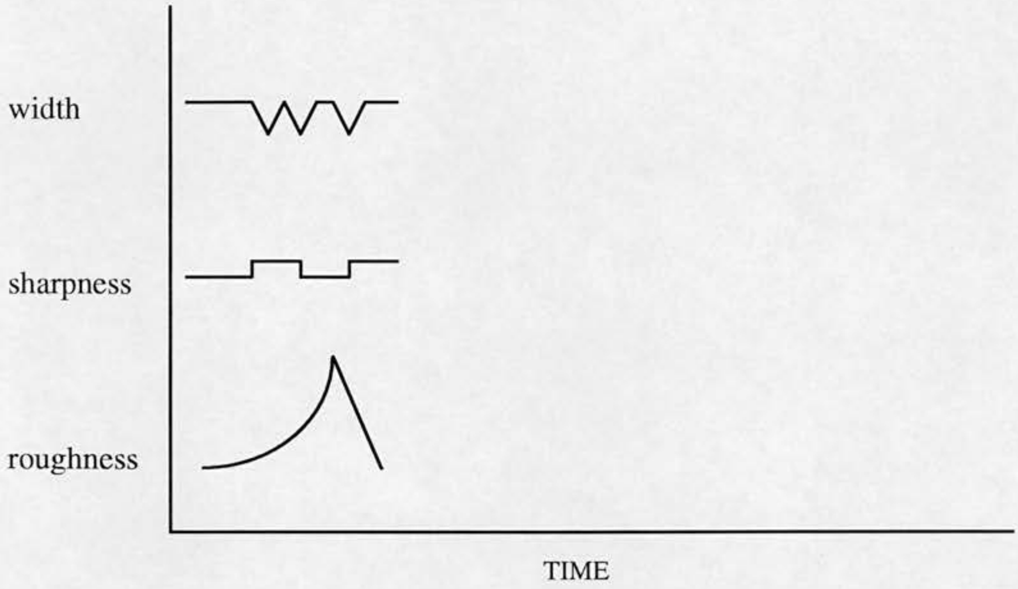

Turning to figures 4.11a and 4.12a, let us now compare the graphs of timbral weight for section [A] of the two performances. In performance 1, we see that weight fluctuates - generally staying around a measure of 0.25, but going as high as 0.45 and as low as -0.05, which is a range of 0.5. In performance 2, weight fluctuates much less - very generally, it lies around 0.3, and spans 0.1 to 0.4, which is a range of 0.2. As can quite readily be seen, the fluctuation in weight in performance 1 is more rapid than that in performance 2. Why is this the case, especially as we saw that changes in timbral width were more pronounced in performance 2 than in performance 1 ? The reason lies partly in the 'closer' style of the recording of performance 2. Because of the closer recording, we hear the comings and goings of the dyadic pairs a little more clearly than in performance 1, and this has already been seen to have an effect on timbral width. However, the closer recording style highlights the fundamental pitches of the dyadic pairs at the expense of the overall timbral effect of shifting areas of energy. This means that timbral weight changes less in performance 2. The loudest band moves less in performance 2 (timbral pitch - figure 4.14a), because the loudest band is generally centred on the area of the fundamentals. For example, about 10 seconds into section [A] in both performances (at 73 seconds in performance 1; 57 seconds in performance 2) the loudest band in performance 1 (see figure 4.13a) is band number 28 - centred roughly on D$^{\#5}$ - and the loudest band in performance 2 is band number 25 - centred roughly on D$^{\#4}$. This would suggest that the loudest band in performance 1 is composed solely of the energy of upper harmonics (the highest fundamental that is present is B^{4}), whereas the loudest band in performance 2 is composed largely of the energy of fundamentals. This means that the position of the loudest band in performance 2 is more stable, as is the measure of timbral weight, and the position of the loudest band in performance 1 is more fluctuating, as is the measure of timbral weight. Perceptually, the greater movement in timbral weight in performance 1 reflects the broader palette of timbral shades that is heard in this section when compared with performance 2.

Section [B] forms the centre section of the first half of *Atmosphères.* It is a climactic section. Section [B] contains the widest pitch cluster of the first half of the piece, and marks the beginning of the symmetrical pitch movement to letter [D] (see figure 4.1); it has the greatest amplitude (see figure 4.6a - section [B] comprises the excerpt labelled 'b'); and it shows a very distinctive rise in roughness and sharpness (figures 4.9a and 4.10a). The section is orchestrated for full orchestra playing a virtually complete chromatic cluster spanning Ab^{2} to E^{7} (only the pitch C^{5} is omitted), with an additional pitch, Eb^{1}. Through the use of selective crescendo-decrescendo markings, Ligeti brings out first a 'white-note' pitch collection, followed by a 'black-note' pitch collection. The section finishes with a decrescendo to *ppp* for strings alone.

As in the discussion of the previous sections, let us begin with a consideration of roughness.

The roughness curve for section [B] in both performances shows a very distinctive double peak. The first of these two peaks represents the sounding of the white-note collection, the second peak represents the sounding of the black-note collection. These two peaks can be seen as the culmination of a build-up of roughness beginning from the high-point of roughness in section [A], to the beginning of section [B] (which begins either a little below or a little above this peak, depending on the performance), to the brief roughness plateau that occurs just before the first of the two large peaks, to the first of the two peaks itself.

Roughness rises at letter [B] because the pitch cluster is so much wider than the cluster we left at the end of section [A]. Note, however, that the level of roughness is only slightly higher than the level at the very beginning of the opening, where the pitch cluster had a similar range to the pitch cluster here. Indeed, the cluster at the opening was played *pp*, while the cluster in section [B] is played *ppp*. The level of roughness is higher because in section [B] the orchestration is fuller - importantly, it includes the full brass section.

Roughness increases from the beginning of section [B] due to a general crescendo for the strings. At the end of bar 16 (the bar numbers of section [B] are marked on both figures 4.9a and 4.10a), the 'black-note' strings begin to decrescendo, while the 'white-note' strings begin to crescendo further. The 'white-note' strings reach the dynamic *ff* at the end of bar 17, and then hold this dynamic while the other 'white-note' instruments crescendo up to their highest dynamic in bar 18. From careful listening, it appears that the plateau in roughness that occurs immediately before the first of the two peaks is caused by the crescendo of the 'white-note' collection. So, although the dynamic continues to increase (the crescendo of the 'white-notes' outweighs the lessening in loudness due to the decrescendo of the 'black-notes), less beating occurs between sound components because we hear a predominantly 'white-note' collection, as compared to a chromatic cluster. Indeed, this roughness plateau in fact shows a small decrease in roughness over its length.

Looking at the graphs of both performances, it can be seen that the two peaks in performance 1 are much sharper than the two peaks in performance 2. This partly reflects the greater degree of shape and control in performance 1, and partly the greater degree of 'noisiness' in performance 2 because the recording style is 'closer.' This extra noisiness adds to the background level of roughness, and so, for example, the approach to the first roughness peak is not as smooth in performance 2 as in performance 1. After the second roughness peak, the roughness level plateaus for a short time, before dropping towards the end of this section. This short roughness plateau is caused by a crescendo for all the strings, which emerges as the 'black-note' collection dies away. All the strings decrescendo towards the end of the section.

Just as there is a climax of roughness in this section, so there is a climax of sharpness. Sharpness has climbed during section [A] and through the first half of section [B], and then falls as we approach the end of the section.

The sharpness graphs of the two performances show the same general contour, but also differ in significant ways. Section [B] of performance 1 begins at band number 28, whereas performance 2 begins at band number 29. This is because of the different instrumental balances in the two performances. Performance 1 emphasises the sound of the low brass during bars 14-15, whereas performance 2 emphasises the sound of the high wind - thus, performance 2 sounds sharper. At around 105 seconds in performance 1 the strings begin to dominate the sound (bar 16), and sharpness rises by one band number. The sound of the strings is also the reason why sharpness rises by one band number in performance 2 at the equivalent moment - at around 90 seconds. For both performances, the next rise in sharpness comes with the strong crescendo for the brass that marks the arrival of the 'white-note' collection (bars 17-18). At the height of this first crescendo, performance 1 reaches a sharpness measure of 32 (a rise of 3 bands from the immediately preceding level of sharpness of 29), while performance 2 reaches a sharpness measure of 31 (a rise of only 1 band from the preceding level of sharpness of 30). As was discussed earlier, performance 2 is generally sharper than performance 1 (as is borne out by the difference in the immediately preceding levels of sharpness) - why does performance 1 momentarily become sharper at this point? The reason lies in a mistake made by the first trumpet player of performance 1. The first trumpet's note is very clearly audible immediately prior to the 'white-note' crescendo. The first trumpet reaches the top of the *crescendo* too early, so that it does not coincide with the rest of the trumpets and trombones.[14] It is the sharpness of this 'solo note' that emerges so clearly from the surrounding texture, which causes the sharpness in performance 1 to rise up to band number 32. A close examination of the sharpness readings shows that at exactly125 seconds, where the general brass crescendo is to the forefront and the single trumpet note has died away, sharpness has dropped to band number 31. Indeed, it is at this moment, when sharpness has dropped to band number 31, that roughness reaches its highest, rather than during the 'solo' trumpet note. Very shortly after, sharpness drops once more as the brass sound dies away. During this same period in performance 2 there is

---

[14] In performance 1, the pitch that occurs at the wrong time is $F^{\#5}$. There is no $F^{\#5}$ shown in the score to be played by a trumpet at this moment - the closest notated pitch for a trumpet is F^5 for trumpet 1. Either the trumpet player not only got the timing wrong, but also the pitch, or the entire recording is a semitone sharp. Checking the pitch from a later section of the piece shows that the overall pitch of the recording is correct - the trumpet player got it wrong.

no such error, and sharpness rises to band number 31, the same level of sharpness as we saw for the brass crescendo after the 'solo note' is no longer apparent.

The following sharpness peak, due to the crescendo on the 'black-note' collection (bars 19-20), is very different between the recordings. In performance 2, the trough after the first peak (at band number 29) marks a moment where the strings sound for a short time by themselves, just before the following crescendo. During this crescendo, the sharpness of performance 2 rises to band number 31 - the same band number that was reached during the previous crescendo. In performance 1, it is very different. The equivalent moment in performance 1 to the sharpness trough mentioned above in performance 2 is from 126 to 130 seconds, where sharpness lies at band number 30 - one band number sharper than performance 1 (the upper strings are more prevalent here). When the next crescendo begins in performance 1, sharpness actually falls to band number 29, there being only a brief upwards peak at 133 seconds to band number 30 at the height of the crescendo. The reason for the differences between the two performances is due to a difference in orchestral balance during the second crescendo. In performance 1, the horns are favoured over the winds. In performance 2, the winds are favoured over the horns. This makes this moment in performance 1 less sharp than the equivalent moment in performance 2.

After the second crescendo, all the strings crescendo together up to *mf* (therefore a chromatic cluster is heard), and then die away towards the end of this section (bars 20 - 22). In performance 1, this cluster crescendo causes sharpness to rise (to band number 30) after the previous section where the horns dominated. In performance 2, the cluster crescendo causes sharpness to fall (to band number 30) after the previous section where the winds dominated. Sharpness falls in both performances with the general decrescendo that ends the section.

The information on sharpness for section [B], and the reasons for the differences between the two performances, is summarised below in example 4.15:

Example 4.15

	Performance 1	Performance 2
bb.14-15	band number 28 low brass favoured	band number 29 high winds favoured
bb.16-18	band number 29 strings heard	band number 30 strings heard
	band number 32 'solo' tr. note heard; band number 31 'white-note' collection	band number 31 'white-note' collection
bb. 19-20	band number 30, section between crescendos (upper strings favoured)	band number 29, section between crescendos (upper strings not as strong)
	band number 29 'black note' collection (horns favoured)	band number 31 'black-note' collection (winds favoured)
bb. 20-22	band number 30 string cluster crescendo	band number 30 string cluster crescendo

Let us now turn our attention to timbral width during section [B].

Comparing the graphs of timbral width for the two performances, we find that the general contour of each is very similar. For the first half of the section, width generally first curves upwards and then downwards, and in the second half of the section, width generally curves back up once more. On the graph for width there has been marked the same bar divisions that were used in the analysis of the roughness graph. During bars 14-15, timbral width remains largely static before rising. The dip in width in performance 2 in bars 14-15 is caused by the sound of a flute breaking through the texture - the moment it merges back into the background, width rises once more. The static width is caused by the static cluster at the beginning of the section, and the slight rise towards the end of bar 15 is caused by the beginning of the crescendo. (The sudden drop in width in performance 1, which is a live recording, just before the beginning of bars 16 - 18, is caused by a single cough that makes the timbre momentarily more focused.) Up to a point, the crescendo causes timbral width to continue to become more diffuse. However, as the crescendo continues, and especially as the brass come more to the fore, width begins to become more focused - the ear starts to focus on the brass sound which emerges out of the general textural fog. Thus, the graph of width starts to curve downwards about half way through bars 16-18. In performance 1, there is a sudden dip in width which coincides with the high point in sharpness in this section. As will be remembered, this is the moment when the 'solo trumpet' breaks through the texture; thus, width becomes momentarily more focused. The sharp dip in width in performance 2 about half-way through bars 16-18 is caused not by an element in the performance, but by a fault in the recording - fidelity is momentarily lost, and the sound becomes muffled. Bars 19-20 contain the second crescendo, and, as would be expected, timbral width lies at around the same level as it did during the first crescendo. In performance 2, the crescendo favours the winds, and the upper winds are heard from the very beginning of bar 19. In performance 1, the *crescendo* favours the horns, and before they are heard distinctly, the strings are the dominant sound. Thus, in performance 1, width becomes more diffuse at the beginning of bar 19 before becoming more focused again during the crescendo. As the crescendo grows, width becomes diffuse once more in performance 1 because of the timbral nature of the horn sound as it becomes louder. Because the balance favours the winds in performance 2, this more diffuse timbre does not occur. In both performances, at the end of bar 20, width has a sudden dip - this is caused by the rapid decrescendo that occurs at this point. In bars 20 - 22, a shimmer of strings, gently becoming louder, and then fading away, takes over from the more focused sounds of the crescendos that have come before, and this is reflected in the more diffuse timbre in these bars.

The information on timbral width is summarised below:

Example 4.16

	Width	Reason
bb. 14-15	static, then rises	a static cluster, followed by a *cresc.*
bb. 16-18	rises briefly, then falls	*cresc.* continues - sound of brass comes to the foreground
bb. 19-20	perf.1: changeable perf. 2: more static	*cresc.*, horns dominate *cresc.*, winds dominate
bb. 20-22	rises	shimmering string sound

In the analysis of section [B], I have looked at roughness, sharpness, and timbral width. This now leaves timbral weight and timbral pitch to consider.

As we have seen before, the graphs of timbral weight and pitch often show a high degree of fluctuation. Let us look first at the graphs of timbral weight and pitch for performance 1 (figures 4.11a and 4.13a respectively).

These, like the previous graphs, have bar numbers marked on them, and the shape of the movement of weight and pitch is summarised graphically in an example that follows. In the middle of the section marked as bars 14-15 we see a sudden drop in weight (from 0.35 to -0.3) and a corresponding rise in timbral pitch (from band number 26, with a centre pitch of G^4, to band number 31, with a centre pitch of $D\#^6$). This is due to the upper strings becoming prominent for a moment, with a corresponding 'deepening' in timbral weight. It is important to remember that this is a perceptual phenomenon that is being measured; we hear two different affects - the coming to prominence of the upper strings (the rise in timbral pitch), and the heavier nature of the timbre. This change in timbral pitch and weight appears to be caused by a change in balance in the strings as the first crescendo begins. The movement upwards in timbral pitch and downwards in timbral weight in the second half of the section labelled as bars 16-18, is caused by the coming to prominence of the brass in the latter stages of the first crescendo. The peak in timbral pitch (to band number 33) is caused by the 'solo' trumpet note. Interestingly, the top of the second crescendo (in the centre of the section marked bars 19-20), which in performance 1 brings the horns to the foreground, causes timbral pitch to fall - from band number 31 to band number 29. This is due to the relatively 'darker' nature of the horn sound when compared with the 'brighter' trumpet sound. In bars 20-22, there is the shimmer of the strings, which causes timbral pitch first to rise (to band number 34) and then to drop (to band number 26), and timbral weight to become first heavier, and then much lighter. The first movement is due to the strings being left to play by themselves after the horns and winds have ceased to be heard, and the second movement is due to the general string crescendo, which brings the entire cluster to the foreground, causing the timbre to become much lighter.

Figures 4.12a and 4.14a show graphs of timbral weight and pitch respectively for performance 2. In the section marked bars 14-15, instead of a sharp peak (to band number 31) which we saw in performance 1, we see here that pitch is constantly high (lying mostly in band number 30), and weight lies constantly below 0 - the balance between the high and low strings is more constant in performance 2, and tends to favour the upper strings which causes timbral pitch to remain at band number 30. As the crescendo begins in bars 16-18, pitch falls and weight rises to levels similar to those shown for performance 1. As the brass comes to the foreground, pitch rises and weight falls - just as in performance 1. On the second crescendo (bb. 19-20), pitch remains static and weight rises in performance 2, as opposed to pitch momentarily falling and the overall movement of weight being downwards which we saw in performance 1. This is because of the different balance between the horns and the winds in performance 2 when compared with performance 1 - this has been discussed earlier. In the section marked bars 20-22, pitch and weight move in completely opposite directions when performance 2 is compared with performance 1. Whereas section [B] in performance 1 ends with a light shimmer of strings, performance 2 ends with a more definite, heavier sound in the strings.

The information on timbral pitch and weight for the two performances is summarised below in example 4.17:

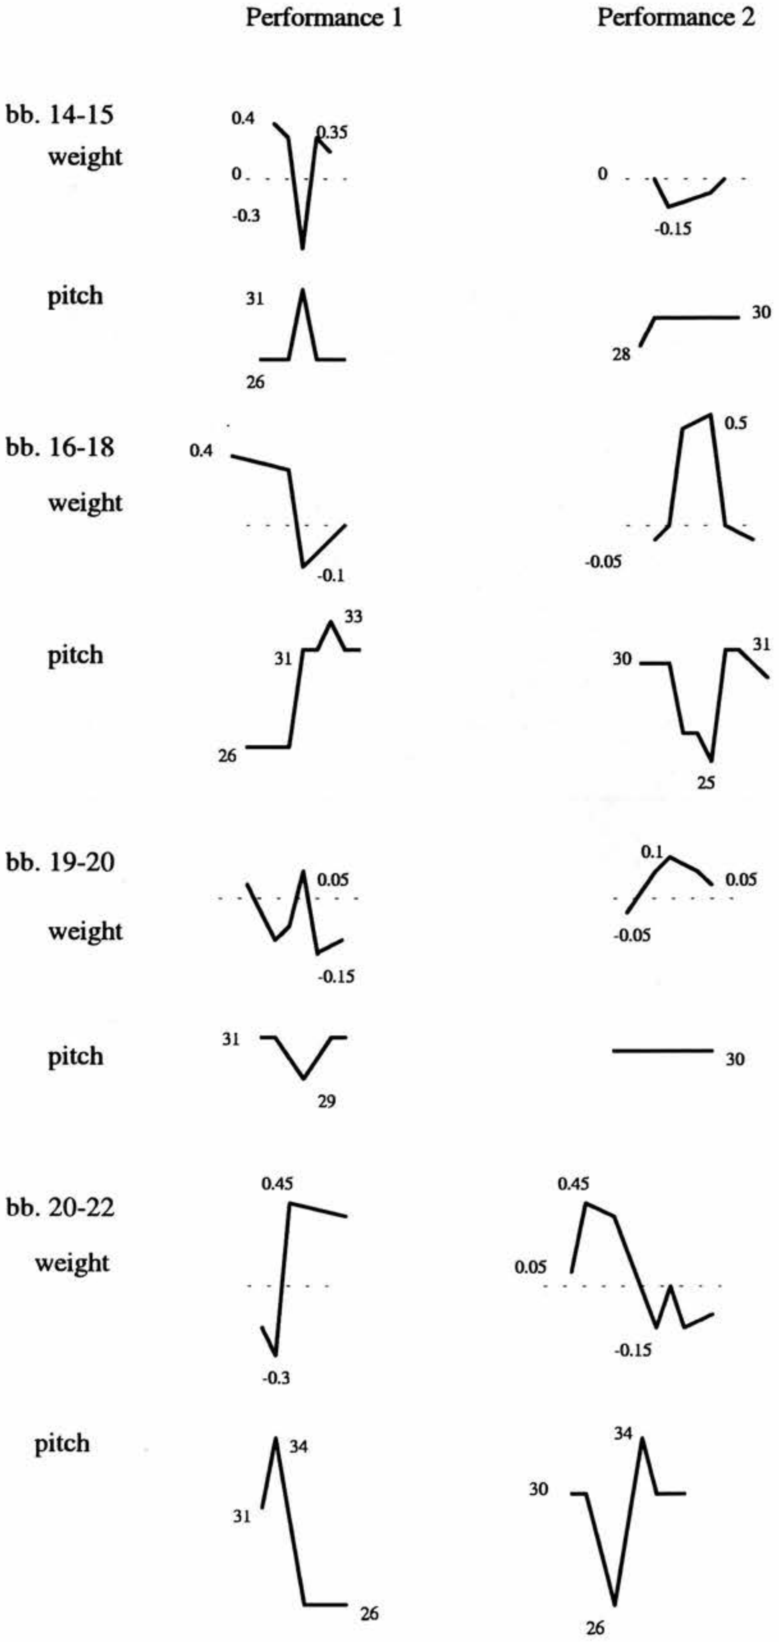
Performance 1
Performance 2
bb. 14-15
weight
0.4
0.35
0
-0.3
0
-0.15
pitch
31
26
30
28
bb. 16-18
weight
0.4
-0.1
0.5
-0.05
pitch
33
31
26
30
31
25
bb. 19-20
weight
0.05
-0.15
0.1
0.05
-0.05
pitch
31
29
30
bb. 20-22
weight
0.45
-0.3
0.45
0.05
-0.15
pitch
34
31
26
34
30
26

Example 4.17

In conclusion, section [B] can be considered to consist of a number of gestures based on increase followed by decrease: increasing followed by decreasing dynamic (which brings to the foreground two distinct subsets of the prevailing pitch cluster - the white-note collection and the black-note collection); and increasing followed by decreasing levels of width, sharpness and roughness. The different 'rhythms' of these gestures, all occurring within a static pitch cluster, give this section its movement, and its power.

After a very detailed look at section [B], let us turn our attention to section [C], which consists of a single gesture, and will be much more straightforward to analyse.

Section [C] is orchestrated for strings and flutes and clarinets; the double-basses and winds enter half way through the section, on the last crotchet of bar 25. Beginning with sustained pitches in the strings, an oscillating and increasingly rapid pitch figure is progressively introduced, so that the texture moves from a completely static state to a saturation of internal movement.

Turning to figures 4.9a and 4.10a, we see that, to a greater extent than in previous sections, the contours of the graphs for roughness, sharpness and width are very similar between the performances. In performance 2, the graph of roughness shows a very high degree of symmetry. The graph begins at a level of 2.5, smoothly curves to a level of 5.5, and then smoothly curves back down to 2.5 once again. The smooth increase in roughness is due to the progressive increase in the density of the musical texture; the decrease in roughness is caused by the maintenance of the same degree of density, but with a gradual decrescendo coupled with the instruction to the strings of *poco a poco sul tasto* (gradually move the bow towards the finger-board). The roughness graph for performance 1 shows the same general contour, but not the same symmetry. This is mostly due to the jump in roughness that occurs at the entry of the winds and double-basses on the last crotchet of bar 25. The moment of this entry is marked on the graphs of both performances (shown as 'ww + cb'). In terms of roughness contribution, the entry of these instruments in performance 2 is much smoother than in performance 1.

The entry of the winds and double basses occurs at the middle-point of the section, and is a decisive event in the timbral progression of this section. In both performances, this mid-point entry is the high-point of the roughness progression, and it is the turning point for the sharpness progression. In performance 1, sharpness has been largely at band number 30. At the entry of the winds and double basses (which play a chromatic cluster spanning B^3 - $A\flat^4$), sharpness drops to band number 28. Sharpness then continues to drop (finally coming to band number 27) because the notes of the winds start to dominate further due to their rate of oscillation decreasing while the rate of pitch oscillation in the strings continues to increase - the slower moving notes are heard more distinctly through the texture. The same movement in sharpness can be seen in performance 2, except that most of the progression is slightly sharper - as we have come to expect of performance 2. Sharpness spends a significant amount of time at band number 31 in the first half of the section. At the entry of the winds and double basses, sharpness drops to band number 29, and completes its descent at band number 27 (the same band number as performance 1 at this moment).

The mid-point instrumental entry also has a significant effect on timbral width. For the first half of the section, width is relatively static, lying at around 0.75 in both performances. At the wind and double-bass entry, as would be expected, timbre becomes more focused, and then width slowly decreases due to the decrescendo, and due to the slower wind notes that come to the foreground.

The details of this movement in width are as follows:

Example 4.18

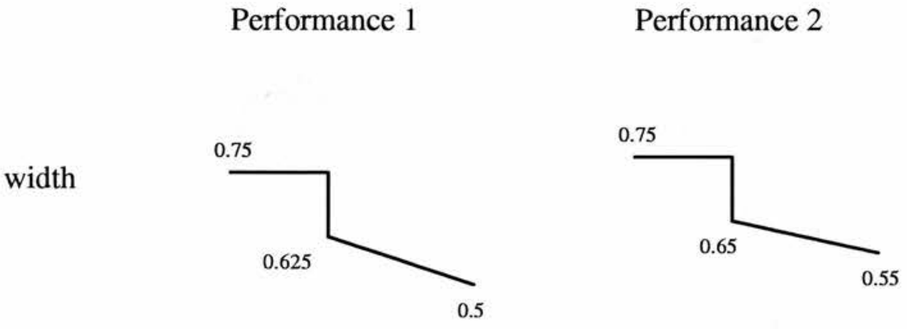

Turning to figures 4.11a, 4.12a, 4.13a and 4.14a, we see that, once again, the general contours of the graphs of timbral weight and pitch are largely different between the two performances, but the graphs do pass through similar levels, and therefore share in similar trends.

The movement of timbral weight and pitch in section [C] is summarised below:

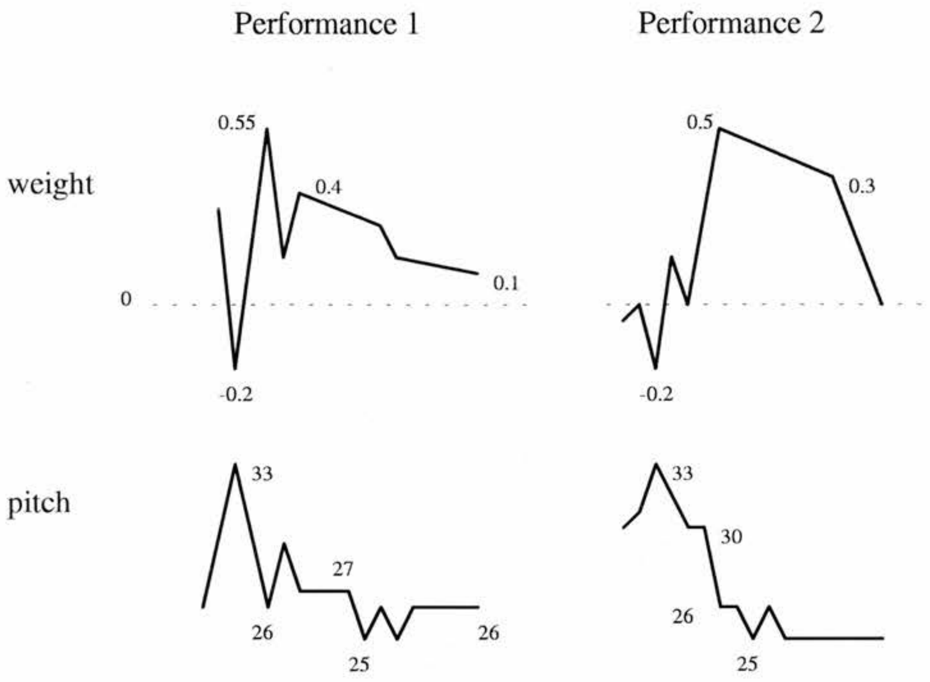

Example 4.19

The graph of pitch in both performances rises from the opening of the section to band number 33, and then falls during the course of the section, coming to rest at band numbers 25/26 (in figures 4.13a and 4.14a note how the red line, which plots data points at an average of 5, oscillates between these two band numbers at the end of the section in both performances). Before the rise to band number 33, performance 1 lies at band number 26, which is centred approximately on G^4. This pitch lies a little below the centre of the pitch cluster. Band number 33 is centred approximately on B^6, which is roughly an octave above the topmost notes of the cluster - in other words, the loudest band lies at the position of the group of second harmonics that belong to the topmost pitches of the cluster. From this, we can see that at the beginning of the section the emphasis shifts from the general sound of the cluster to the highest notes of the cluster. This shift occurs when movement is introduced to the highest notes. A similar situation applies to performance 2, except that the section begins on band number 30, centred approximately on B^5. So, in performance 2, the instrumental balance already favours the topmost notes, but as they begin to move, the emphasis shifts still higher. As movement progressively spreads to the lower pitches, timbral pitch falls, and the entry of the winds and double basses (marked on the graphs) takes timbral pitch still lower (though note that this entry does not play as decisive a role in lowering timbral pitch as it did in lowering sharpness - see above).

Like timbral pitch, timbral weight in both performances follows the same general trends. Weight moves to -0.2 when the topmost pitches begin to oscillate (the timbre is quite 'solid'), and then rises to around 0.5 as lower pitches begin to oscillate as well (the timbre becomes more 'ephemeral'). Timbral weight continues to become heavier as more of the lower pitches join in the oscillation, and as the pitches of the wind and double bass cluster come to the forefront as their notes begin to move more slowly. This tendency for the weight to become heavier as the section progresses reflects the perception that the sound becomes increasingly 'grounded' as we move towards the end of this particular gesture.

Let us now move on to sections [D], [E] and [F], which will be discussed as one section. The reason for this will become apparent if we look once more at figures 4.9a and 4.10a. From letter [D] up until, but not including, letter [G], the graphs show basically a single sweep of movement - a single timbral gesture. Roughness curves upwards, sharpness climbs progressively higher, and timbral width curves downwards (except for a brief section just before letter [G]). Turning to figures 4.13a and 4.14a, we see that timbral pitch also climbs progressively higher during these sections, and in figures 4.11a and 4.12a, we can see a general trend for timbral weight first to climb, and then to fall during section [F].

Section [D] consists of a static chromatic pitch cluster stretching from B^3 to Ab^4, played by the double basses and two cellos. All timbral measures are relatively static during this time, as would be expected. In performance1, however, there are two small peaks in the roughness graph, and the graph of timbral width shows quite a lot of movement. This movement is caused by the sound of the musicians turning the pages of their music - remember that this is a live recording. The levels of these measures for the two performances are compared below:

	Performance 1	Performance 2
roughness	1.75	2
sharpness	28	28
timbral width	0.575	0.65
timbral weight	0.25	0.1
timbral pitch	26	27

Example 4.20

The levels of roughness and sharpness are the same, or very similar. The difference in the measure for timbral width is due to the pitch G^4 dominating the cluster in performance 1, whereas performance 2 presents a more diffuse, evenly distributed sound. This is the same reason for timbral weight being heavier in performance 2 than in performance 1 - without the high single pitch dominating the texture, the timbre of performance 2 is heard to rest more 'solidly'. Timbral pitch is different by one band number. In performance 2, the loudest band, number 27, has a central pitch of approximately B^4, which means that this band must consist solely of energy from the upper harmonics of the lower notes of the cluster. In performance 1, the loudest band is number 26, whose central pitch is approximately G^4, the fundamental of the 'solo' note. Thus, bad balance (Ligeti asks repeatedly in the performance instructions to *Atmosphères* that no individual pitches should stand out) has a marked impact on the nature of the timbre.

In the plan of the movement suggested at the end of the structural analysis (see example 4.7), section [D] was labelled an 'interlude' prior to the second section of the first part of the piece (the second section consists of the music from letters [E] to[G]). After this interlude the 'sweep' of the timbral parameters begins in earnest. Section [E] consists of a rapid pitch expansion. The top section of this expansion goes on to move upwards during section [F] (see figure 4.1). Section [E] is orchestrated for strings, piccolos, oboes, clarinets and trumpets. By section [F], only the violins remain from the strings, which play with the wind instruments that were mentioned previously. By the end of section [F], only the four piccolos remain.

The overall rise in roughness from letter [E] is due to the increasing level of activity and the increasing dynamic, even though instruments drop-out along the way. Looking at the graph of roughness for the two performances during sections [E] and [F], we see that, in both performances, the rise in roughness is not smooth. In performance 1, there is a sudden rise at around 212 seconds. This is due to the entry of the violas and cellos in bar 31 (marked on the graphs of both performances), the cellos playing *sul ponticello tremmolo*. This same event causes timbral width to become momentarily more focused. It does not stay focused for very long, for the entry of the winds, which coincides with the first violins moving to higher pitches, causes width to become more diffuse, and also causes sharpness momentarily to rise (marked on figure 4.9a). However, this entry of the winds also signals a tendency for the pitch cluster to become more compact, so that from here on, timbral width decreases as we approach letter [G]. In performance 1, roughness rises relatively smoothly until about the 243rd second. A peak in roughness at the 243rd second, which coincides with a peak in width, and a momentary rise in sharpness, is due to the final *f* of the crescendo for the three trumpets (marked on the graph). After the 243rd second, roughness begins to fluctuate around a level of 9. These fluctuations are due to the texture now being effectively reduced to just four piccolos (there are still violins playing a sustained cluster, inaudible in both performances). As the players briefly stop playing in order to breathe, and as they articulate changes of pitch, roughness also changes, though there is an impression of an overall level of roughness, which is high.

These changes in roughness do not cause width to fluctuate. As the texture thins to just the four piccolos, width falls quite steeply. It increases once more just prior to letter [G] due to a last crescendo on the part of the piccolos.

In performance 2 there is a strong dip in timbral width at the 177th second. This is due to a note played by a cello suddenly coming to the forefront of the texture, so that timbre briefly becomes much more focused (marked on the graph). This event also causes a small rise in roughness and a momentary rise in sharpness. The rise in roughness at the 183rd second occurs at the same moment in the score as the rise in roughness which was discussed above for performance 1 at the 212th second - due to the entry of the violas and cellos. This entry is interpreted in a more 'dramatic' way than in performance 1, so timbral width does not decrease here as it did in performance 1. As in performance 1, sharpness increases on the entry of the violas and cellos, and remains higher.

With the entry of the winds, width begins to fall. As in performance 1, when the piccolos start to come to the foreground (at around the 203rd second), roughness rises steeply. The peak in roughness at 209 seconds, as was the case in performance 1 at this same moment in the score, is due to the trumpets reaching the top of their crescendo. This moment also causes sharpness to rise and width to increase briefly. From this moment, roughness fluctuates around a level of 13. This is considerably rougher than this moment in performance 1, where a roughness level of 9 was reached. The higher reading in performance 2 is due to the piccolos being recorded more closely, and because the players articulate the change of pitch more strongly, and re-articulate more strongly after pausing for breath. Just as in performance 1, there is a brief section just prior to letter [G], where width increases once again, along with increasing roughness. As in performance 1, this is due to the piccolos pushing through a final crescendo before they come to the end of their passage.

The movement of roughness, sharpness and timbral width in this section of the music is summarised in the following diagram:

Example 4.21

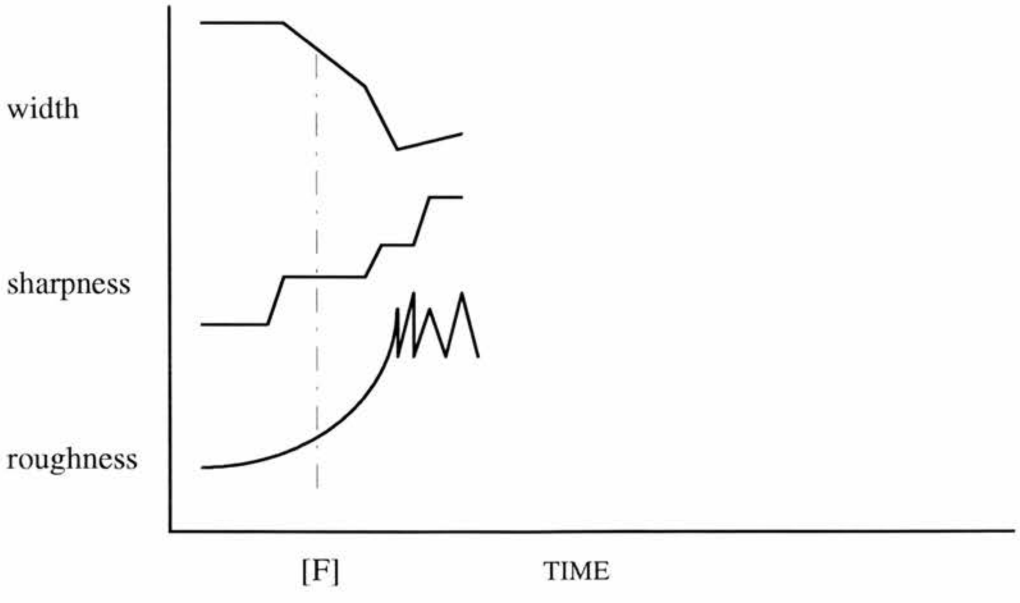

To conclude this look at the second section of the first part of *Atmosphères*, let us examine the graphs of timbral weight and pitch for sections [E] and [F]. The movement of these graphs is summarised below:

Example 4.22

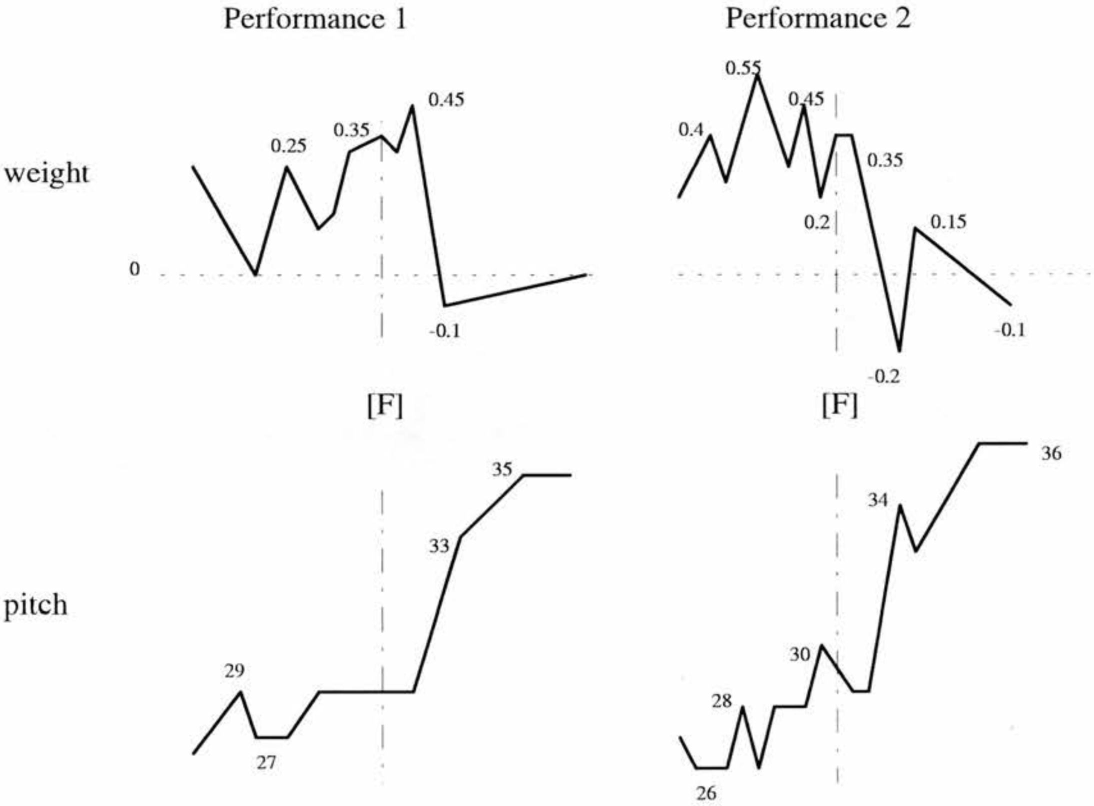

The common feature in the two performances is the sudden fall in timbral weight a little after [F], along with the corresponding rise in timbral pitch. This is caused by the piccolos coming to the forefront of the texture - the position of the loudest band rises to a level equivalent to the frequency of the fundamentals (the centre of band number 33 is approximately B^6; the centre of band number 34 is approximately D#^{7}), and the weight of the sound, the 'solidity' of the sound, becomes greater. Prior to this, the two performances show different tendencies in the movement of timbral pitch and weight. In performance 2, weight becomes lighter on the entry of the violas and cellos - this entry makes the sound more amorphous. But on the entry of the winds (where weight drops to 0.2), the sound begins to become heavier, more 'solid,' and this tendency continues as the piccolos rise and *crescendo*. In performance 1, however, the entry of the violas and cellos, and the entry of the winds both cause weight to become lighter. In performance 1, the entry of the winds sounds much gentler than in performance 2 - in performance 2 the winds play more 'insistently' and they are recorded more 'closely.' In both performances, timbral pitch follows the fundamentals upwards. The differences in timbral pitch between the two performances are not significant.

The first half of *Atmosphères* has been shown to consist of a number of musical gestures - pitch, textural, and timbral. These gestures move one into the other, so that we feel the sections of the piece to be very much a part of a flowing continuity. The second half of the piece, however, is built from contrasting parts. Here, the emphasis is not on continuity, but on disruption; that is, disruption within the context of the piece, which provides cohesion.

Figures 4.9b and 4.10b show the combined graphs for performances 1 and 2 respectively of the second half of *Atmosphères* - the graphs begin just after the beginning of section [G]. Like the first half, if we compare the overall shapes that appear in each graph, we find that the two performances are basically similar to each other. It is in the detail that we find the significant differences.

Although the sections in the second part of *Atmosphères* contrast with each other much more than the sections in the first part, the sections in the second part (that is, from letter [G] onwards) can still be grouped together into coherent musical gestures.

The first gesture we will consider is that consisting of sections [G], [H] and [I].[15] Section [G] consists of a static chromatic pitch cluster spanning C#^{1} to Ab^1, played *ffff, tutta la forza*, by the double basses. Sections [H] and [I] consist of a texture that collapses into itself. There is constant movement made up of individual descending lines. Overall, this texture begins as a very wide pitch cluster (C^3 to G^6), and then narrows until, just before letter [J], it consists of the very narrow pitch cluster Bb^3 to C#^{4}. Sections [H] and [I] are orchestrated for all the strings.

In performance 1, roughness is seen to drop during section [G] - from a level of 6 to a level of 5. In performance 2, roughness measures 4.5 at the beginning of the graph, and drops to around 4 just prior to letter [H]. These differences are due to section [G] being played much more aggressively by the double bass players in performance 1 than by those in performance 2. The section is also recorded more 'closely' in performance 1 - surprisingly, as performance 2 has, up until now, been the 'closer' recording. The decline in roughness during section [G], seen in both performances, is due to a small decrescendo, despite there being no decrescendo marked in the score. Sharpness in performance 1 lies at band numbers 25/26. In performance 2, sharpness lies in band number 25. The centre pitch of band number 25 is approximately D#^{4}, which means the loudness centroid lies well into the upper harmonics of the double basses fundamentals. Timbral width shows a relatively high degree of fluctuation, which would be expected due to the timbral changes caused by the constantly alternating changes of bow. The level of timbral width for both performances lies at around 0.625.

---

[15] [I] is not marked on the graph because its position cannot be confidently identified. Section [I] does not distinguish itself by audible instrumental or textural change.

Timbral weight and pitch are shown in figures 4.11b, 4.12b, 4.13b and 4.14b. In performance 1, pitch lies at band number 21 (approximate central pitch: B^2), which means it occupies the space of the second and third harmonics of the double bass fundamentals. Perhaps surprisingly, weight for section [G] in performance 1 is relatively light - around 0.45. Our initial reaction when thinking of a double bass sound is to imagine it as heavy. However, this is its *low pitch* we are thinking of. The sound of a double bass is somewhat 'woolly' and unfocussed when compared with the timbre of, say, a piccolo or a trumpet, evidence for whose relatively heavier weight we have seen in the first part of *Atmosphères*. In performance 2, timbral pitch lies initially at band number 21 and timbral weight initially at 0.45 (the same levels for this section which we saw in performance 1). However, before we reach letter [H], pitch rises to band numbers 25/26, and weight drops to -0.15, before rising a little to 0.05. The reason for this is difficult to hear. It is certain that the music *does* sound heavier at this moment, and then becomes somewhat lighter soon afterwards. It is likely that one of the double bass players begins to play more loudly, thus putting more energy in the upper harmonics, thus raising the position of the loudest band - that is, timbral pitch - thus causing timbral weight to become heavier.

The information on timbre for section [G] is summarised below:

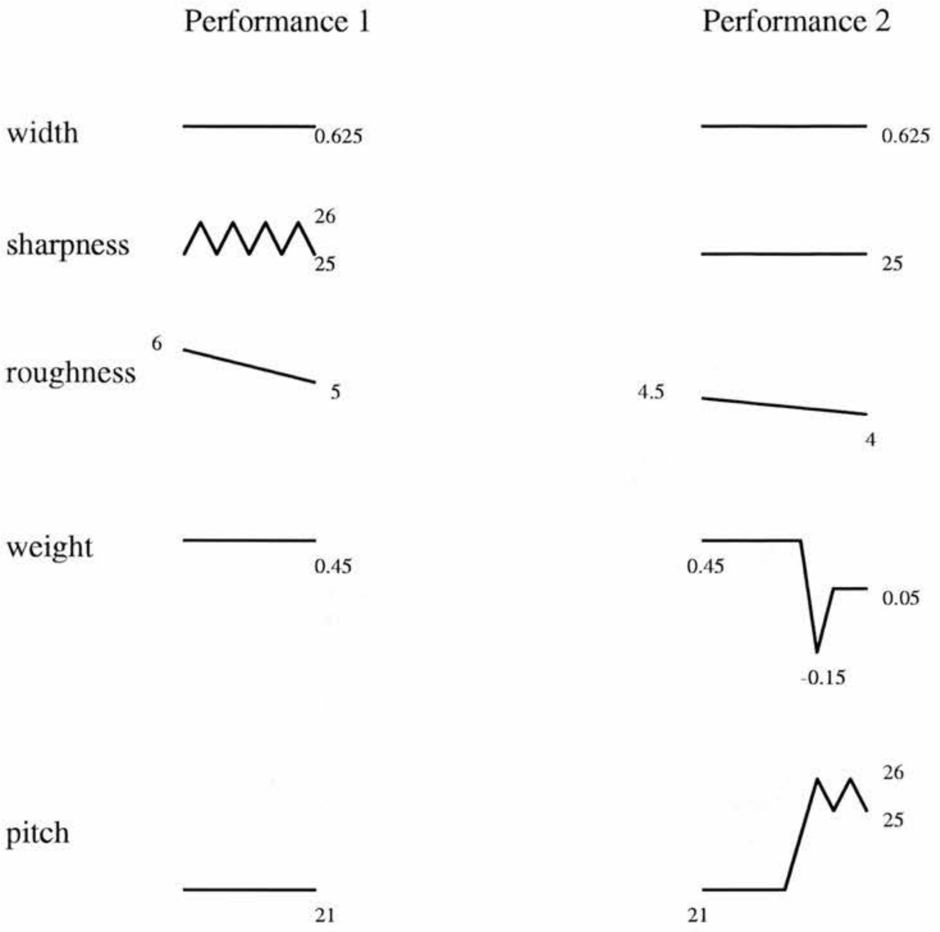

Example 4.23

Sections [H] and [I] show a sweep of movement in the timbral graphs. In both performances, roughness curves upwards to a relatively high level, sharpness rises during the first half of the section, and then plateaus, and timbral width rises a little, and then gently falls.

In performance 1, the level of roughness continues at a level just a little below the level reached at the conclusion of section [G], and then continues to fall for approximately 10 seconds, before rising. In performance 2, roughness begins to climb from the very beginning of section [H]. In performance 1, sharpness does not begin to climb until a little before 10 seconds into the section, whereas in performance 2, sharpness leaps upwards immediately at the beginning of the section. It can also be seen that timbral width begins its descent earlier in performance 2 than in performance 1. The reason for all of these differences between the performances is that the double basses play more loudly and more forcefully in performance 1 than in performance 2, as has already been noted above. Because the double basses are playing very loudly, when the other strings enter at letter [H], the double basses are still clearly audible, therefore it is the masking affect of the other strings which causes roughness to initially decrease in performance 1 at the beginning of section [H]. As the double basses decrescendo, the decrease in roughness continues. When the double basses stop playing (marked on figure 4.9b), the degree of internal movement in the strings begins to increase, and thus roughness starts to climb. Also when the double basses decrescendo and then stop playing, sharpness begins to climb because the affect of the new higher string cluster starts to be heard more distinctly. As for width, it climbs during the first 10 seconds of section [H] because of the additive effect of the strings playing their texture along with the sound of the double basses. When the double basses stop playing, width decreases, and it then continues to decrease because the pitch cluster folds in on itself, as was described above. In contrast, in performance 2, these movements in the timbral graphs occur from the very beginning of the section because the double basses are not heard nearly so strongly.

Levels of sharpness and timbral width, after taking into account the 'delay' in performance 1 due to the double basses, are the same between the performances. Sharpness, however, is higher in performance 1. This is due to stronger, more 'aggressive' playing in the strings, a style of playing carried over from the playing style of the double basses in section [G].

Below is a summary of timbral width, sharpness and roughness for sections [H] and [I]:

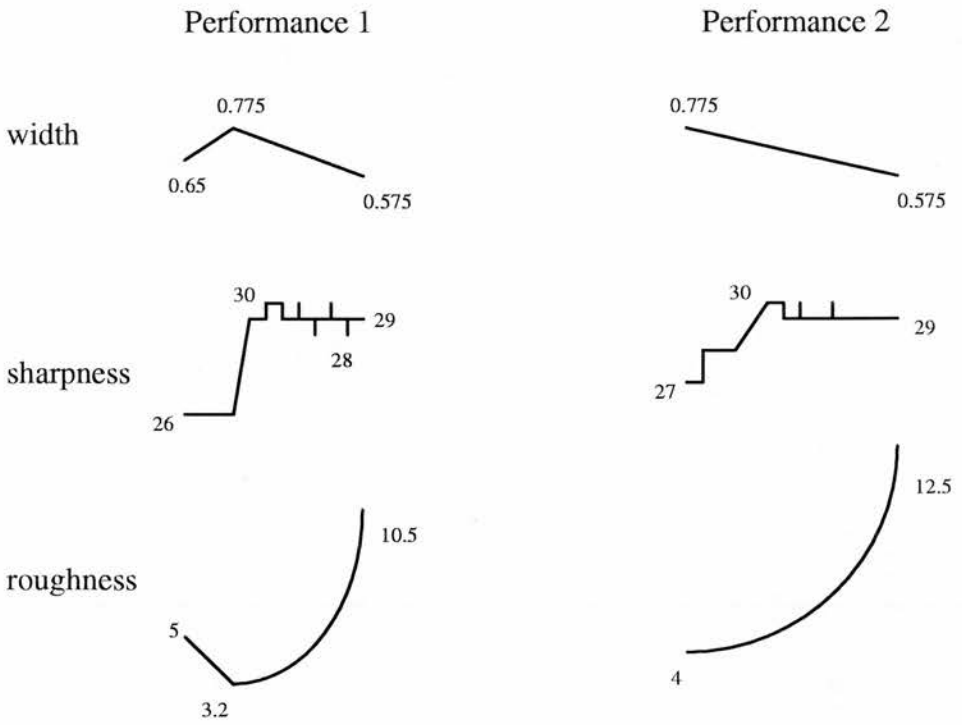

Example 4.24

The 'delay' seen in performance 1 in the measures of width, sharpness and roughness, is also seen in the measures of timbral weight and pitch. At the start of section [H] in performance 1 (see figure 4.11b) weight is still light (0.5) because the sound of the double basses predominates. When the sound of the upper strings comes to the foreground of the texture, but the double basses are still heard, weight becomes heavier (0.05), as the sound is now more 'solid.' It is near to this level of weight (0.1) at which section [H] of performance 2 begins, as the double basses do not dominate as they do in performance 1. For both performances, as the sound of the double basses fades still further, and then ceases, the activity of the strings becomes more and more busy, and weight becomes lighter once again. This lighter, less 'well-defined' sound seems to be due to the pitches of the cluster coming closer together coupled with the rapid movement in the individual parts. Note that in figure 4.11b (performance 1) during section [H], there is a lot of movement in the trace which plots data at an average of 5 (data points are an interval of 0.232 seconds apart) from 25 seconds to 33 seconds. Thus, the smooth movement towards a lighter weight which the trace at an average of 50 suggests during these eight seconds is not, in fact, heard this way. It is only at around 35 seconds that the lighter timbre becomes clear.

In both performances, after taking account of the 'delay,' timbral pitch for sections [H] and [I] moves between band numbers 25, 26 and 27.

The information on timbral weight and pitch for sections [H] and [I] is summarised below. A dotted line indicates the general direction of movement, but that there is much fluctuation either side of this line (this convention is also used in later summaries - the degree of the fluctuation is often indicated by figures in brackets).

Example 4.25

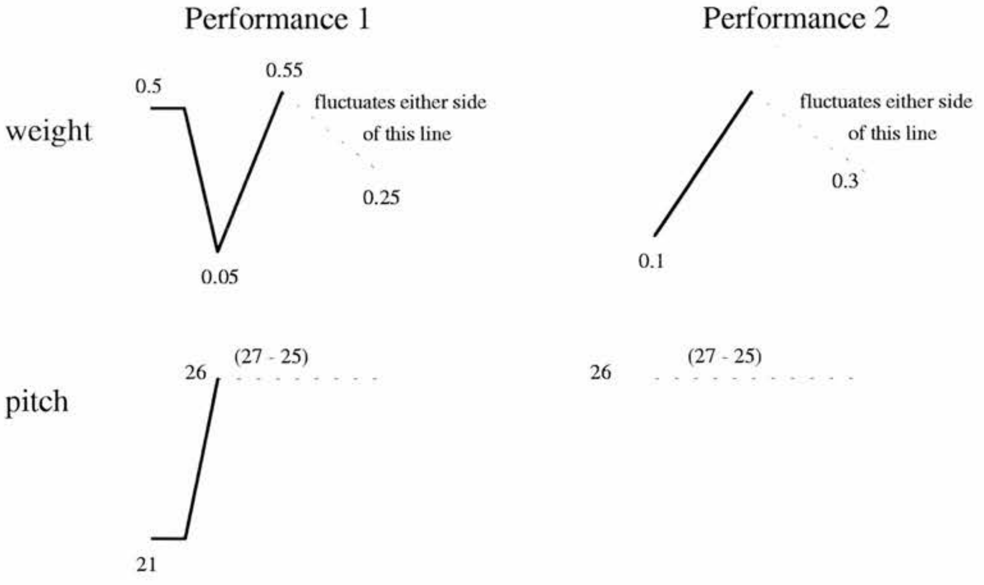

Sections [J], [K], [L] and [M] offer a great contrast to the textures and timbres heard in the previous two sections, [H] and [I]. From the full sound of the entire string section at the end of section [I] (discussed above), we move to the gentle, murmuring texture presented by the clarinets and flute, plus a single horn, in section [J]. Section [K] offers more of a 'pointillist' texture, played *pppp* by flutes, horns, trumpets and just a very few players from the string section. Section [L], scored for flutes, clarinets and horns, begins to build towards section [M] by beginning to expand the pitch cluster, and by introducing crescendos to *forte*. Section [M] itself shifts to a different pitch cluster twice as wide as the pitch cluster at the end of section [L], and introduces very strong individual crescendos to *ffff* which involve the entire brass section.

Turning to figures 4.9b and 4.10b once again, we see that during sections [J], [K], [L] and [M] there is distinctive increase in roughness and sharpness. In performance 1, roughness increases gradually during sections [J], [K] and [L], and then leaps upwards at letter [M]. In performance 2, sharpness remains relatively static during sections [J] to [L], but becomes much higher at letter [M]. In performance 1, sharpness lifts during section [L] before the further increase at letter [M]. Timbral width in both performances tends to fall during sections [J], [K] and [L], and then rises again during section [M].

The measurements for width, sharpness, roughness, weight and pitch for section [J] in the two performances are summarised below:

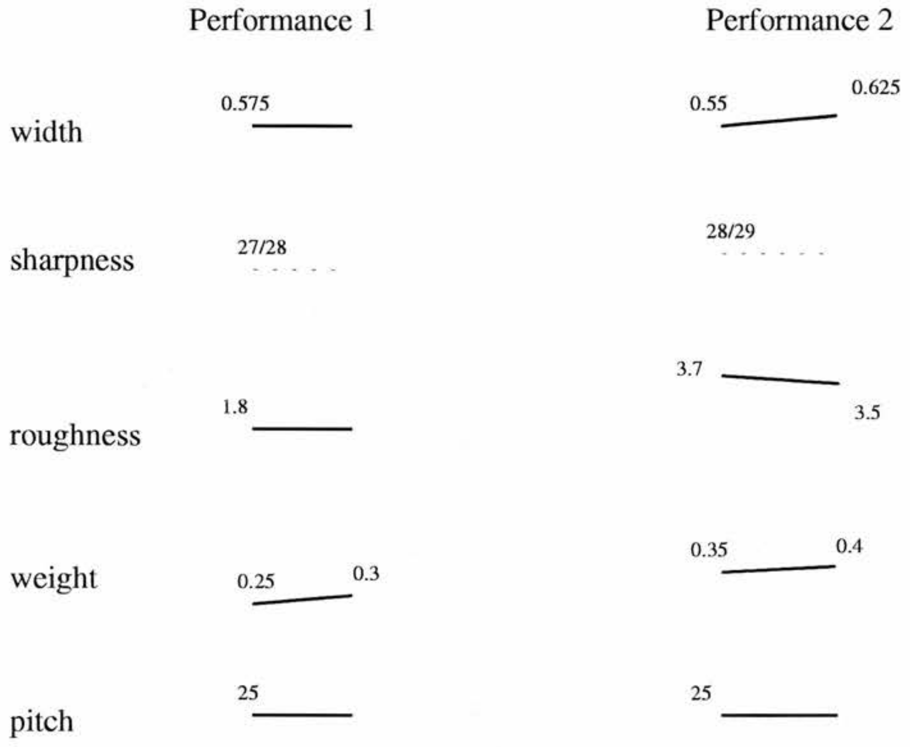

Example 4.26

Section [J] only lasts for approximately 5 seconds. Over this brief length of time, the fact that weight slopes upwards slightly in both performances, and that in performance 2 width increases and roughness decreases a little, is insignificant. The more significant differences between the performances are that performance 2 is rougher and sharper. This would appear to be caused by performance 2 being recorded more closely, and a little more brightly. There is very little difference in actual playing style between the two performances.

A summary of the timbral parameters of section [K] is shown below:

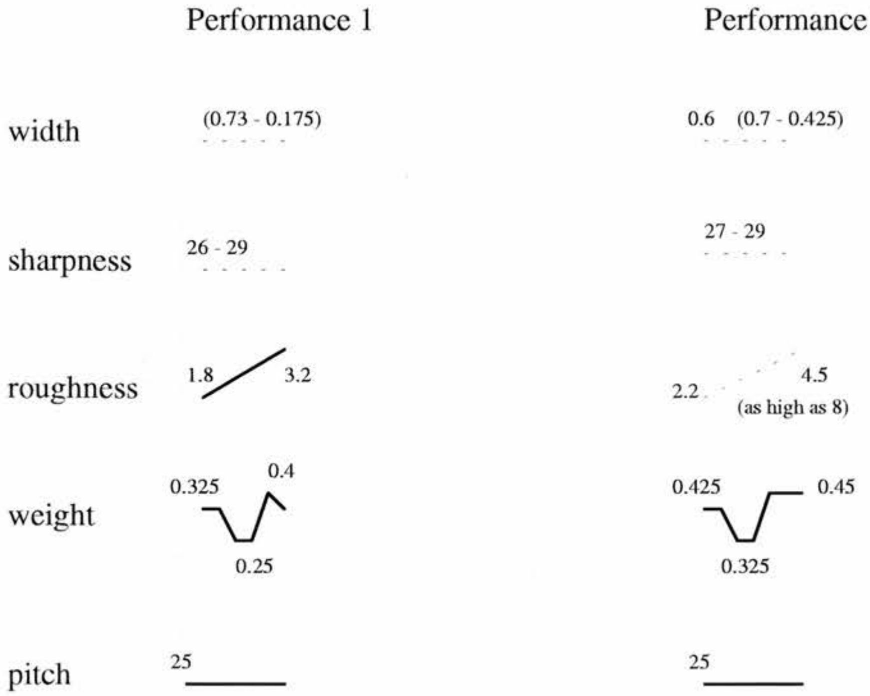

Example 4.27

In comparison to section [J], section [K] shows a decrease in timbral width (evident in the level of fluctuation), but much the same level of sharpness (though the fluctuation of sharpness is a little greater); it begins with a similar level of roughness, and then roughness increases during the section; weight is similar, and timbral pitch is the same. Thus, when comparing the timbre between sections [J] and [K], we find it is similar, but with decreased width and an increasing level of roughness. Section [K] consists of staggered entries, of brief duration. A crescendo occurs during each note, and the entries occur more densely as the section progresses. Section [J] consisted of a soft murmur of clarinets and flutes. The articulations of the entries along with the crescendos, the sparser texture, and the increasing density of entries are what cause section [K] to have decreased width (the crescendos and the sparser texture) and increasing roughness (the articulations, the crescendos, and the increasing textural density).

If we compare the two different performances of section [K], we find various differences. The two timbral measures which are most obviously different are timbral width and roughness. In performance 1, width shows a much greater degree of variation, and in performance 2, roughness is higher, and also shows much more variation.

Why does the measure of timbral width show so much movement in performance 1? Width is a measure of how focused or diffuse a sound is: the measure of timbral width therefore shows the timbre rapidly alternating between a more focused and a more diffuse sound in performance 1. Because the entries of the individual instruments and their individual crescendos are distinctly heard in the sparseness of the texture, an entry or a crescendo can cause the timbre to become suddenly more focused. In performance 2 we also see a moderately high degree of movement in timbral width. As has been observed before, performance 2 is often recorded more closely than performance 1. Here, this closer recording is coupled with a slightly more 'aggressive' style of playing - the crescendos are played with more of a bite in performance 2. These two factors - the closer recording style and the more 'aggressive' crescendos (which are the reasons that roughness is higher and more varied in performance 2) - mean that there is more 'noise' present during the sounding of the pitches, which is an important factor that causes a note to sound more diffuse. In performance 1 these factors are absent - the recording is a little more 'distant' and the crescendos are played more gently. The factor that contributes especially to the tendency of the music to become repeatedly suddenly more focused is the manner in which the horns are played in performance 1. They enter quite strongly, and their sound cuts through the texture. The sound they produce is quite 'pure,' and it is these notes in particular that cause width to become much more strongly focused.

In section [K], performance 2 is just a little sharper than performance 1. In performance 2, sharpness is at the topmost band number for this section, 29, for a greater percentage of the time than in performance 1, and in performance 1 sharpness goes as low as band number 26, whereas it only goes as low as band number 27 in performance 2. The reason for this slightly greater degree of sharpness lies in the sharper overall sound in this section in performance 2 (due partly to the 'closer' recording style) and also in the more forceful style of playing. The same reasons explain why performance 2 is slightly lighter in timbral weight than performance 1.

A summary of the timbral parameters of section [L] is shown below:

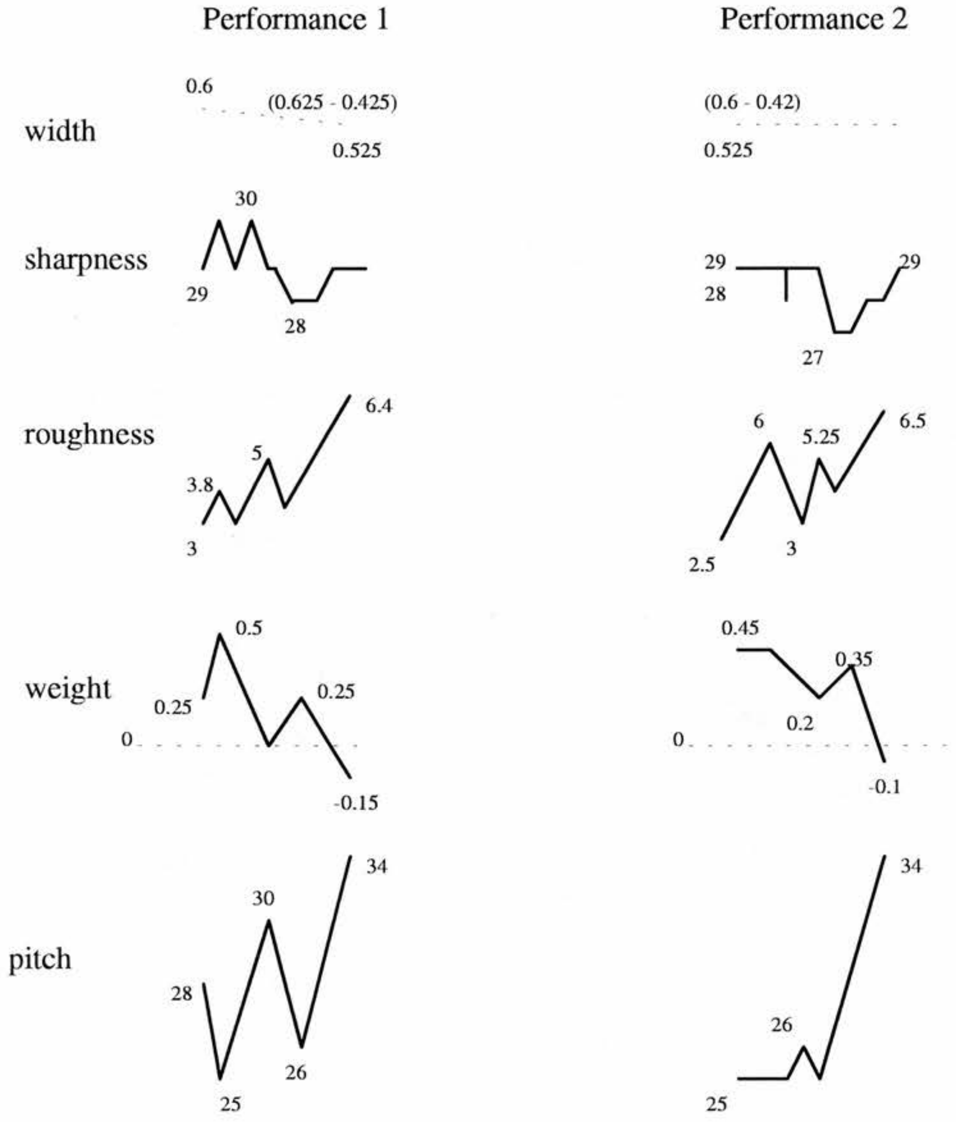

Example 4.28

In sections [J] and [K] we saw that the two performances were quite similar in their timbral measures. In section [L], however, we sometimes see instances of very different tendencies in timbre when we compare the performances.

Looking first at roughness, we see that, overall, section [L] has a higher level of roughness than section [K]. In section [L] in both performances, roughness rises three times, the third time taking us into the next section. These three roughness peaks (marked on the graphs of both performances) correspond to the three principal groups of instruments of this section as each group crescendos independently. Thus, the first roughness peak is caused by a crescendo by the flutes, the second by a crescendo by the clarinets, and the third by a crescendo by the horns, this last crescendo moving the listener into the much more active section [M]. In performance 1, these three roughness peaks lie in ascending order of magnitude - flutes are the least rough, followed by clarinets, followed by horns. This ascending level of roughness is not surprising given the type of sound of these three different types of instrument. In performance 2, however, the level of roughness for the flutes is surprisingly high - up to a level of 6, compared with a level of 5.25 for the clarinets. The primary reason is the closeness of the recording of the flutes in performance 2. A contributing factor may well be that the flutes in performance 2 are played with a very wide vibrato, so that moments of maximum roughness between two frequencies may occur more often in performance 2 as the different frequencies associated with the four flutes approach and part from each other.

Sharpness lies in much the same region in section [L] as in section [K]. Comparing performances, performance 1 moves between sharpness levels of band numbers 28 to 30, whereas performance 2 moves between sharpness levels of band numbers 27 to 29. So, performance 1 is a little sharper than performance 2. As in previous differences in timbre between performances in this part of the piece, this difference in sharpness is due to the way the instruments have been recorded. Performance 2 sounds more 'mellow' in this section than performance 1. This 'mellowness' is the reason width remains more focused (after becoming more focused in both performances at the entry of the horns) in performance 1 towards the end of this section and performance 2 becomes more diffuse (see figures) - the sound of the horns as they reach the end of their crescendo towards the end of section [L] in performance 1 creates a more concentrated, more 'nasal' timbre than the sound we hear in performance 2. Otherwise, timbral width covers virtually the same range in both performances.

If we examine the graphs of timbral width, sharpness and roughness together, we find that the three crescendos that occur during section [L] which were shown to affect roughness, also affect the other timbral parameters. The first two peaks in roughness coincide with sharpness rising from band number 29 to 30 in performance 1, and from band number 28 to 29 in performance 2 (these points are marked on the figures). During the beginning of the horn crescendo (the third crescendo), sharpness drops to band number 28 in performance 1, before rising to band number 29; and drops to band number 27 in performance 2, before also rising to band number 29. Although rises in timbral width (i.e., moving towards diffuse) are not exclusive to changes in roughness and sharpness in this section, the rises in roughness associated with the first two crescendos do coincide with rises in timbral width in both performances, while the beginning of the third (horn) crescendo coincides in both performances with a clear fall in timbral width. From these observations, we can summarise as follows:

Example 4.29

Section [L]

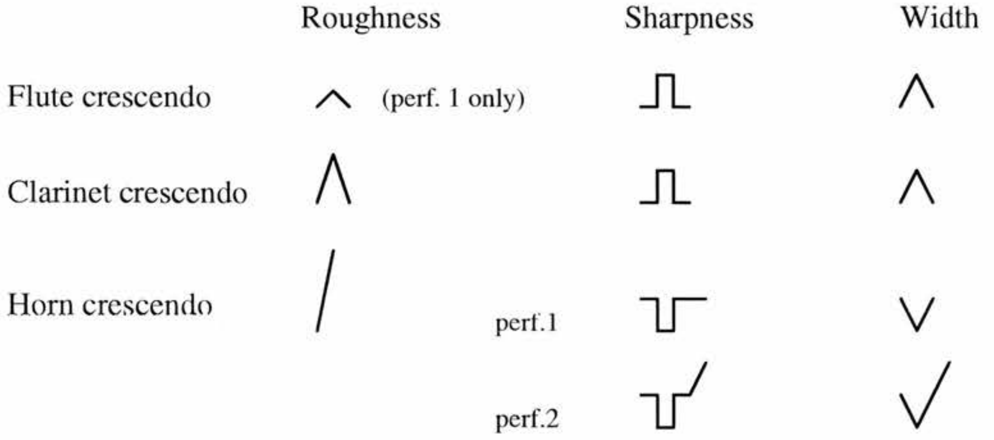

Thus, on comparing the flute and clarinet crescendos, we find that roughness is higher during the clarinet crescendo (see the reasons above why in performance 2 the flute crecendo is rougher), and sharpness and width both increase to similar degrees - as the volume increases the sound of the flute and of the clarinet becomes sharper and more diffuse. At the beginning of the horn crescendo, sharpness and width drop - the tendency for the horns to cause width to become more focused has already been discussed in association with section [K] in performance 1. As the horn crescendo continues (we do not hear the horn decrescendo as it is covered by the entry of the trumpets at [M]) sharpness and width both rise once more, rising further in performance 2 than in performance 1.

Timbral weight and pitch move in similar ways in both performances. Just as with the other timbral parameters, movements in weight and pitch coincides with the three crescendos. The movement is summarised below:

Example 4.30

Section [L]	Weight
Flute crescendo	∧
Clarinet crescendo	∨
Horn crescendo	∧

Thus, during the flute crescendo, weight becomes lighter; during the clarinet crescendo, weight becomes heavier; and during the horn crescendo, weight first becomes lighter at the horns' entry, and then heavier as the crescendo continues. The movement in pitch is simply the inverse of the movement in weight.

The timbral shifts in section [L] are very subtle, comprising both the change in timbre due to change in volume, and the change in timbre due to change in instrument. These timbral changes are organised so that the timbre has a sense of continuous movement, and importantly there is control of the rate of timbral change: a greater number of timbral parameters move in an opposite direction when comparing the clarinets' and the horns' crescendos than when comparing the flutes' and clarinets' - the rate of timbral change increases in order to prepare us for the dramatic entry of section [M].

A summary of the timbral parameters of section [M] is shown below:

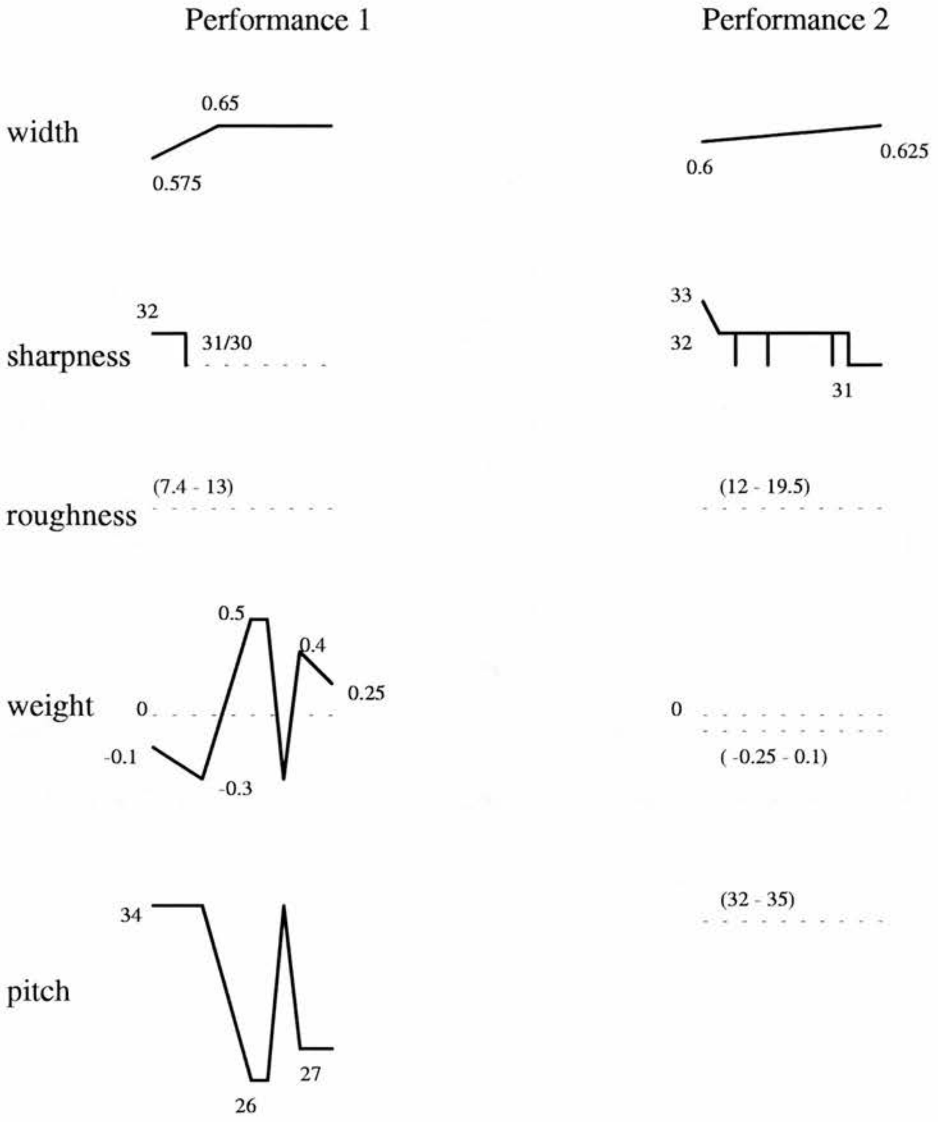

Example 4.31

Section [M] is the climax of the part of *Atmosphères* labelled as Episode 2 in the plan of the work proposed in example 4.7. Roughness in section [M] reaches the highest level for the second half of the work in both performances, and in performance 2 roughness reaches the highest level for the entire piece.

Section [M] consists of a series of re-articulated notes played by the horns and trumpets at *ffff*, and by the trombones and tuba at *ff*. This results in a constant, loud 'stabbing' texture. This is why roughness fluctuates to such a degree, and it contributes to the fluctuation seen in the degree of sharpness. Because of the nature of this texture, timbre will be discussed in terms of overall levels, rather than isolating individual timbral events.

On comparing the performances, we see that performance 2 is, as it has been in the past, one band number sharper, on average, when compared with performance 1. It is also much rougher. As has been stated above, these two differences are due to recording style, and recording closeness, rather than differences in the way instruments are played. Both performances show sharpness to be at its highest at the start of the section, and roughness to be at its lowest. This is because the entries of the different groups of brass instruments are staggered - the trumpets enter first (sharpness is at its highest), followed by the trombones and then horns (sharpness drops and roughness increases).

The measure of width in the two performances is very similar. Both performances show width rising during the section. In performance 1, the rise is due to the brief time it takes for all of the brass to begin their re-articulations - a texture consisting of held notes has a timbre which is more focused than a texture consisting of many loud articulated notes. In performance 2, although this same time exists when the brass are sustaining their notes, the closer style of recording means that the sound has a more diffuse timbre, so we do not see the same initial rise as we saw in performance 1 in this section. However, in performance 2, the brass play more forcefully as the section progresses (not indicated in the score), so timbral width slowly rises during the section.

In the balance between trumpets, horns, trombones and tuba, performance 2 seems to achieve a reasonably even balance, favouring the higher brass a little. In performance

1, the balance very much favours the low brass - the trombones and especially the tuba. This is reflected in the different levels of timbral pitch. In performance 1, pitch moves between band numbers 34 and 26, reflecting the swinging between the sounds of the high and low brass. The movement between these two extremes is very clear if the red graph (average of 5) is examined in figure 4.13b. In performance 2, however, pitch moves between band numbers 35 and 32, reflecting the balance more in the favour of the high brass. For these same reasons, although timbral weight in the two performances is similar, performance 1 contains episodes of much lighter weight (less 'solid' sound) than performance 2, whose weight lies predominantly below 0. (Lowness of pitch and loudness is no guarantee of heaviness of weight, as was seen in the graph of weight for section [G], which consists of solo double basses.)

Section [N] begins the third episode in the second half of the work (see example 4.7). After the loud and aggressive section [M], section [N] allows the energy to dissipate. Moving into section [N], the brass instruments, now joined by the clarinets, sustain their pitches from section [M], and then decrescendo to *morendo* a couple of bars later. At the start of section [N], a cluster enters, over five octaves in height, played by the string section. At bar 70, elements of this cluster begin to drop out, so that the sound becomes thinner and quieter, until only three cellos remain. During this process, section [O] has begun. There is, however, no audible sign to identify the beginning of section [O], so sections [N] and [O] will be considered as just one section, which, for convenience, will be referred to simply as section [N]. At the end of this section, just prior to [P], a new five note cluster enters, played by the violas. These pitches are marked to be played *sul ponticello*, with the additional instruction: "do not put the finger of the left hand all the way down; draw the bow without pressure. Almost without tone (more of a bowing noise)." The idea of playing 'almost without tone' is developed in the two-bar-long section [P], where the brass instruments are asked to play a series of articulations, with the instruction: "very soft blowing, without production of tone."

A summary of the timbral parameters for sections [N] and [P] is shown below:

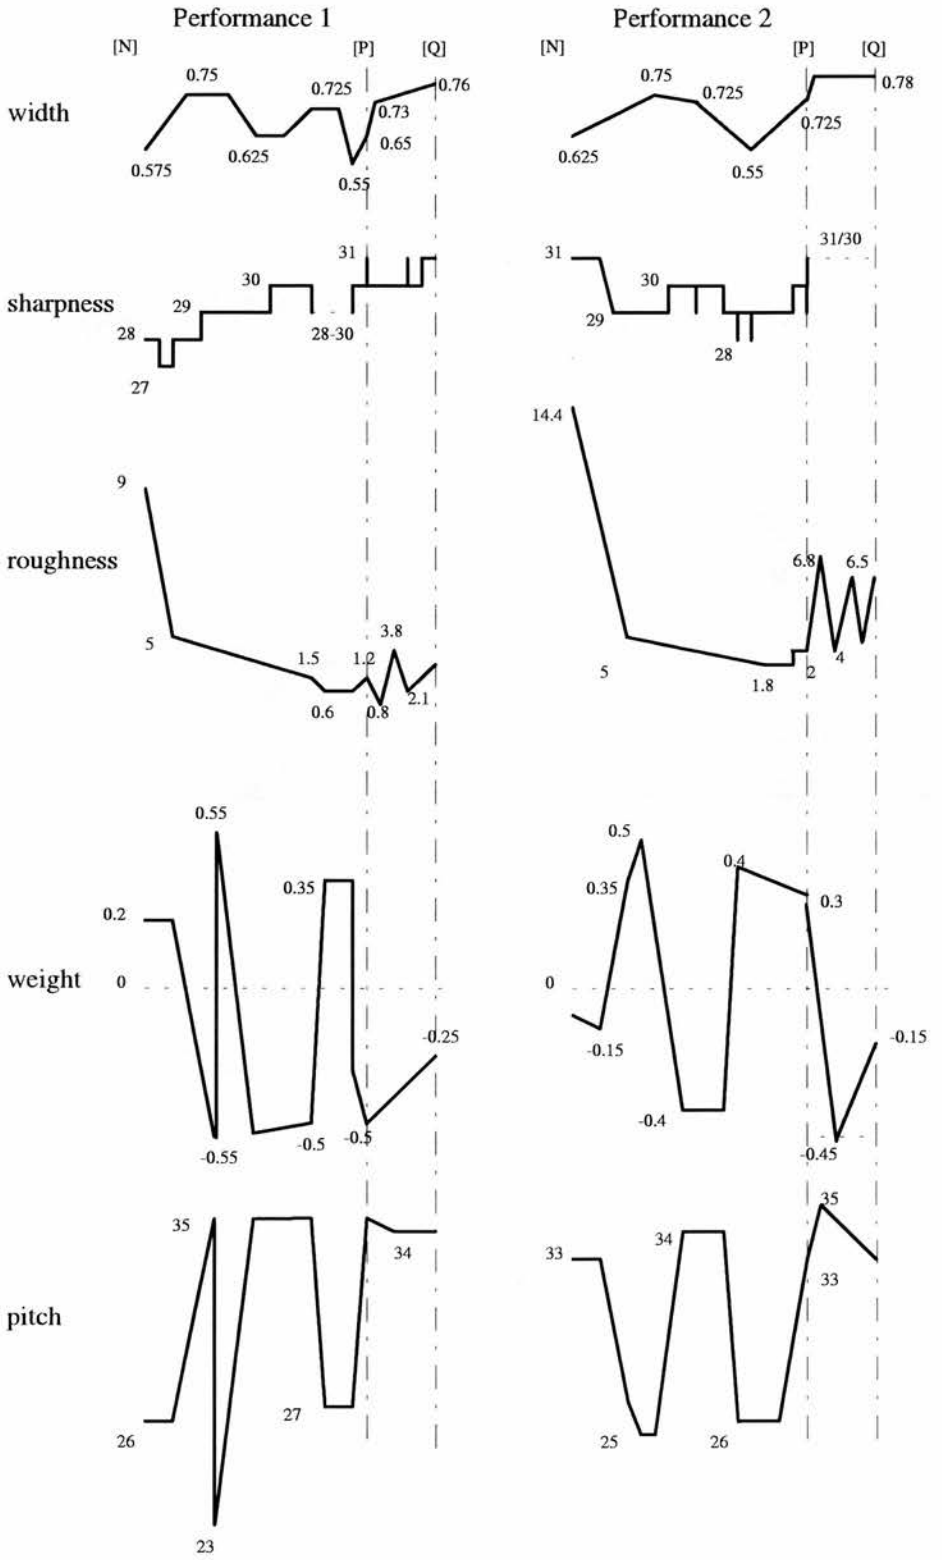
Performance 1
Performance 2
[N]
[P]
[Q]
width
0.575
0.75
0.625
0.725
0.55
0.73
0.65
0.76
0.625
0.75
0.725
0.55
0.725
0.78
sharpness
28
27
29
30
28-30
31
31
29
30
28
31/30
roughness
9
5
1.5
0.6
1.2
0.8
3.8
2.1
14.4
5
1.8
2
6.8
4
6.5
weight
0.2
0
0.55
-0.55
0.35
-0.5
-0.5
-0.25
0
-0.15
0.35
0.5
-0.4
0.4
0.3
-0.45
-0.15
pitch
26
35
23
27
34
33
25
34
26
35
33

Example 4.32

We can see that in both performances width becomes more diffuse at the start of the section. In performance 1, this increase in diffuseness (from 0.575 to 0.75) is due to the gradual dropping-out of the brass instruments - the presence of the string cluster plus the brass causes the timbre to be more focused than the string cluster alone (the brass cluster is less than an octave in height). When the brass are no longer heard, the level of timbral width ceases to increase. In performance 2, this same factor is responsible for the increase in width for the first half of the length of the increase. The second half of the increase in width is due to changes in the balance between the strings.

After remaining relatively stable for about 20 seconds following the increase in diffuseness discussed above, the level of width in performance 2 begins to decrease. This is due to individual string parts beginning to drop-out, so that the texture becomes thinner, and the timbre more focused. After reaching its lowest point, width increases again at the entry of the 'bowing noise' cluster - the sound is now more diffuse. At the entry of the brass at [P], playing a 'soft blowing' noise, width becomes still more diffuse. Looking at figure 4.10b, the shape of the curve of timbral width in sections [N] and [P] can be seen to be part of a much larger structure beginning half way through section [L], where a steady increase in width began. This low point in turn derives from a steady decrease in width that began at the beginning of section [H].

Width in section [N] of performance 1 does not show the same smooth-flowing lines as seen in performance 2. The pronounced dip at 125 seconds is caused by a fault in the recording, rather than by a musical event. Momentarily, there is a concentration of sound at low frequencies, as the upper frequencies drop-out. Thus, just for a moment, the timbre is much more concentrated. Whereas the overall movement in width in performance 2 was for it to rise, fall, and then rise again, the overall movement in performance 1, as can be seen in the summary above, is for it to rise, fall, rise, and fall. The reason for the first rise has already been discussed. The beginning of the fall is, like performance 2, due to instruments dropping-out of the texture, and the fall continues due to individual pitches moving into focus within the texture. When only a few strings are left playing, the balance changes, and width becomes more diffuse once more. The final drop in width (to 0.55) occurs at the entry of the 'almost without tone' *sul ponticello* section. In performance 2, the entry of this section occasions a progressive rise in width to 0.725. The reason for this difference is in the different ways of playing the *sul ponticello* section. In performance 2, the playing is quite light, so that the level of noise in relationship to tone is quite high - as the instructions say, "almost without tone." In performance 1, however, the strings play the *sul ponticello* section in a much harder, more 'determined' manner, so that the level of noise in relationship to tone is lower. Thus, in performance 1, the *sul ponticello* section has a more focused timbre when compared with the same section in performance 2. Because of this manner of playing the *sul ponticello* section in performance 1, the transition to section [P] involves more of a leap in timbral width than the smooth transition to section [P] seen in performance 2. I regard performance 2 as the performance that follows more closely the directions in the score in this section, and because of this, performance 2 shows a smoother movement in timbral width.

Turning now to sharpness, we see that the shape of the movement of sharpness in the two performances is quite different. Sharpness in performance 1 rises from a low of band number 27 near to the beginning of the section, to a high of 31 at and during letter [P]. In contrast, in performance 2 we see sharpness at band number 31 at the opening of section [N], then falling to band number 29. For most of section [N], sharpness lies at either band number 29 or 30. Then in section [P], sharpness lies at band numbers 30/31. Thus, section [P] shows the same level of sharpness between the recordings, while in sections [N] and [P], performance 1 shows a rising level of sharpness and performance 2 shows a much more static level of sharpness. The reason for this difference is that performance 1 in this section favours the sound of the lower strings, while performance 2 represents the sound of the large string cluster more evenly. To understand why this causes the difference in the sharpness levels, we need to understand the pitch structure of this section.

The pitches of the string cluster drop-out in a very ordered manner. Two processes happen, one after the other. First, the 'black' notes begin to drop-out one after the other, beginning from the bottom of the cluster. Second, when the top of the cluster is reached, the 'white' notes begin to drop-out, starting from the top and bottom of the cluster simultaneously, and funnelling towards the centre. So it is the four pitches near to the centre of the original cluster that finally remain. This process is seen very clearly in the graph of sharpness of section [N] in performance 2. Sharpness begins at band number 31 where the brass and strings are sounding at the same time. Because the brass are rich in upper harmonics, when they stop sounding, sharpness drops to band number 29. The 'black notes' then begin to drop-out from the bottom of the cluster moving towards the top. About half-way through this process, sharpness rises to band number 30 because there are more high pitches sounding than low pitches. As the process nears completion, sharpness drops back again to band number 29 because the gaps in the cluster are now evenly distributed, and the pitches now drop-out from both the top and bottom simultaneously. Sharpness rises at the end of section [N] on the entrance of the *sul ponticello* section. Sharpness rises yet again at [P] for the 'soft blowing' section for the brass.

The sharpness graph for performance 2 follows very well what we see happening in the pitch structure. The graph for performance 1 does not, because, given that the recording favours the lower strings, as the lower strings drop-out sharpness moves upwards, and the absence of the lower strings has a greater affect on sharpness than the absence of the upper strings. Thus, in performance 1, sharpness rises during the pitch process described above.

In both performances, roughness shows a very rapid decrease at the beginning of section [N]. This is due to the great change in texture from the 'stabbing' motion of section [M] to the sustained pitches in section [N]. The decrease in roughness is continued at the beginning of section [N] by the brass dropping-out after approximately two bars. The moment in both graphs when the strings are left playing alone can be clearly seen because the graph of roughness suddenly shows a far slower decrease in roughness level. Once more, in both performances, the effect of the *sul ponticello* section can be seen in the small increase in roughness at the end of section [N]. The roughness spike in performance 1 at 135 seconds occurs at the same time as, and is due to the same fault in the recording, as the downwards timbral width spike discussed above.

After this decrease in roughness over a relatively long period of time, section [P] brings a sudden increase in roughness with the section of 'soft blowing' for the brass. The score shows somewhat staggered entries at the start of section [P], and crescendo / decrescendo markings begin a few moments into the section. From the roughness graphs, it appears that performance 1 interprets this section well, showing increasing roughness during the section, whereas performance 2 does not show the increase in roughness associated with the staggered entries and dynamic markings. There is no audible increase in roughness or loudness during section [P] in recording 2.

Looking at the summary graphs of timbral weight and pitch above, we see that the two performances are similar in these respects. There is a fall in weight during the middle of the section (to -0.55 in performance 1; to -0.4 in performance 2) , followed by a rise in weight (0.35 - performance 1; 0.4 - performance 2), followed by another fall in weight at the start of section [P]. What causes these movements, and the dissimilarities between the performances? Section [N] in performance 1 starts with a much lighter timbral weight than performance 2. This is primarily because the brass is recorded much more closely in performance 2, and the sound is more 'solid' than that heard in performance 1 - timbrally heavier. After the brass cease playing, weight in performance 1 drops immediately to -0.55. It then rises immediately to 0.55, before dropping once more to around -0.5. In performance 2, the brass drops-out at the point in the above summary where weight reaches 0.35. Weight then keeps rising for about five seconds, until dropping to -0.4. Looking at the graphs of timbral weight (figures 4.11b and 4.12b), we can observe that when the strings begin to play alone, timbral weight is very unstable. The summary above of weight and pitch is based on an averaging factor of 50 (data points every 2.32 seconds). On figures 4.11b and 4.12b, we can see this graph along with the graph of points at a distance of 0.232 seconds (an averaging factor of 5). Thus, between 127 and 137 seconds in performance 1, and 117 and 123 seconds in performance 2, weight, as shown by the data points connected by the red line, is oscillating rapidly either side of 0. The peak in weight up to 0.55 in the summary of performance 1 is representative of this rapid, and wide oscillation. This rapid oscillation in weight reflects a rapid oscillation in timbral pitch - in performance 1, between band numbers 25 and 35, and in performance 2, between band numbers 25 and 34. Thus, the widely spaced pitch cluster creates widely spaced bands which compete for the status of the loudest band. Towards the end of the section, there is the passage played *sul ponticello*. It will be recalled that in performance 1 this passage was played quite strongly so that timbral width became more focused, while in performance 2, this passage was played less forcefully, so that timbral width became more diffuse. This difference in playing styles is also borne out in timbral weight. In performance 1, weight is heavy - the sound is quite 'solid' - in performance 2, weight is light - the sound is more 'indistinct.'

In section [P], both performances show the sound of the brass (played without tone) to be heavy. Both performances also show weight becoming lighter during this brief section. In performance 1, this is due to the greater forcefulness of the playing as the section progresses (this tendency was noted in the discussion of roughness). In performance 2, the rise shown in the summary presents the rise of the black-line graph in figure 4.12b. However, if the red-line graph is examined, it will be seen that this black-line rise reflects a rapid movement in weight, rather than a rise in its own right. Thus, it is concluded that weight in performance 2 of section [P] remains, overall, static. This fact is reflected by the placement of a dotted line in the summary, which does not slope.

The last segment of *Atmosphères* has been designated as a coda in the previous sectional analysis. This segment consists of sections [Q], [R], [S] and [T]. The coda introduces the previously unheard sound of piano strings played with brushes. It also introduces two similar types of string textures/timbres that have not been previously heard. One of these, heard during section [Q], consists of a 'pointillist cloud' of sound - a combination of *staccato*, *sul ponticello*, *sul tasto*, and *col legno* bowing techniques applied to *tremolo* and slurred notes. The other consists of delicate waves of *glissando* harmonics played very rapidly.

Section [Q] begins with brushed piano strings, emerging as the previous section for 'toneless' brass decrescendos to nothing. Then begins the first of the string textures described above. Four flutes then enter ([R]), playing a quiet pitch cluster, first over the string texture, then by themselves ([S]). At [T] the second string texture enters while the flutes sustain their pitches over a decrescendo to nothing. During section [T], low pitches enter played by the trombones and tuba, as well as low pitches played on the brushed strings of the piano. *Atmosphères* concludes as this string texture dies away, and all that is left is the soft sound of the piano strings.

The timbral parameters of sections [Q], [R], [S] and [T] are summarised below:

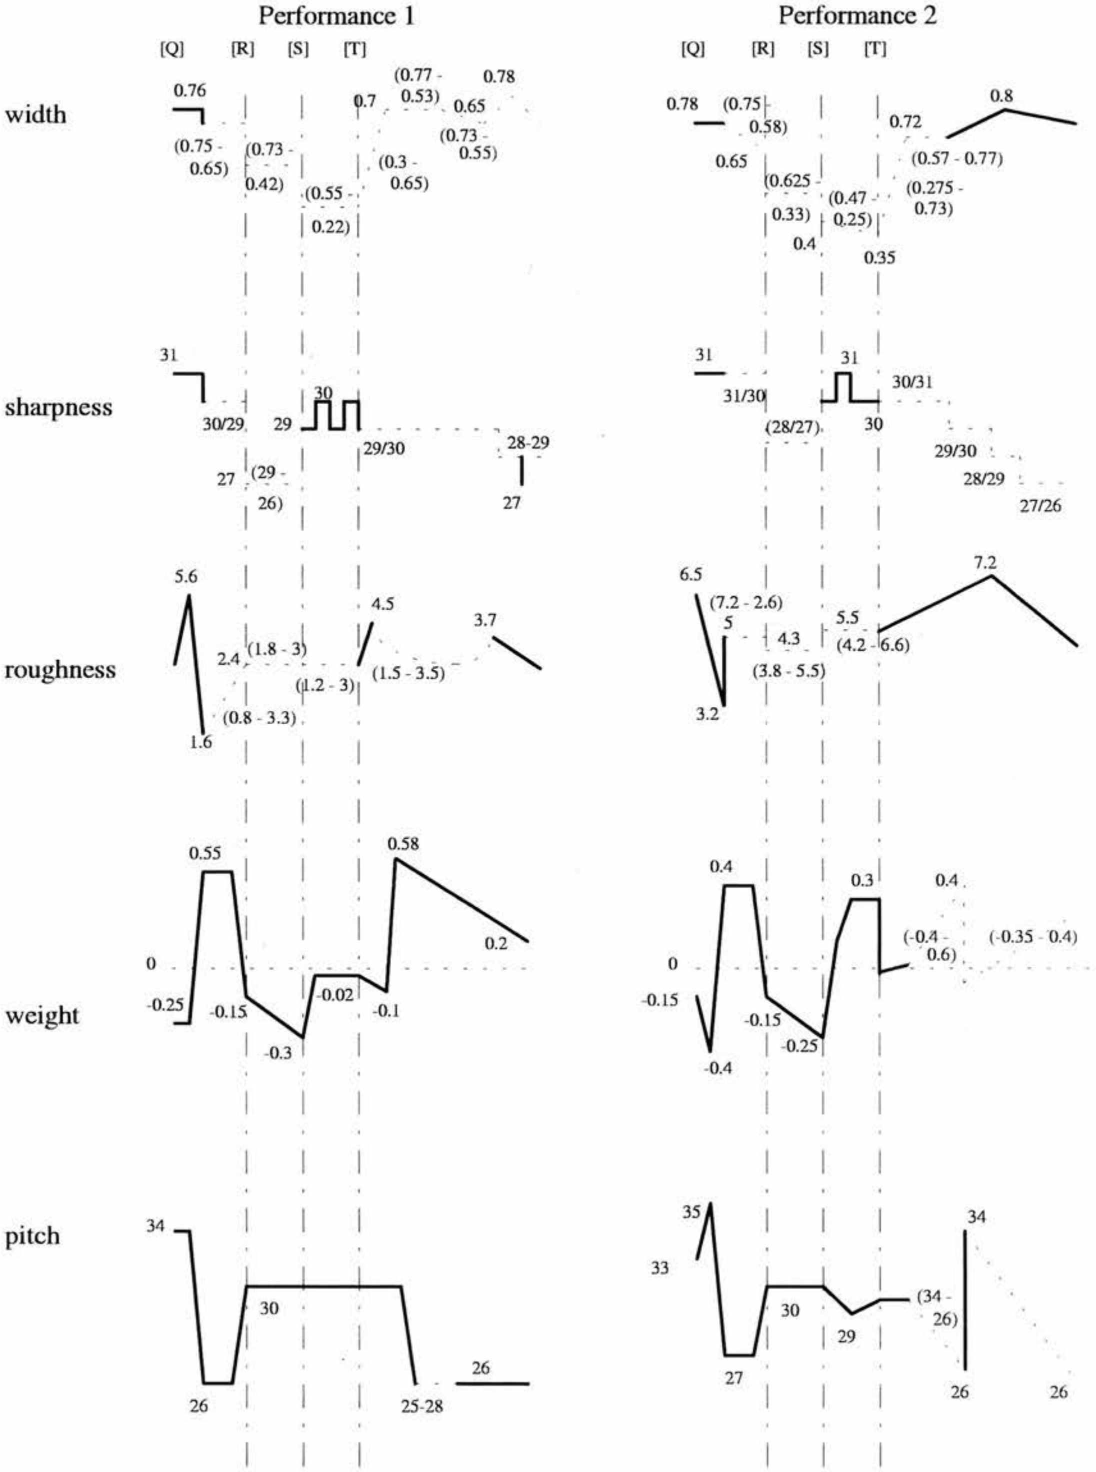

Example 4.33

Looking at the above summary, we see that during sections [Q], [R] and [S], in both performances, width becomes progressively more focused, though there is often quite a wide range of values for width in these sections. In performance 1, the sound of the brushed piano strings is recorded very closely, and the player is quite 'deliberate' in the way it is played. In performance 2, the sound of the piano blends more effectively. Therefore, when we look at the measurements for width and roughness at the very beginning of section [Q], we see that in performance 1 they rise relative to the levels at the end of section [P], whereas, in performance 2, they fall relative to the levels at the end of section [P]. In absolute terms, however, performance 2 begins more roughly than performance1.

In sections [Q], [R] and [S] roughness often varies considerably, and the degree of variability is shown on the above summary. Overall, however, after the strings have entered in the second bar of section [Q], roughness remains at the same average level until section [T].

Looking at sharpness in both performances as presented in the above timbral summary, we see that in section [Q] sharpness begins one bandwidth higher than it was in the previous section. During section [Q], sharpness drops by one bandwidth. At the start of section [R], sharpness drops by another three bandwidths, and in section [S] sharpness rises by three bandwidths again. Note that performance 2 is, on average, one bandwidth sharper than performance1, as we have come to expect.

Timbral weight moves very similarly in both performances. Section [Q] begins with a heavy timbre, and on the entry of the strings moves to a much lighter timbre. Section [R] brings a heavier timbre once again, and timbre continues to become heavier during section [R]. Section [S] brings a lighter timbre relative to section [R]. In performance 1, this timbre still lies below 0 (-0.02), whereas in performance 2, the timbre in section [S] is much lighter than this, moving to 0.3 a little into section [S]. As would be expected from the above observations on weight, timbral pitch is also very similar between the two performances. The slight variations in timbral pitch between the two performances in sections [Q], [R] and [S] can be readily observed from the above graphical summary.

Let us now compare these observations on timbre with what we know is happening in the score.

Section [Q] begins with the entry of the brushed piano strings that occurs as the previous timbre (softly blown brass) gradually decrescendos. The effect of the entry of the piano on width and roughness has already been discussed. Upon the entry of the strings, the piano soon after ceases to play. The entry of the strings and the cessation of the piano and softly blown brass marks the point where timbre becomes more focused, even though the *texture* of this string gesture has a very diffuse feel to it. At this point, timbre also becomes less sharp and less rough (sounding pitches are not as sharp or rough as the sound of 'toneless' blown brass or brushed piano strings), weight becomes lighter and timbral pitch drops (the timbre is now less 'solid,' and the 'centre of gravity' of the sound is lower).

At letter [R], the flutes enter with a four-note cluster while the previous string texture (which now includes double-basses) continues. Width decreases - a four note chromatic cluster played by flutes has a more focused timbre than a fifteen note chromatic cluster played with such techniques as *sul ponticello* and *col legno*. Sharpness also decreases. Roughness remains close to the level it reached at the end of section [Q] - remember that roughness measurements incorporate normalisation for amplitude, so although a greater number of sinusoidal components will be present when the flutes play as well as the strings, the result is divided by the increased amplitude. Weight becomes heavier, as the timbre of the flutes along with the timbre of the strings is more 'solid' than the timbre of the previous string texture.

Section [S] sees the flutes playing by themselves. Width becomes more focused, as would be expected, and sharpness rises, because the low string fundamentals are no longer present. Roughness remains very similar in performance 1, and rises in performance 2. Closely spaced flute notes produce a lot of beating, so, due to the normalisation in the roughness calculation, mentioned above, roughness does not decrease when the strings stop playing. Timbral pitch is very similar between the performances. Timbral weight, however, shows a significant difference - section [S] is lighter in performance 2 than in performance 1. The reason is due to the closeness of the recording of the flutes in performance 2. This close recording brings out a lot of ringing upper harmonics in the sound, causing the sound to be lighter. These upper harmonics are not nearly as prominent in the more distant recording of performance 1, therefore the timbre is heavier. The stronger presence or absence of these upper harmonics is the reason why the flutes in performance 2 have a higher roughness level than the flutes in performance 1.

We have tracked through how and why timbre changes during the first half of the Coda. Let us now examine the second half of the Coda - section [T].

Both recordings show width to increase at the beginning of section [T], beginning from the level at the end of section [S]. This increase is due to the string texture becoming 'thicker.' Surprisingly, the entry of the brushed piano strings in bar 93 makes no difference to timbral width in either recording. The clearest contribution to width is to be seen in performance 2 at 212 and 216 seconds, where the downward spikes show moments where a piano string is brushed harder than its neighbours, and the timbre momentarily becomes more focused. Indeed, surveying the other timbral parameters at the point of entry for the piano (performance 1 - 223 seconds; performance 2 - 208 seconds) we find no change in timbre. The sound of the piano *does* blend well with the sound of the strings, but it is still audibly a different timbre from the strings. It appears that, although we can audibly pick out the different timbre of the piano from its surroundings at the beginning of its entry, the contribution of the piano at this point to the *overall* timbre is insignificant as far as the timbre analysis computer programs are concerned. The entry of the brass at 245 seconds in performance 1, and at 222 seconds in performance 2, also makes very little impact on timbral width. However, we can see that there is an overall rise in timbral width in performance 2 (from 0.72 to 0.8 - see above summary). This rise is due to the accumulative effect of the presence of the piano and brass. This rise is not seen in performance 1 due to the recording being more distant, and the timbres therefore blending with each other more completely. Towards the end of performance 1, there is a 'hump' shape to the graph of timbral width. This is due to changes in timbre as the low piano strings are brushed. No other instrument plays at the end of the Coda. The sudden drop in width at the end of performance 2, as well as the sudden rise in sharpness, are both due to the sound of the piano becoming inaudible, and the graph plotting the timbre of the background noise of the recording.

At the beginning of section [T], sharpness in performance 1 lies at bands 29/30, and in performance 2 at bands 30/31. As has been pointed out before, performance 2 is, overall, sharper than performance 1. The entry of the piano makes little difference to the sharpness, though the two points in performance 2 when a piano string is brushed harder than it should be stand out as moments when the red graph of sharpness moves up to band number 32 (at 212 and 216 seconds). In both performances, sharpness drops by one band number at the entry of the low brass and low piano notes at bar 98 (performance 1: 28/29; performance 2: 29/30). At the end of the piece, the sound of the brushed low notes on the piano is all that remains. In performance 1, the sound of the live performance fades-out soon after the piano is left by itself. In performance 2, the sound of the solo piano is sustained for a much longer time - indeed, the recording finishes with the sound of the piano strings ringing after the last brush across them has been made. This is why the sharpness graph for performance 2 keeps descending (finally to band number 26/27), whereas the sharpness graph for performance 1 does not (this can be seen clearly on the summary above). Sharpness in performance 2 keeps getting lower as the sound of the brushed piano strings fades away.

In performance 1, roughness follows an approximate bowl-shaped arc for the first three-quarters of section [T], and then trails away at the conclusion of the section. The high-point at the beginning of the arc is due to beating between the flute notes. As we progress along this arc, the flutes drop out (217 seconds), the piano enters (220 seconds) and the low brass and low piano enter (245 seconds). Roughness trails away as the strings and brass cease playing, and the sound of the piano fades.

In performance 2, instead of an arc shape for roughness before roughness starts trailing away, we see a rise in roughness. The roughness spikes at the beginning of section [T] are, as in performance 1, due to the beating of the flute notes. At 205 seconds, the flutes cease playing, and a few seconds later, the piano enters. The low brass and low piano enter at 221 seconds. Because the strings are recorded more closely in performance 2 than in performance 1, the level of roughness is sustained through this first part of section [T], rather than describing the arc shape that we saw in performance 1 - falling as the flutes drop out and rising as the brass and low piano enter. The long tail of roughness that drops away in magnitude in performance 2 is due to the piano continuing for some time after the other instruments have ceased playing, and then roughness continuing to decrease as the piano strings are allowed to keep ringing after the brushing has stopped.

Weight in section [T], in both performances, begins with a short period of relatively little movement. This is the period where the flutes play with the strings. Once the flutes drop-out of the texture, and the strings play by themselves, weight either immediately becomes lighter (as in performance 1), or it starts to fluctuate over a wide range (performance 2). In performance 1, once weight has become lighter, weight gradually becomes heavier with the addition of the piano, and then of the low brass and piano notes - the timbre becomes more 'solid' once again. In performance 2, while weight fluctuates over a wide range, it can be seen to make two moves at becoming lighter. The first rise in the graph of weight (in the above summary shown reaching a measure of 0.4) takes us to the point just prior to where the piano is heard by itself at the very end of the piece. The second rise covers the brushing of the solo piano strings, and the sound of the strings left to ring as the piece comes to an end.

Although the summaries of the graphs of timbral weight shown above look quite different for the two performances, the graphs of timbral weight are, in fact, more similar than at first appears. To understand this, we need to examine figures 4.11b and 4.12b. We need to know that we are comparing the same sections of the music when we compare the graphs. As has been mentioned, performance 2 contains much more of the sound of the piano at the end of the work than does performance 1. Indeed the sound of the piano at the conclusion of the work is very different in the two performances. In performance 1, the sound is very indistinct, it being hard to make out what the sound actually is. In performance 2, the sound is recorded much more closely, without the distractions of the 'live' noises of performance 1. The point where the piano is left to sound on its own is marked on both graphs. Because of the very different natures of these solo piano sections between the performances, it is not meaningful to compare them for information on differences in interpretation. We are merely comparing the very different recording situations, and the inequality in the lengths of these sections.

If we now compare the parts of section [T] up to the point of the solo piano, we find that there are more similarities than at first seen. As has already been observed, both performances begin with a short period of relatively static weight near to 0. Performance 2 then contains many more fluctuations in weight than performance 1. However, the lightest moments of these fluctuations lie very close to the general level of weight in performance 1 for the equivalent section of music. In other words, weight in performance 2 does show similar readings as performance 1 if the upper limits of its fluctuations are read. In fact, if the upper fluctuations of both performances are compared, we find that they both move from around 0.6 at the point where the flutes drop out, down to 0.45 just prior to the piano playing by itself. Why does performance 2 show such rapid and wide fluctuations at this point in comparison to performance 1? Actually, performance 1 also shows some equally wide fluctuations in the section under consideration. The nature of the music in this part of the piece is to show wide variations in weight, because the position of the loudest band is shifting (as can be seen in figures 4.13b and 4.14b). Because performance 2 is recorded more closely than performance 1 (a fact that seems especially evident in this section of the piece), performance 2 shows more intense variation in weight and timbral pitch. The effect is for performance 1 to have a more unified timbre in this section in terms of weight and pitch, whereas performance 2 demonstrates a much more lively and changeable style in weight and pitch.

That completes the detailed examination of timbral width, sharpness, roughness, timbral weight and timbral pitch in *Atmosphères*. The timbral parameter yet to be considered is that of harmonicity - in other words, cepstrum analysis.

Figures 4.21a and 4.22a show cepstrum analyses of the first half of *Atmosphères*. The opening section (from the beginning to letter [A]) shows very little evidence of harmonicity. Performance 2 shows a slightly higher level of harmonicity than performance 1 in this section. This is largely due to a stronger sense of pitch coming through in performance 2 from the flutes, which are recorded more closely than in performance 1. In section [A], both performances show increased levels of harmonicity, due to individual pitches coming to the foreground of the texture as they crescendo. In section [B] we see much more energy represented in the cepstrum graph. At times, this energy shows evidence of harmonicity. Looking first at performance 1, we see clearly the presence of the 'solo' trumpet note, marked by the first asterisk, which has already been discussed in detail. Prior to the trumpet note, there is some evidence of harmonicity due to particular pitches momentarily coming to the foreground. At the second asterisk, we again see evidence of harmonicity, though less than at the first asterisk. This marks the second brass crescendo (on 'black' notes), where, although no 'solo' note is heard, the ringing brass sounds still increase the degree of harmonicity to a marked degree. In performance 2, where no 'solo' trumpet note is heard, we do not see the sudden increase in harmonicity in the first brass crescendo marked by the first asterisk (on 'white' notes), though we still see the increase in harmonicity on the 'black' note crescendo (marked by the second asterisk). This increase in harmonicity in both performances at the 'black' note crescendo is most probably due to the brass playing the relatively concordant 'black-note' pitch collection. The cepstral lines near to the start of section [B] in performance 2 are due to notes in the wind instruments, in particular the flutes, being more closely recorded, and therefore more prominent, than the other instruments playing at this time. In both performances, section [C] shows some degree of harmonicity in its first half. This is due to individual notes briefly emerging from the texture prior to the movement of the pitches becoming much faster and individual pitches becoming submerged into a shimmer of sound. The decrease in the level of harmonicity seen near the beginning of section [C], when compared with that seen in section [B], is continued into the second half of section [C] and into section [D]. A decisive increase in the level of harmonicity is seen in the second half of section [E].

This marks the beginning of the climb upwards in pitch towards the highest point in pitch of the first half of the piece, played by the piccolos, just prior to letter [G]. At this increase in harmonicity, the texture becomes more sparse, and the winds are the most prominent instruments. Where the piccolos are left to play largely by themselves, the many cepstral lines are seen to lie very close together, showing firstly that the sound is harmonic (the pitches lie very close together and there is very little beating - see the graph of roughness), and secondly, that the pitches are very high. For further discussion of this point, see Chapter 2, section 3.5.

Figures 4.21b and 4.22b show cepstrum analyses of the second half of *Atmosphères*. Starting from the beginning of the graphs, the first difference we see between the performances is a strong rahmonic line prior to letter [H] in performance 1 - there is no such line found in performance 2. This rahmonic line reflects the presence of a strong Ab^1 in performance 1 during the section for double basses. The sound of the double basses in performance 2 is much more blended as regards pitch. Moving towards the end of section [H], we see rahmonic lines beginning to form in both performances (the clear set of rahmonics in performance 1 before this is due to a cello pitch, played as an harmonic, suddenly coming to the foreground of the texture). These rahmonic lines just prior to [J] are due to the pitches funnelling towards the three-note cluster B^3 - C^4 - $C\#^4$ - as the pitch cluster narrows, so harmonicity increases. Sections [J], [K] and [L] all show strong harmonic tendencies. The texture in these sections is quite sparse, so that individual notes, or small groups of notes, are easily heard. Section [M] consists of sudden jabs of brass sound. The effect of this is for individual pitches to be clearly heard as they are re-articulated at *ffff*. This playing style, coupled with the 'ringing' nature of the brass sound, gives section [M] a relatively high degree of harmonicity. Section [N] introduces a very diffuse chromatic cluster, and any sign of families of rahmonics disappears. In performance 2, the brass is heard for a longer period of time into section [N] than in performance 1, so the rahmonic lines persist into the beginning of section [N] in performance 2. As was the case at the end of section [H], near the end of section [N] rahmonic families begin to re-appear as the pitch cluster narrows. Sections [P], [Q], [R] and [S] all show very little cepstral activity. They all consist of very diffuse textures - section [P] consists of soft blowing, without tone, for brass; section [Q] consists of a two-and-a-half octave chromatic cluster, where no individual pitches come to the foreground; and sections [R] and [S] consist largely of a softly played cluster for four flutes. Harmonicity increases significantly only once more before the end of the piece - at the beginning of section [T], where, although the pitches consist of an almost complete chromatic cluster of three-and-a-half octaves, the notes are produced as string harmonics, thus making the texture very light and almost transparent, and increasing the degree of harmonicity.

To summarise, cepstrum analysis shows that, during the first half of *Atmosphères*, there are two principal moments of increase in harmonicity - during the middle of section [B], and during section [F], the increase beginning towards the end of section [E]. In other words, harmonicity increases during the middle of the first half of the piece, and at the end of this first half. In the second half of *Atmosphères*, the degree of harmonicity often changes quite suddenly from section to section - for example, there is a high degree of harmonicity in section [M], but low harmonicity in section [N]. It appears that harmonicity helps to delineate the middle and end of this first half. Harmonicity in the second half is used as a method of helping to delineate abrupt changes between sections.

This completes the detailed timbral analysis of *Atmosphères*, which incorporated a comparative timbral analysis of two different performances. In the next section, I will summarise my findings on the timbral structure of this work, draw the reader's attention to the large scale timbral gestures, and discuss how these gestures interrelate with the other structural parameters of the work. Observations on the large-scale timbral gestures will apply equally to both performances, unless otherwise stated.

6. Timbral Gestures and Musical Structure

The musical gestures in the first half of *Atmosphères* (to a little after letter [G]) create directed movement - flowing towards moments of climax, and ebbing to moments of rest. The gestures in the second half of *Atmosphères* (from a little after letter [G] to the end) form oppositional units - the energy of the music is created through these oppositions, rather than through directed motion. The music between letters [G] and [H] is both a climax to the first half and a bridge to the second half of the piece, forming an interlude between these two different types of music. The point of division between the first and second halves is situated between these two letters.

The first half of the piece begins with two largely static pitch clusters, the second pitch cluster appearing at letter [A]. From letter [B], pitch is organised symmetrically around the pitch-class dyad E/Eb, and the height of the pitch clusters contracts around this dyad. Thus, the pitch organisation suggests that a new structural section begins at letter [B]. This description is slightly at odds with the sectional description given in example 4.7, which is reproduced below as example 4.34.

Analysis of Atmosphères

Opening ; [A] [B] [C] ; [D] ; [E] [F] [G] ;
Introduction ; Section 1 ; Interlude ; Section 2 (bridge) ;

[H] [I] ; [J] [K] [L] [M] ; [N] [O] [P] ; [Q] - end
Episode 1 ; Episode 2 ; Episode 3 ; Coda

Example 4.34

Here the introduction consists only of the Opening of the piece to letter [A], rather than to letter [B] as the pitch structure suggests. This structural ambiguity of section [A] is reflected in the timbral domain of the work. We see a low, relatively even level of roughness, sharpness (figures 4.9a and 4.10a), and harmonicity (figures 4.21a and 4.22a) during the opening of the piece up to letter [A]. During section [A], which has been described above as both part of the introduction and part of Section 1, roughness becomes greater, but then moves back down to a level very close to that seen at the end of the Opening. Sharpness moves in a similar way - it rises during section [A], but then either returns to the sharpness level of the Opening (performance 2) or returns to a level one band higher than the Opening at letter [B] (performance 1). Harmonicity increases about half-way through section [A], but then decreases once more at letter [B]. These three timbral measures suggest that section [A] contains an impetus for musical movement, but the movement does not quite 'get off the ground.' It is with section [B] that these three timbral measures do move decisively. Thus, section [A] contains aspects of both introductory gestures, and intimations of the movement to come in Section 1. The remaining timbral parameters also contribute to this ambiguity of section [A]. Timbral width in the Opening is relatively smooth. In section [A], width begins to become much more variable, which remains a characteristic throughout the first half of the piece. The graphs of timbral weight (figures 4.11a and 4.12a) show section [A] as part of a gradual lightening of weight from the opening to letter [B], before the changes from light to heavy weight that occur in section [B]. Thus, width shows section [A] to be different from the Opening, whereas weight connects section [A] to the Opening. Timbral pitch (figures 4.13a and 4.14a) also connects section [A] to the Opening - either by remaining largely in a band number the same as that predominant in the Opening (performance 2), or by first rising and then falling back to the predominant band number of the Opening (performance 1). Thus, the timbral parameters, along with the pitch and sectional structures, very strongly reinforce the ambiguity of section [A], and its role in moving the listener from the relatively static, pregnant quality of the Opening, to the much more fluidly moving quality of Sections 1 and 2.

The overall pitch structure (and Sectional structure) of the first half of the piece can be described as follows: introductory material from the opening to letter [B] (Introduction, and first part of Section 1); large-scale symmetrical movement around pc-dyad E/Eb from letters [B] to [D] (remainder of Section 1); static pitch cluster from letter [D] to letter [E] (Interlude); concluding pitch material to the first half of the piece, which moves rapidly upwards, away from the central E/Eb pc-dyad - letter [E] to a little after letter [G] (Section 2); this material also propels us into the second half of the work.

The timbral measures support the division of the first part of *Atmosphères* in the manner described above. Overall, roughness and sharpness form a sweep of movement from letter [A] to letter [D]. As was discussed above, in section [A] roughness makes a movement upwards, but then falls once more. In section [B], this upwards impetus is allowed to rise fully to the twin peaks of roughness at the moments of the 'white note' and 'black note' collections, dominated by the brass. Roughness then rapidly falls towards letter [C]. In a gesture reminiscent of that seen in section [A], roughness rises once more in section [C] and then falls again, falling to a level close to the level at which roughness began its rise at letter [A]. The shape of the roughness gesture from letters [A] to [D] is reflected in the sweep of the sharpness gesture from letters [A] to [D]. Sharpness begins its ascent at letter [A] at a level the same as that predominant during the Opening (performance 2), or one band higher than the level predominant at the Opening (performance 1). It reaches its highest point thus far half-way through section [B], at the moment of the 'white note' collection, which is also the moment of the first of the twin peaks of roughness. For the remainder of section [B] and for the first half of section [C], the level of sharpness repeatedly drops and then moves back to this high level of sharpness. At the beginning of the second half of section [C] (where the wood-wind and double-basses enter), at the same moment as roughness begins to drop, sharpness descends, reaching a level before the end of this section the same as that at the Opening (performance 1), or one band lower than the Opening (performance 2). Thus, sharpness, from letters [A] to [D] (from the beginning of Section 1 to a point just prior to the beginning of Section 2) can be characterised by a movement upwards, followed by the same (or almost the same) sized movement downwards. The degree of harmonicity can also be characterised in this same way. There is an initial increase in the second half of section [A], and then a large increase during the middle of section [B]. Harmonicity then declines as we approach letter [D]. Thus, roughness, sharpness, and harmonicity all move in a sweep of increase followed by decrease between letters [A] and [D]. The overall movement of timbral width shows a trend of gradually increasing diffuseness from the very beginning of the work to a point half-way through section [C] - the same mid-way point where roughness and sharpness begin their downwards movement. Thus, width unites all sections up until this mid-way point in a gesture of gradually increasing diffuseness. From this mid-point to letter [D], width becomes more focused. Timbral pitch shows a very similar gestural shape as that described for sharpness. Timbral weight moves towards lightness from the Opening to letter [B], and fluctuates between lightness and heaviness between letters [B] and [D].

Timbre contributes to the unity and direction of the music from the opening to letter [D] - unity and direction also seen in pitch organisation and sectional organisation. Section [D] itself is a static Interlude. From letters [E] to [G] (which comprises Section 2), tension grows as pitch moves more rapidly, and becomes higher. Timbral parameters too move more rapidly, and with a strong sense of direction. Roughness rises smoothly until half-way through section [F], where roughness fluctuations become much greater, though still within an overall increasing level of roughness moving towards letter [G]. Sharpness rises steadily until letter [G]. Width, on the other hand, shows an overall movement towards becoming more focused. Thus, while roughness and sharpness both show a gesture of steady increase, width shows a gesture of steady decrease. Harmonicity, while its increase cannot be described as steady over the length of the music from letters [E] to [G], does increase in degree from letter [E] to letter [F], and it then remains high till letter [G]. Weight, while harder to characterise because of its tendency to fluctuate, does show a tendency to move from lighter to heavier, especially in performance 2. Timbral pitch, like sharpness, moves steadily upwards from letters [E] to [G]. Importantly, all these timbral parameters move, overall, in a gesture consisting of a single direction, just as pitch shows an overall tendency to move in a single direction - upwards.

The music either side of letter [G] forms the emotional climax of the work, as well as the turning point for compositional style. At letter [G], pitch suddenly drops, and timbre suddenly changes. Roughness, sharpness and the degree of harmonicity all drop, while width suddenly rises; weight suddenly becomes lighter, and timbral pitch suddenly falls. This suddenness marks the musical way for the second half.

To summarise the timbral structure of the first half of *Atmosphères* (in which I include the beginning of section [G]), we can say that the gestures of roughness, sharpness, harmonicity and timbral pitch all consist broadly of three stages (see example 4.35 for a summary of all major gestures to be found in the piece). A relatively static introduction (the Opening); an increase followed by a corresponding decrease (letters [A] to [D]); and an increase followed by a very sudden fall (from letters [D] to just after [G]). The shape of the gesture for weight is similar. Overall, it consists of a movement towards lightness from the Opening to letter [B]; fluctuation between lightness and heaviness from letters [B] to [E]; and a movement towards heaviness from letters [E] to [G], followed by a sudden increase in lightness just after [G]. Unlike the three-part division of the timbral parameters discussed above, the gesture of width falls into two sections - a rise towards diffuseness from the Opening to a point mid-way between letters [C] and [D], followed by an overall fall towards being more focused, followed by a sudden increase in diffuseness just after letter [G]. These are the timbral, pitch and sectional gestures whose movement and interaction give shape and energy to the first half of *Atmosphères*.

Let us now look at the pitch movement in the second half of the piece. Pitch contracts in on itself during Episode 1 (see sections [H] and [I] in figure 4.1). The last stages of this contraction (letters [J] to [L]) continue into the beginning of Episode 2 (which consists of sections [J] to [M]), but with a change of instrumentation from string to wind instruments. This pitch contraction contrasts to the rapid pitch expansion just before it during Section 2 (letters [E] to [G]), and echoes the large-scale pitch contraction found from letters [B] to [D]. Episode 2 is concerned with the contrast of two pitch areas - the pitch movement between letters [J] and [M], and the pitch movement between letters [M] and [N]. This Episode, consisting of relatively small pitch clusters, acts as a 'buffer' between the large cluster at [H] (which was the beginning of the first Episode) and the very large cluster at [N] (the beginning of the third Episode). Episode 3 is about the dissolution of this cluster into noise (bowing noise from the strings moving on to soft, toneless blowing by the brass). The Coda (letter [Q] to the end) opens out the pitch-space once again, re-asserts the p-c dyad which was central to the first half of the work (E/Eb), and allows the energy of the work to dissipate into low, rumbling pitches.

The steady roughness rise during Episode 1 (letters [H] to [J]) is one of a whole series of steady rises in roughness that occur during the piece. Its immediate predecessor is the rise that moves from letter [D] to letter [G]. That rise was followed by a sudden fall. The rise in Episode 1 is also followed by a sudden fall at letter [J]. This fall is followed by yet another rise, as the pitch cluster expands, and there is a subsequent plateau of roughness during section [M] (though with much fluctuation). This rise and plateau comprise the gesture of roughness during Episode 2. The rise and plateau is the last large movement in roughness of the work. The drop in roughness following the plateau, which occurs at the beginning of Episode 3 (along with a large increase in the size of the pitch cluster) is not as sudden as the previous drops. The movement of roughness is becoming less energetic. Indeed, the drop after letter [N] is followed by a gentle downward gradient of roughness as pitches progressively drop-out from the pitch cluster, and the drop in roughness reaches the end of its movement just prior to letter [P]. The end of Episode 3 (section [P]) overlaps with a new movement in roughness - a rise (caused by the introduction of toneless noise), which continues very gradually (after a drop in performance 1) during most of the Coda as the pitch-space opens out once more, until the drop in roughness that accompanies the gradual dying away of sound at the conclusion of the piece. This very gentle rise and fall is a distant echo of the much larger rises and falls in roughness that have occurred through this piece.

We have seen three types of sharpness gesture in the first part of the piece: a static section, a section which progressively rises and then falls, and a section which progressively rises and then suddenly falls. The second part of the piece introduces two more types of sharpness gesture. The first, which begins after the fall at letter [G], comprises a rise in two stages, with a plateau after each stage. There is a rise from letter [H] to half way through section [H], followed by a plateau of sharpness to letter [J]. This is Episode 1. There is then a small drop in sharpness (at letter [J]), in preparation for the next rise in sharpness that carries through to letter [M]. From letters [M] to [N] there is another plateau in sharpness. This second rise and subsequent plateau is Episode 2. Thus, this first new sharpness gesture carries across Episodes 1 and 2. The second new sharpness gesture consists of sharpness moving below and above a central point of sharpness. This is how I characterise the sharpness during Episode 3 and the Coda. This gives five different types of sharpness gesture during the course of *Atmosphères* - five sharpness gestures which, because of their long durations in the second half of the piece, help to unite different sections of the work, and which show us five complementary methods of manipulating the timbral parameter of sharpness.

Width falls into four gestures when the entire piece is considered, the beginnings and endings of these gestures not coinciding (except around letter [G]) with the beginnings and endings of any other timbral gestures. Thus, we begin to see evidence of the complex web of gestures that form this work. As was said in the discussion of the first half of the piece, width becomes more diffuse from the Opening to half-way through section [C], and then becomes more focused from this point to letter [G]. Just after letter [G], width suddenly becomes more diffuse again. All this I regard as the first gesture. From after letter [G] to letter [M] there is a rise followed by a fall in width. This is the second gesture - it comprises the end of section [G], Episode 1 and all but the final section (section [M]) of Episode 2. From letter [M] to letter [T] there is another gesture of a rise followed by a fall, which comprises the third gesture of width. This third gesture comprises the last section of Episode 2, Episode 3 and half of the Coda. The fourth gesture consists only of a rise in width, stretching from letter [T] to the end of the piece (the final fall on the graph is due to the measurement of background noise). This gesture comprises the second half of the Coda. Thus, we see that the gestures of width are not coincidental with the other gestures. The gesture of width contributes to the 'seamless' quality of *Atmosphères* .

Because the degree of harmonicity that is present is shown only in a general way in the cepstrum analyses, I shall not attempt to define overall gestures of harmonicity. Nor shall I attempt to define overall gestures of weight, as the measure of weight fluctuates too widely and differs too greatly between the performances.

The gestures of timbral pitch are, as would be expected, similar to those of sharpness. I find five gestures of timbral pitch. The first is a gesture where timbral pitch is relatively static (the Opening); the second is a gesture of timbral pitch rising and then falling (letter [A] to letter [D]); the third is a gesture consisting of a steep rise, followed by a sudden fall (letter [D] to a little after letter [G]). The fourth gesture, which is the first of the second half of the piece, demonstrates yet a different way in which sharpness can rise. It consists of a rise at letter [H], and then a period of relative stasis; and then another rise at letter [M], and then another period of relative stasis (until letter [N] in performance 2, but only until half-way through section [M] in performance 1). The fifth timbral pitch gesture is different yet again from the previous gestures - it consists firstly of sudden movements in timbral pitch, oscillating over a distance of approximately 10 bands, followed by a fall in timbral pitch - from letter [N] to the end. This fall is more prolonged in performance 2 than in performance 1.

The overlapping gestures are shown below in example 4.35. Breaks in the lines indicate the ending of one gesture and the beginning of another. All sections break during section [G].

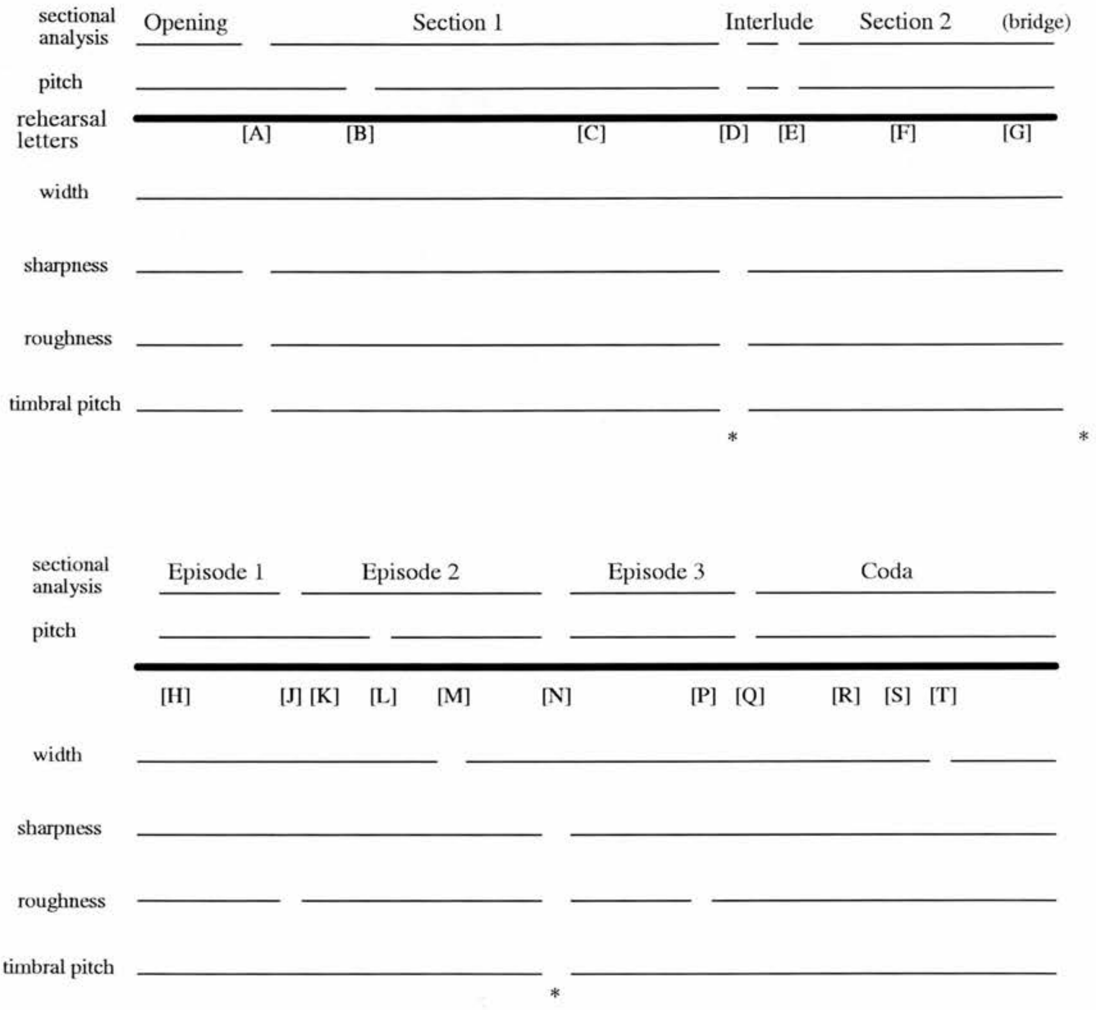

Example 4.35

From this diagram, we can see that it is possible to divide *Atmosphères* into large units based upon the number of breaks in gestures that occur at the same moment. The greater the number of simultaneous breaks, the stronger the segmentation. The strongest break (consisting of six simultaneous breaks) occurs between letters [G] and

[H]. - the mid-point of the work. Points consisting of five simultaneous breaks occur between Section 1 and the Interlude, and between Episode 2 and Episode 3. That is, the next-strongest breaks occur at the sectional mid-points of the first and second halves of the work. Thus, the breaks in gestures divide the work first into half, and then divide each of these halves in half again. These points are marked by an asterisk in the above example.

This concludes the analysis of *Atmosphères*.

***

CHAPTER 5

Conclusions

CHAPTER 5

Conclusions

---

1. Timbre and Analytical Tools

Jeux vénitiens and *Atmosphères* are two outstanding examples of 20th-century music in which timbre is elevated to a high level of structural significance. Timbre is utterly intrinsic to the identity of these two pieces. One needs only to consider the effect of playing either of the pieces in piano reduction - their identity would be lost. Yet, while we can hear that the movement of the music 'makes sense', we have not, until now, been able to answer *why* the music makes sense.

In order to 'see into' the structure of these works, we need analytical tools that are able to extract information on timbre in a way that is sensitive to the disciplines of both psychoacoustics and music analysis. The first attempt to analyse a piece acoustically, broadly within the discipline of music analysis, was that by Robert Cogan (1984) (see Chapter 1, section 3.2). Cogan looked at the pattern of the spectrographic image, and, through his own system of oppositional dyads, described what he saw. While he used the term 'tone color', he did not attempt to define it - it was viewed as an undifferentiated aspect of 'sonic design.' Cogan made no attempt to interpret the spectrographic images in the light of perceptual models. He treated the spectrographs as objects in themselves, rather than as representations of data that can be further processed.

Loudness-Distribution Analysis (width, weight and timbral pitch), and measures of Roughness, Sharpness and Harmonicity take the data as represented in a spectrograph, weight it according to perceptual models, and then reduce it so that just one perceptual attribute is represented. This takes us much further than Cogan's study, for by weighting and reducing the data, we can identify and plot specific timbral attributes in a piece.

In this thesis information on timbre has been represented in a series of graphs. In the interpretation of these graphs, we have investigated the 'gestures' of the graph. These represent 'gestures' of perception. It is the gestures of.the timbral graphs of the A- and B-textures of *Jeux vénitiens* that demonstrate the perceptual timbral gestures that create a progression from one A-texture and from one B-texture to the next. In *Atmosphères*, the timbral gestures form an interlocking net of gestures that both help to give a sense of form to the piece, and create a feeling of constant metamorphosis. The gestures of the timbral graphs form the basis for the uncovering of timbral structures that help to move a piece through time. Indeed, the notion of a timbral trajectory or gesture is not too far from Cogan's theory of sonic design, except that here we are dealing with specific perceptual aspects, whereas Cogan was dealing with spectrographic pattern. Both Cogan and I are concerned with how sound is 'sculpted' through time.

2. The Methodology of Timbre Analysis

Having established a set of analytical tools, a methodology is needed whereby the tools can be applied in a way that will maximise their music analytical worth.

In timbre analysis, we are concerned with an aspect of the piece for which there exists no notation. A notation is made (the timbral plots) using the timbre analysis tools. The first stage of the analytical methodology consists of a broad survey of the shape of the timbral plots of performances. We might see broad areas of contrasting roughness (as in *Jeux vénitiens*), or an overall shape to the progression of timbral width (as in the first half of *Atmosphères*). When all the timbral plots have been considered, a 'macro-survey' may suggest an overall analytical strategy, that will either be validated or need to be reconsidered when the strategy is tried (for example, examining the A- and B-textures of *Jeux vénitiens* separately, which was validated by the discovery of timbral progressions across these textures). The second stage of analysis consists of micro-analysis - the matching of the individual timbral events of the graph with the individual musical events of the performance. For example, the positions of the individual 'exclamations' of the B-textures in *Jeux vénitiens* were carefully timed and then plotted onto the timbre graphs. This task is vital for the interpretation of the timbre graphs, and, in this thesis, for the further investigation and subsequent validation that the analytical tools are measuring the particular perceptual attributes that is claimed for them. The micro-analysis, which forms the bulk of the analytical work, is performed for two different recordings. The findings of the micro-analyses make possible a detailed comparison of the timbre graphs of the different recordings. This comparison is essential, for it allows a decision to be made as to when a particular timbral event belongs to one realisation of the piece, and when it belongs to both of the recordings. If, when considering a section of the music, the timbre graphs of the two recordings largely agree with one another (timbre graphs of two different performances will never precisely agree, because two performances are never exactly the same), then the shape of the graphs is considered not only to be showing an invariant quality between two performances, but also an aspect of the timbre design of the piece of music. If the relevant sections of the timbre graphs are largely inconsistent with one another, the audible reasons why the graphs are different are located, and a decision made as to which realisation of the piece is closer to the analyst's notion of 1) a 'realistic' performance (this relates to the properties of the recording); and 2) the interpretation that makes the most musical sense (based on the analyst's knowledge of the piece and its genre). An example of the process of making the decision regarding (1) was seen in the analysis of *Jeux vénitiens* where the properties of the different recording styles of the two performances had to be considered, and a decision reached as to which was closer to a 'realistic' hearing of the piece. Depending on the timbral attribute being considered, one or other of the recordings was judged to be closer to 'reality.' An example of dealing with (2) was seen in the analysis of *Atmosphères*, where a decision was reached that the timbre change due to a particular pitch coming to the foreground was not part of the timbre design (Ligeti specifically requests that individual notes should not be heard to emerge from the 'sound mass').

There are many different ways in which a piece of music can be performed, and we have seen that timbre analysis is a powerful method whereby similarities and differences in interpretation and recording styles can be illuminated. And following the investigation of the causes of these similarities and differences in timbre data from two different performances, we are in a position to decide on a representative graph for a particular aspect of timbre for the piece. Is the comparison of two performances sufficient in order to come to an understanding of a representative timbral plot? While it is not argued that the analysis of a greater number of different recordings might not give further insight into the timbral structure of the piece, it is clear from the analyses presented here that, although two performances might be quite different in many respects, invariants of timbre exist between them. This discovery is very important, for it speaks of the robust nature of timbre structures to withstand the vagaries of different interpretations and of different recording styles, and it supports the assertion that the timbre structures presented here are intrinsic to the piece, and not simply a product of two particular recordings.

More than two recordings could, of course, be used in the timbre analysis of a work. Two recordings were used here because of limits on the number of CD recordings that were available for the pieces that were analysed, and because of constraints of time.

Once the comparison of the two recordings is concluded, the analysis moves on to the third stage: arriving at an overall timbral design for the piece. It is the larger-scale timbral movements and shapes that are of interest, for we are now looking for timbral trajectories and gestures. These trajectories and gestures comprise the timbral design of the piece of music. The timbral trajectories and gestures inter-relate with one another to create overall timbral contrast and progression. The last step in the analytical process is to consider how the movement in timbre inter-relates with pitch and sectional structures, and thus to examine the overall musical design of the work.

The steps in the timbre analytical process are illustrated below:

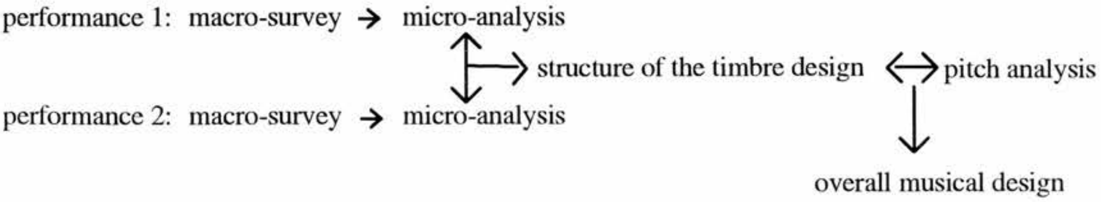

3. The Future

This thesis demonstrates that the successful understanding of timbre in music relies on the bringing together of the disciplines of psychoacoustics and music analysis. Without the discipline of psychoacoustics, the analysis of timbre remains either vague and anecdotal, because of the lack of an effective language, or, if physical analysis is employed, 'scientific' but unreliable, due to the failure to extract the relevant data and to weight the data so that it accords more strongly with auditory perception. If the discipline of music analysis is not rigorously employed, then timbre analysis can occur in isolation from other well-established music analytic methods. Our understanding of how timbre functions in a piece of music is greatly enhanced if we are aware of the pitch and sectional 'structural environments' in which the timbre progressions occur. We thus gain insight into the entire musical structure.

3.1 Aspects for improvement

While the tools and methods presented here appear to have been effective, there are, of course, aspects that could be improved.

During the course of this study, no perceptual tests were carried out to establish the effectiveness of the analytical tools. The measures of roughness, sharpness and harmonicity have been investigated by a number of researchers, though the measures have not been used in the field of the analysis of ensemble music before. Loudness-Distribution analysis, while derived from a well-tested technique (the Tristimulus

Method), has never been subjected to psychoacoustic testing. Thus, while the analytical tools are grounded in established psychoacoustic techniques, their effectiveness in the type of auditory situations in which they have been used here has not been psychoacoustically tested. The decision not to move into the area of psychoacoustic testing was made at the very start of this research - the primary thrust of the thesis is a music analytic one, and to move into psychoacoustic testing would have been another piece of research all of its own. Instead of formal psychoacoustic testing, the analytical tools have been demonstrated on simple examples (Chapter 2), and their measures of *Jeux vénitiens* and *Atmosphères* have been shown to be consistent with the 'ordinary' experience of such events as pizzicato, the playing of a note as an harmonic, and the many other small-scale musical situations that were described. While it is hoped that, through the examples given, and through the detailed description during the course of the analyses, the reader will be satisfied that the analytical tools are measuring valid timbral parameters, there is no doubt that methodical psychoacoustic testing of the analytical tools would assist in their validation, as well as possibly showing ways in which the tools could be improved.

In Chapter 1, it was decided not to consider separately the perceptual consequences of the time varying aspects of a sound - in particular, the starting transient. It was made clear that it was the premise of this thesis that most of the information on timbre in a musical setting, particularly an ensemble musical setting, is carried by the steady-state portion of sounds. Depending on the *type* of musical setting, however, more or less timbre information will be held in the time-varying aspect of the sound - particularly in the starting transient. Thus, a fuller understanding of music timbre would come with an understanding of how our perception of the time-varying aspect of sound is affected by musical context. More research needs to be carried out to determine this, as well as to discover if the starting transient affects our judgement of the subsequent roughness, sharpness, and so forth of the steady state.

Timbre and musical context is also an issue for this study in that it is quite possible for a listener to distinguish both an overall timbre during a section of music and a number of different timbres. For example, we can experience an overall timbre 'shade' during a *tutti* section of a symphony, and we can also distinguish the different timbres of the flutes, strings, percussion, and so forth. The timbre graphs only give a measure that reflects an *overall* timbral impression - they cannot account for the existence of different but simultaneous timbral strands, nor can they account for the way in which we can shift our attention between different timbres. The type of music selected for analysis minimised this problem, for in these pieces it is strongly evident that the object of both Lutoslawski and Ligeti is to work with overall timbral shades, and to work very little with simultaneous timbral strands. However, in music where simultaneous timbral strands did appear to be more structurally important (for example, in most western tonal music) it would be very useful if it were possible to seek-out computationally different timbral strands in a recording, and then view their affect both separately and in combination. This would be a major improvement in the perceptual representation of timbre. A way forward in this task might be to consider the role of phase in the differentiation of separate but simultaneous notes, and to incorporate in the computer programme the known spectral patterns of particular instruments at particular registers so that they might be recognised, or at least differentiated.

Lastly, improvements could also be made in LD-analysis and in the measurement of harmonicity. We saw in the measurement of timbral pitch and weight that sometimes the degree of movement in the graph was not proportional to the perceived degree of change in that aspect of timbre (see especially Section 3.2.1, Chapter 2). A damping technique was tried (so that the position of the loudest band would change on the graph only if the alternative band was a certain degree louder than the current band), but it was not found to be effective. An effective damping method would require specific perceptual research into the minimum degree of change that is necessary for a physical change in the position of the loudest band to be perceived.

Cepstrum analysis, while it was found to be a useful method for determining the degree of harmonicity, was limited by the absence of an actual measure for the amount of harmonicity present. A successful measure might be found in the signal to noise ratio of the cepstrum measurements.

3.2 Long-term development

At the opening of this thesis, timbre was defined as consisting of the primary aural information that is used in the perceptual task of assigning an identity to a sound. This aural information has been represented in a series of graphs. These graphs form the analytical timbre vocabulary of this study. Having established a language of analysis, a possible development would be to use these results as the basis for a language of synthesis.

The combined timbral graphs, seen at the conclusion of the analysis of *Jeux vénitiens* (figures 3.41 - 3.44), and near the beginning of the analysis of *Atmosphères* (figures 4.9 and 4.10) were likened to a timbral 'score'. It is possible to read these graphs, and hear the change in timbral quality, independently of imagining changes in pitch or loudness. While there are possibly further acoustical parameters to be discovered that help to generate a particular timbral quality, if it is possible to 'read' a score of timbre, it should be possible to re-synthesise from this same information: a composer draws a line on a computer screen, and the timbral progression is played back. Music analysis and composition have a rich history of interaction. We have seen how useful for analysis the beginnings of a timbre notation can be - we have seen that composers can organise the sound-world of timbre with great care and delicacy. In the future, composers may find in this type of analysis the inspiration to explore further the terrain of timbral structures.

****

Appendix

1/3 Octave Bands

Band no.	Ctr. frq.	Pitch (approx.)	Bandwidth (approx.)
17	50	G^1	11.6
18	63	B^1	14.6
19	80	$D^{\#2}$	18.3
20	100	G^2	23
21	125	B^2	29
22	160	$D^{\#3}$	37
23	200	G^3	46
24	250	B^3	58
25	315	$D^{\#4}$	73
26	400	G^4	92
27	500	B^4	116
28	630	$D^{\#5}$	145
29	800	G^5	183
30	1000	B^5	230
31	1250	$D^{\#6}$	290
32	1600	G^6	370
33	2000	B^6	460
34	2500	$D^{\#7}$	580
35	3150	G^7	730
36	4000	B^7	920
37	5000	$D^{\#8}$	1160
38	6300	G^8	1450
39	8000	B^8	1830
40	10000	$D^{\#9}$	2300
41	12500	G^9	2900

Bibliography

BIBLIOGRAPHY

American Standards Association: (1960). *Acoustical Terminology* SI, 1-1960. New York: American Standards Association.

Balzano, Gerald: (1986). "What are Musical Pitch and Timbre?" *Music Perception* 3/3: 297-314.

Belkin, Alan: (1988). "Orchestration, Perception, and Musical Time: A Composer's View." *Computer Music Journal* 12/2: 47-53.

Bent, Ian: (1987). *The New Grove Handbooks in Music: Analysis*. London: Macmillan Press.

Bernard, Jonathan: (1987). "Inaudible Structures, Audible Music: Ligeti's Problem, and His Solution.." *Music Analysis* 6/3: 207 - 236.

Bismarck, G. von: (1974a). "Timbre of Steady Sounds: A Factorial Investigation of its Verbal Attributes." *Acustica* 30: 146-159.

----: (1974b). "Sharpness as an Attribute of the Timbre of Steady Sounds." *Acustica* 30: 159-172.

Boulez, Pierre: (1987). "Timbre and Composition - Timbre and Language." *Contemporary Music Review* 2: 161-71.

Björk, E.: (1985). "The Perceived Quality of Natural Tones." *Acustica* 57: 185-88.

Bruijn, A. de: (1978). "Timbre Classification of Complex Tones." *Acustica* 40: 108-114.Brunner, C. "The Perception of Contemporary Pitch Structures." *Music Perception* 2/1 (1984): 25-40.

Brunner, Lance: (1985). Review of Cogan (1984). *Notes* 42: 283-86.

Burkhart, Charles: (1974). "Schoenberg's *Farben*: An Analysis of Op.16, No.3." *Perspectives of New Music* 12: 141-72.

Campbell, Murray, and Greated, Clive: (1987). *The Musician's Guide to Acoustics*. London: J.M. Dent & Sons.

Charbonneau, Gérard: (1981). "Timbre and the Perceptual Effects of Three Types of Data Reduction." *Computer Music Journal* 5/2: 10-19.

Chou Wen-Chung: (1979). "*Ionisation*: The Function of Timbre in its Formal and Temporal organisation." In *The New Worlds of Edgard Varèse: A Symposium*. Edited by Sherman von Solkema. Brooklyn: Institute for Studies in American Music.

Cogan, Robert: (1969). "Towards a Theory of Timbre: Verbal Timbre and Musical Line in Purcell, Sessions and Stravinsky." *Perspectives of New Music* 8/1: 75-81.

----: (1975). "Reconceiving Theory: The Analysis of Tone Color." *College Music Symposium* 15: 52-69.

----: (1980). "Tone Color: The New Understanding." *Sonus* 1/1: 3-24.

----: (1984). *New Images of Musical Sound.* Cambridge, Mass.: Harvard University Press.

REVIEWS:

Brunner, L.: (1985). *Notes* 42: 283-86.

Strawn, J.: (1986a). *Computer Music Journal* 10/1: 97-9.

----: (1986b). *Computer Music Journal* 10/4: 9-13 (includes reply from Cogan).

Swift, R.: (1986). *Journal of Music Theory* 30: 277-83.

Whittall, A.: (1986). *Music & Letters* 67: 187-88.

----: (1987 & 1988). "Science, Music Values and the Study of Music: Mozart, Tibetan Chant, Boulez (Parts 1 & 2). *Sonus* 7/2: 10-28, and 8/2: 16-36.

----: (1990). "Hildegard's Fractal Antiphon." *Sonus* 11/1: 1-19.

----: (1991). "Varèse: An Oppositional Sonic Poets." *Sonus* 11/2: 26-35.

Cogan, Robert, and Escot, Pozzi: (1976). *Sonic Design - The Nature of Sound and Music*. New Jersey: Prentice-Hall.

REVIEWS

Geissler, F.: (1977). *Notes* 34: 330-34.

Thomson, W.: (1979). *Journal of Music Theory* 23: 125-34.

Danner, Gregory. "The Use of Acoustic Measures of Dissonance to Characterise Pitch-Class Sets. *Music Perception* 3/1 (1985): 103-122.

Erickson, Robert: (1975). *Sound Structure in Music*. California: University of California Press.

REVIEWS

Rahn, J.: (1976). *Notes* 32: 754-5.

Slawson, W.: (1978). *Journal of Music Theory* 22: 105-09.

Swift, R.: (1975). *Perspectives of New Music* 14: 148-58.

Whittall, A.: (1976). *Music and Letters*. 57: 73-4.

Fletcher, Harvey: (1940). "Auditory Patterns." *Review of Modern Physics* 12: 47-65.

Floros, Constantin: (1996). *Gyorgy Ligeti: jenseits von Avantgarde und Postmoderne.* Vienna: Wien Osterreiches Musikzeit Edition.

Francois, Jean-Charles: (1991). "Organisation of Scattered Timbral Qualities: A Look at Edgard Varese's *Ionisation*." *Perspectives of New Music* 29/1:

Fry, Laura: (1984). "Stockhausen's *You are a Spirit*: Musical-Mystical Unities." *Sonus* 4/1: 46-57.

Gebura, Alice: (1980). "Spectral Shape as a Reflection of the Text in Stockhausen's *In the Sky* from Indianlieder. *Sonus* 1/1: 48-58.

Geissler, Fred: (1977). Review of Cogan & Escot (1974). *Notes* 34: 330-34.

Giesse, Gregory: (1984). "Timbre Dimensions for Musical Control." Ph.D. dissertation, University of Minnesota.

Goad, Pamela: (1992). *Loudness and Sharpness Measurements Relevant to Musical Instrument Timbres.* Technical Report Series no. 9202. Systematic Musicology Programme, University of Washington.

Grey, John: (1975). "An Exploration of Musical Timbre using computer based techniques for analysis, synthesis and perceptual scaling." Ph.D. dissertation, Stanford University.

----: (1977). "Multidimensional Perceptual Scaling of Musical Timbres." *Journal of the Acoustical Society of America* 61/5: 1270-77.

----: (1978). "Timbre Discrimination in Musical Pattern." *Journal of the Acoustical Society of America* 64/2: 467-72.

Grey, John, and Gordon, J.M.: (1978). "Perception of Spectral Modifications on Orchestral Instrument Tones." *Computer Music Journal* 2/1: 24-31.

Grey, John, and Moorer, James: (1977). "Perceptual Evaluations of Synthesised Musical Instrument Tones." *Journal of the Acoustical Society of America* 62/2: 454-62.

Gunden, H.C. von: (1977). "Timbre as Symbol in Selected Works of Olivier Messiaen." Ph.D. dissertation, University of California

Hatmaker, J.E.: (1985). "A Theory of Timbre in the Late Classical Symphony." Ph.D dissertation, University of Iowa.

Helmholtz, Herman von: (1863). *Die Lehre von dem Tonempfindung als physiologische Grundlage für die Theorie der Musik.* Braunschweig: F. Vieweg & Sohn. Translated by A.J. Ellis (1954) as *On the Sensations of Tone as a Physiological Basis for the Theory of Music*. New York: Dover Publications.

Huron, D. "Tonal Consonance versus Tonal Fusion in Polyphonic Sonorities." *Music Perception* 9/2 (1991): 135-154.

Hutchinson, W, and Knopoff, L.: (1978). "The Acoustic Component of Western Consonance." *Interface* 7: 1

----: (1979). "The Acoustic Measurement of Consonance and Dissonance for Common Practice Triads." *The Journal of Musicological Research* 3: 6

Jaffe, David: (1987). "Spectrum Analysis Tutorial, Parts 1 & 2." *Computer Music Journal* 11/2: 9-24, and 11/3: 17-35.

Kameoka, Akio and Kuriyagawa, Manoru: (1969a). "Consonance Theory Part I: Consonance of Dyads." *The Journal of the Acoustical Society of America* 45/6: 1451-1459.

Kameoka, Akio and Kuriyagawa, Manoru: (1969b). "Consonance Theory Part II: Consonance of Complex Tones and Its Calculation Method." *The Journal of the Acoustical Society of America* 45/6: 1460-1469.

Kendall, Roger: (1986). "The Role of Acoustic Signal Partitions in Listener Categorisation of Musical Phrases." *Music Perception* 4: 185-214.

Lerdahl, Fred: (1987). "Timbral Hierarchies." *Contemporary Music Review* 2: 135-60.

Lichte, W.H.: (1941). "Attributes of Complex Tones." *Journal of Experimental Psychology* 49: 455.

Ligeti, György: (1960). German text, trans. 1965 as "Metamorphoses of Musical Form." *Die Reihe* 7: 5-19.

----: (1983). *György Ligeti in Conversation.* With various interviewers. London: Eulenburg Books, 1983

McAdams, Stephen: (1987). "Music: A Science of the Mind?" *Contemporary Music Review* 2: 1-61.

McAdams, Stephen, and Bregman, A.: (1979). "Hearing Musical Streams." *Computer Music Journal* 3: 26-43.

McGee, William: (1982). "An Expanded Concept of Timbre and its Structural Significance, with a timbral analysis of George Crumb's *Night of the Four Moons*." Ph.D dissertation, University of Arizona.

Moore, Brian: (1989). *An Introduction to the Psycholohy of Hearing.* 3rd ed. London: Academic Press.

Nowalis, Susan M.: (1975). "Timbre as a Structural Device in Berlioz' Symphonies." Ph.D. dissertation, Cape Western Reserve University.

Padgham, C.: (1986). "The Scaling of the Timbre of the Pipe Organ." *Acustica* 60: 189-204.

Parks, Raymond: (1988). "Spectral Analysis." Notes privately supplied.

Pellman, Samuel: (1979). "An Examination of the Role of Timbre in a Musical Composition, as exemplified by an analysis of *Sequenza V* by Luciano Berio." Ph.D dissertation, Cornell University.

Pierce, John: (1983). *The Science of Musical Sound.* New York: Scientific American Books.

Plomp, Reiner: (1976). *Aspects of Tone Sensation.* New York: Academic Press.

Plomp, R, and de Laat, J.A.P.M.: (1984). Research Notes - "Comparison of Organs is a Spectrum Space." *Acustica* 55: 193-194.

Plomp, R, and Levelt, J.M.: (1965). "Tonal Consonance and Critical Bandwidth." *Journal of the Acoustical Society of America* 38: 548-560.

Plomp, Reiner, and Steeneken, H.J.M.: (1973). "Place Dependence of Timbre in Reverberent Sound Fields." *Acustica* 28: 50-59.

Pollard, Howard: (1988). "Feature Analysis of Musical Sound." *Acustica* 65: 232-244.

----: (1990a). "Timbre Assessment." *Acoustics Australia* 18/1: 19-23.

----: (1990b). ""Timbre Measurement." *Acoustics Australia* 18/3: 65-69.

Pollard, H., and Jansson, E.V.: (1982a). "A Tristimulus Method for the Specification of Musical Timbre." *Acustica* 51: 162-71.

----: (1982b). "Analysis and Assessment of Musical Starting Transients." *Acustica* 51: 249-262.

Potter, R., and Kopp, G., and Kopp, H.: (1947). *Visible Speech.* N.p.: Dover Publications.

Preis, A.: (1984). "An Attempt to Describe the Parameter Determining the Timbre of Steady State Harmonic Complex Tones." *Acustica* 55: 1-13.

Rae, Charles Bodman: (1994). *The Music of Lutoslawski.* London: Faber and Faber Limited.

Rahn, John: (1976). Review of Erickson (1975). *Notes* 32: 754-55.

Randall, J.K.: (1965). "A Report from Princeton." *Perspectives of New Music* 3/2: 84-92.

----: (1967). "Three Lectures to Scientists." *Perspectives of New Music* 5/2: 124-40.

Rigden, John: (1985). *Physics and the Sound of Music.* 2nd ed. New York: John Wiley & Sons.

Risset, Jean-Claude: (1985). "Computer Music Experiments 1964-...". *Computer Music Journal* 9/1: 11-18.

Risset, Jean-Claude, and Wessel, David: (1982). "Exploration of Timbre by Analysis and Synthesis." In *Psychology of Music*, pp.26-58. Edited by Diana Deutsch. New York: Academic Press.

Roederer, Juan G.: (1973). *Introduction to the Physics and Phychophysics of Music.* London: The English Universities Press Ltd.

Schaeffer, Pierre: (1968). *Traité des objets musicaux.* Paris: Editions du Seuil.

Scharf, B.: (1981). "Loudness Adaptation." In *Hearing Research and Theory.* Edited by J.V Tobias and E.D. Schubert. New York: Academic Press.

Schoenberg, Arnold: (1978). *Theory of Harmony*. Translated by Roy Carter. London: Faber and Faber. Originally published 1911.

Searle, Campbell: (1982). "Some Perceptual Aspects of Timbre." *Canadian University Music Review* 3: 80-101.

Seashore, Carl: (1938). *Psychology of Music*. New York: McGraw Hill.

Sethares, Willaim A.: (1993). "Local Consonance and the Relationship between Timbre and Scale." *Journal of the Acoustical Society of America* 94/3: 1218-1228.

Sethares, Willaim A.: (1994). "Adaptive Tunings for Musical Scales." *Journal of the Acoustical Society of America* 96/1: 10-18.

Slawson, A. Wayne: (1968). "Vowel Quality and Musical Timbre as Functions of Spectrum Envelope and Fundamental Frequency." *Journal of the Acoustical Society of America* 43/1: 87-101.

----: (1978). Review of Erickson (1975). *Journal of Music Theory* 22: 105-09.

----: (1981). "The Color of Sound: A Theoretical Study in Musical Timbre." *Music Theory Spectrum* 3: 132-41.

----: (1982). "The Musical Control of Sound Color." *Canadian University Music Review* 3: 67-79.

----: (1985). *Sound Color*. California: University of California Press.

----: (1987). "Sound-Color Dynamics." *Perspectives of New Music* 25: 156-181.

Stevens, S.S.: (1972). "Perceived Level of Noise by Mark VII and Decibels (*E*)." *Journal of the Acoustical Society of America* 51/2: 575-601.

Strawn, John: (1986a). Review of Cogan (1984). *Computer Music Journal* 10/1: 97-9.

----: (1986b). Review of Cogan (1984). *Computer Music Journal* 10/4: 9-13 (includes reply from Cogan).

Strong, William, and Clark, Melvile: (1967). "Perturbations of Synthetic Orchestral Wind-Instrument Tones." *Journal of the Acoustical Society of America* 41/2: 277-285.

Swift, Richard: (1975). Review of Erickson (1975). *Perspectives of New Music* 14: 148-58.

----: (1986). Review of Cogan (1984). *Journal of Music Theory* 30: 277-83.

Taylor, C.A.: (1980). "Sound." In *The New Grove Dictionary of Music and Musicians*, pp. 545-63. Edited by Stanley Sadie. London: Macmillan and Company.

Terhardt, E.: (1974). "On the Perception of Periodic Sound Fluctuations (Roughness)." *Acustica* 30: 201-213.

Thomson, Wiliam: (1979). Review of Cogan & & Escot (1976). *Journal of Music Theory* 23: 125-34.

Trubetzkoy, N.S.: (1969). *Principles of Phonology*. Translated by C. Baltaxe. Berkely: Unversity of California Press.

Warner, David: (1983). "Notes from the Timbre Space." *Perspectives of New Music* 21: 15-22.

Wessel, David: (1979). "Timbre Space as a Musical Control Space." *Computer Music Journal* 3: 45-52.

Whittall, Arnold: (1976). Review of Erickson (1975). *Music and Letters* 57: 73-4.

----: (1986). Review of Cogan (1984). *Music and Letters* 67: 187-88.

Yee On Lo: (1987). "Toward a Theory of Timbre." Ph.D. dissertation, Stanford University. Dept. of Music Report no. STAN-M-42.

Yilmaz, H.: (1967 & 1968). "A Theory of Speech Perception, 1 and 2." *Bulletin of Mathematical Biophysics* 29/4: 793-824, and 30/3: 455-79.

******

Figures for

chapter 2

Tools for Exploring Timbre

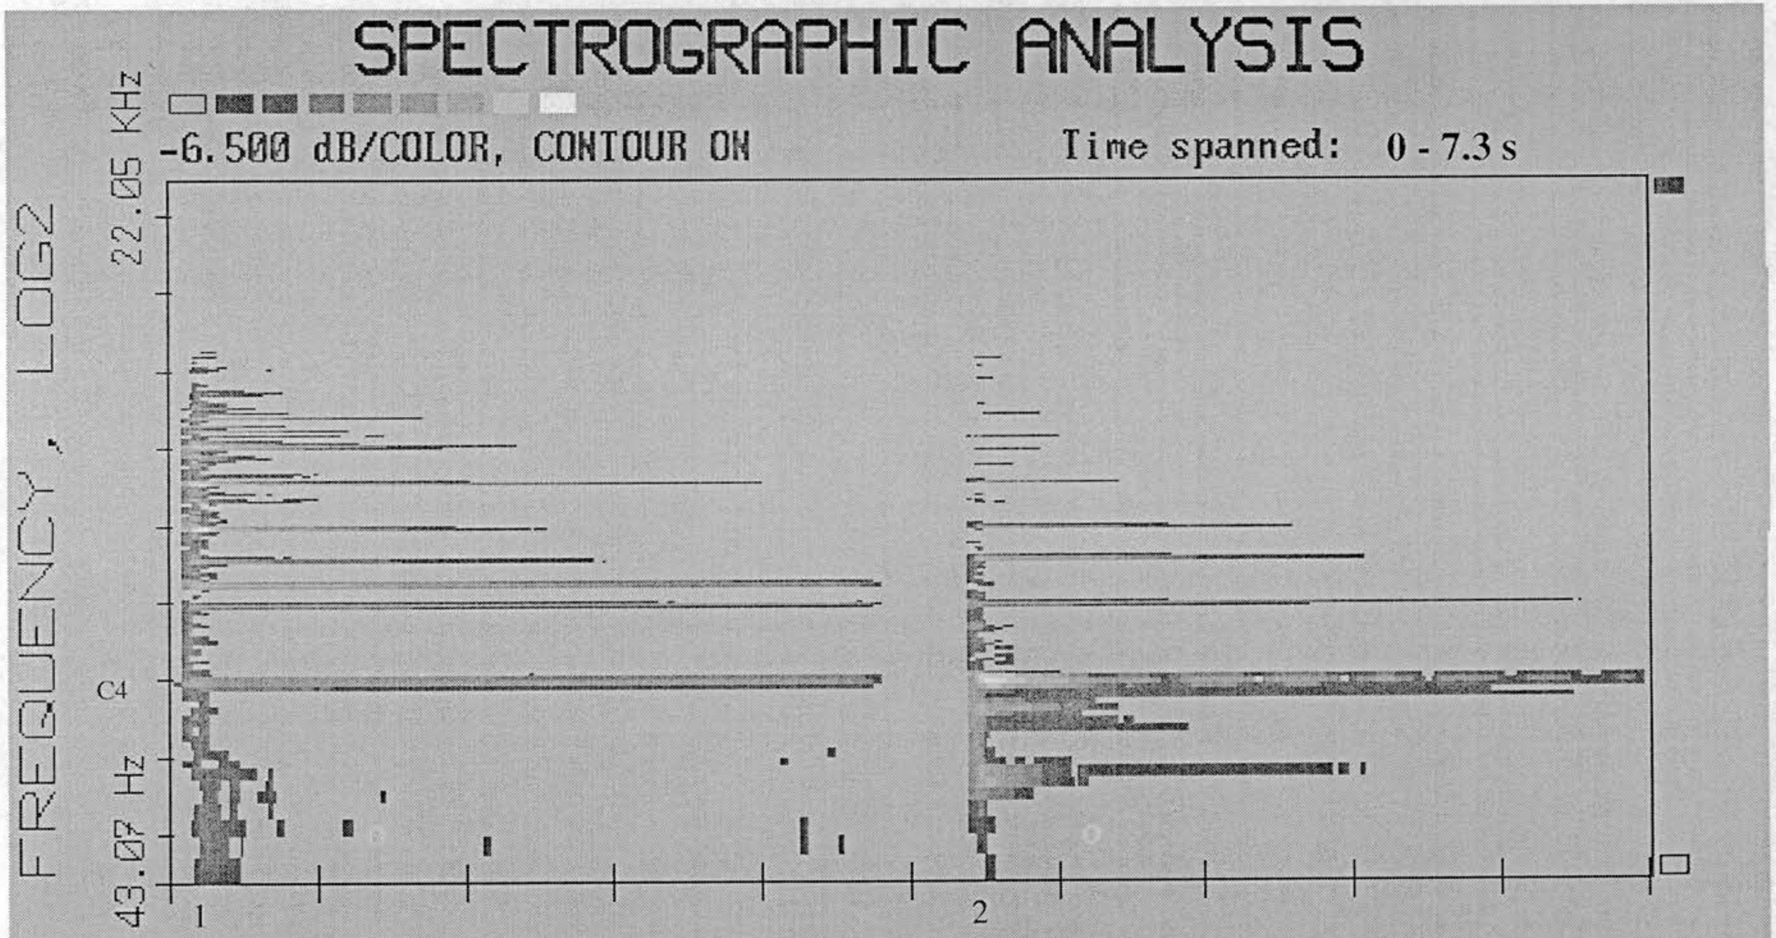

Figure 2.1

1. Spectrum of carillon bell

2. Spectrum of guitar tone

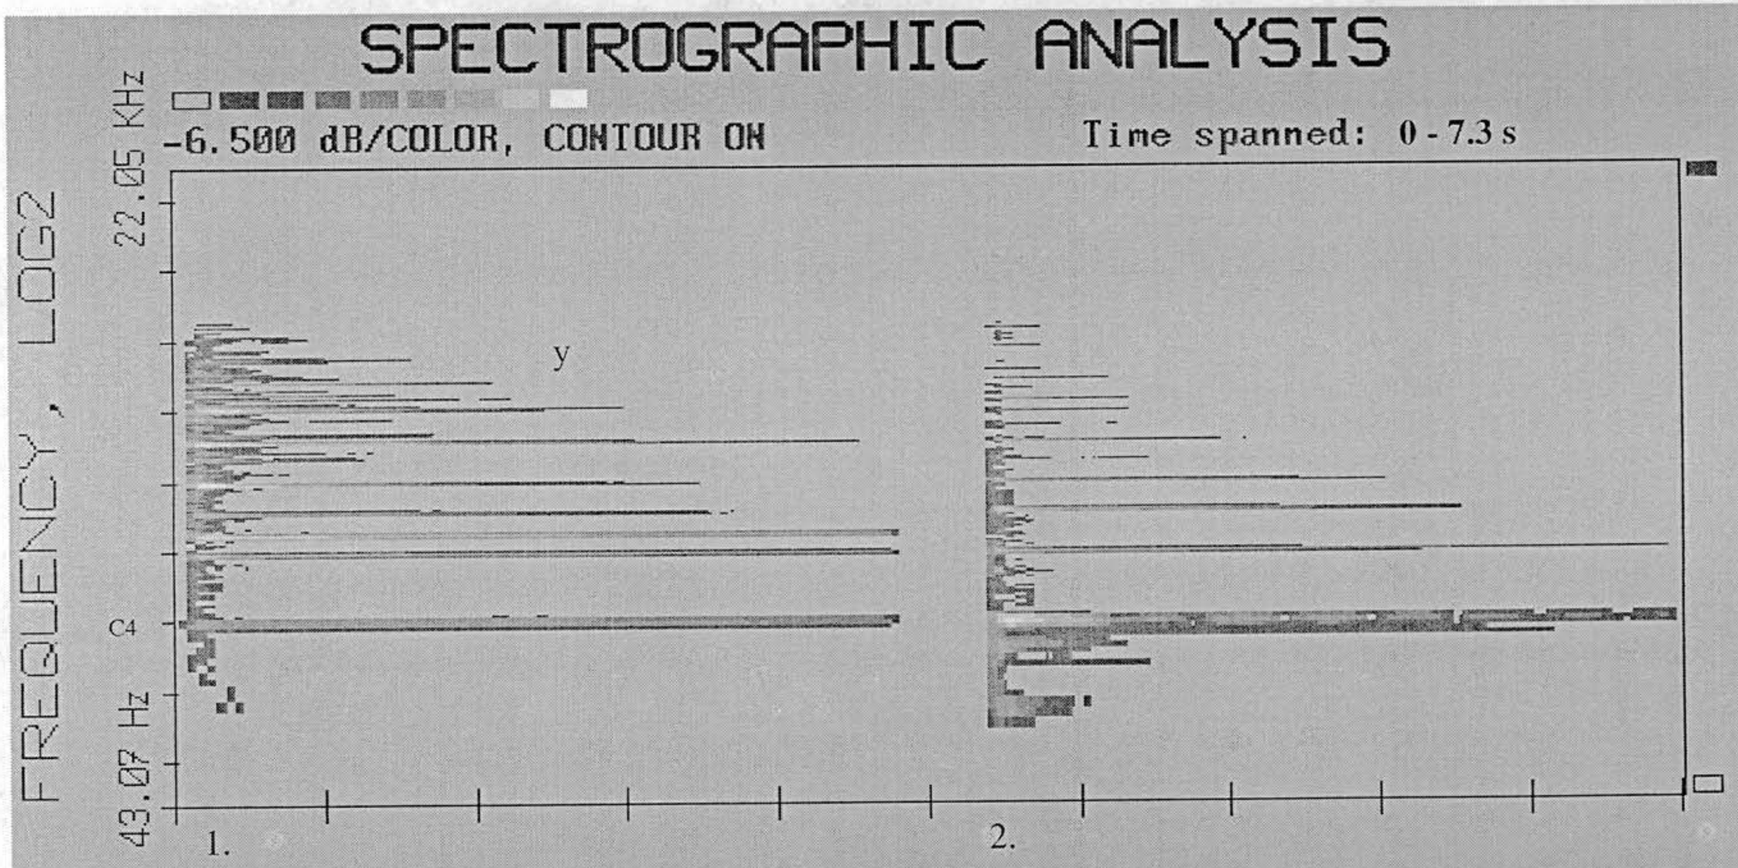

Figure 2.2 A-weighted version of fig. 2.1

1. Spectrum of carillon bell

2. Spectrum of guitar tone

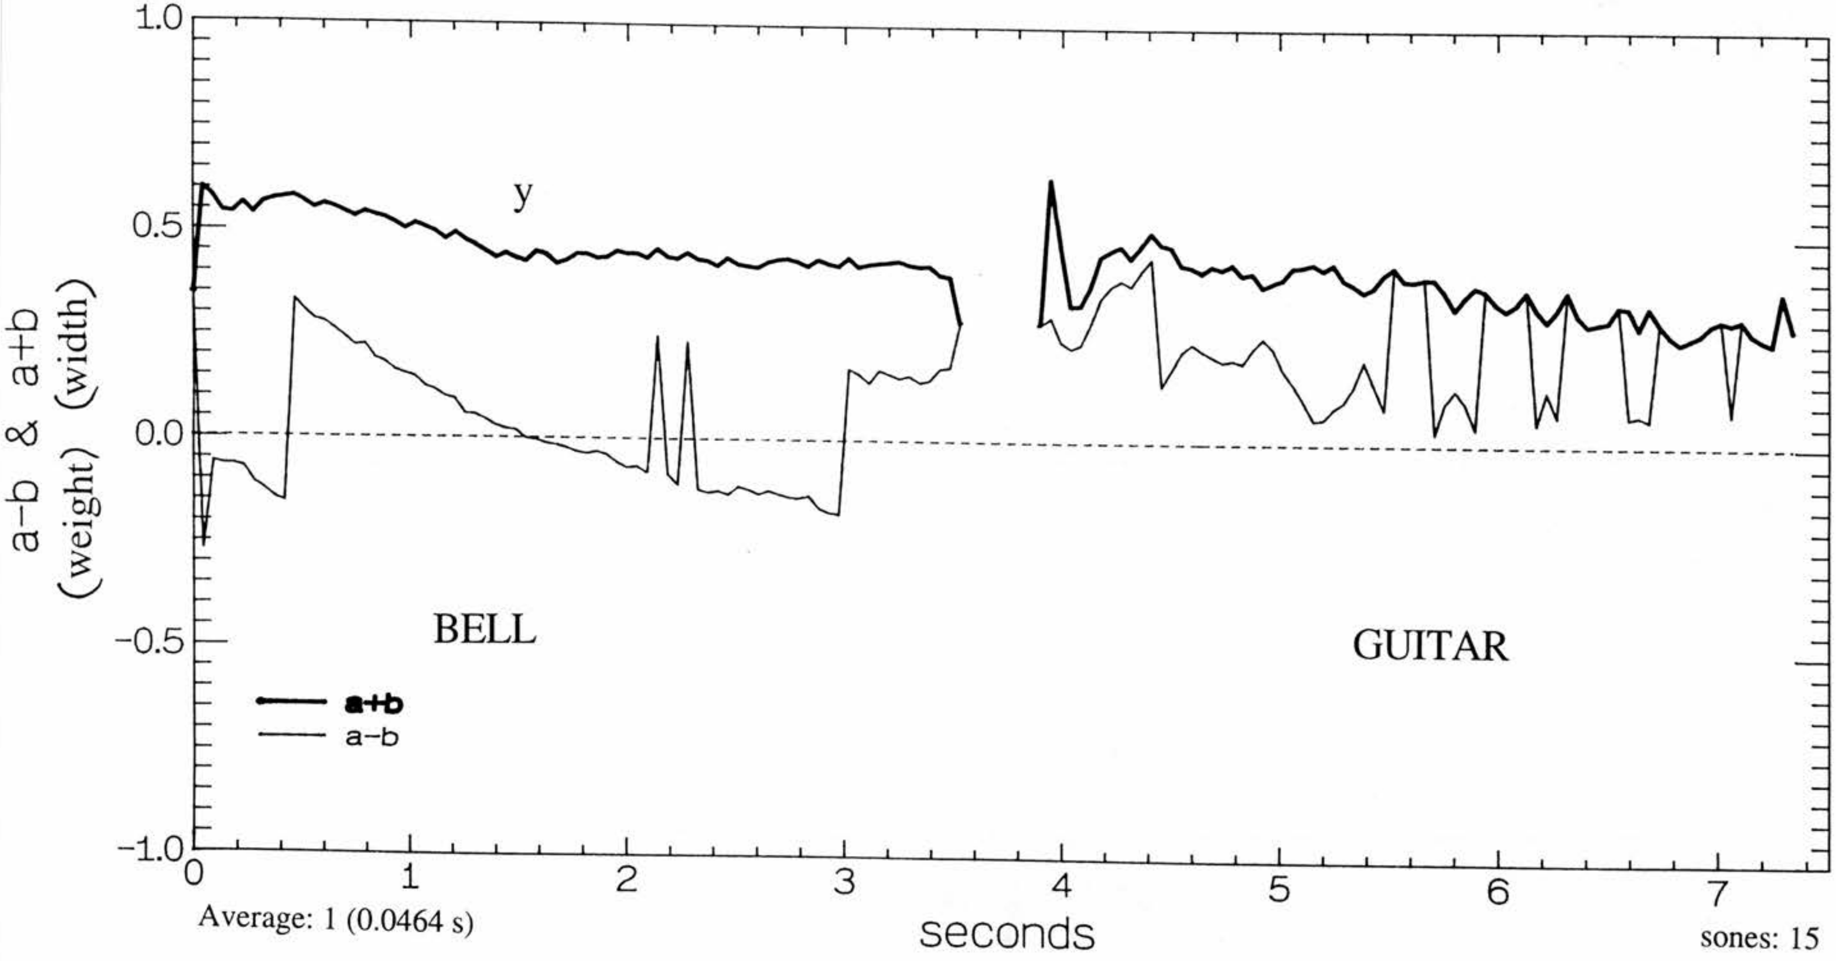

Figure 2.3 LD-graph of Bell tone followed by Guitar tone

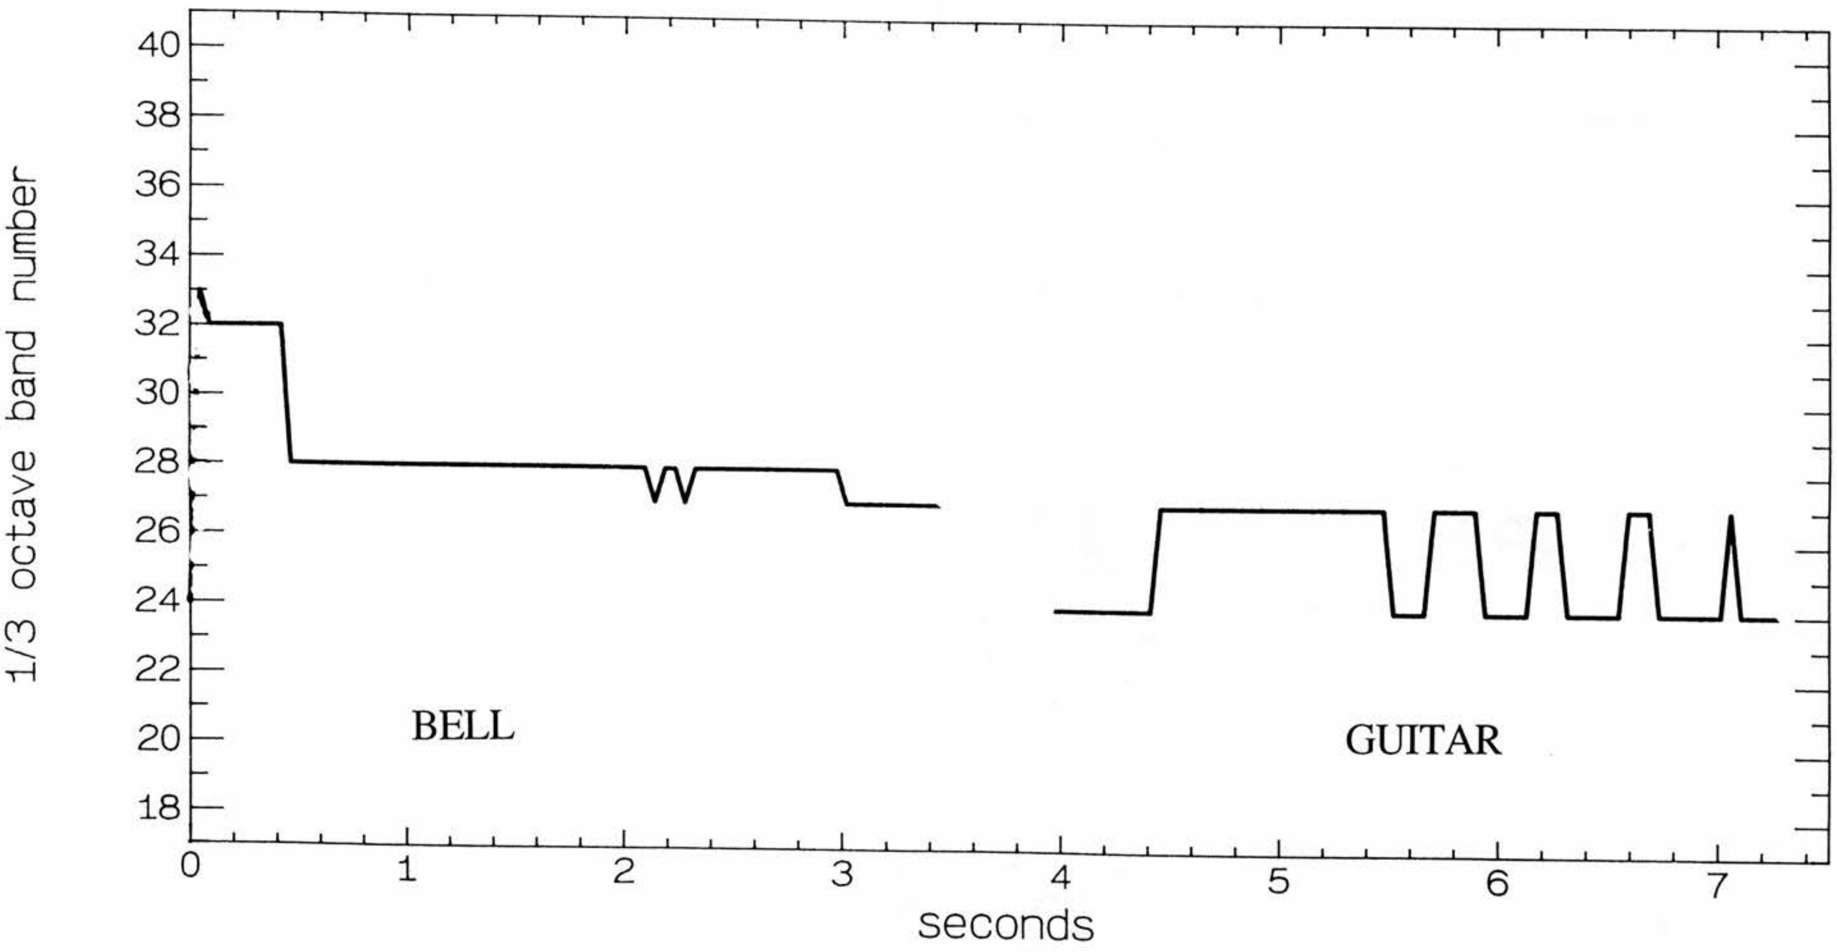

Figure 2.4 Timbral Pitch of Bell and Guitar tones

Figure 2.5a

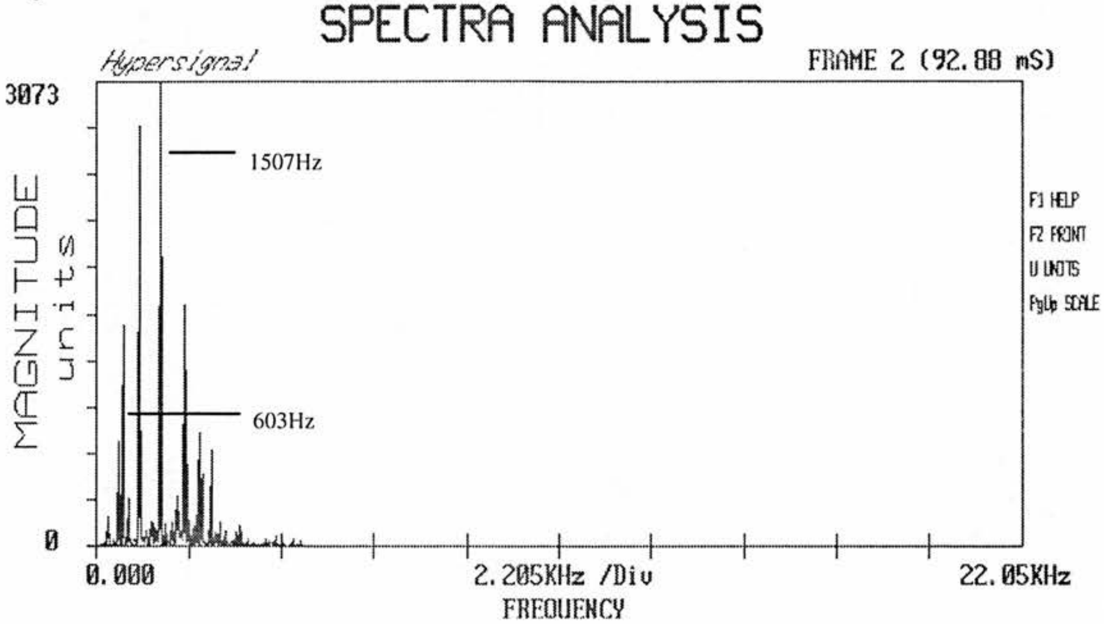

Stevens: band 28 - 17.3 sones
band 32 - 29.6 sones

Figure 2.5b

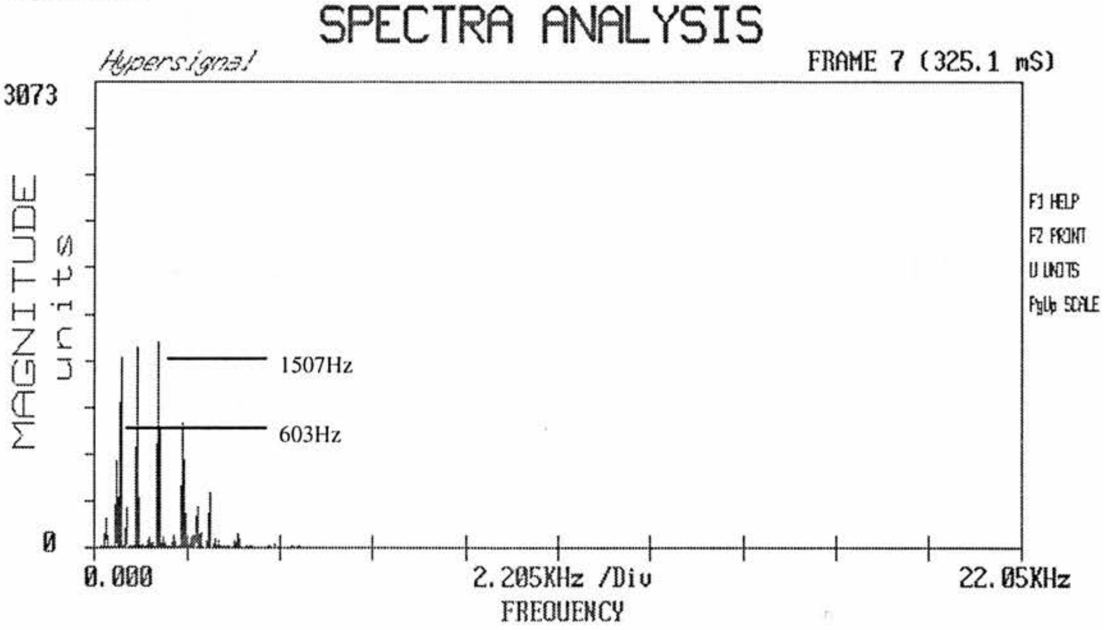

Stevens: band 28 - 17.3 sones
band 32 - 20.2 sones

Figure 2.5c

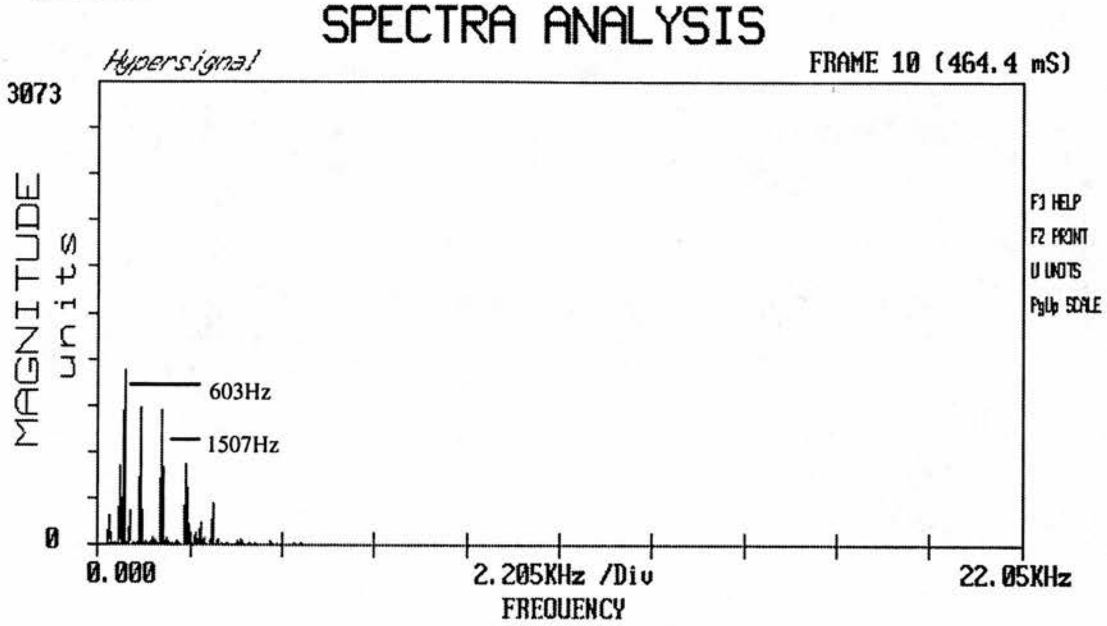

Stevens: band 28 - 17.3 sones

band 32 - 13.7 sones

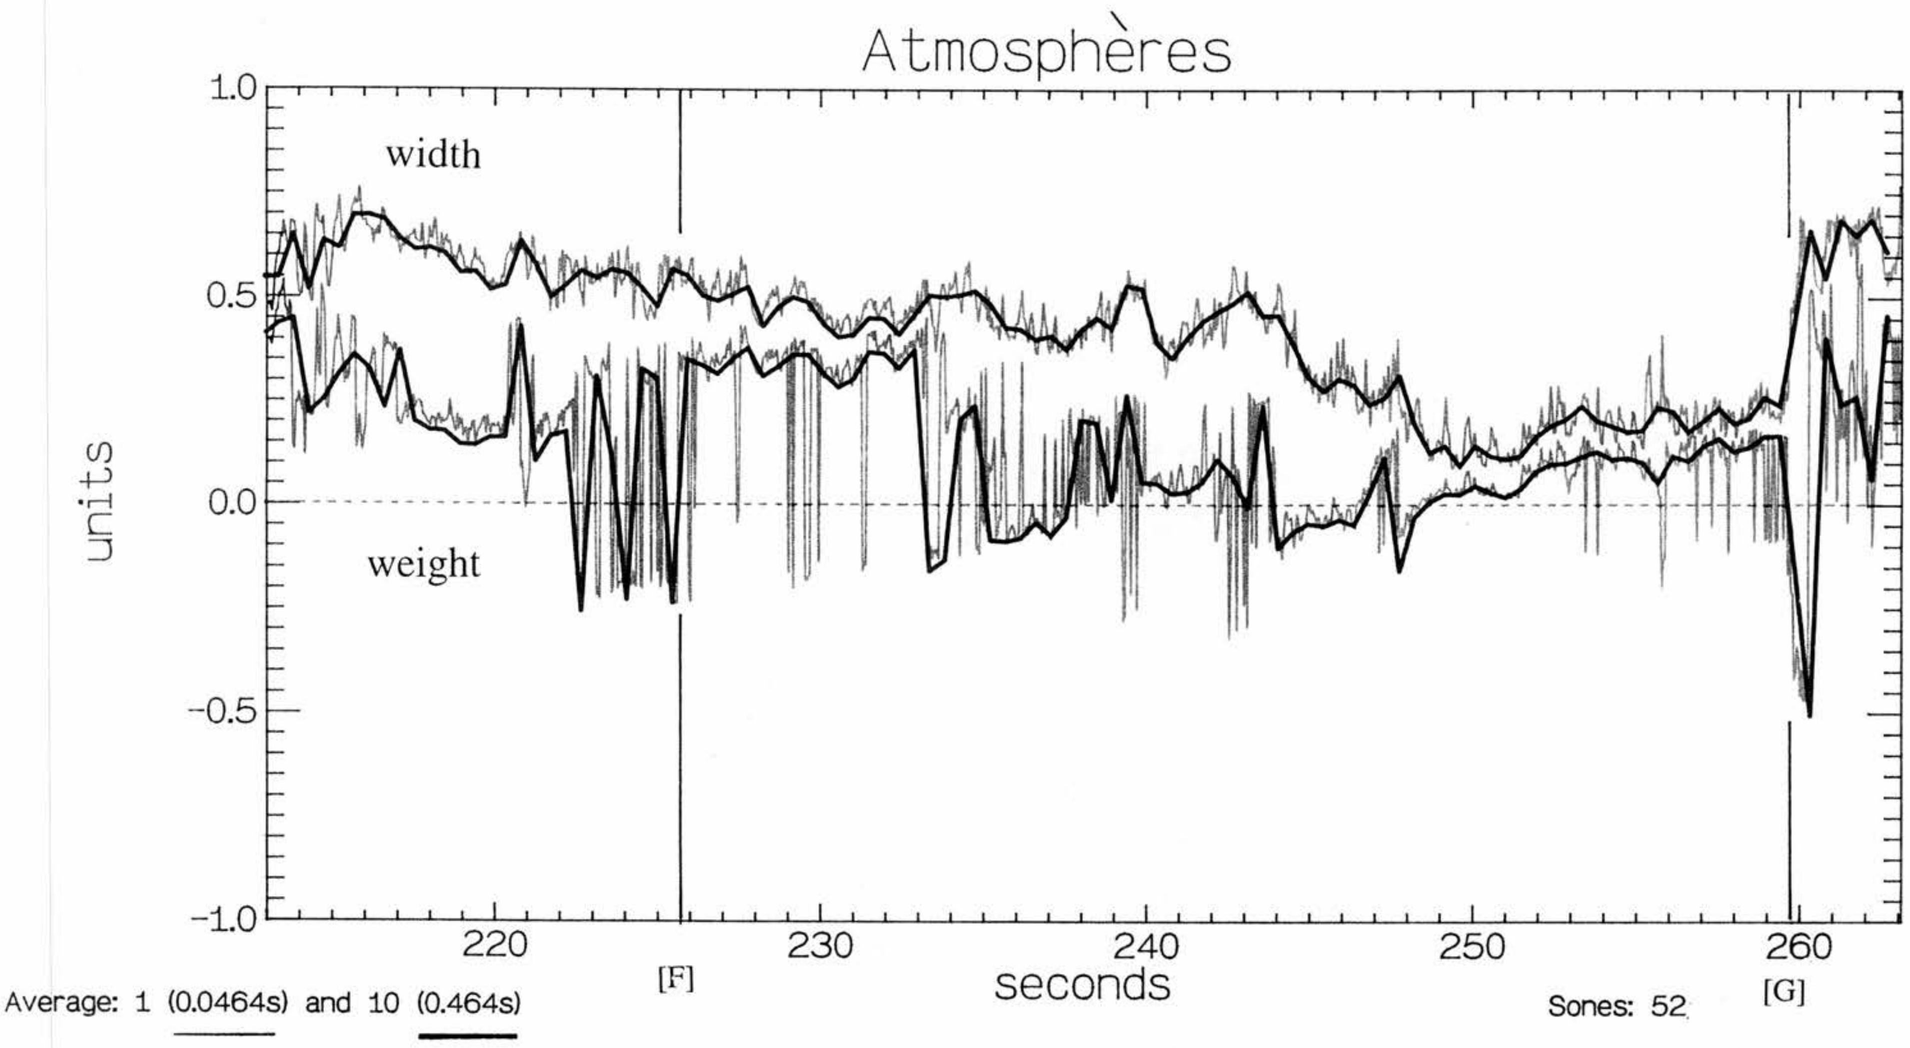

Figure 2.6 LD-graph of a section from Atmosphères

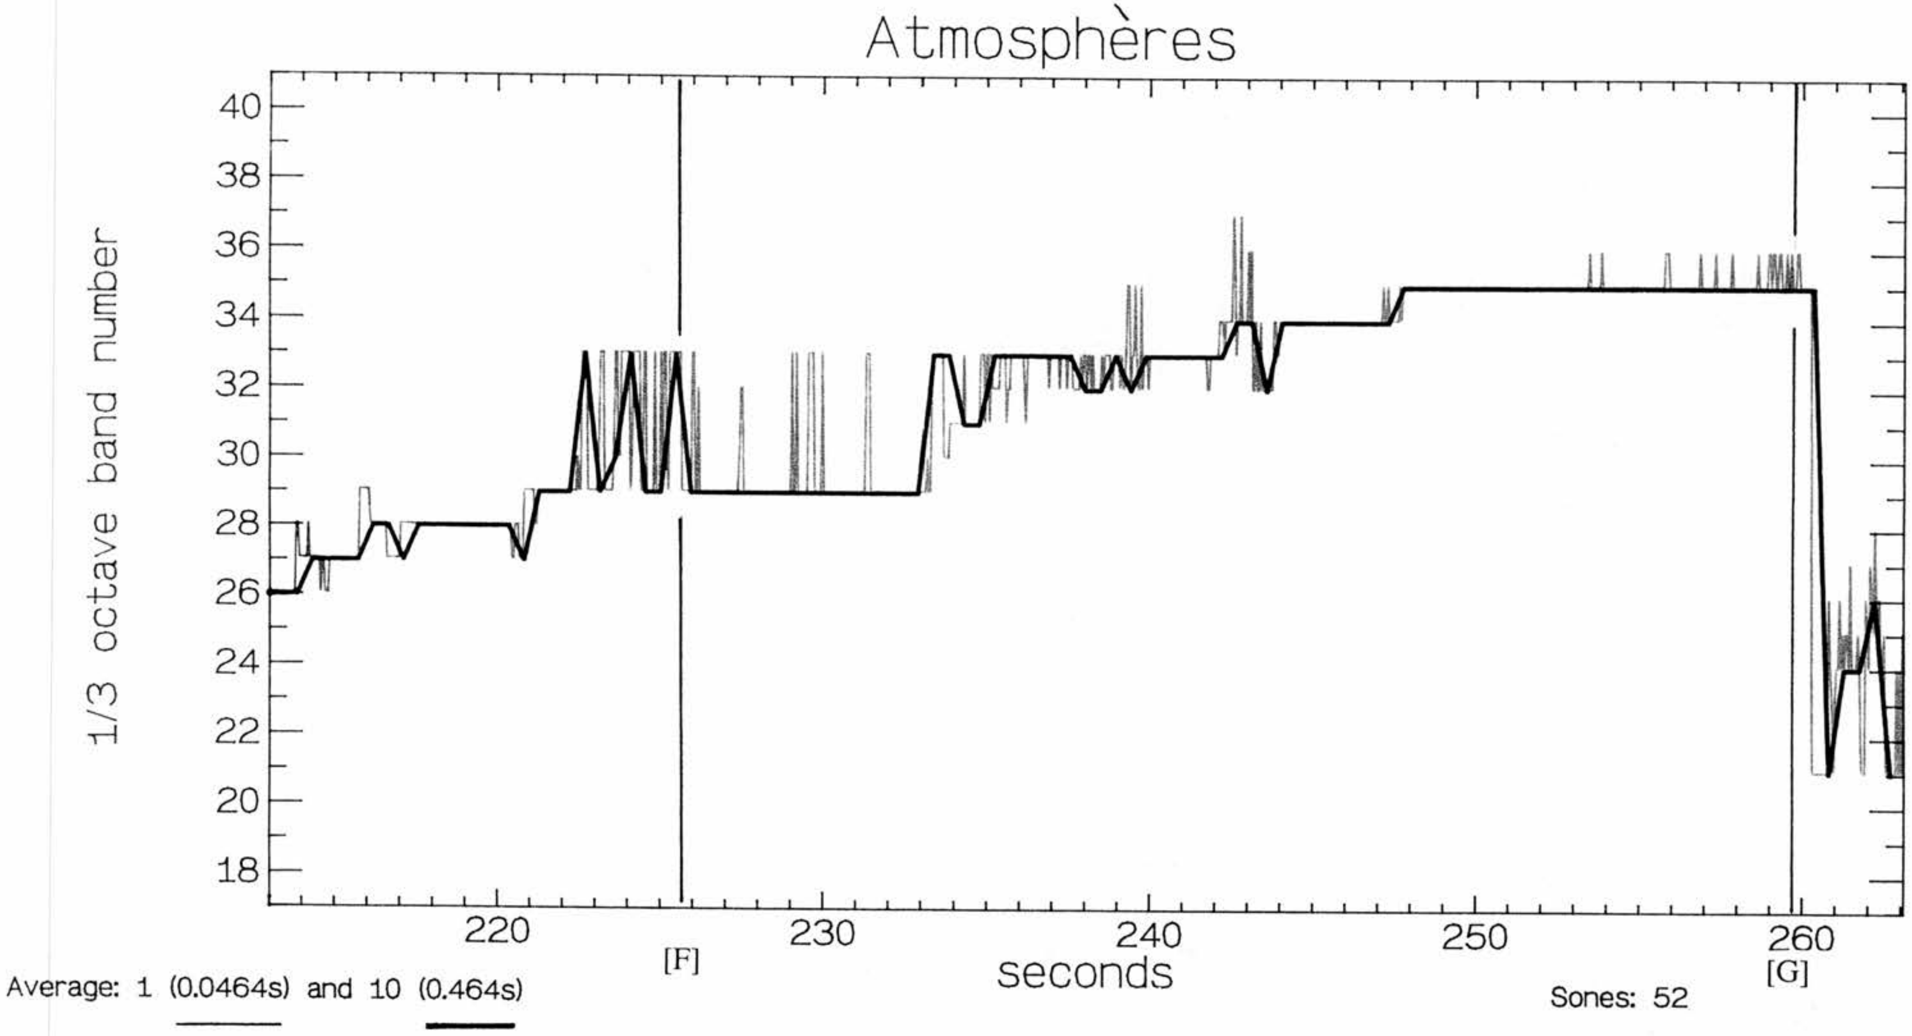

Figure 2.7 Timbral pitch from a section of Atmosphères

Figure 2.8

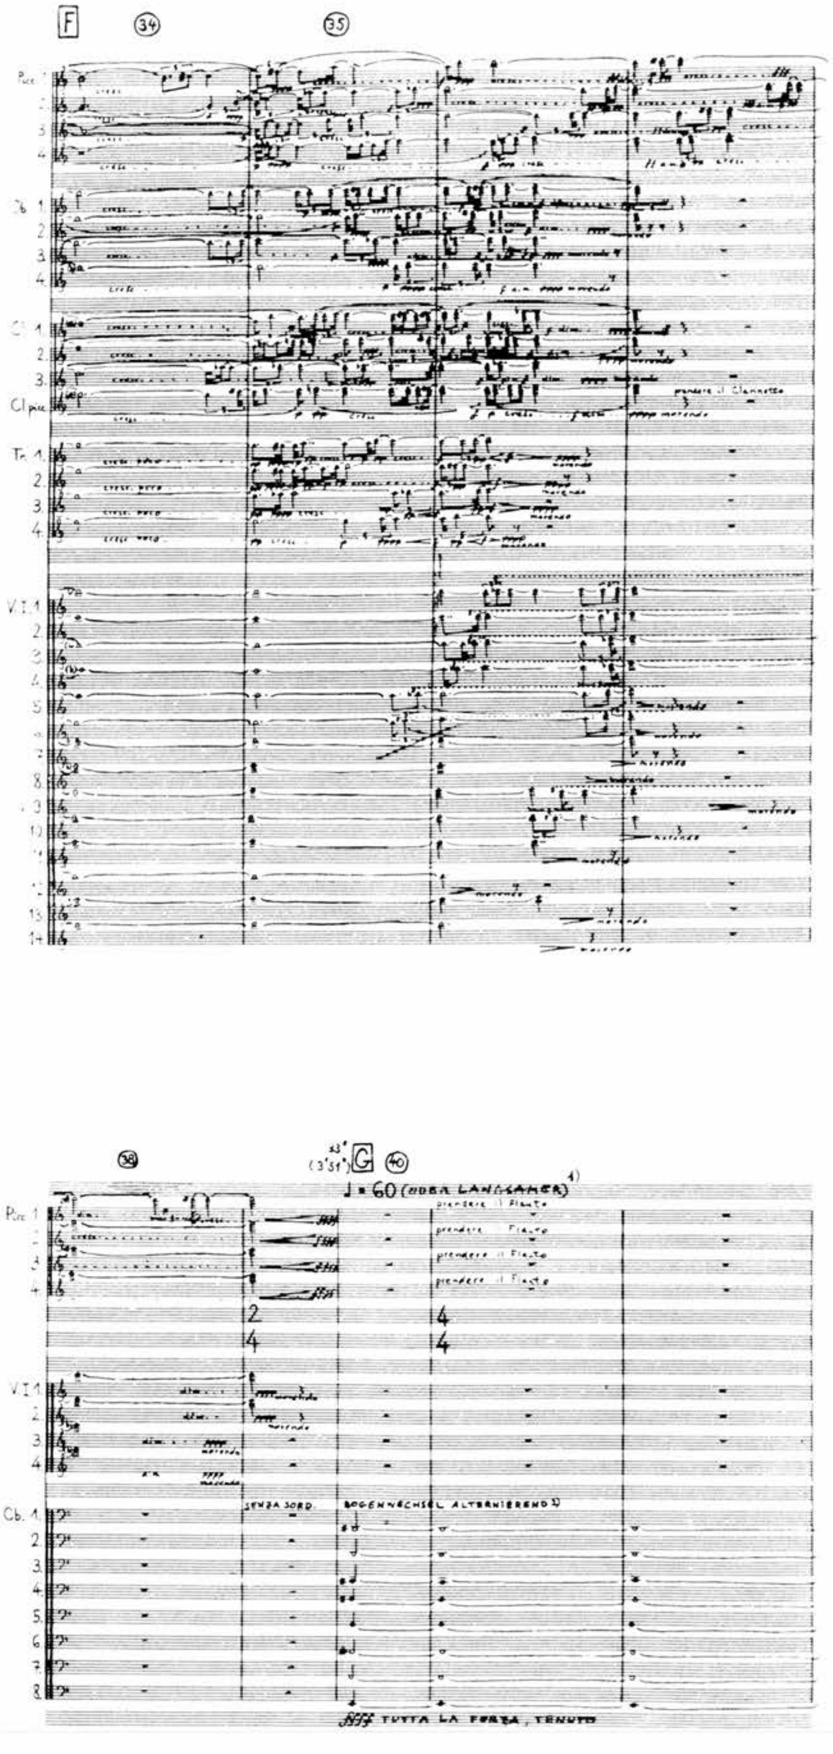

or slower

2) alternating change of bow

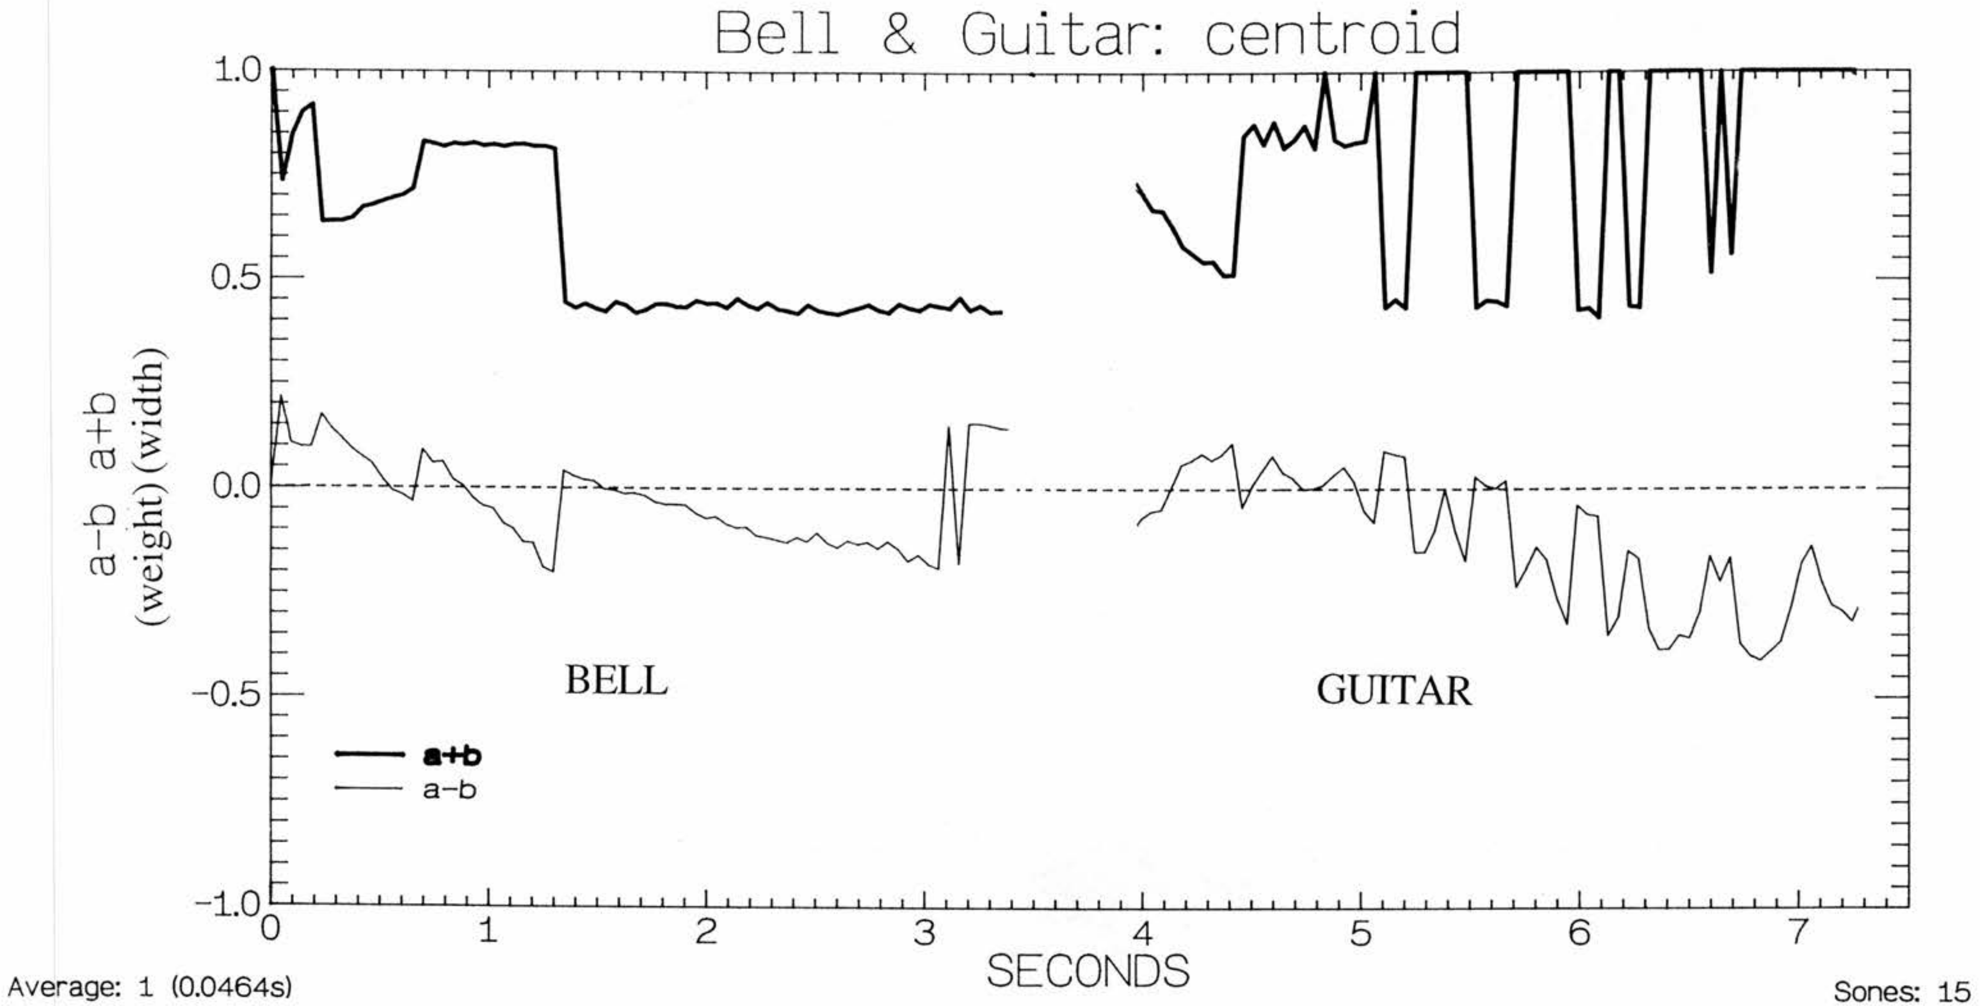

Figure 2.9 LD-graph calculated using the position of the loudness centroid

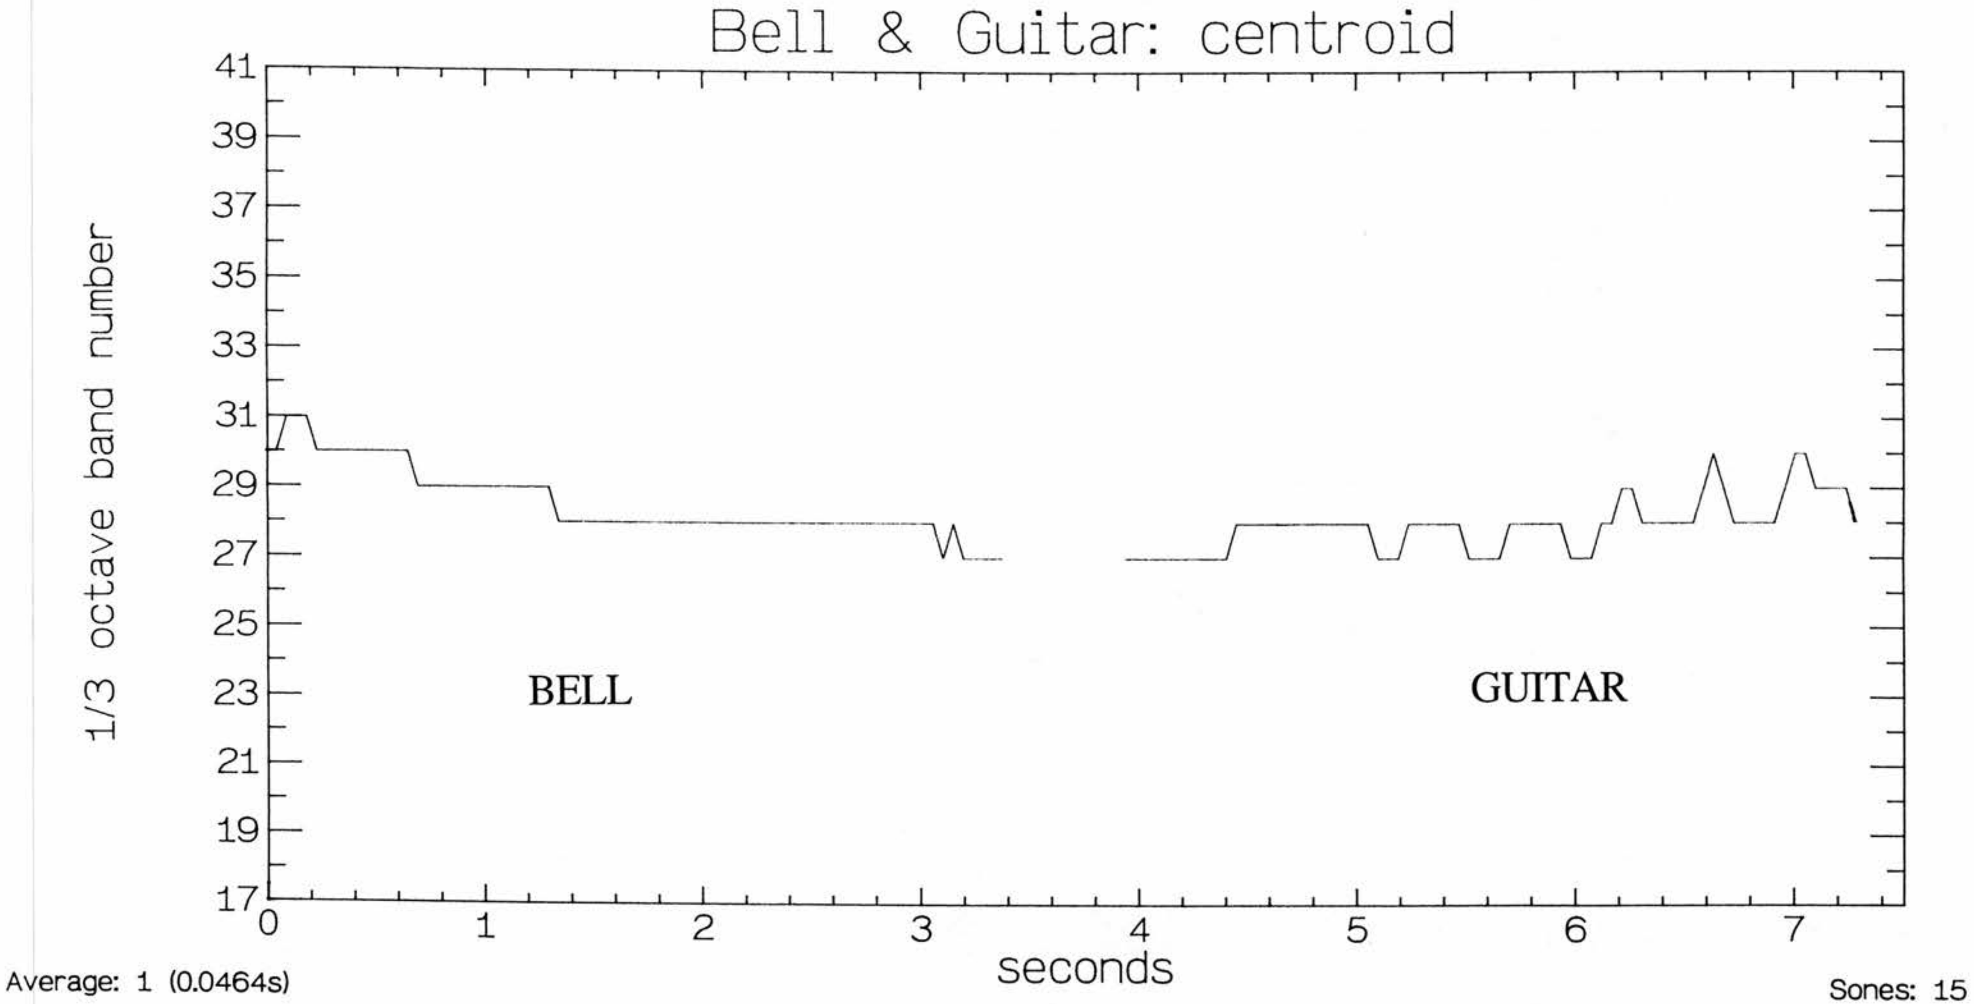

Figure 2.10 Position of the loudness centroid used in the calculation of the data dislpayed in figure 2.9

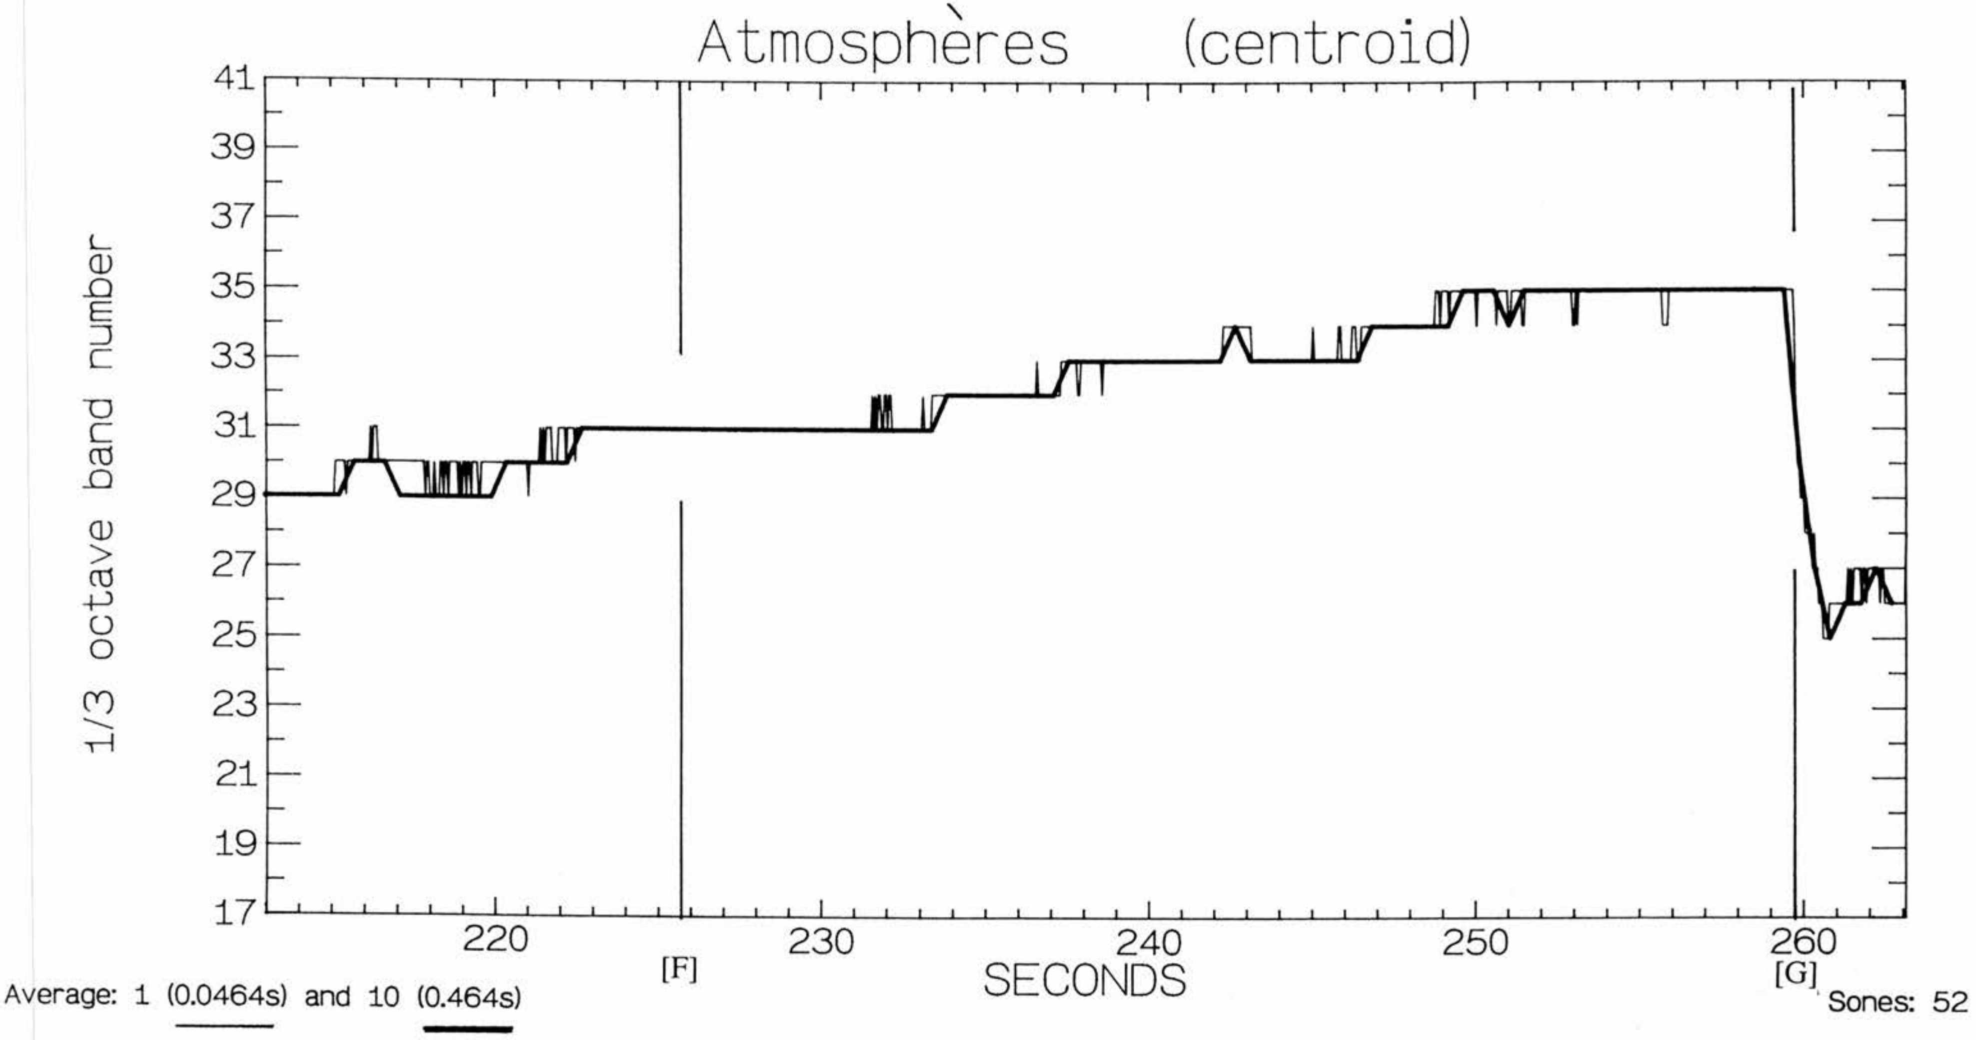

Figure 2.11 Loudness centroid of a section from Atmosphères

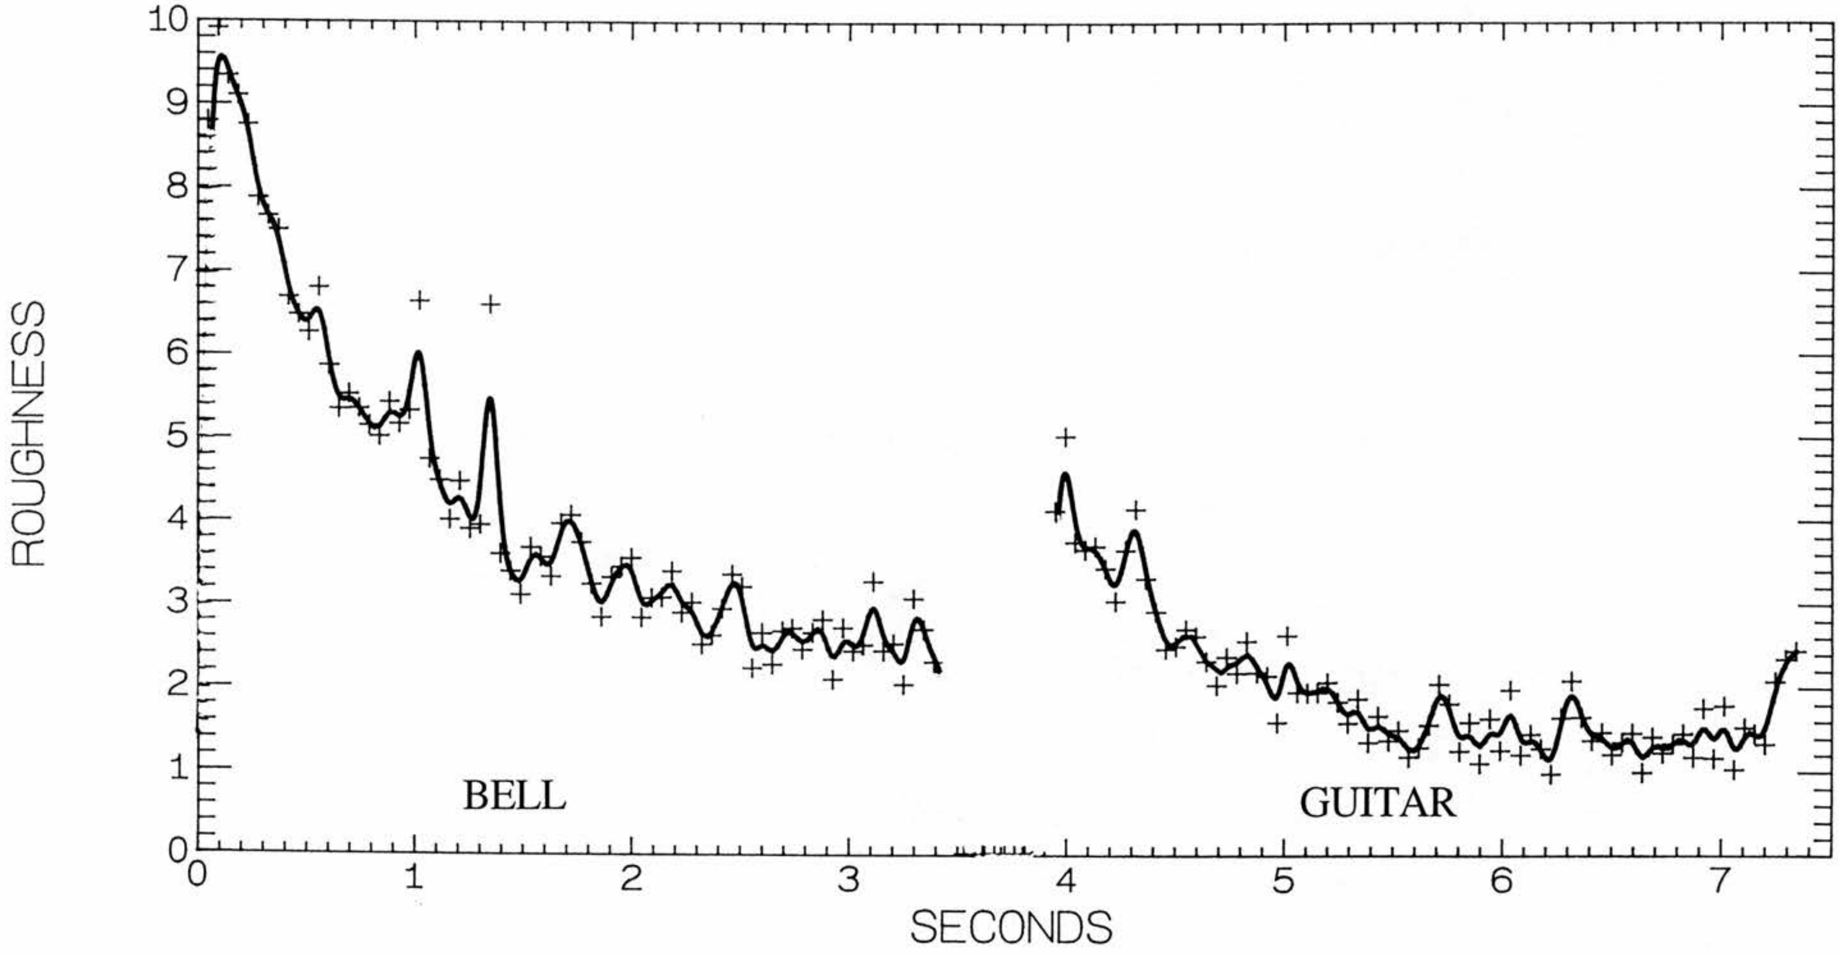

Figure 2.12 Roughness of Bell and Guitar tones

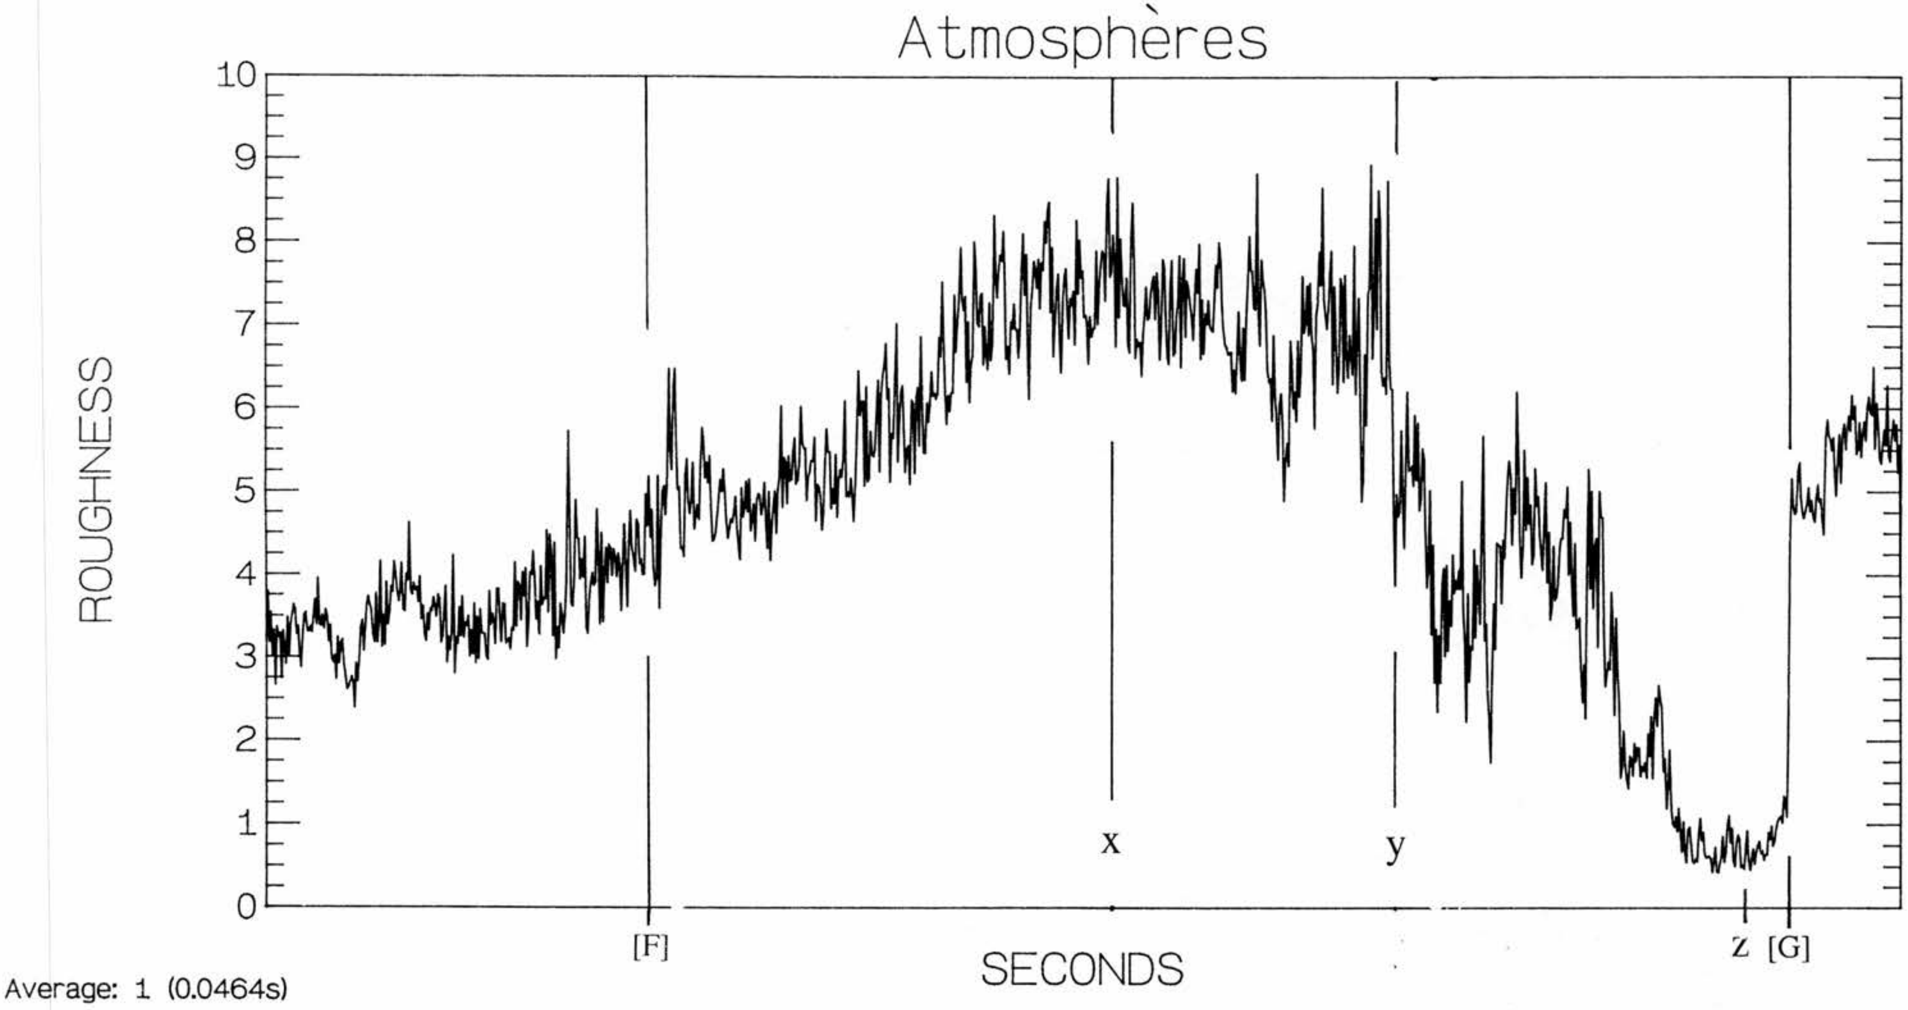

Figure 2.13 Roughness of a section from Atmosphères

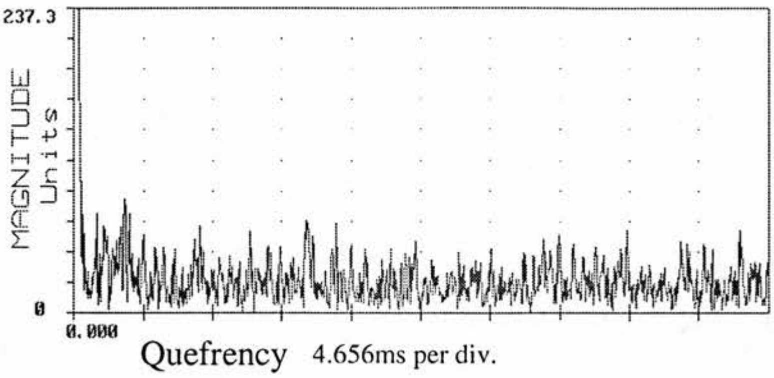

noise

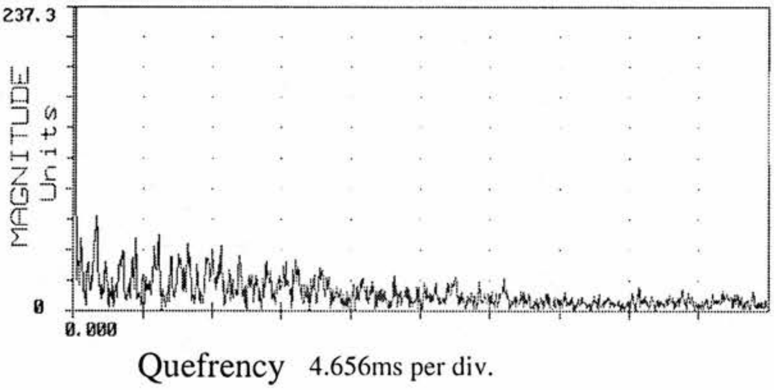

carillon bell

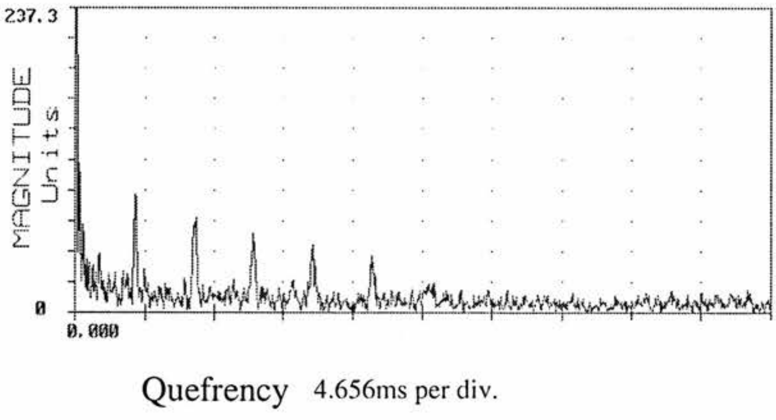

guitar

Figure 2.14 Cepstrum analysis: comparison of different degrees of harmonicity

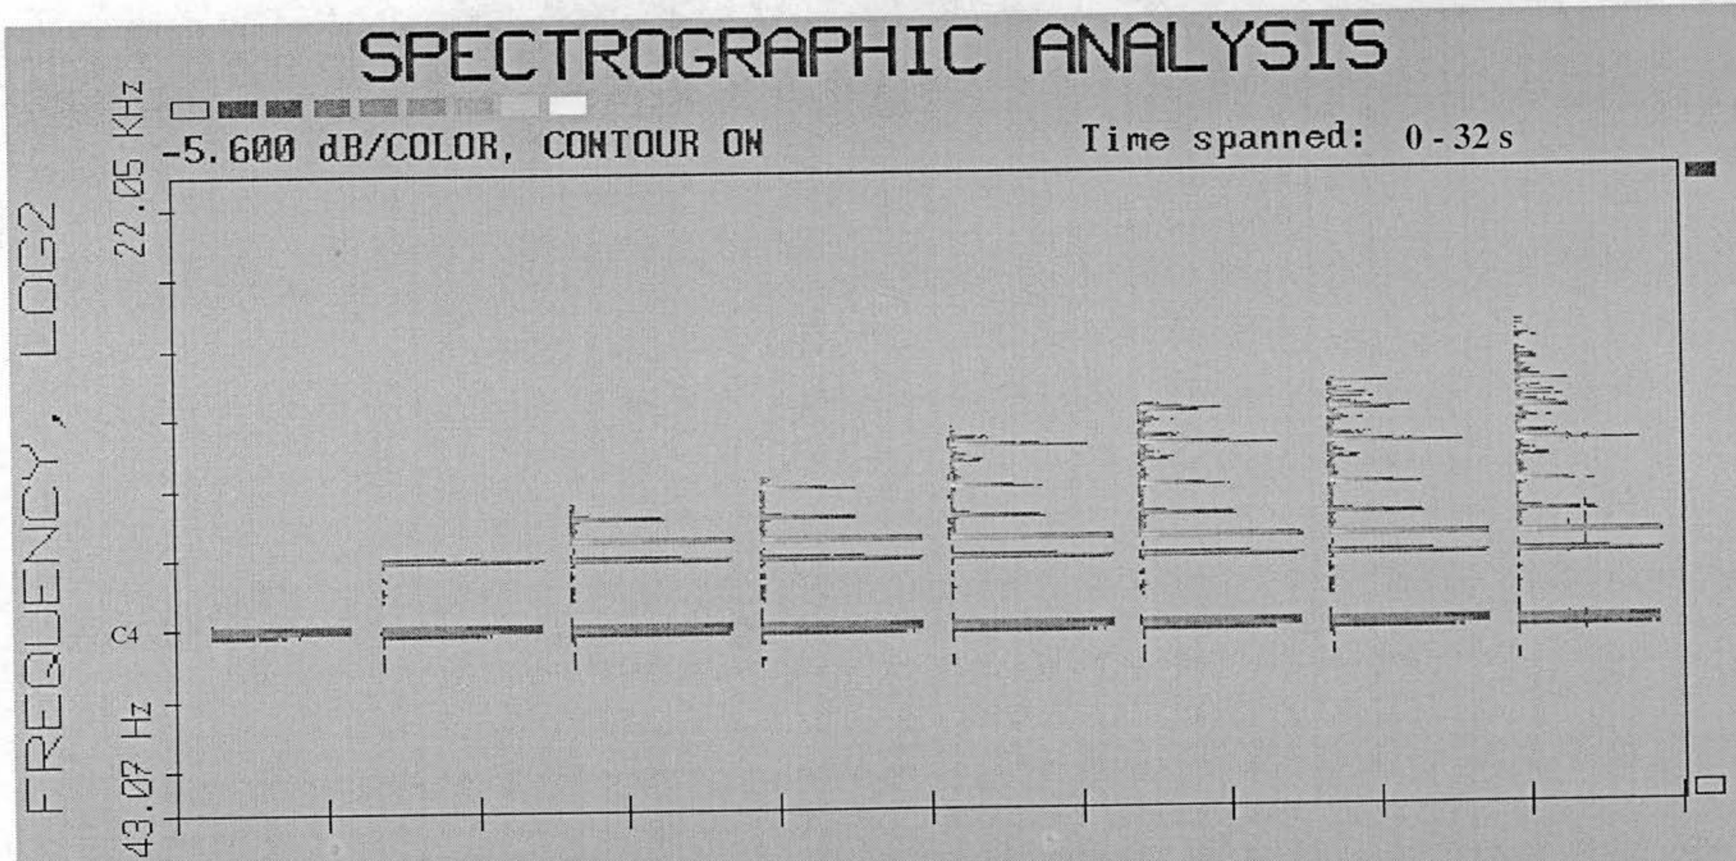

Figure 2.15 Eight stages in the re-synthesis of a carillon bell

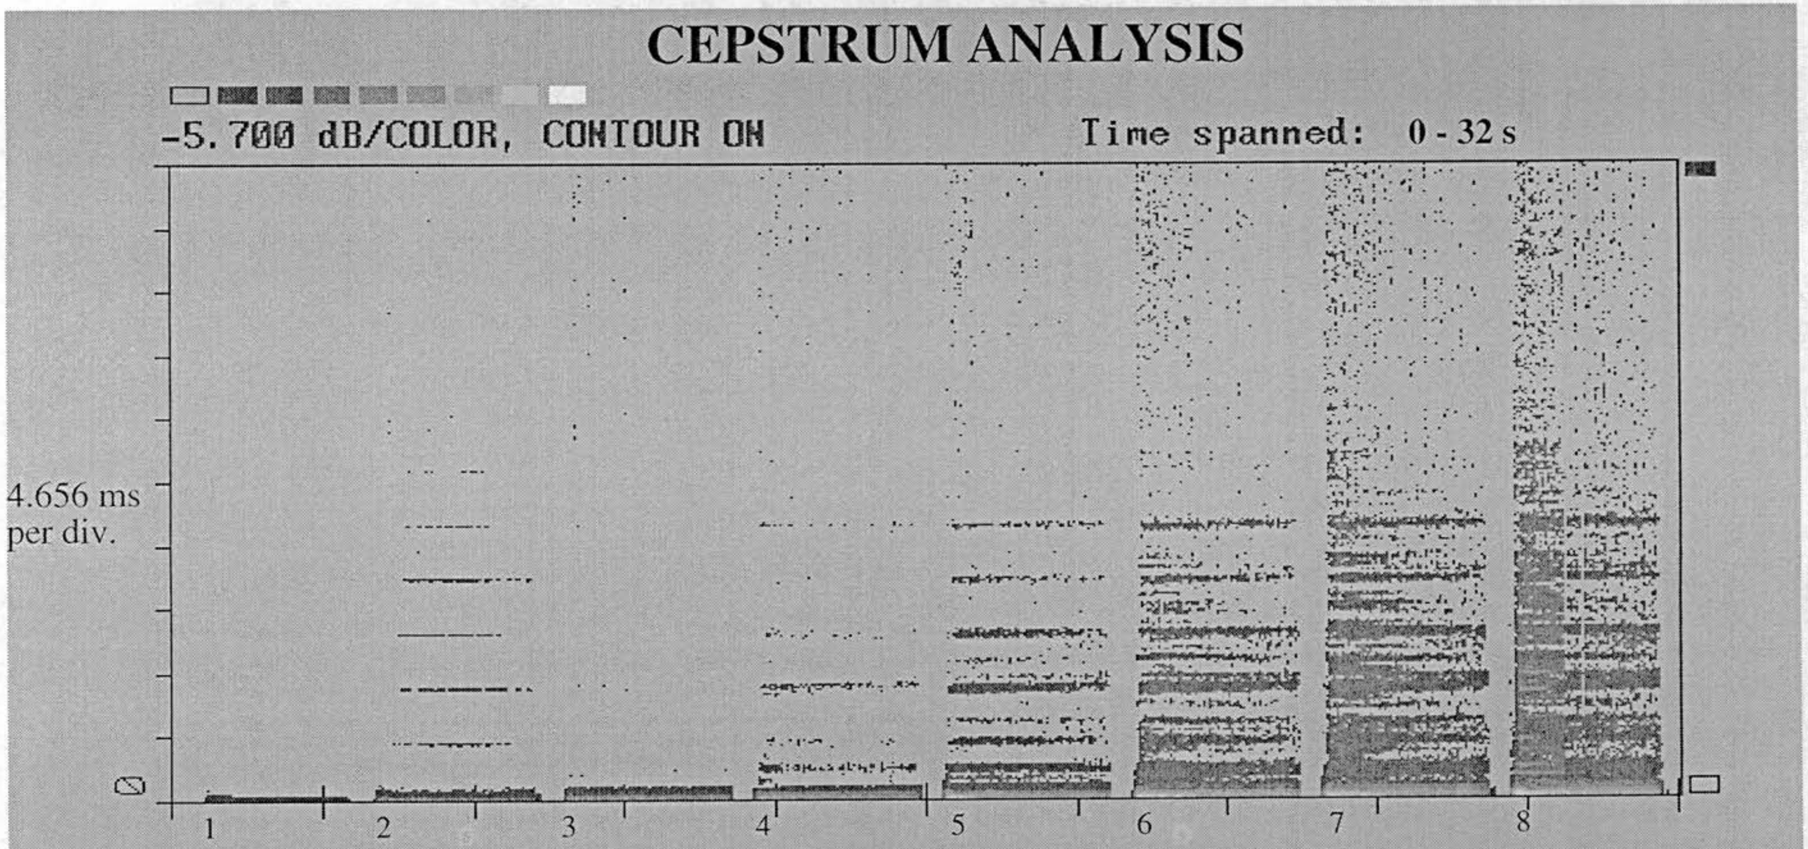

Figure 2.16 Eight stages in the re-synthesis of a carrillon bell tone

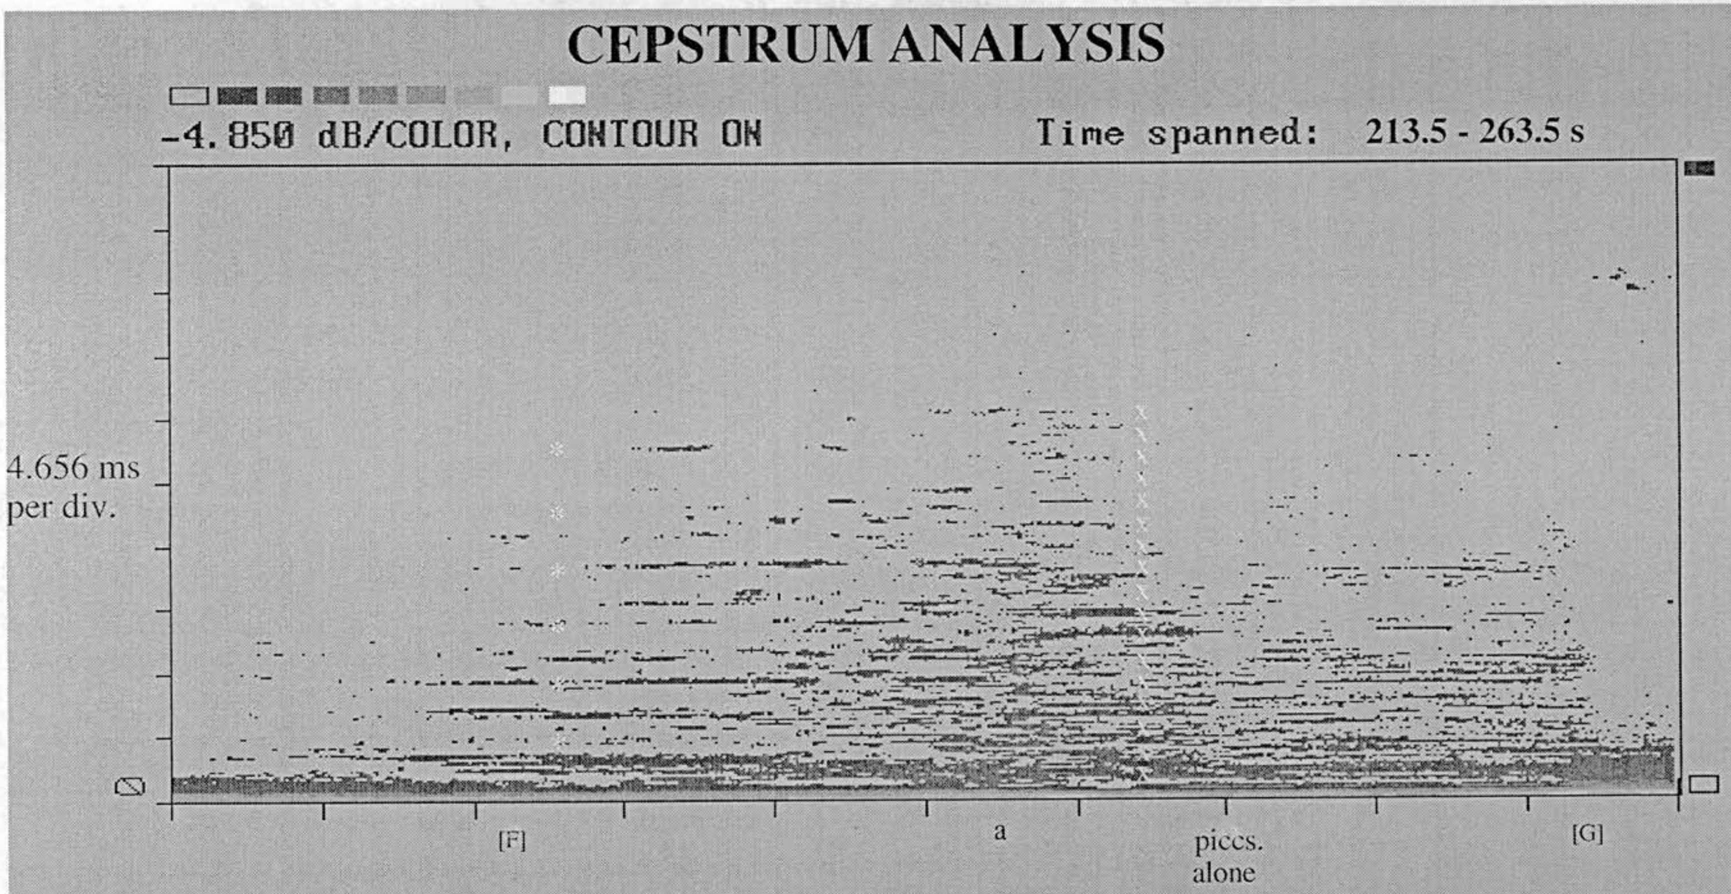

Figure 2.17 Analysis of a section of Atmospheres

Figures for

chapter 3

Analysis of *Jeux vénitiens*: movement 1

Figure 3.1

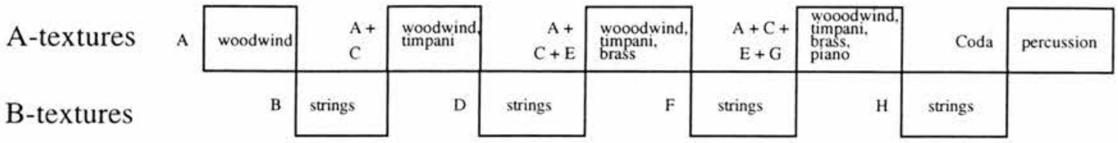

Structure of the first movement

Figure 3.2

Pitch structure of the first movement:

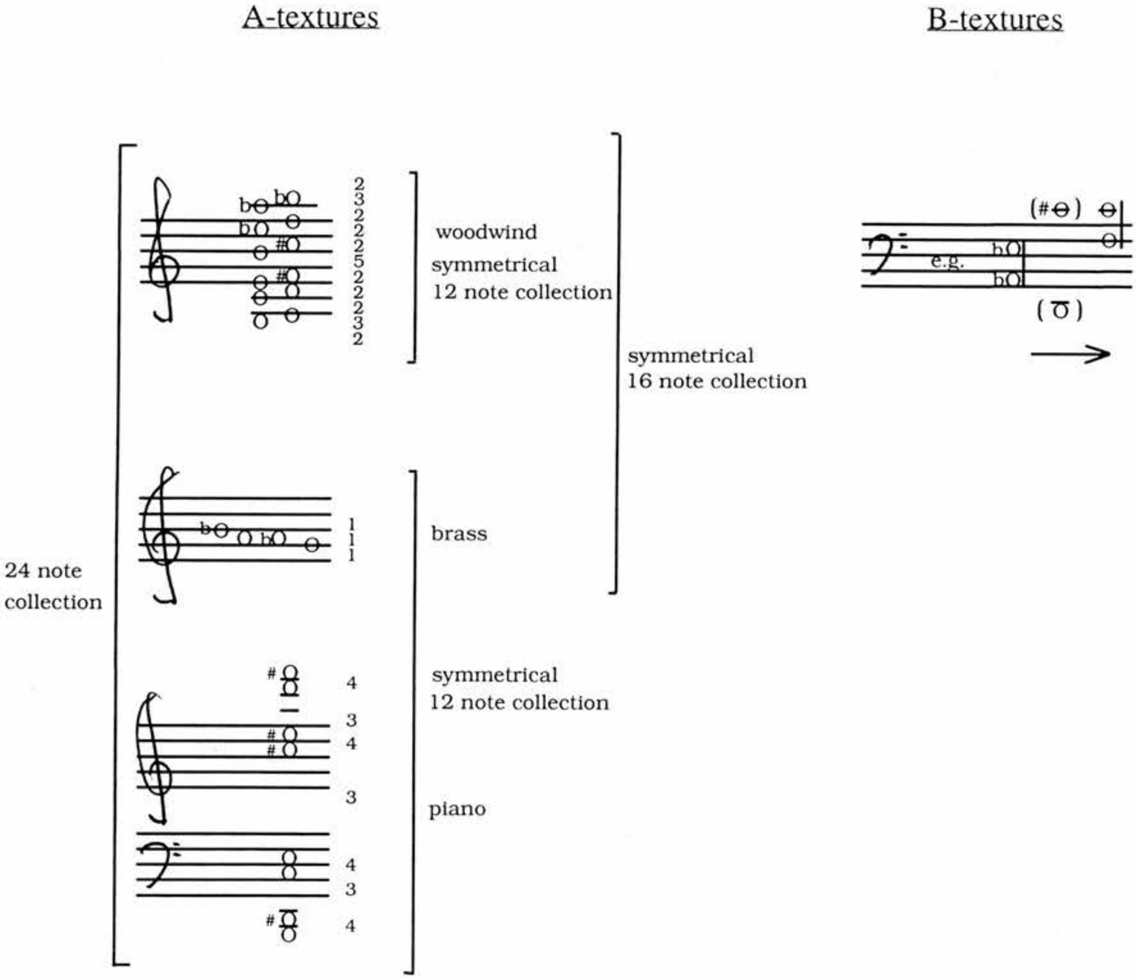

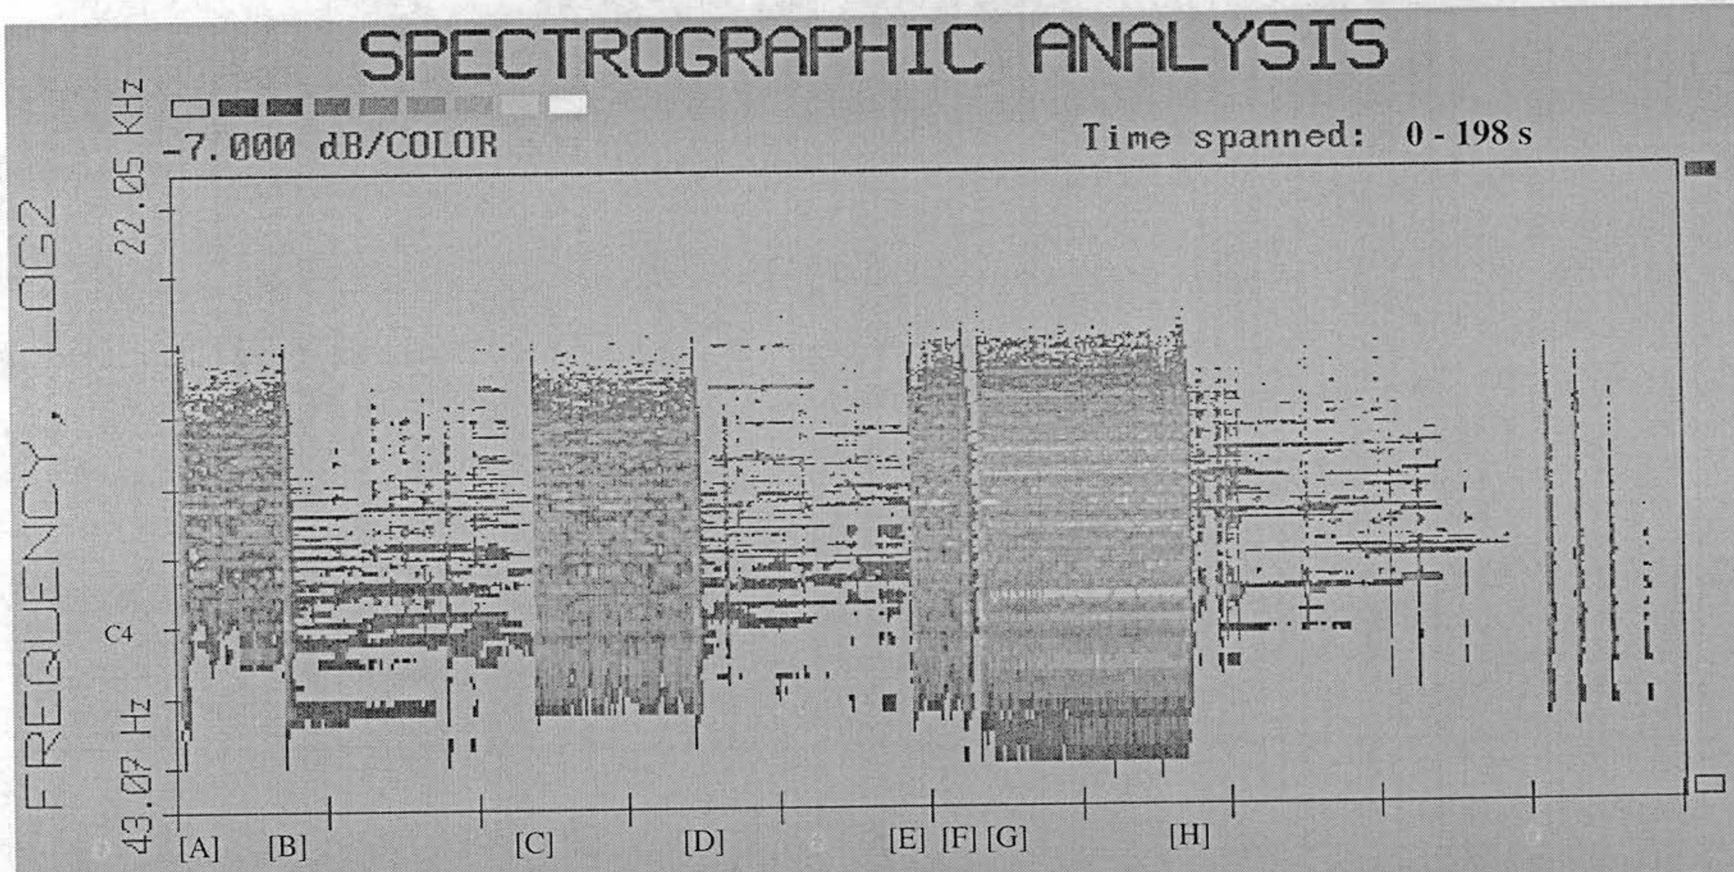

Jeux venitiens: movement 1

Figure 3.3

Figure 3.4
Timbral weight and width

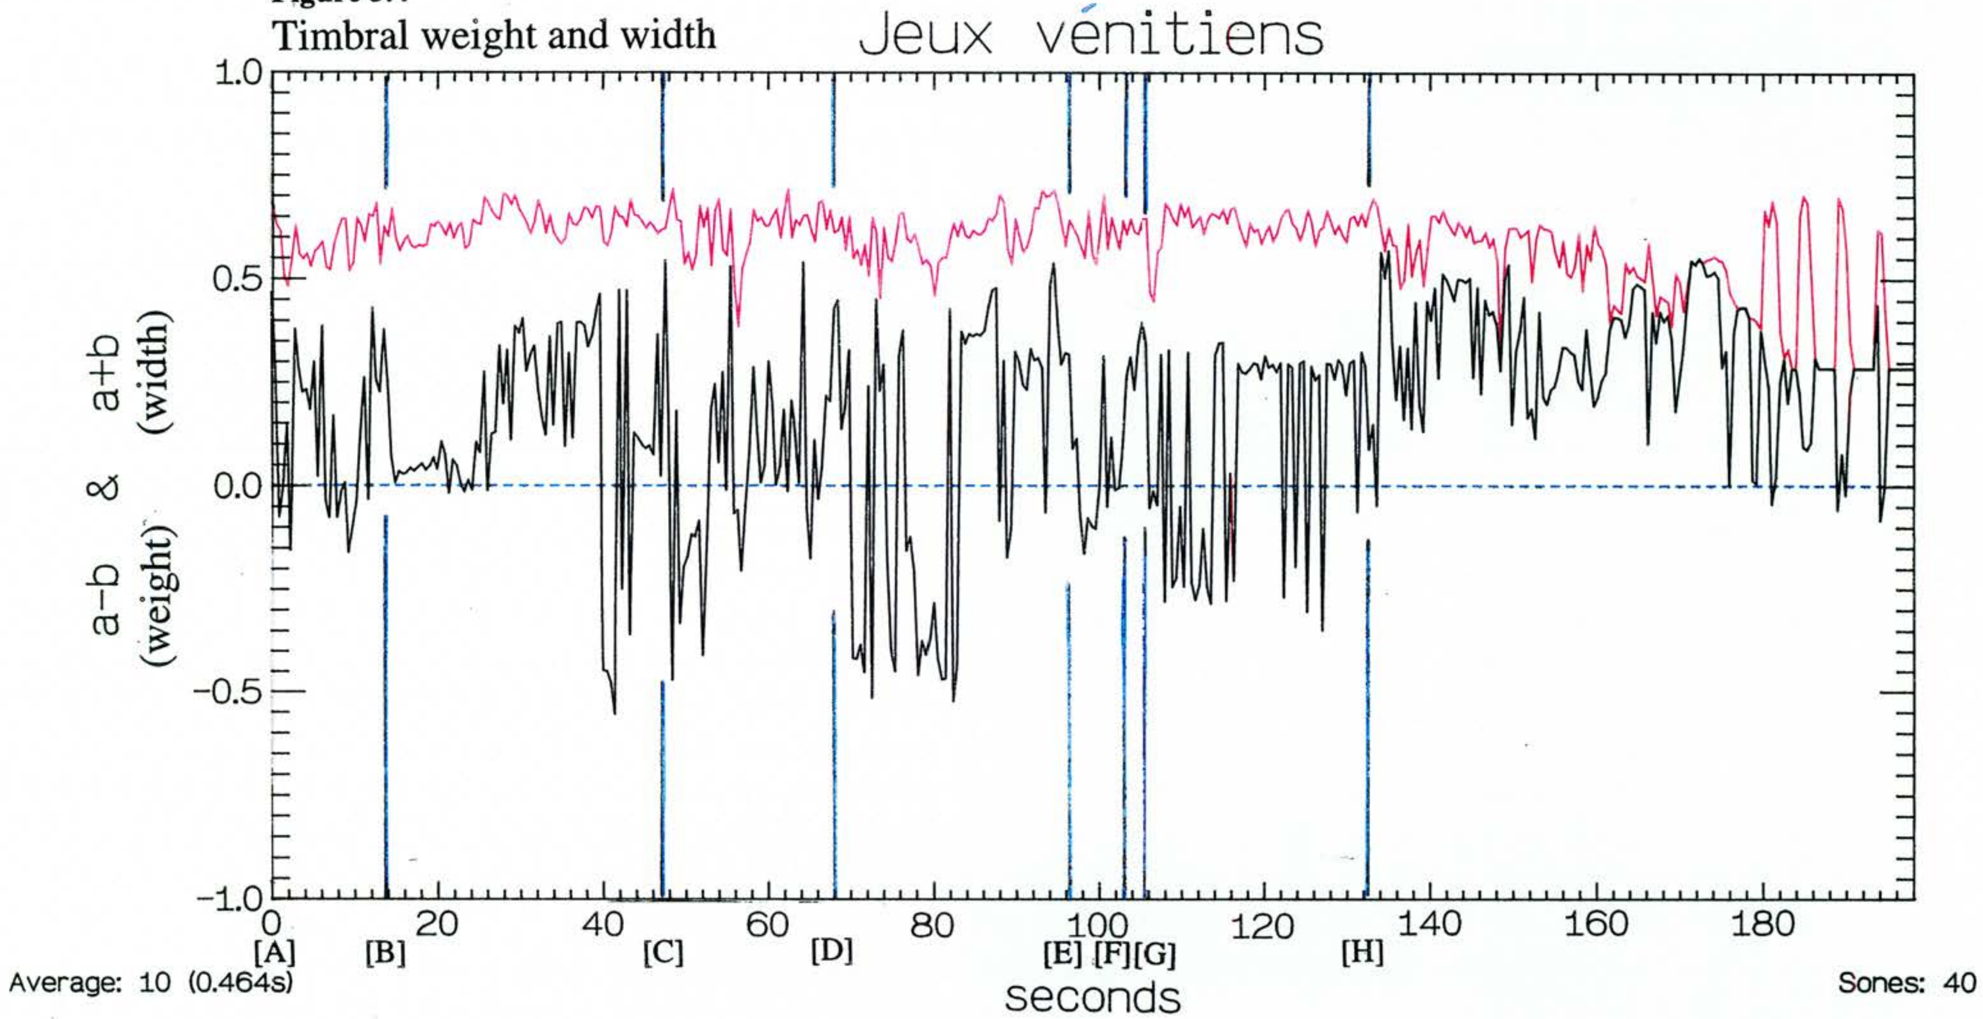

Figure 3.5

Timbral pitch

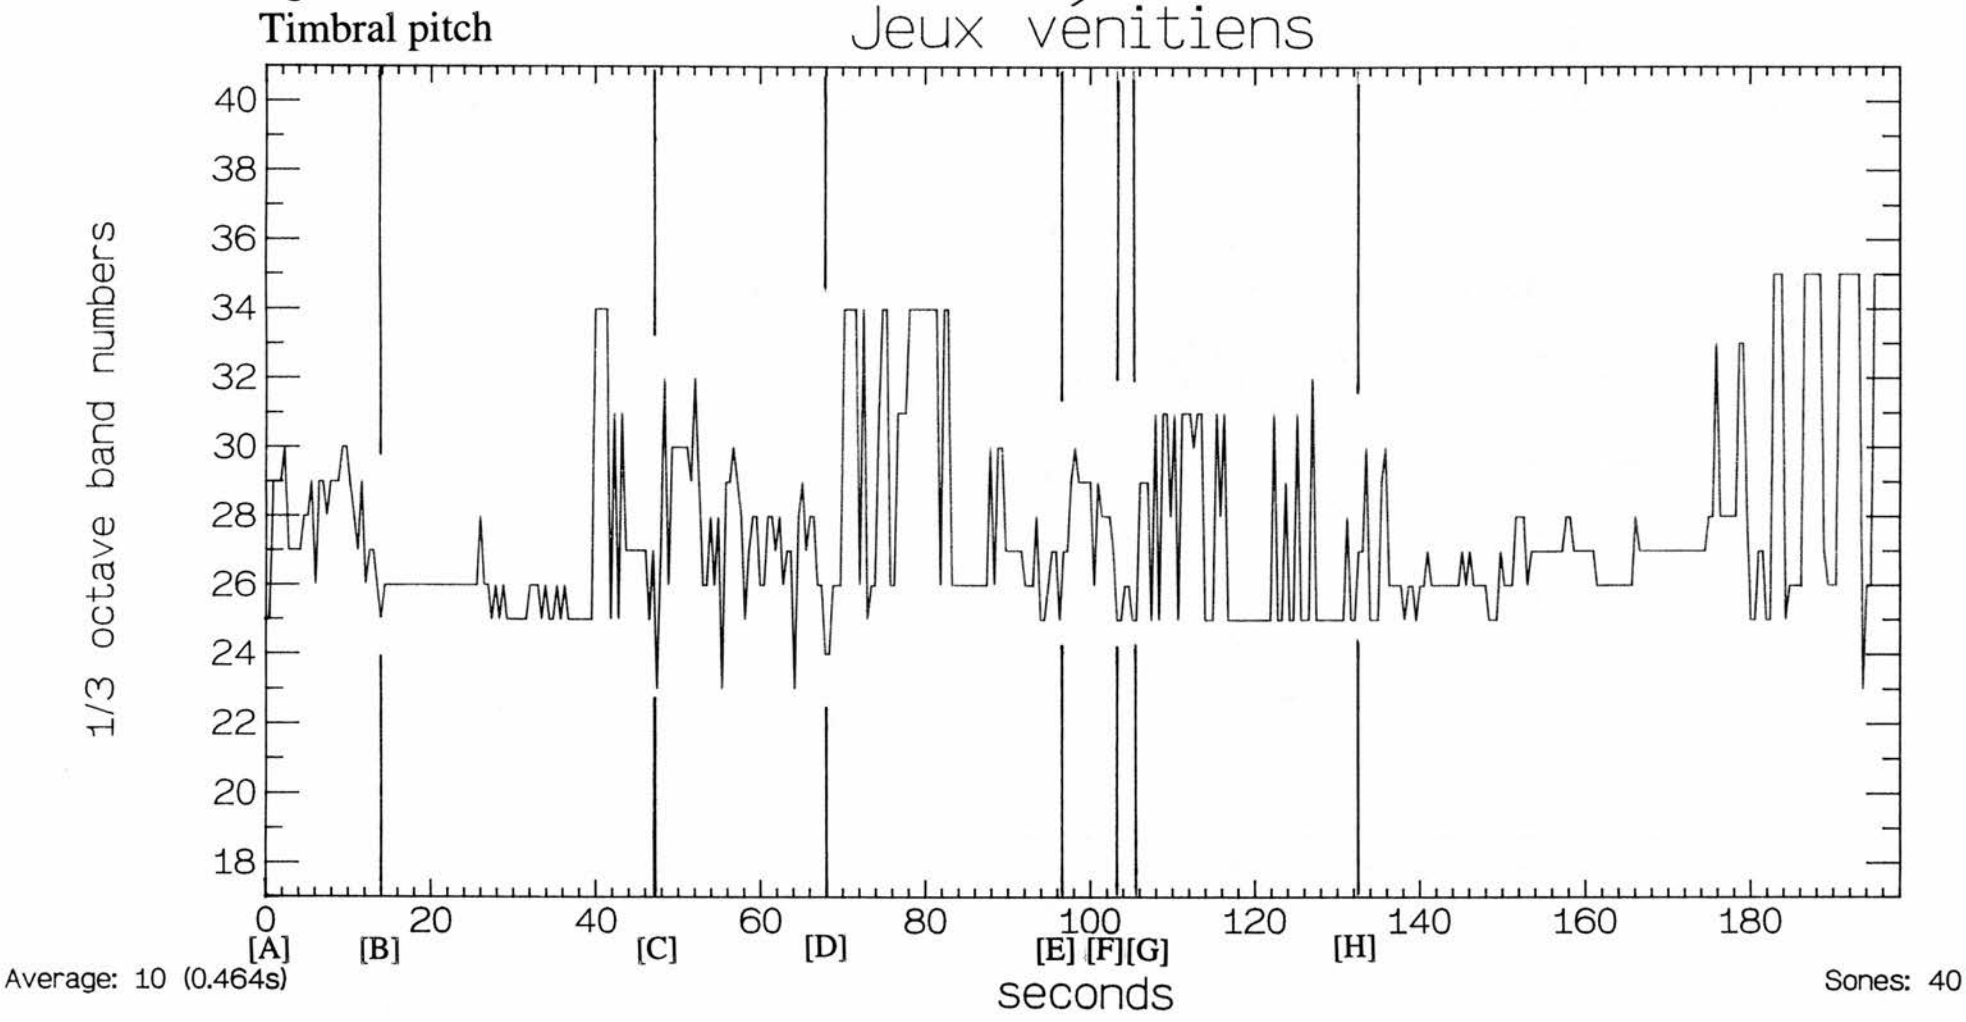

Figure 3.6
Roughness

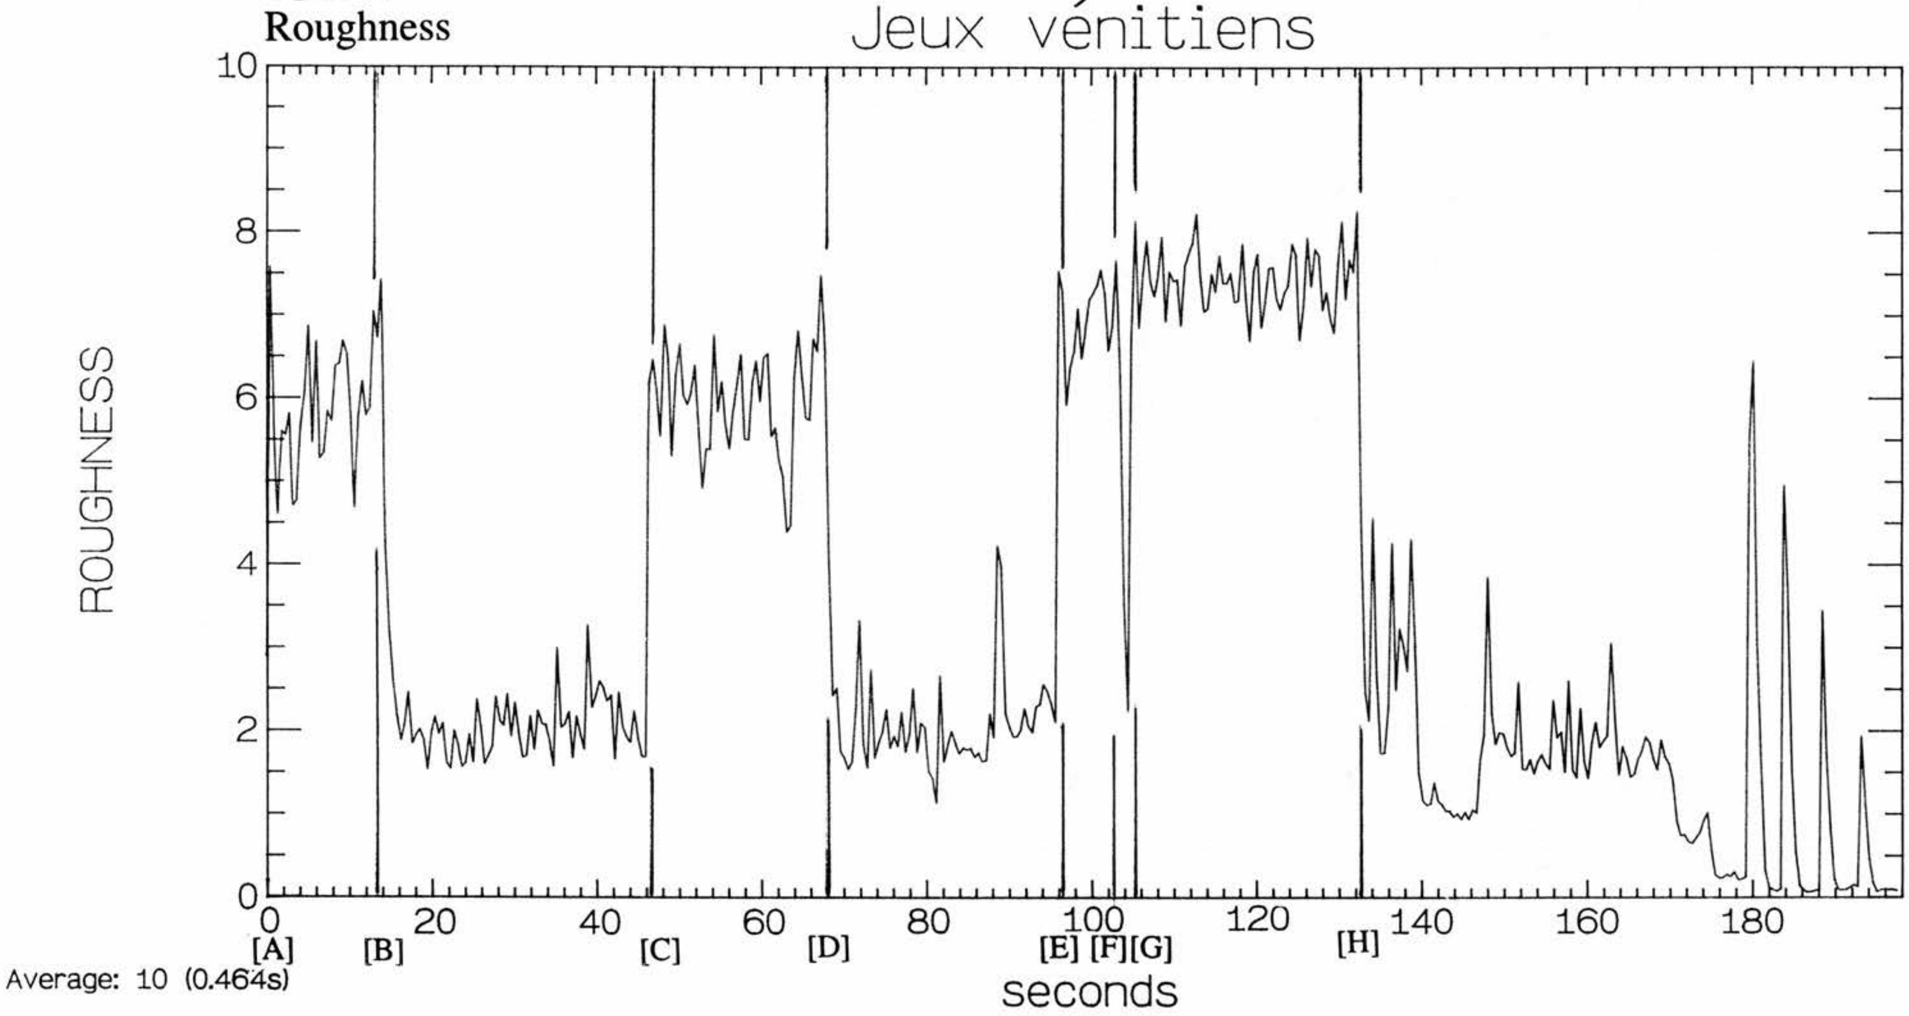

Figure 3.7
Roughness

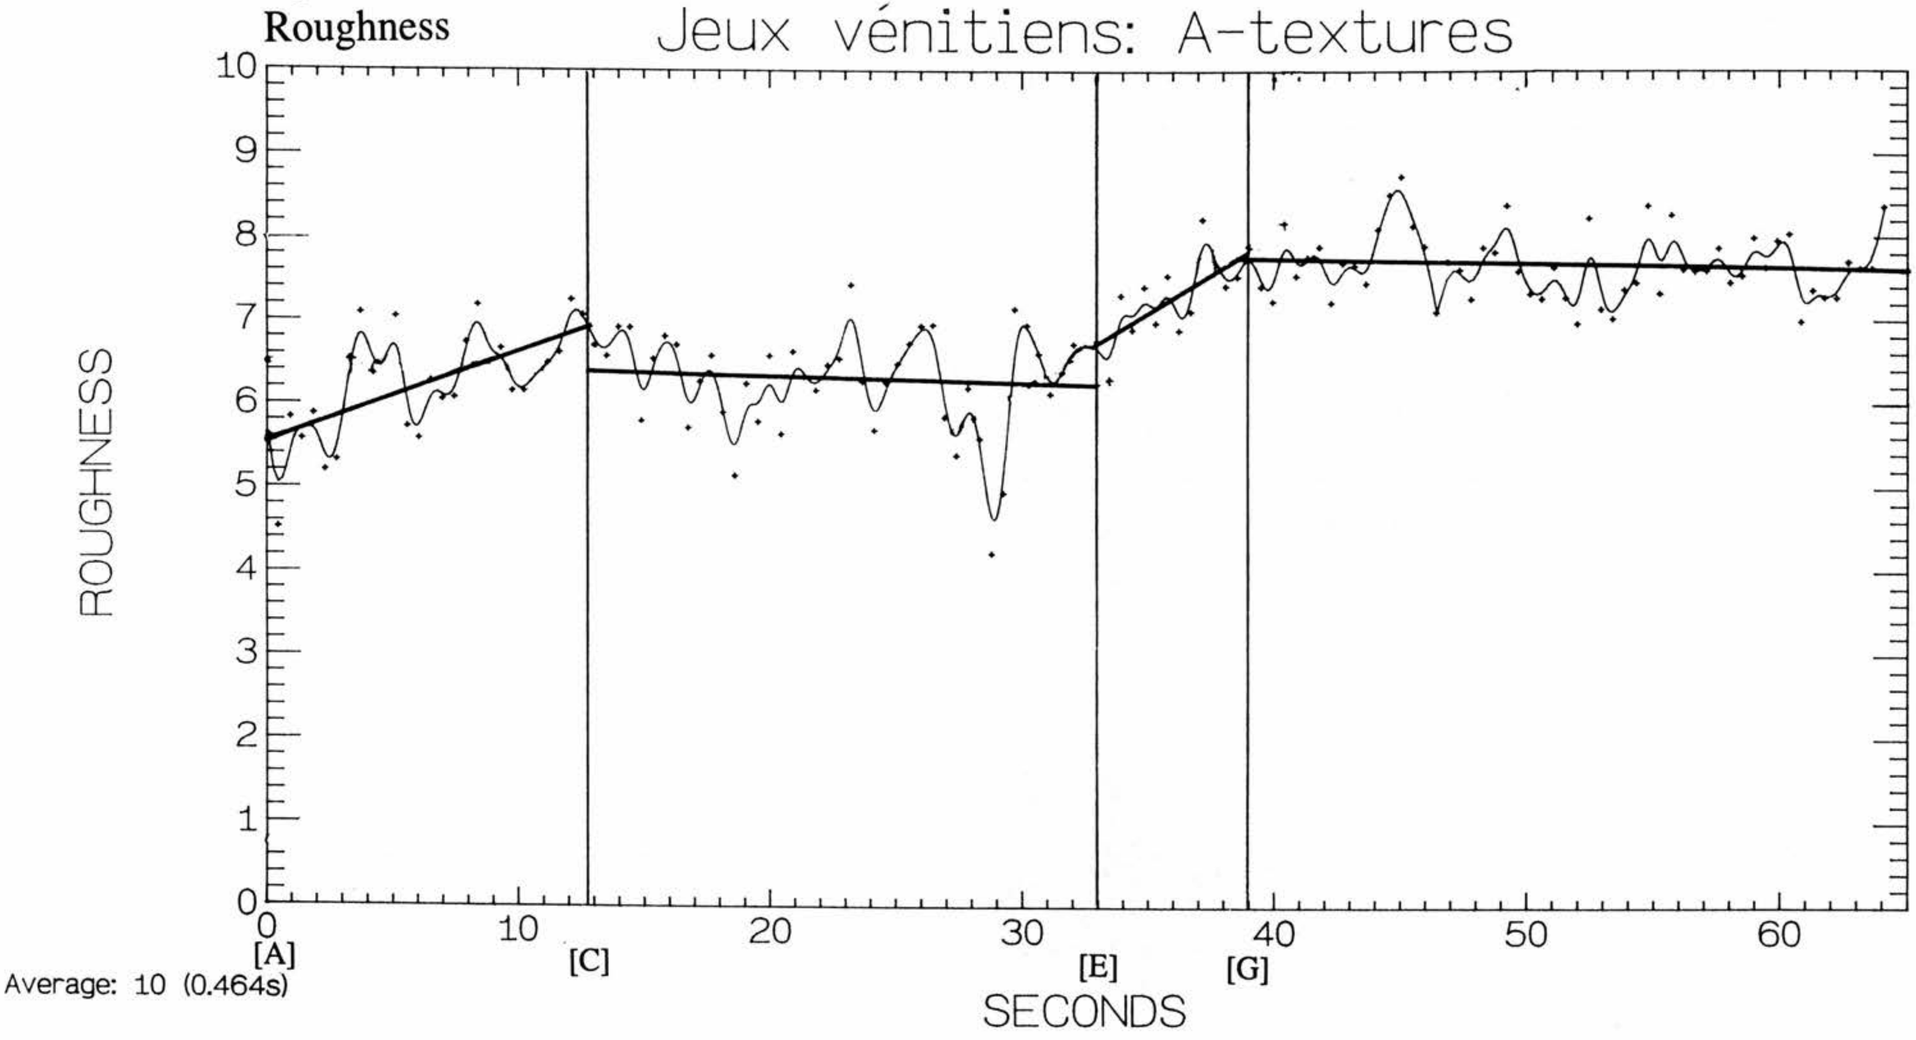

Figure 3.8

JEUX VÉNITIENS

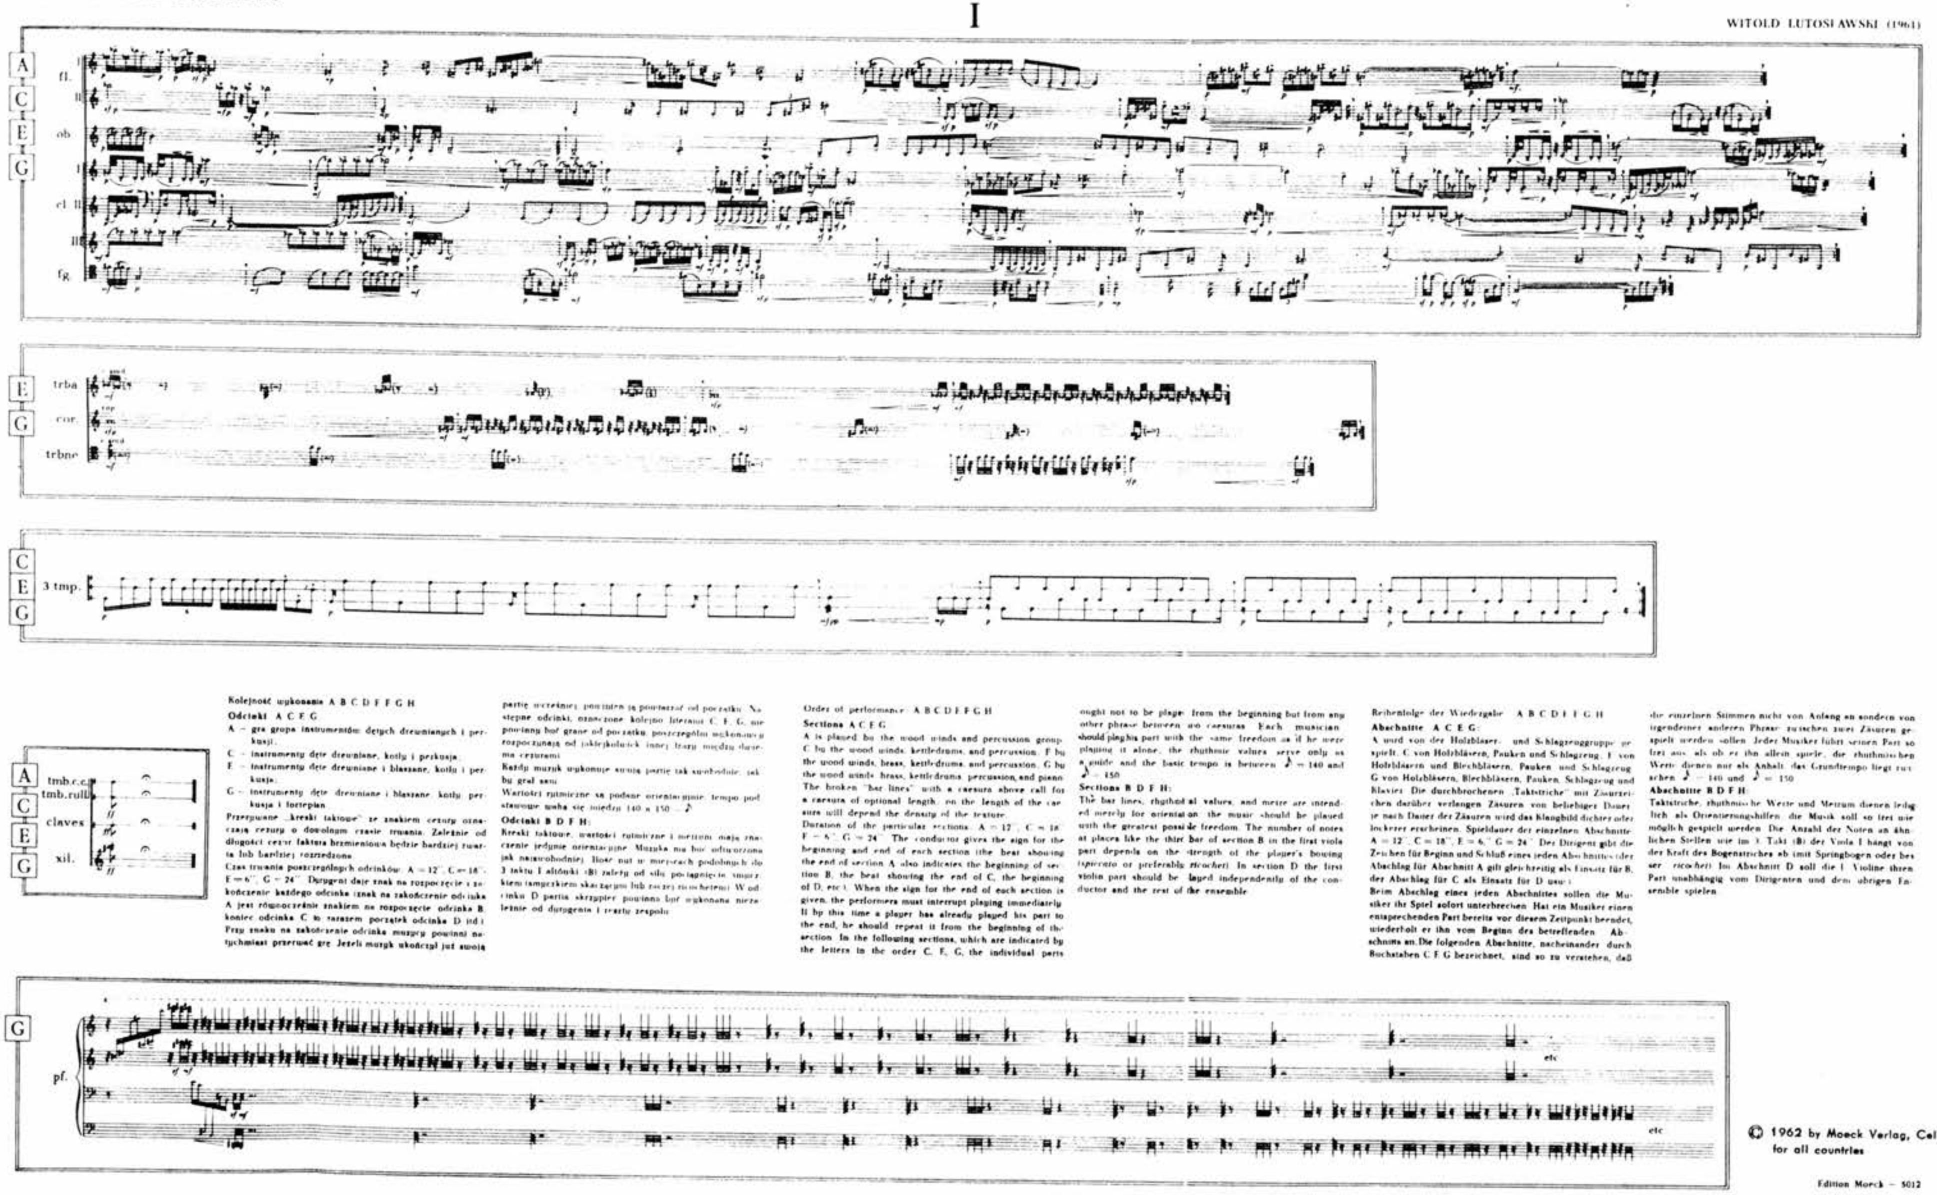

Kolejność wykonania A B C D E F G H

Odcinki A C E G

A – gra grupa instrumentów dętych drewnianych i perkusji;
C – instrumenty dęte drewniane, kotły i perkusja;
E – instrumenty dęte drewniane i blaszane, kotły i perkusja;
G – instrumenty dęte drewniane i blaszane, kotły, perkusja i fortepian.

Przerywane „kreski taktowe" ze znakiem cezury oznaczają cezury o dowolnym czasie trwania. Zależnie od długości cezur faktura brzmieniowa będzie bardziej zwarta lub bardziej rozrzedzona.

Czas trwania poszczególnych odcinków: A = 12", C = 18", E = 6", G = 24". Dyrygent daje znak na rozpoczęcie i zakończenie każdego odcinka (znak na zakończenie odcinka A jest równocześnie znakiem na rozpoczęcie odcinka B, koniec odcinka C to zarazem początek odcinka D itd.). Przy znaku na zakończenie odcinka muzycy powinni natychmiast przerwać grę. Jeżeli muzyk ukończył już swoją partię wcześniej, powinien ją powtarzać od początku. Następne odcinki, oznaczone kolejno literami C, E, G, nie powinny być grane od początku, poszczególni wykonawcy rozpoczynają od jakiejkolwiek innej frazy między dwiema cezurami.

Każdy muzyk wykonuje swoją partię tak swobodnie, jak by grał sam.

Wartości rytmiczne są podane orientacyjnie, tempo podstawowe waha się między 140 a 150 – ♪

Odcinki B D F H:

Kreski taktowe, wartości rytmiczne i metrum mają znaczenie jedynie orientacyjne. Muzyka ma być odtworzona jak najswobodniej. Ilość nut w miejscach podobnych do 3 taktu I altówki (B) zależy od siły pociągnięcia smyczkiem (smyczkiem skaczącym lub raczej ricochetem). W odcinku D partia skrzypiec powinna być wykonana niezależnie od dyrygenta i reszty zespołu.

Order of performance: A B C D E F G H

Sections A C E G

A is played by the wood winds and percussion group, C by the wood winds, kettledrums, and percussion, E by the wood winds, brass, kettledrums, and percussion, G by the wood winds, brass, kettledrums, percussion, and piano. The broken "bar lines" with a caesura above call for a caesura of optional length; on the length of the caesura will depend the density of the texture.

Duration of the particular sections: A = 12", C = 18", E = 6", G = 24". The conductor gives the sign for the beginning and end of each section (the beat showing the end of section A also indicates the beginning of section B, the beat showing the end of C, the beginning of D, etc.). When the sign for the end of each section is given, the performers must interrupt playing immediately. If by this time a player has already played his part to the end, he should repeat it from the beginning of the section. In the following sections, which are indicated by the letters in the order C, E, G, the individual parts ought not to be played from the beginning but from any other phrase between two caesuras. Each musician should play his part with the same freedom as if he were playing it alone; the rhythmic values serve only as a guide and the basic tempo is between ♪ = 140 and ♪ = 150.

Sections B D F H:

The bar lines, rhythmical values, and metre are intended merely for orientation; the music should be played with the greatest possible freedom. The number of notes at places like the third bar of section B in the first viola part depends on the strength of the player's bowing (spiccato or preferably ricochet). In section D the first violin part should be played independently of the conductor and the rest of the ensemble.

Reihenfolge der Wiedergabe: A B C D E F G H

Abschnitte A C E G:

A wird von der Holzbläser- und Schlagzeuggruppe gespielt, C von Holzbläsern, Pauken und Schlagzeug, E von Holzbläsern und Blechbläsern, Pauken und Schlagzeug, G von Holzbläsern, Blechbläsern, Pauken, Schlagzeug und Klavier. Die durchbrochenen „Taktstriche" mit Zäsurzeichen darüber verlangen Zäsuren von beliebiger Dauer; je nach Dauer der Zäsuren wird das Klangbild dichter oder lockerer erscheinen. Spieldauer der einzelnen Abschnitte: A = 12", C = 18", E = 6", G = 24". Der Dirigent gibt die Zeichen für Beginn und Schluß eines jeden Abschnittes (der Abschlag für Abschnitt A gilt gleichzeitig als Einsatz für B, der Abschlag für C als Einsatz für D usw.).

Beim Abschlag eines jeden Abschnittes sollen die Musiker ihr Spiel sofort unterbrechen. Hat ein Musiker einen entsprechenden Part bereits vor diesem Zeitpunkt beendet, wiederholt er ihn vom Beginn des betreffenden Abschnitts an. Die folgenden Abschnitte, nacheinander durch Buchstaben C E G bezeichnet, sind so zu verstehen, daß die einzelnen Stimmen nicht von Anfang an sondern von irgendeiner anderen Phrase zwischen zwei Zäsuren gespielt werden sollen. Jeder Musiker führt seinen Part so frei aus, als ob er ihn allein spiele; die rhythmischen Werte dienen nur als Anhalt; das Grundtempo liegt zwischen ♪ = 140 und ♪ = 150.

Abschnitte B D F H:

Taktstriche, rhythmische Werte und Metrum dienen lediglich als Orientierungshilfen; die Musik soll so frei wie möglich gespielt werden. Die Anzahl der Noten an ähnlichen Stellen wie im 3. Takt (B) der Viola I hängt von der Kraft des Bogenstriches ab (mit Springbogen oder besser ricochet). Im Abschnitt D soll die I. Violine ihren Part unabhängig vom Dirigenten und dem übrigen Ensemble spielen.

© 1962 by Moeck Verlag, Celle/FRG
for all countries

Edition Moeck – 5012

Figure 3.9
Roughness: recording 2

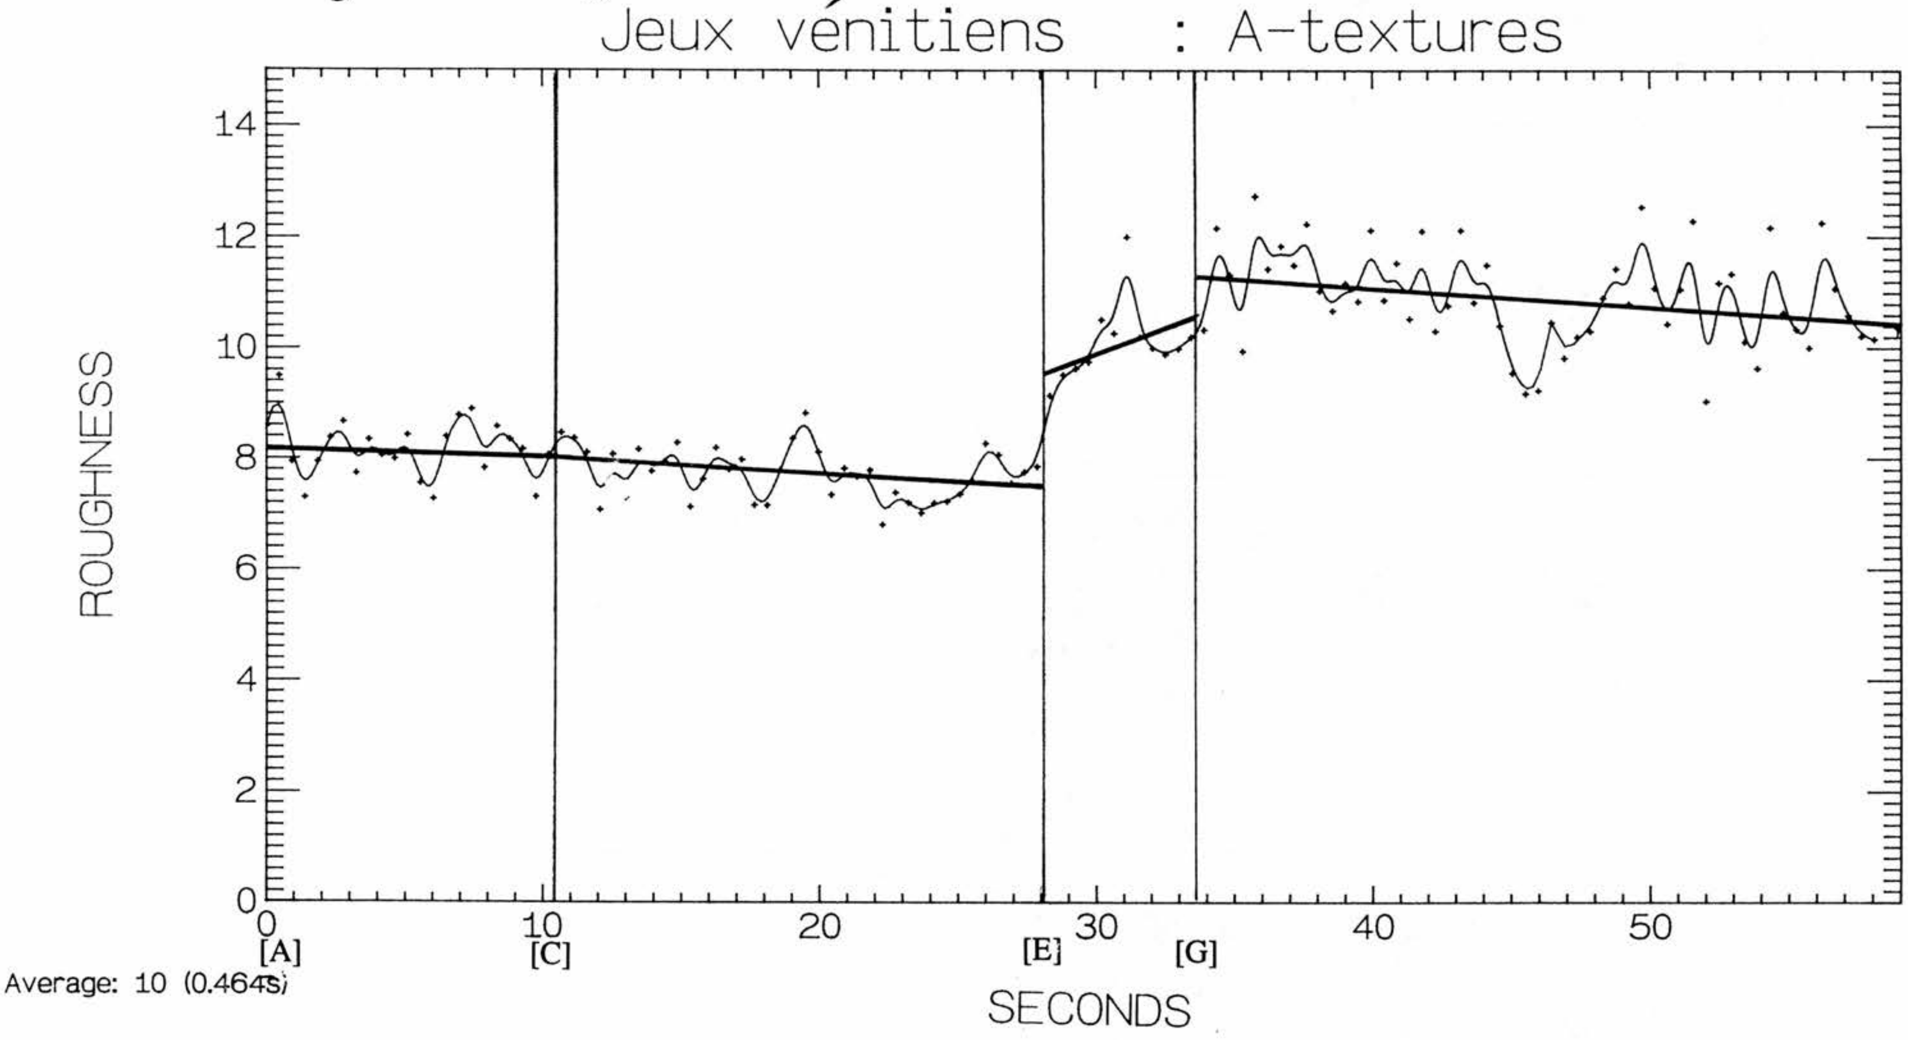

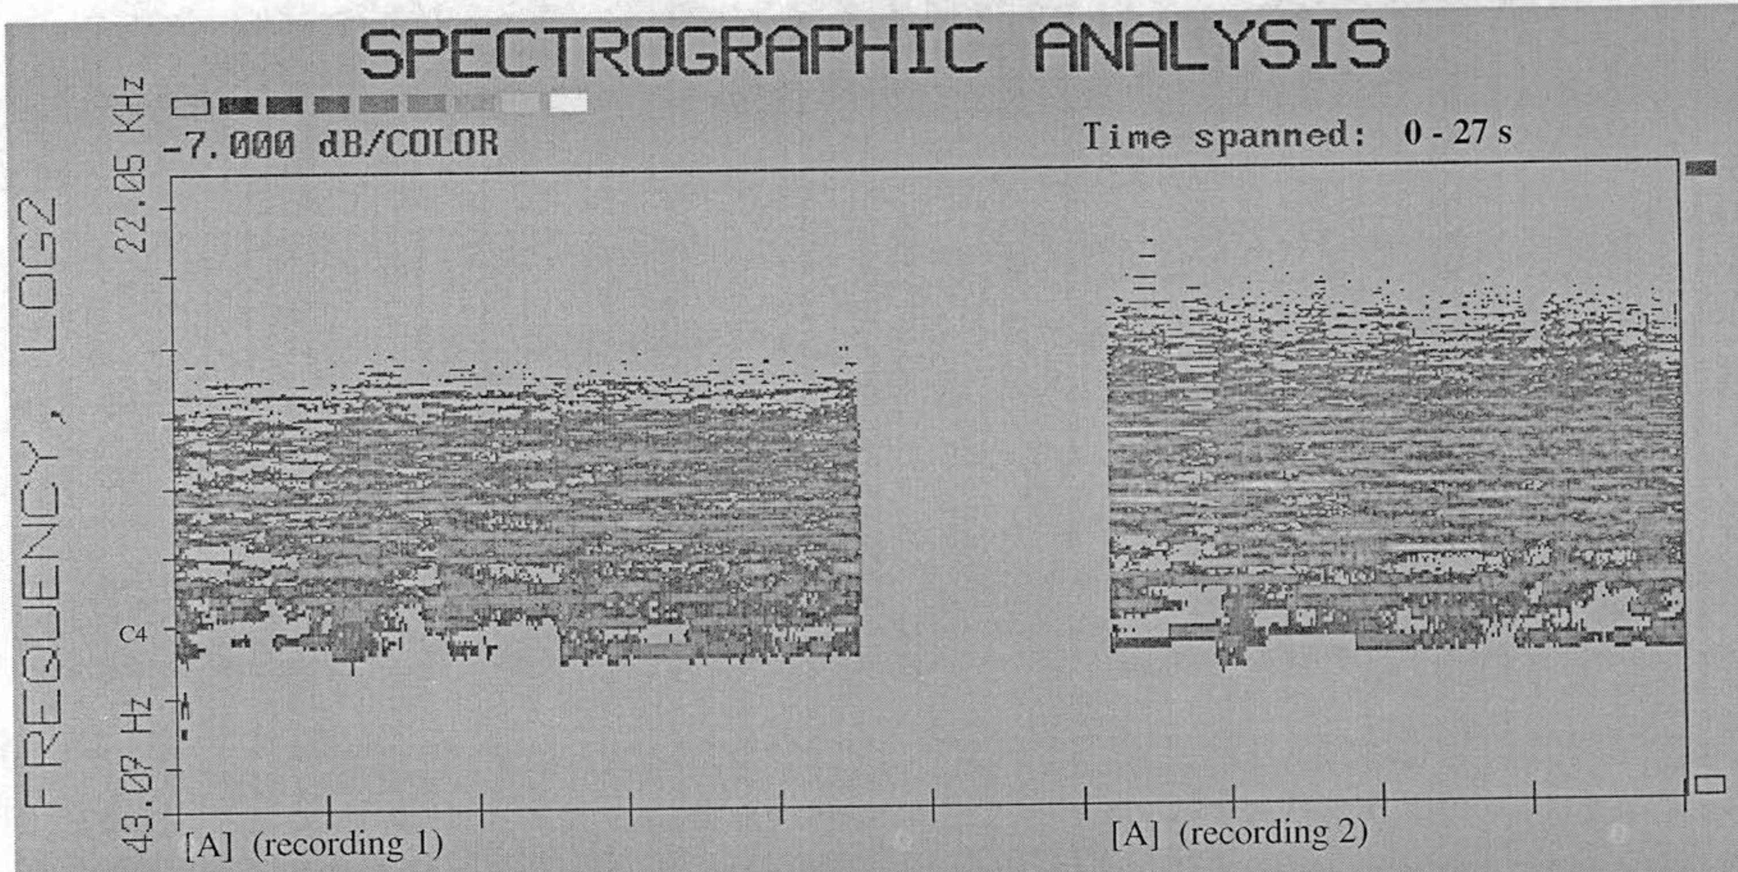

Jeux venitiens: Comparison of section [A] from recordings 1 and 2

Figure 3.10

Figure 3.11
Roughness: recording 1

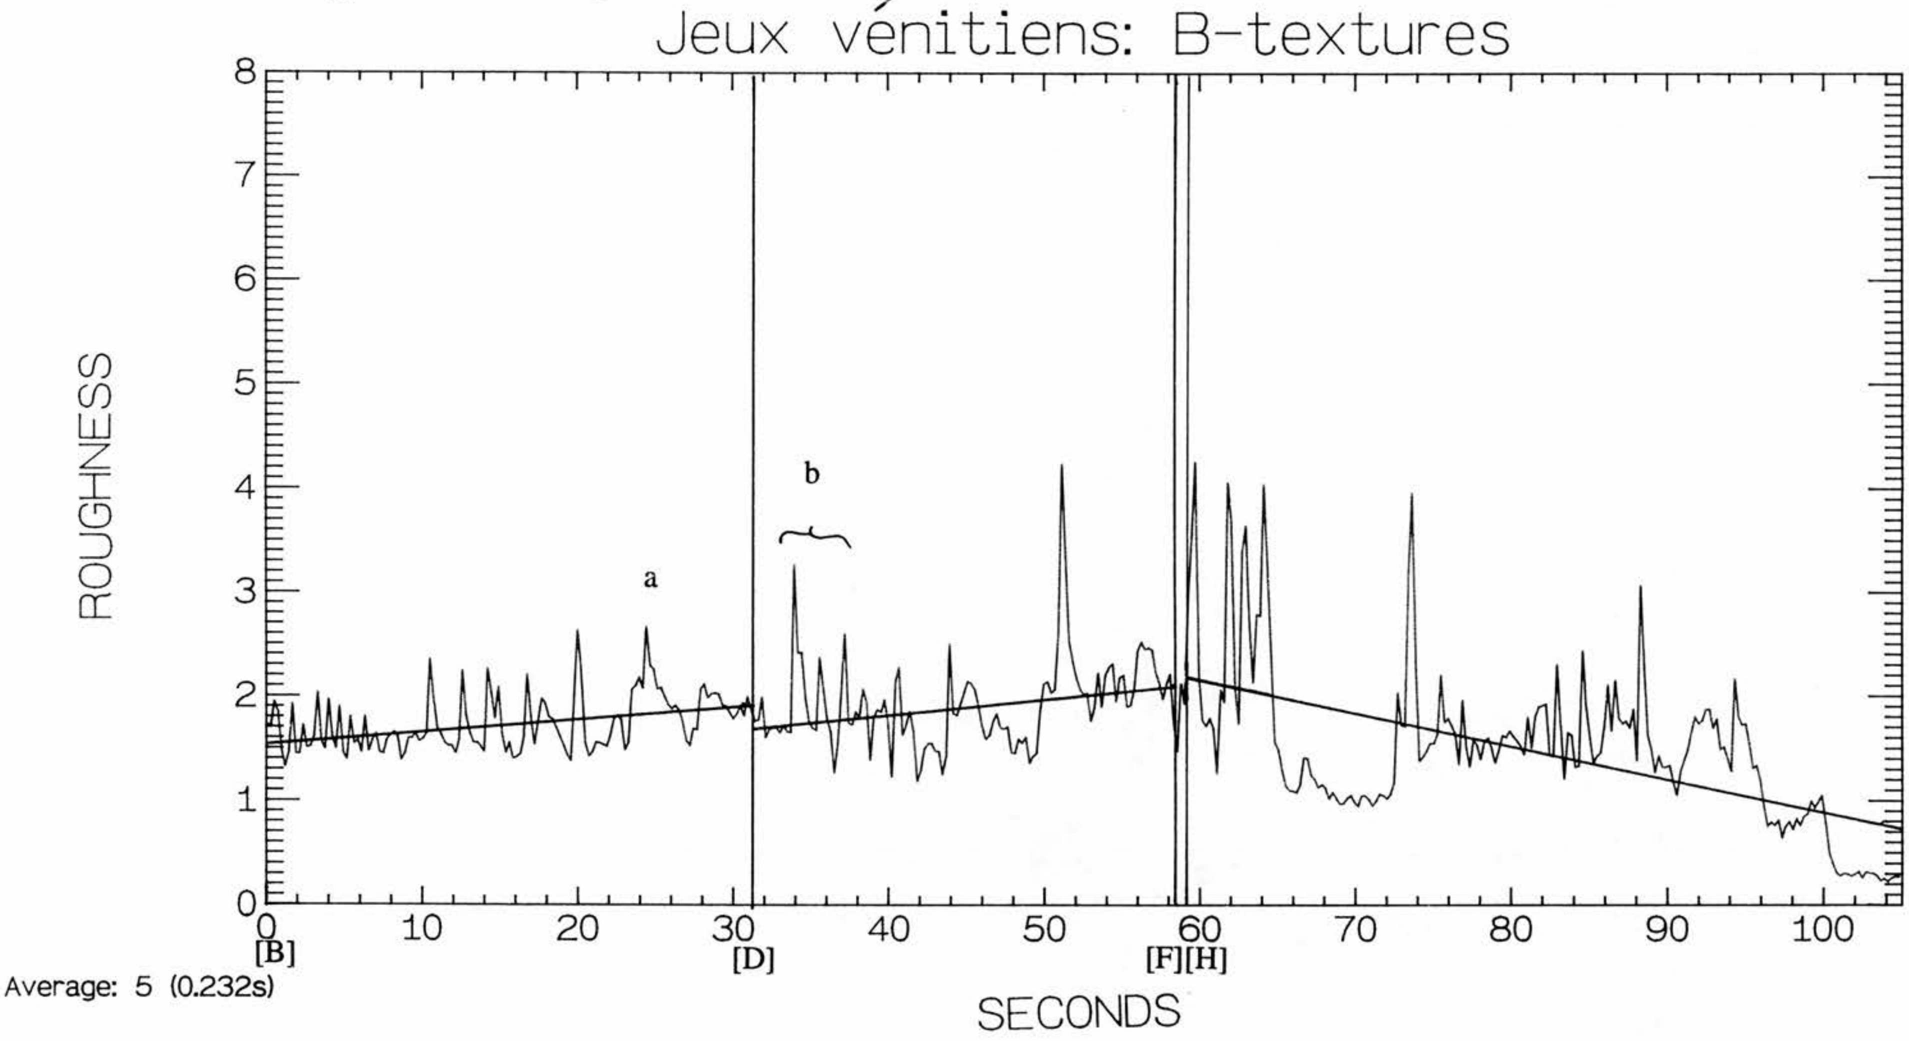

Jeux venitiens: Section [B] - pitch analysis Figure 3.12

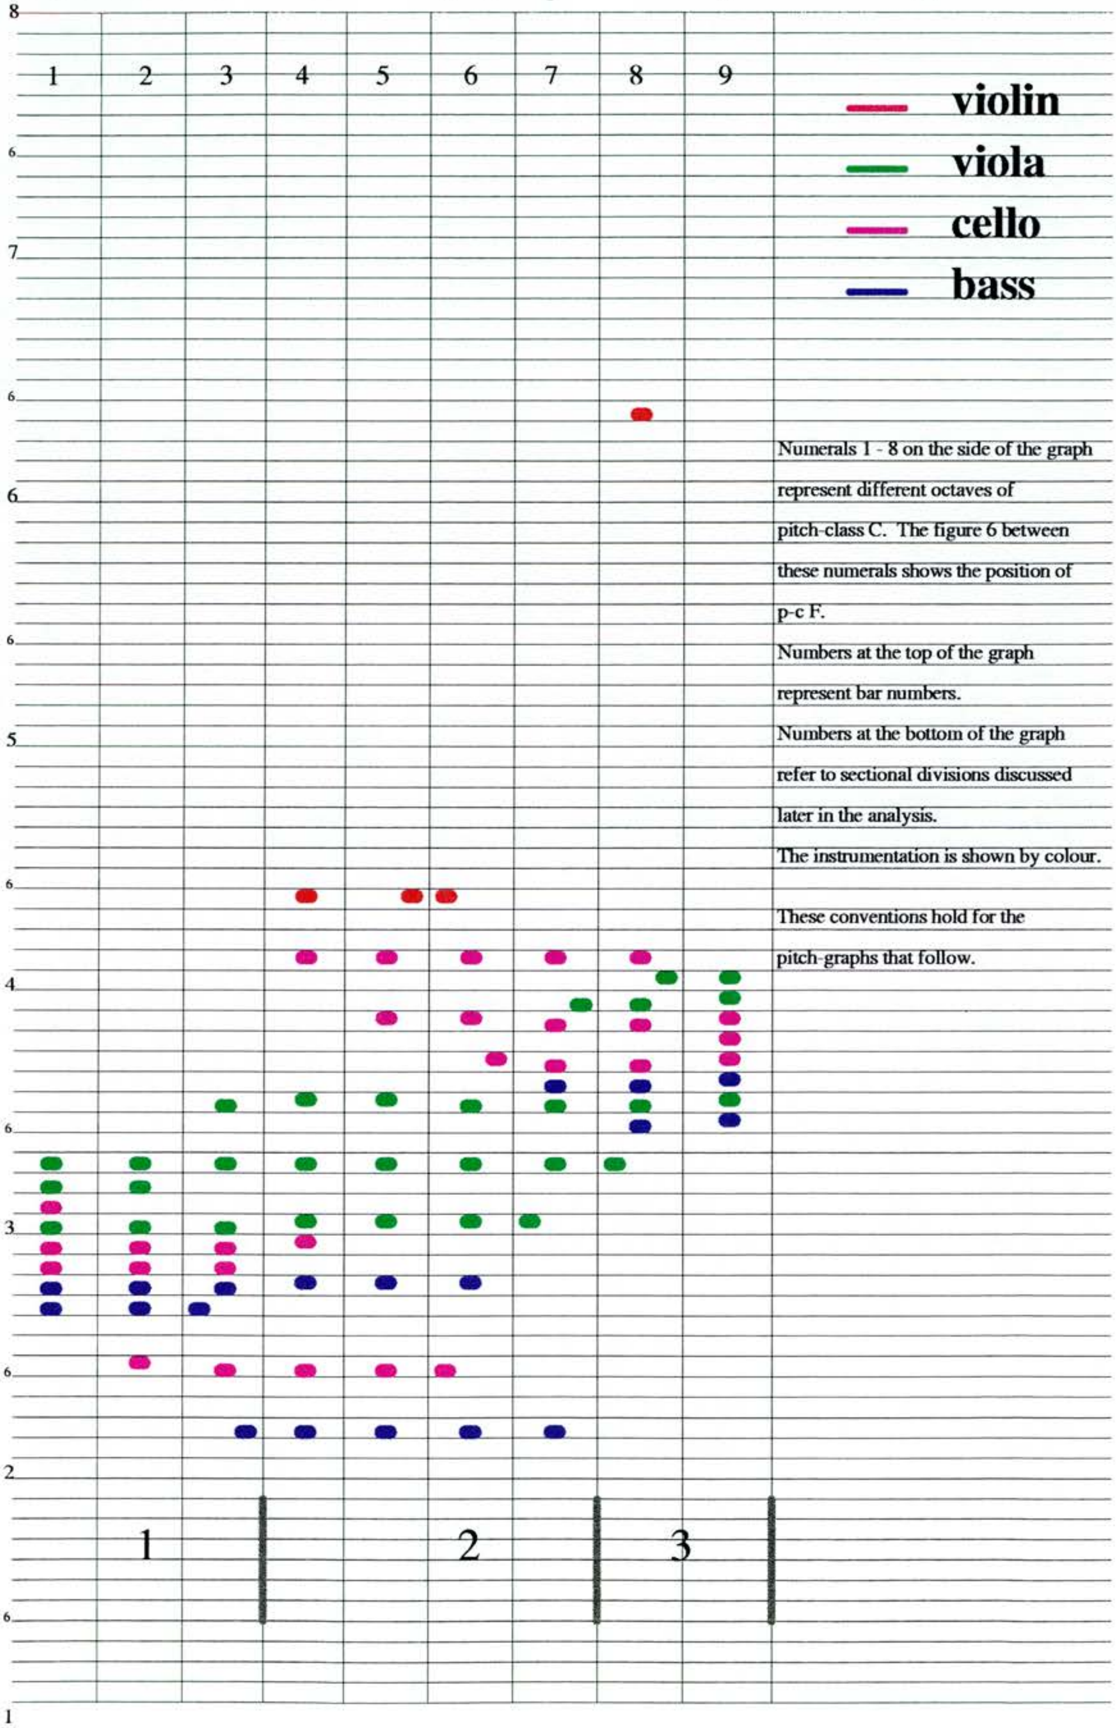

Numerals 1 - 8 on the side of the graph represent different octaves of pitch-class C. The figure 6 between these numerals shows the position of p-c F.

Numbers at the top of the graph represent bar numbers.

Numbers at the bottom of the graph refer to sectional divisions discussed later in the analysis.

The instrumentation is shown by colour.

These conventions hold for the pitch-graphs that follow.

Jeux venitiens: Section [D] [F]

Figure 3.13

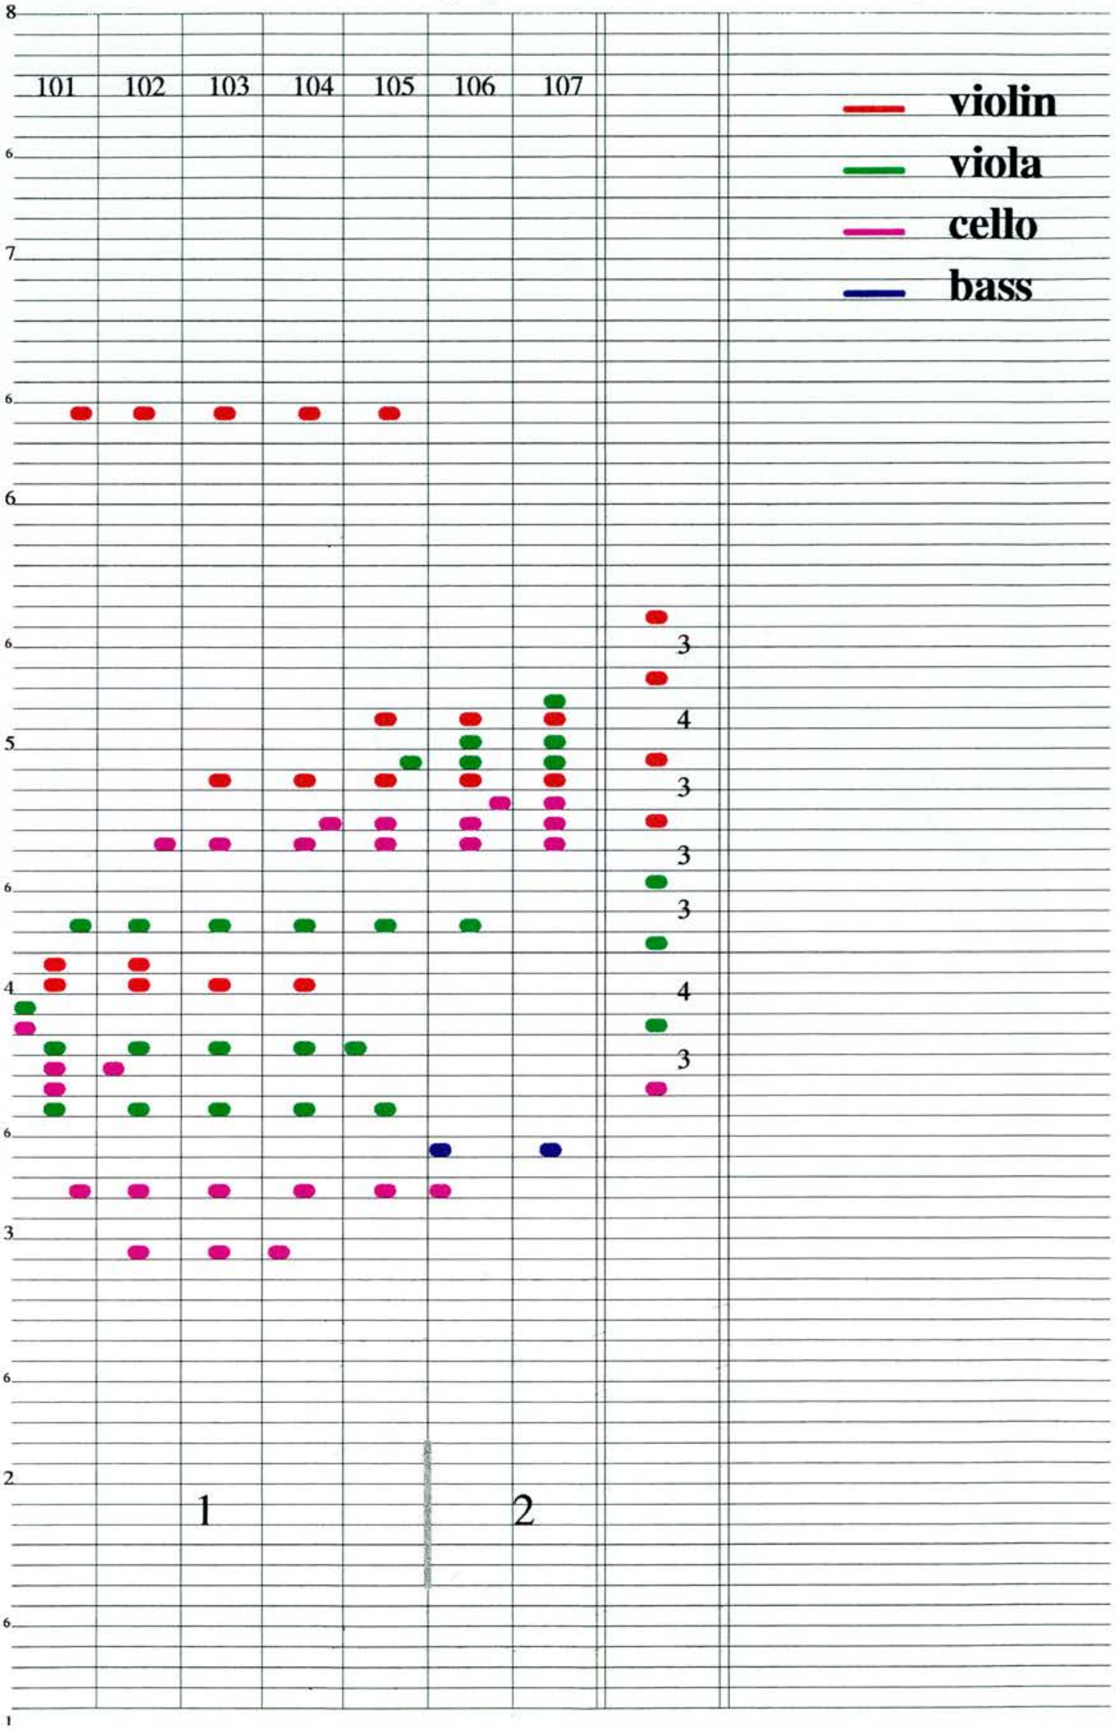

Jeux venitiens: Section [H]

Figure 3.14

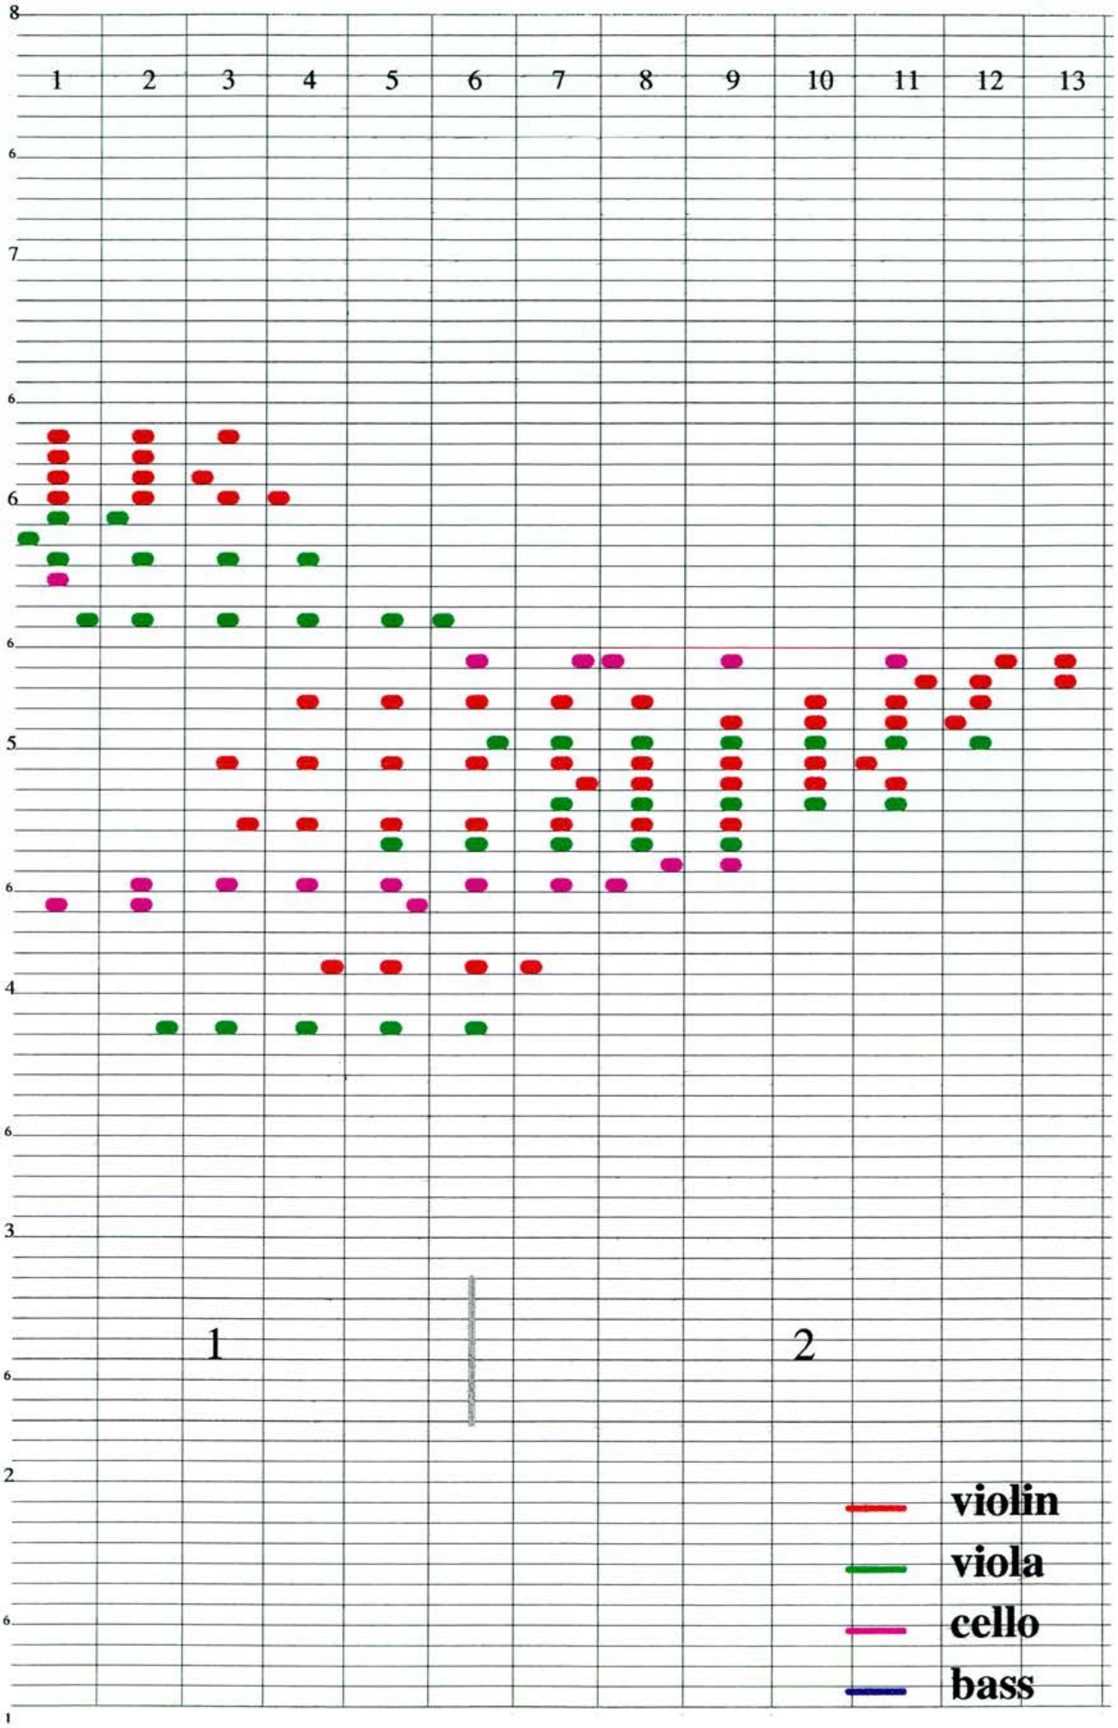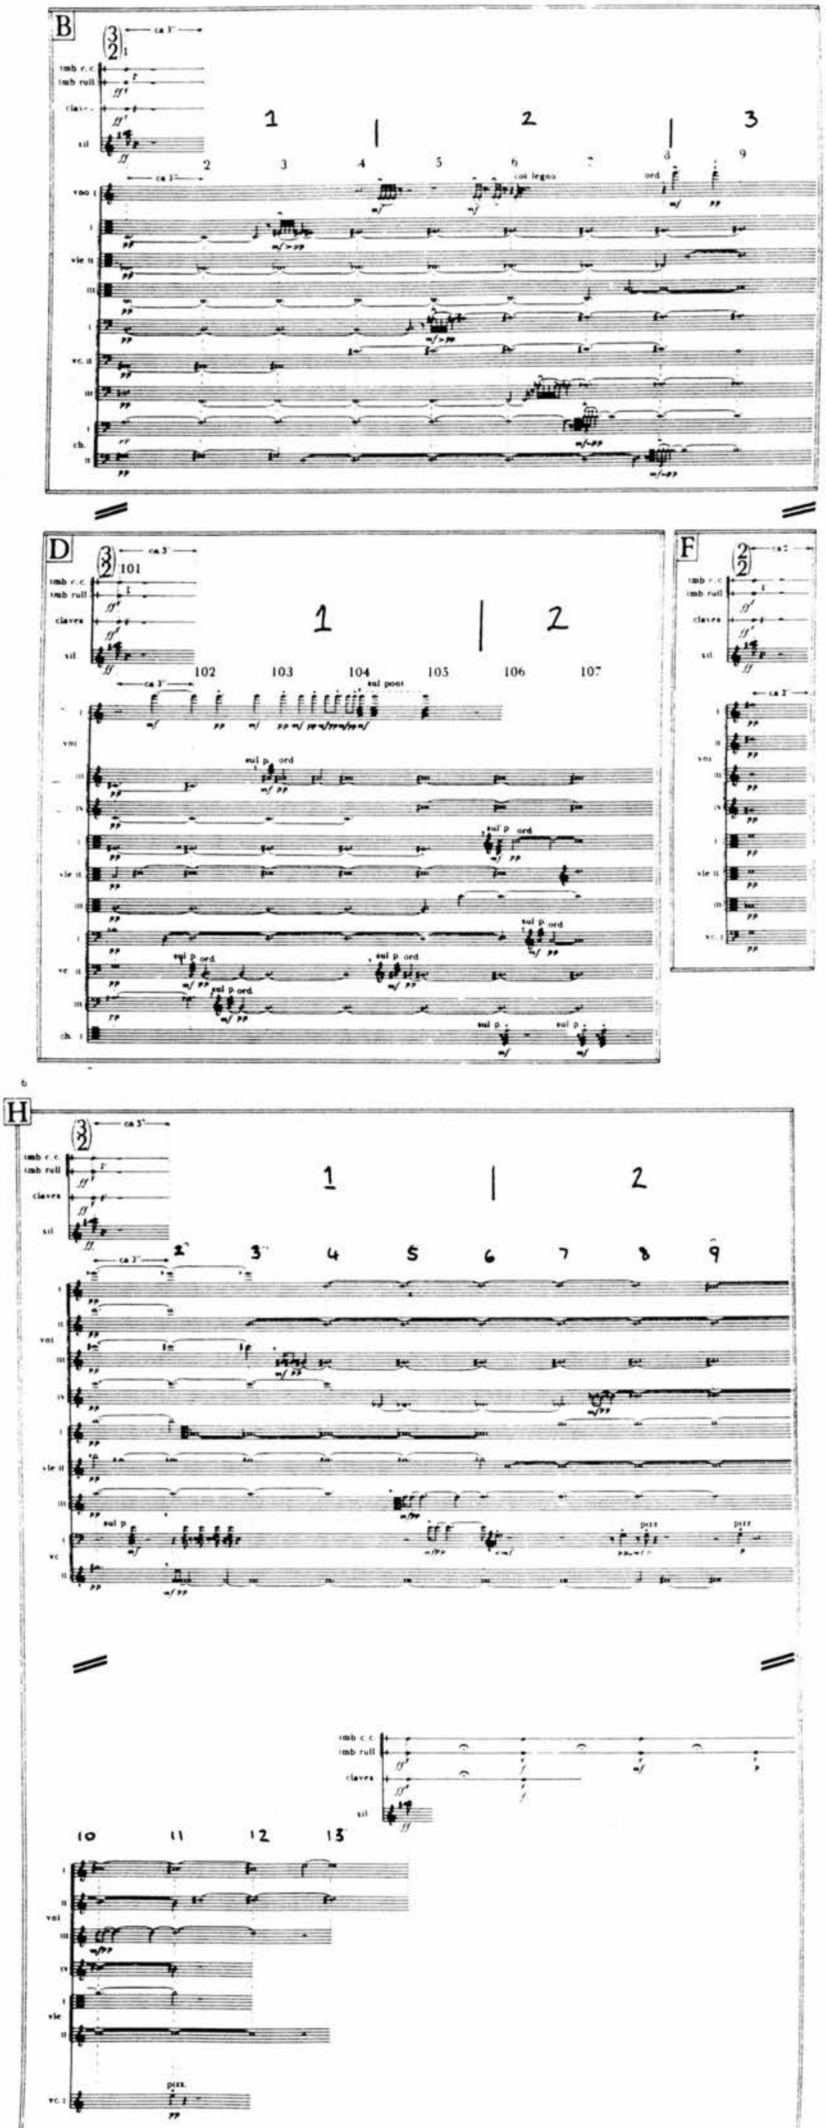


B
ca 3"
timb c.c.
timb rull
claves
xil
1 | 2 | 3
2 3 4 5 6 7 8 9
col legno
ord
vni I
vle II
vc. II
cb.
D
101
1 | 2
102 103 104 105 106 107
sul pont
sul p
ord
vni
vle II
vc II
cb. I
F
ca 2"
H
1 | 2
2 3 4 5 6 7 8 9
vni
vle II
vc
pizz
10 11 12 13
vle
vc I
pizz.

Figure 3.15

Figure 3.16
Roughness: recording 2

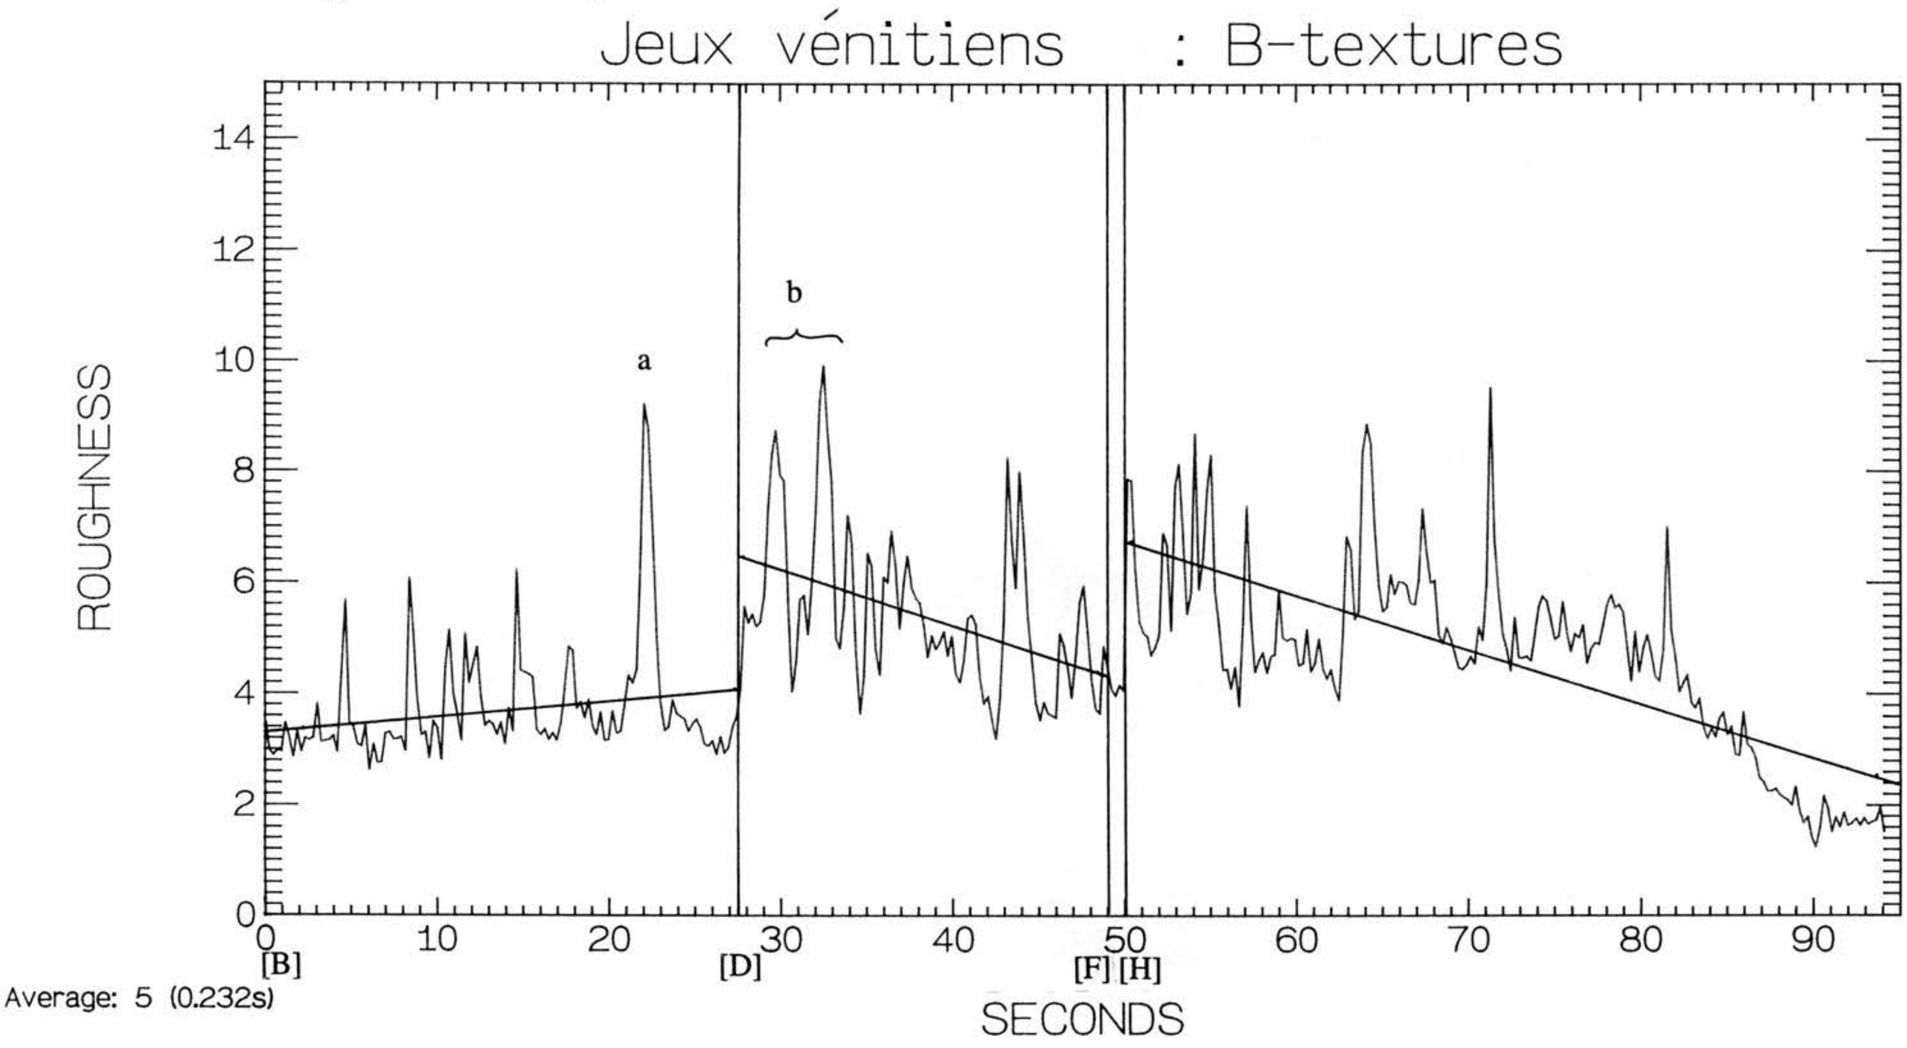

Figure 3.17a

Roughness: recording 1

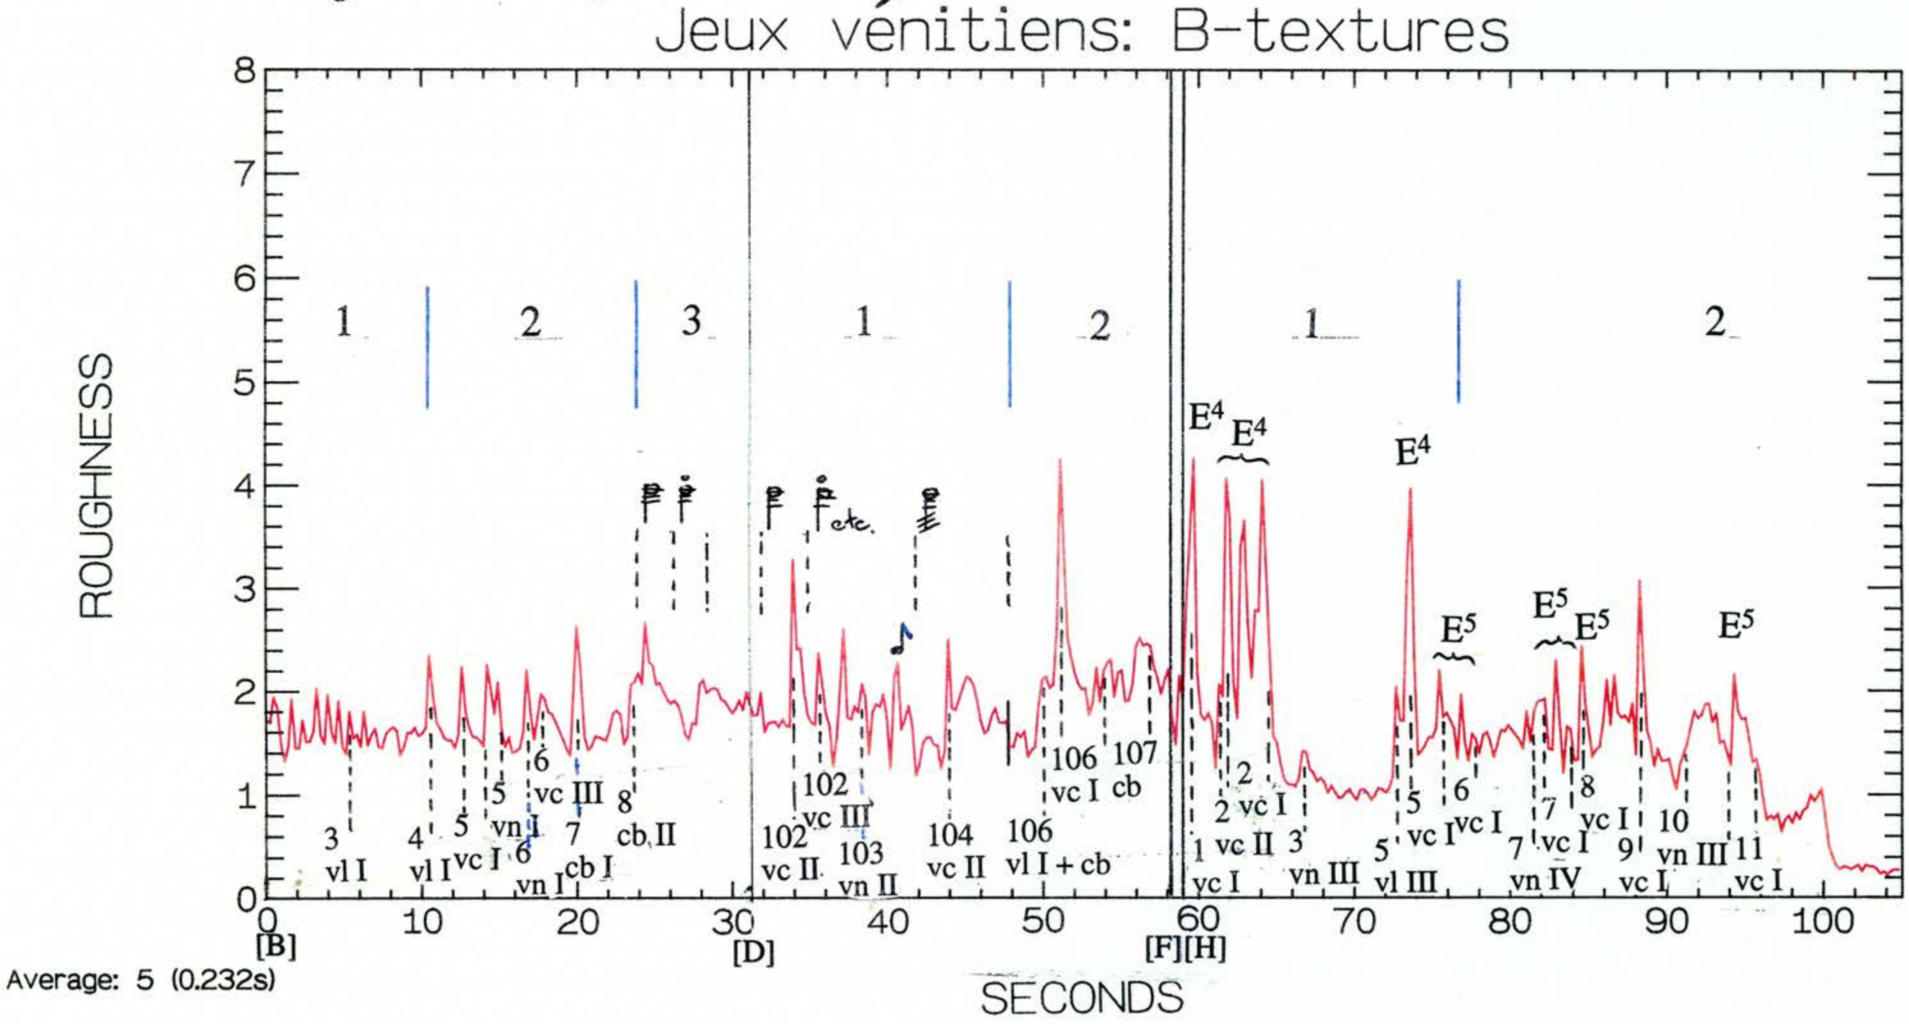

Figure 3.17b

Roughness: recording 2

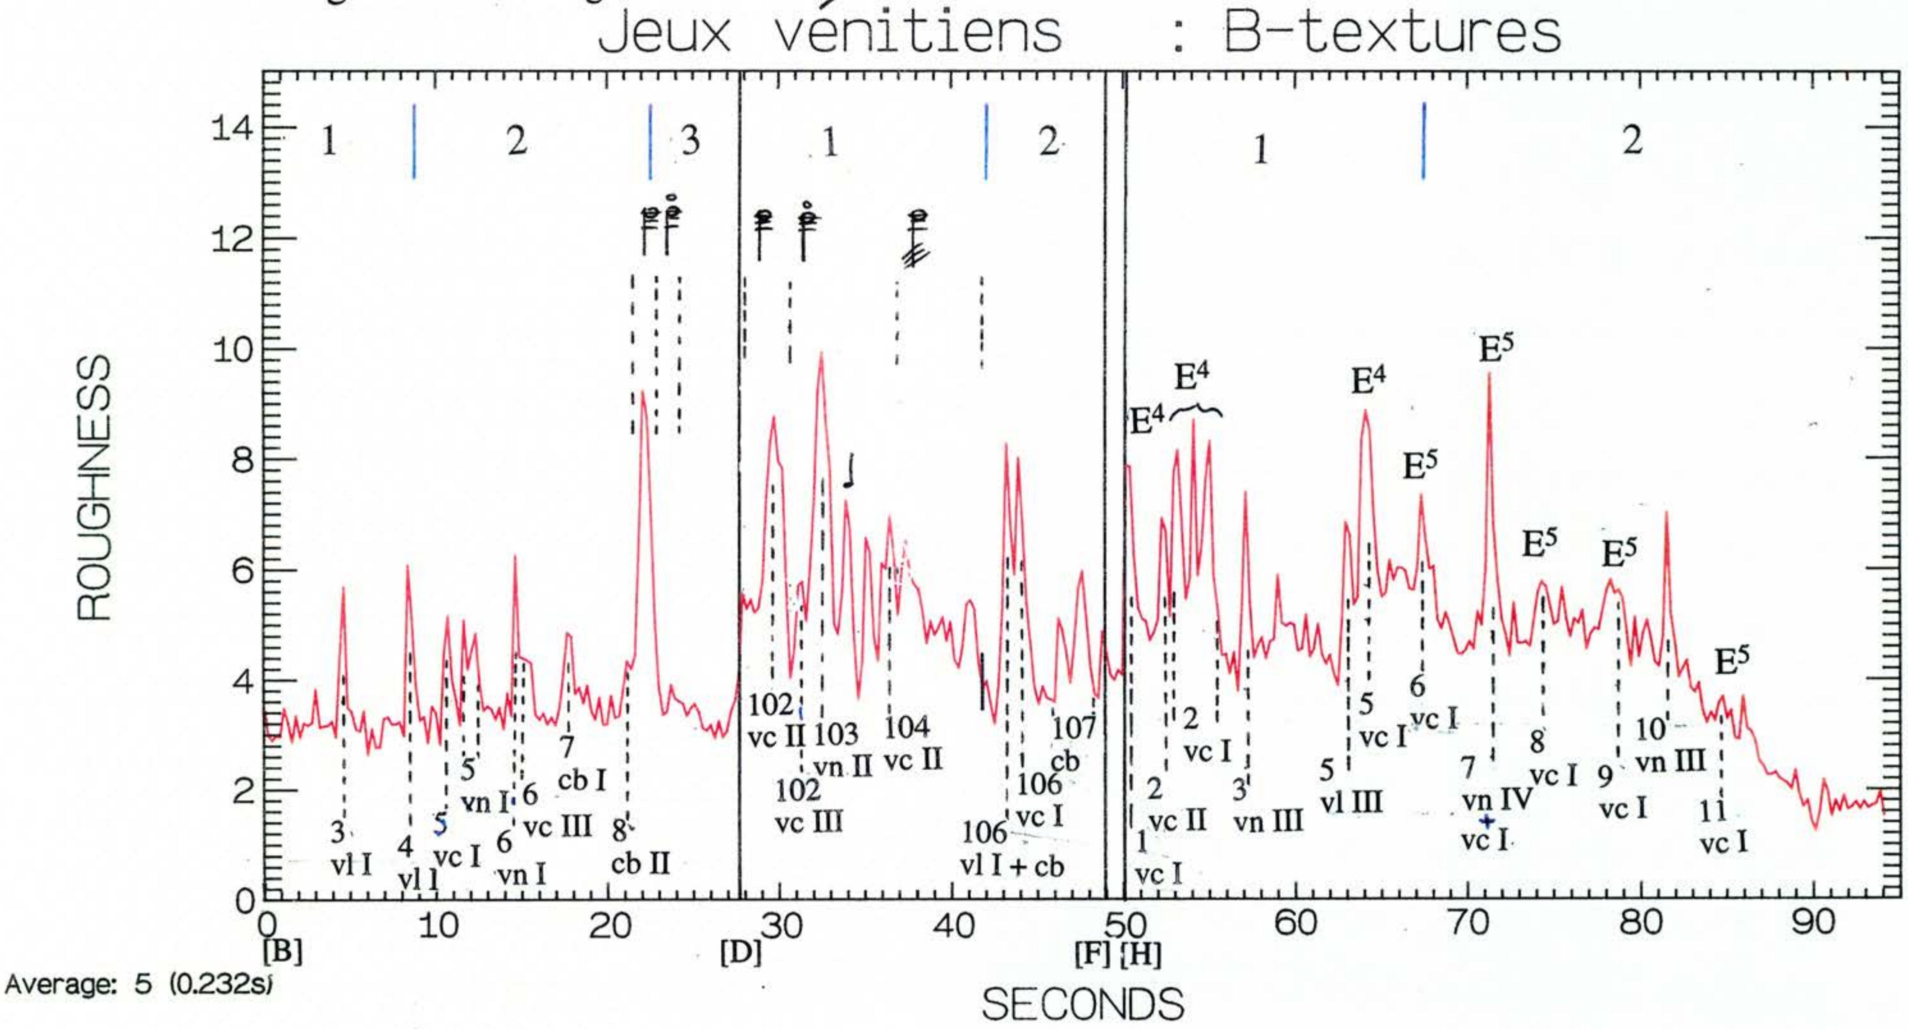

Figure 3.18
Sharpness: recording 1

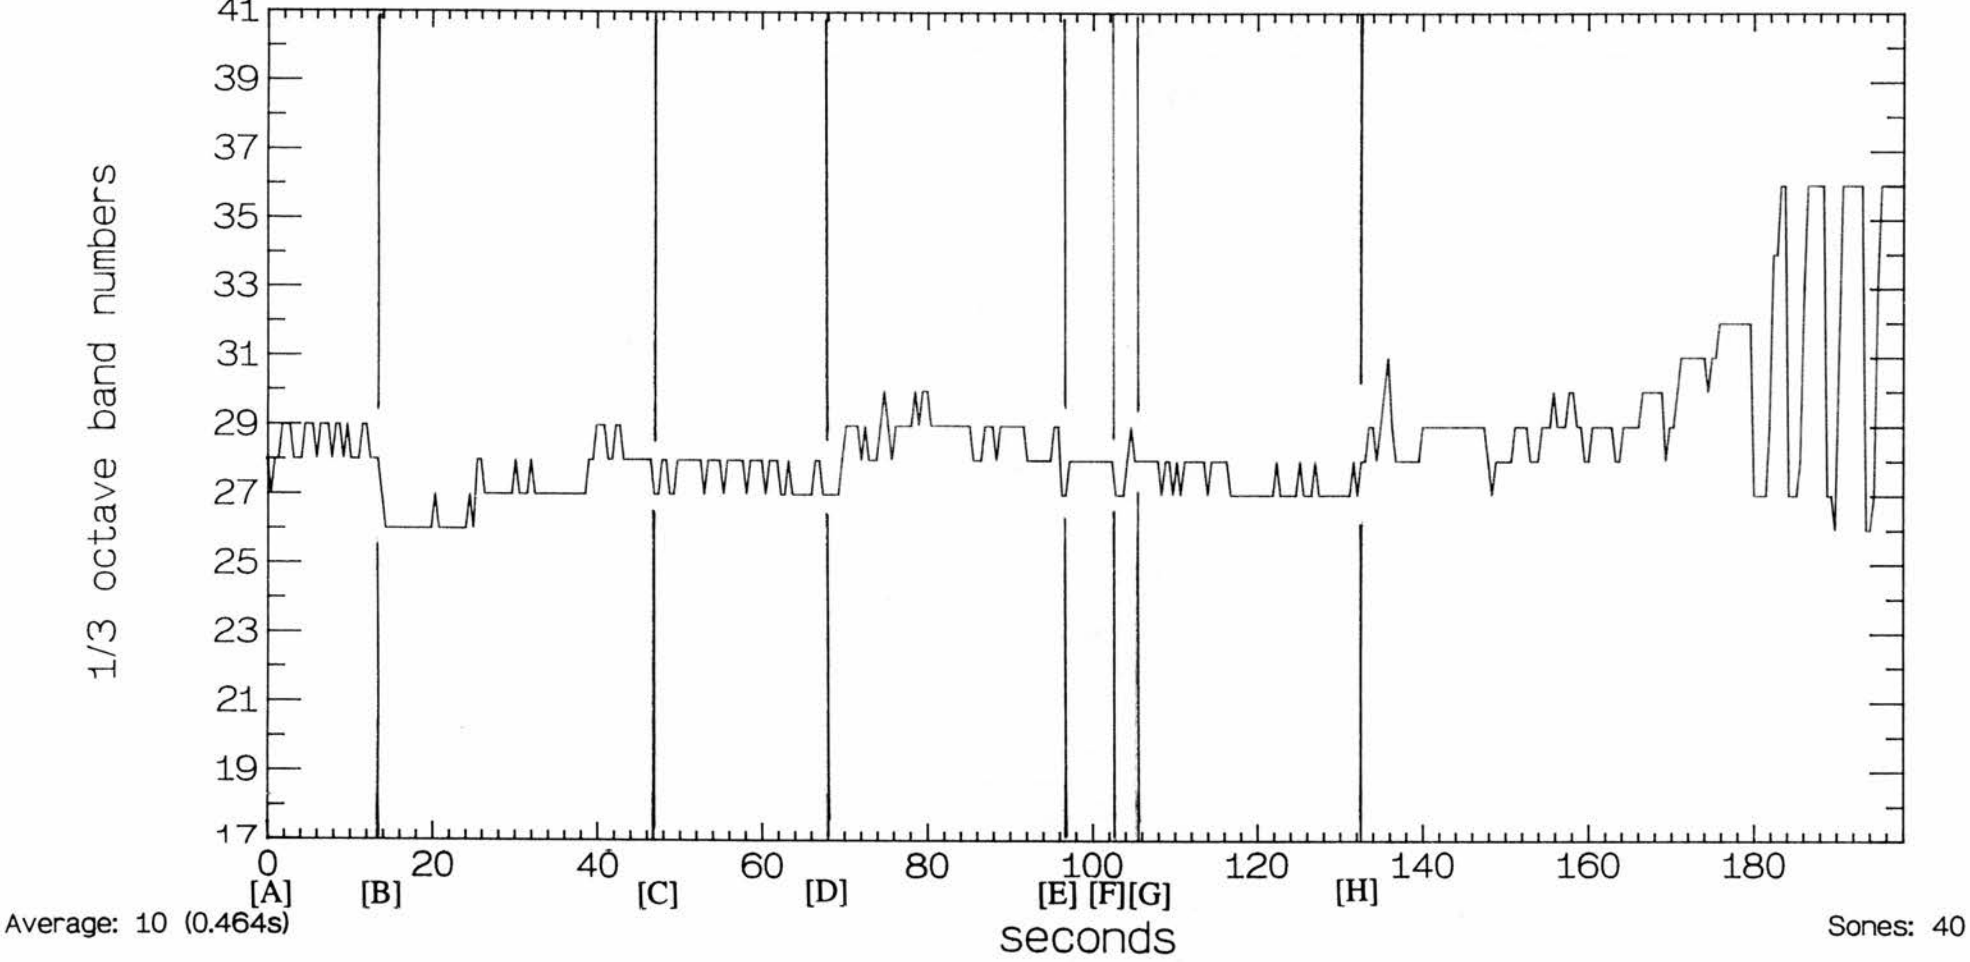

Figure 3.19
Sharpness: recording 1

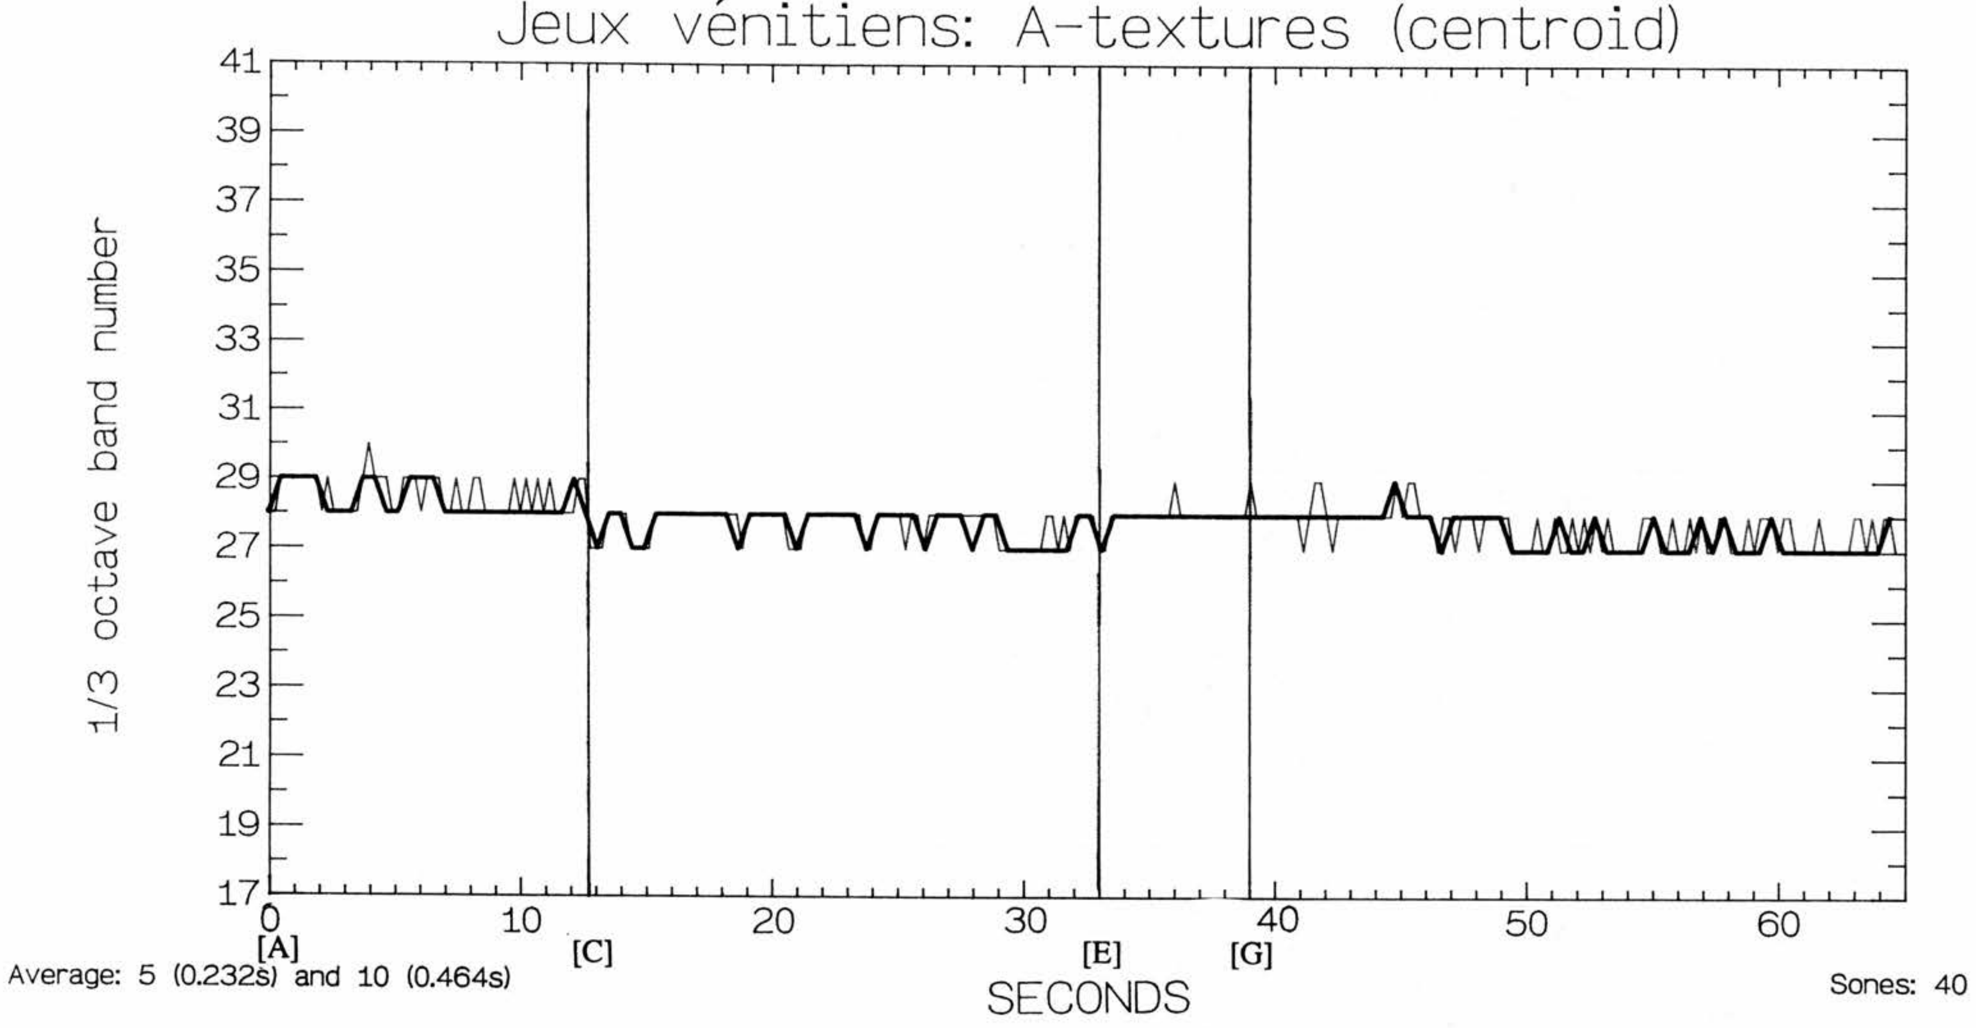

Figure 3.20
Sharpness: recording 2

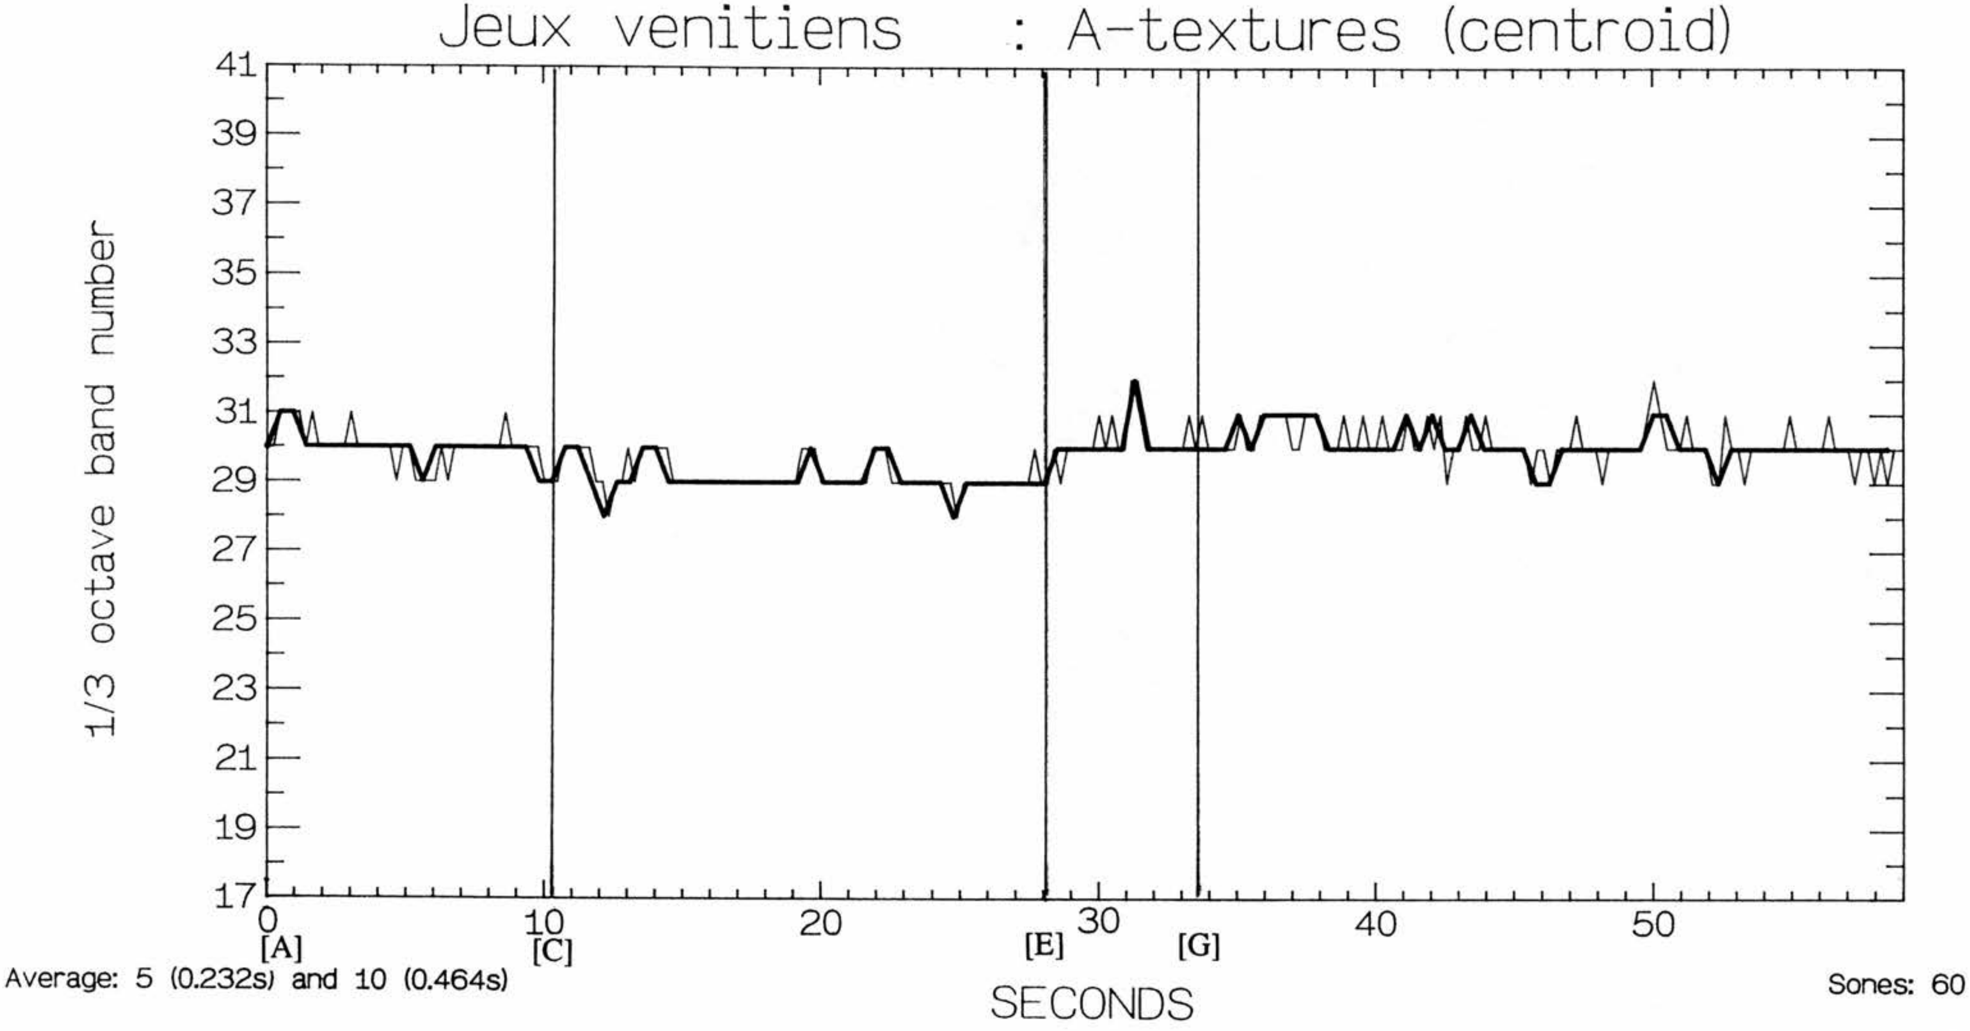

Figure 3.21
Sharpness: recording 1

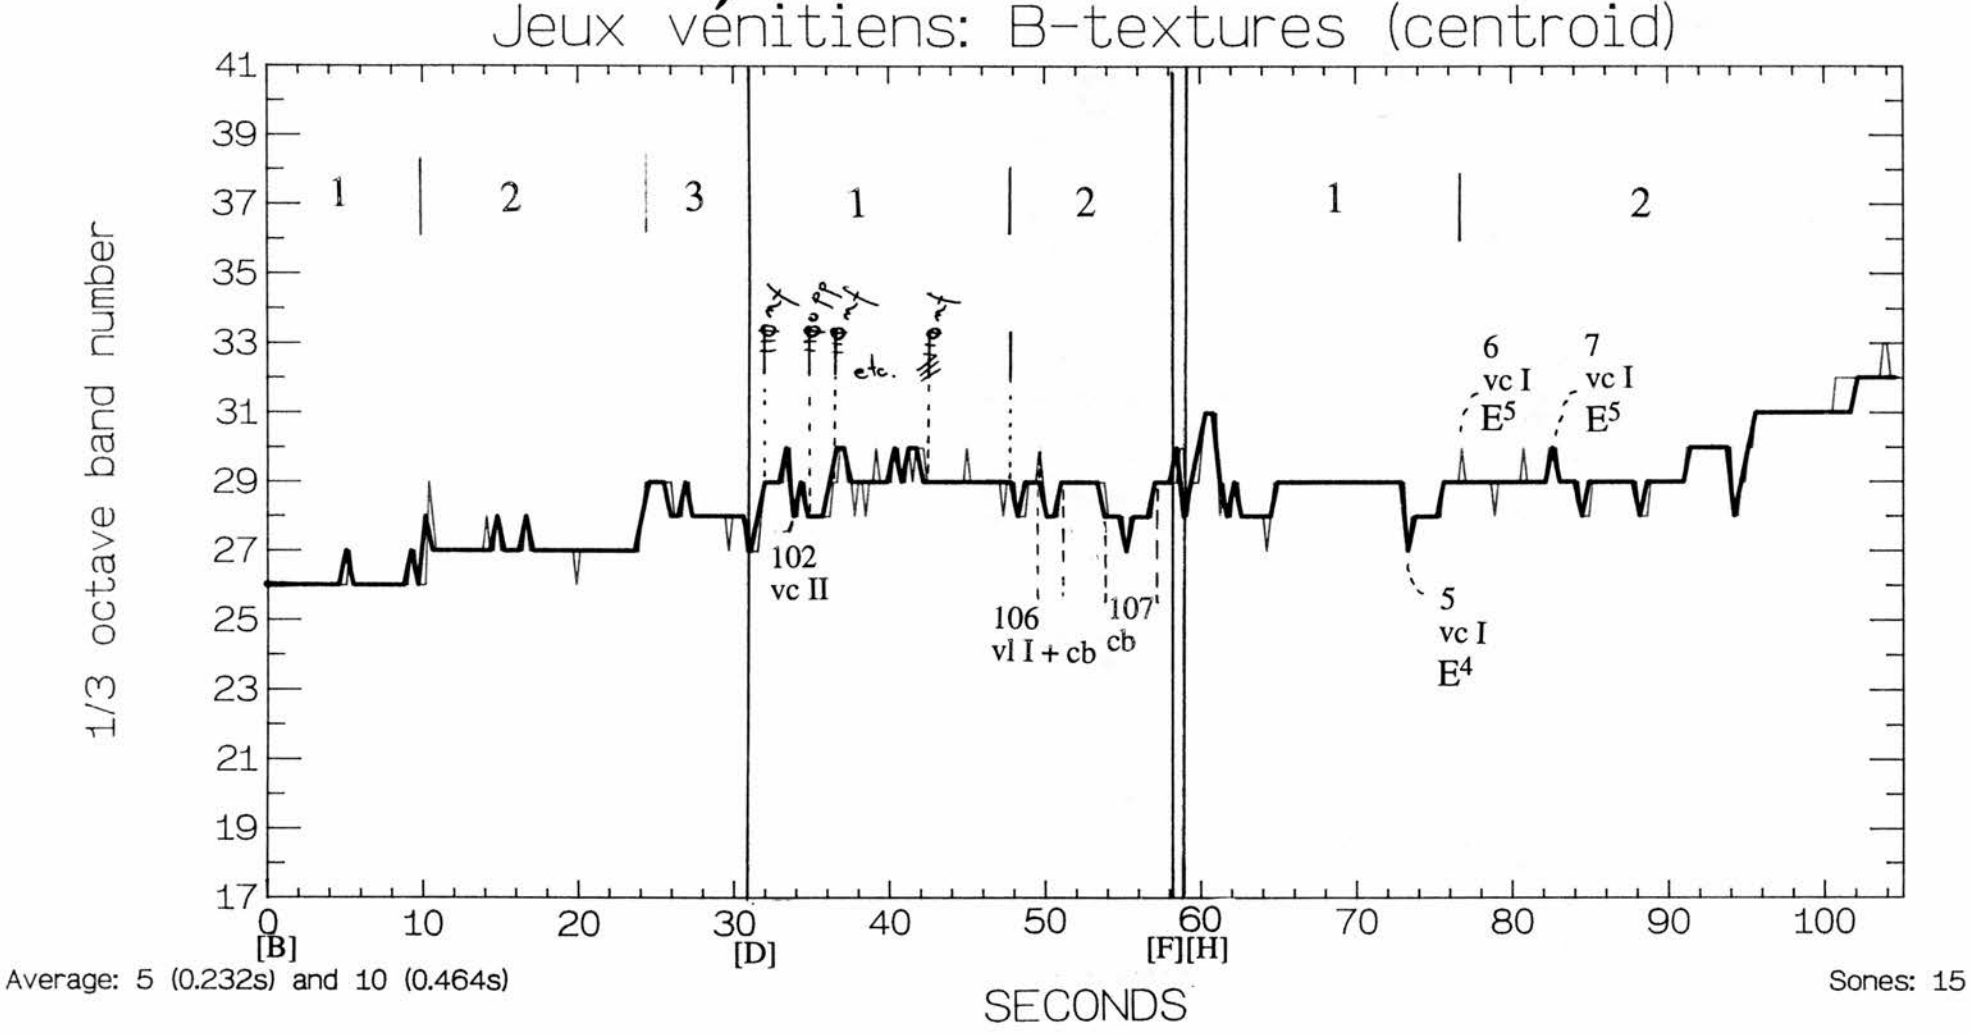

Figure 3.22
Sharpness: recording 2

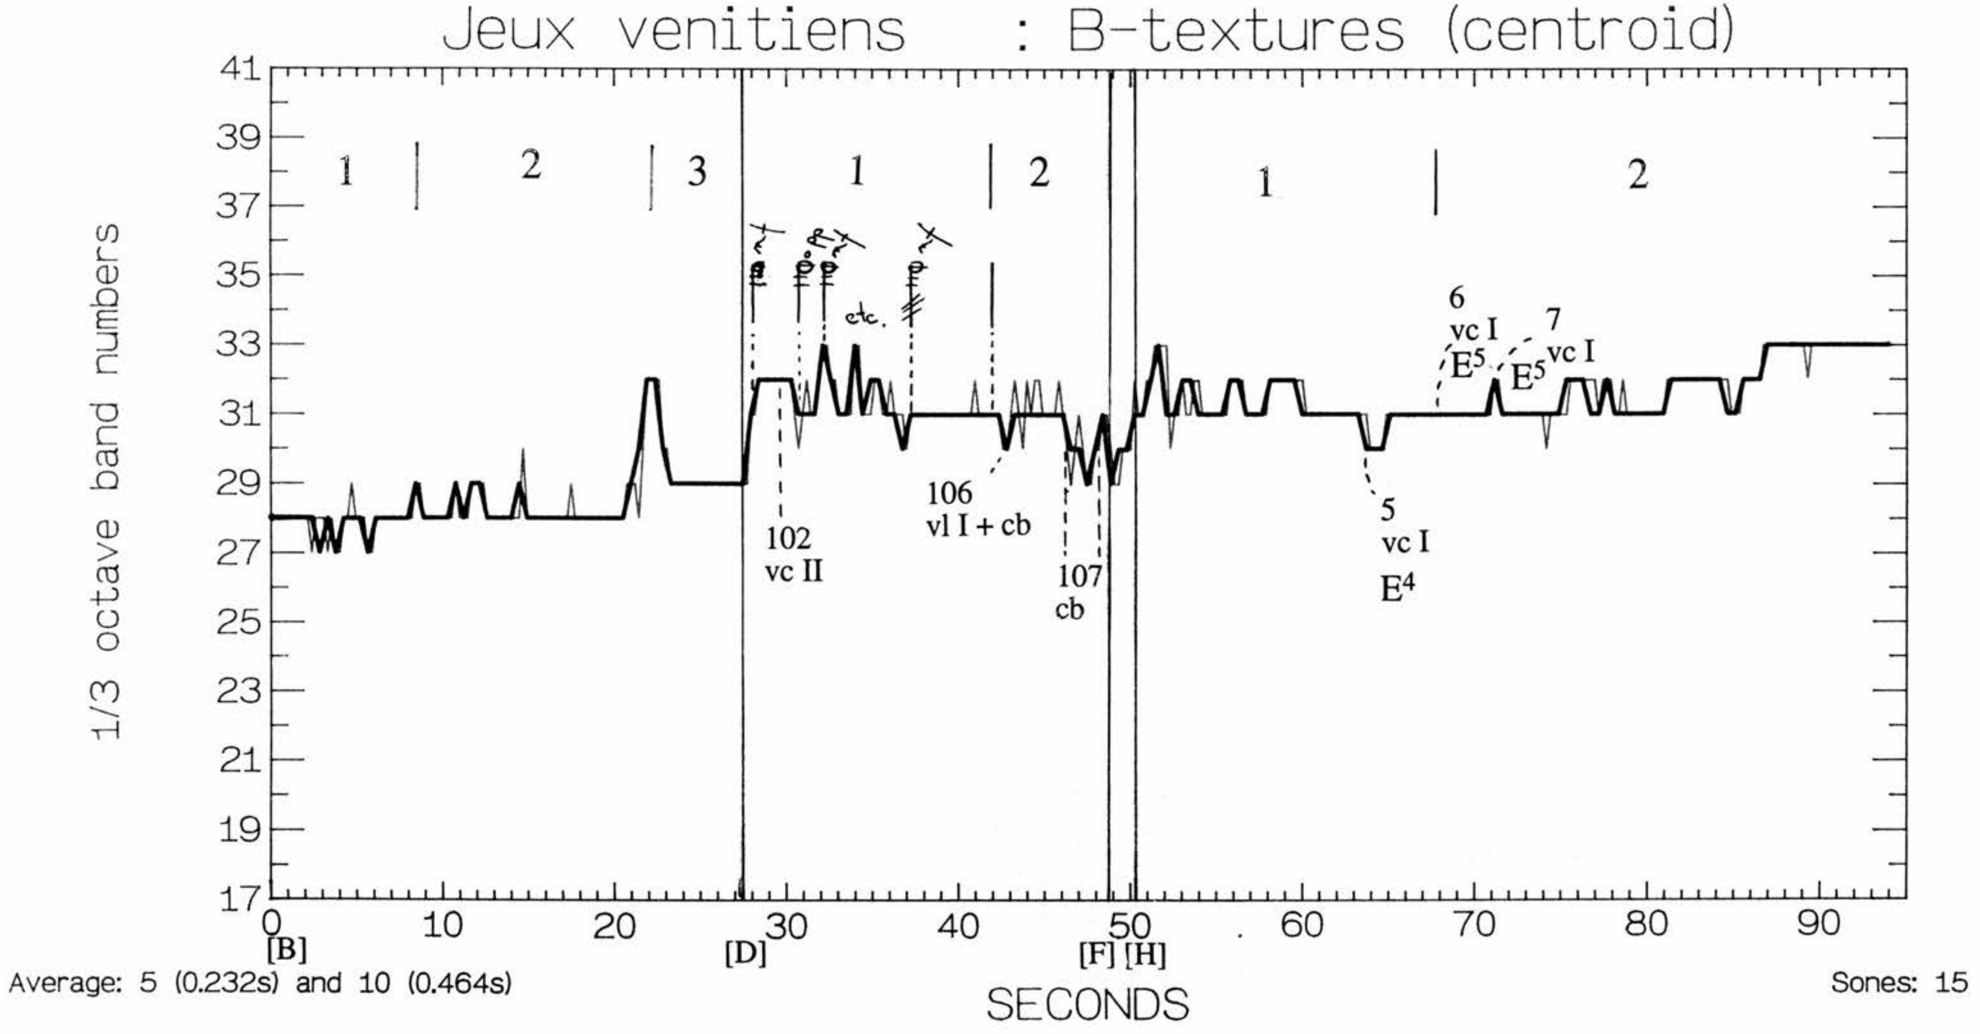

CEPSTRUM ANALYSIS

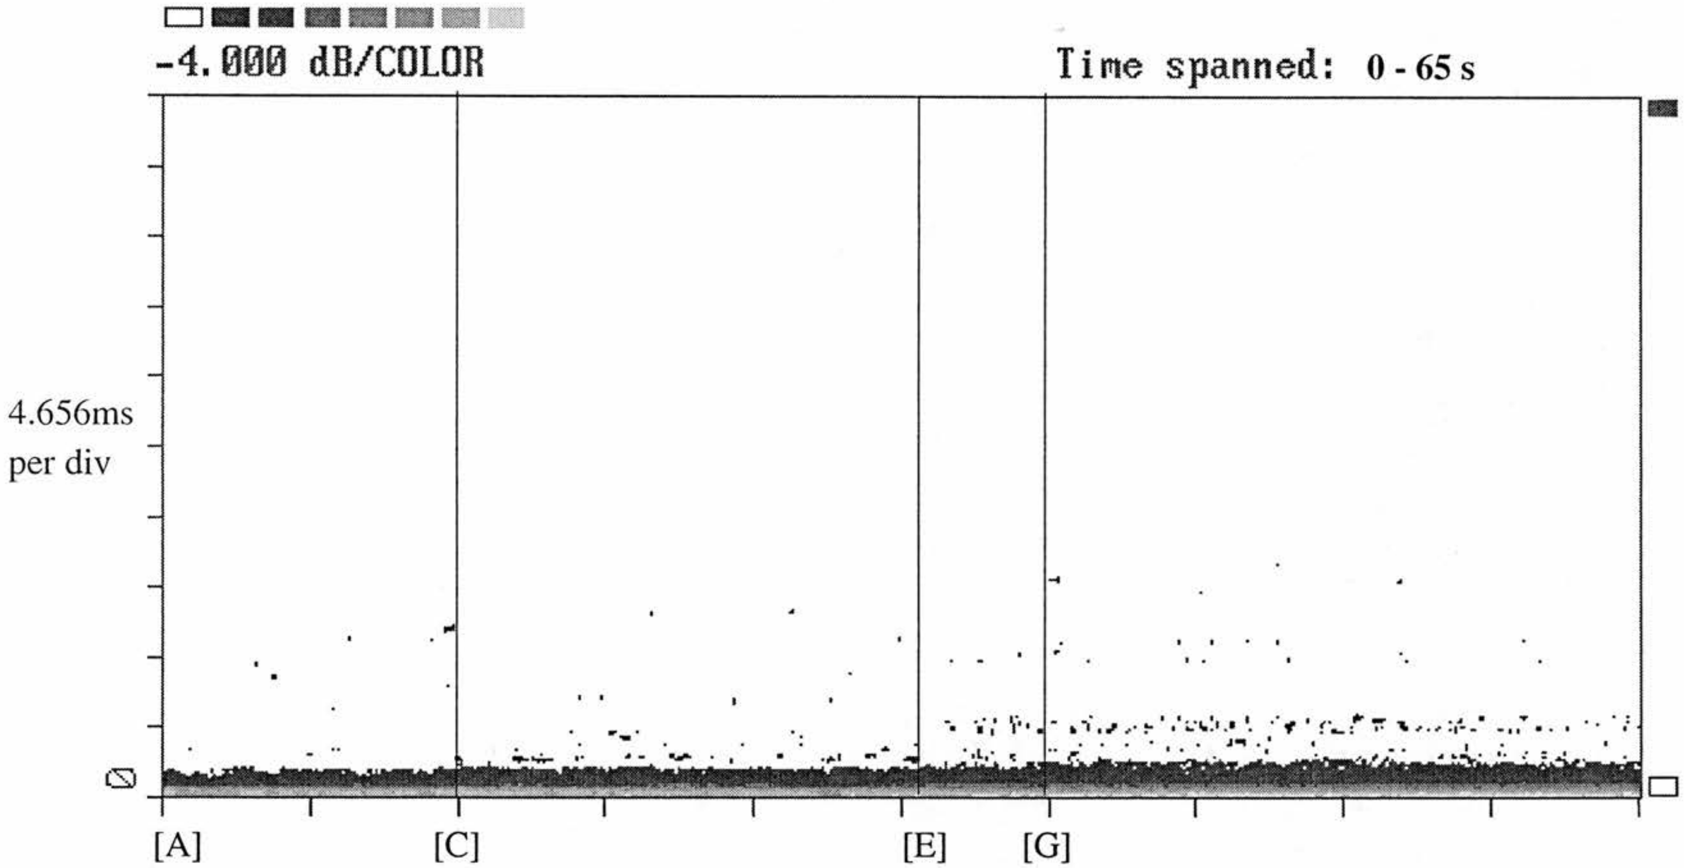

Figure 3.23 A-textures - recording 1

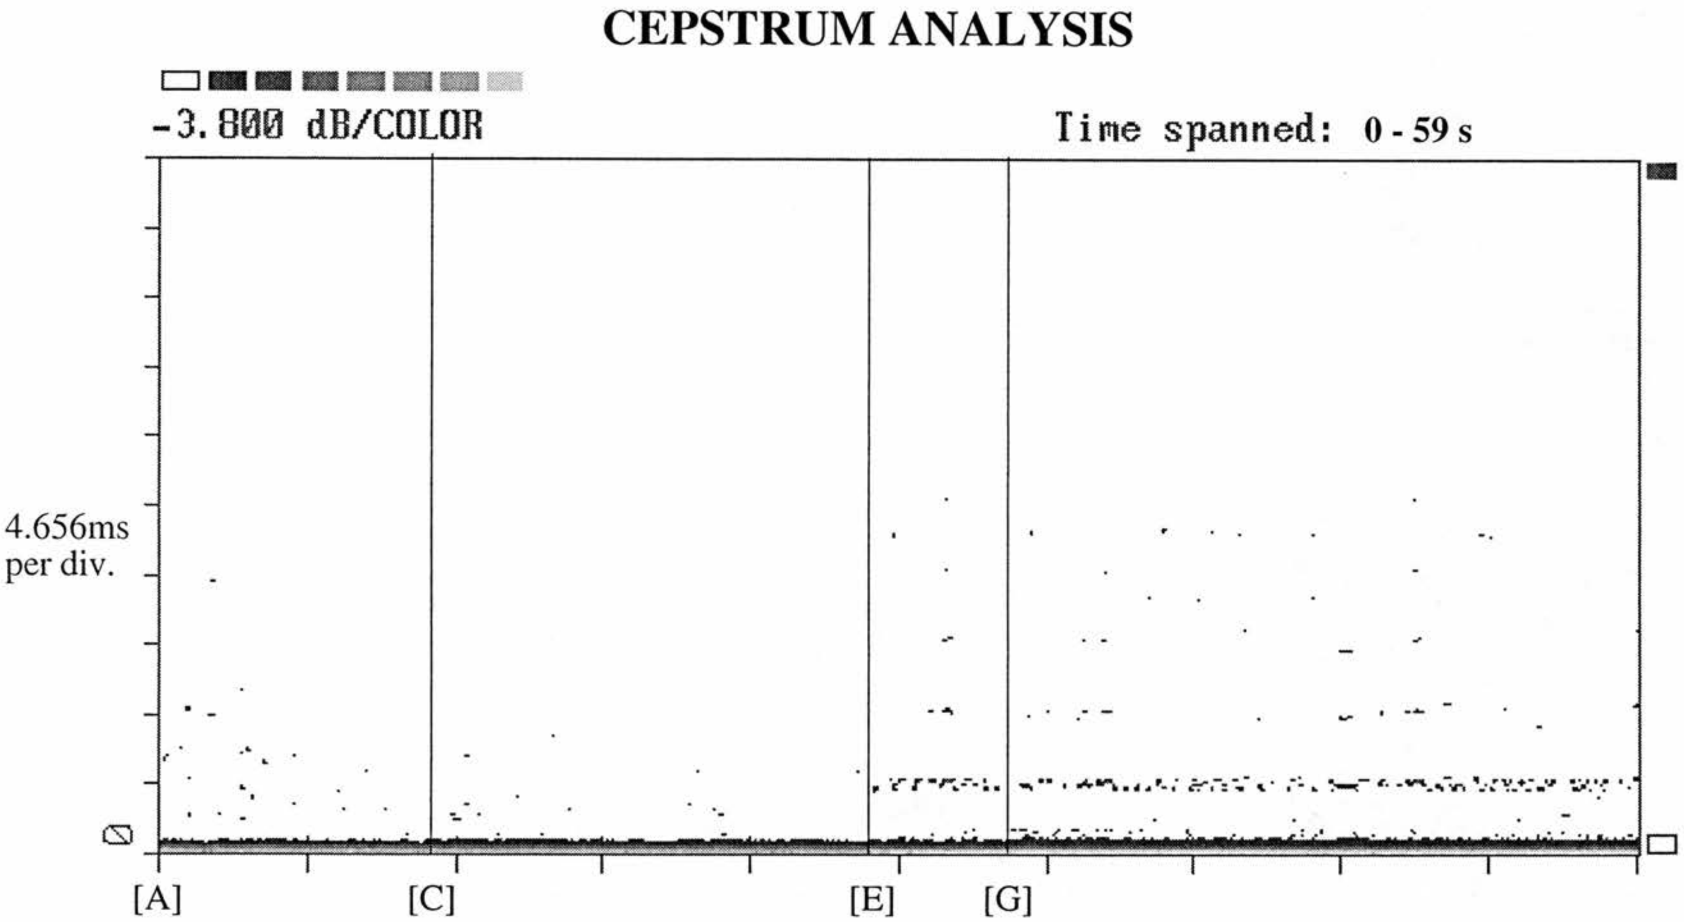

Figure 3.24 A- textures - recording 2

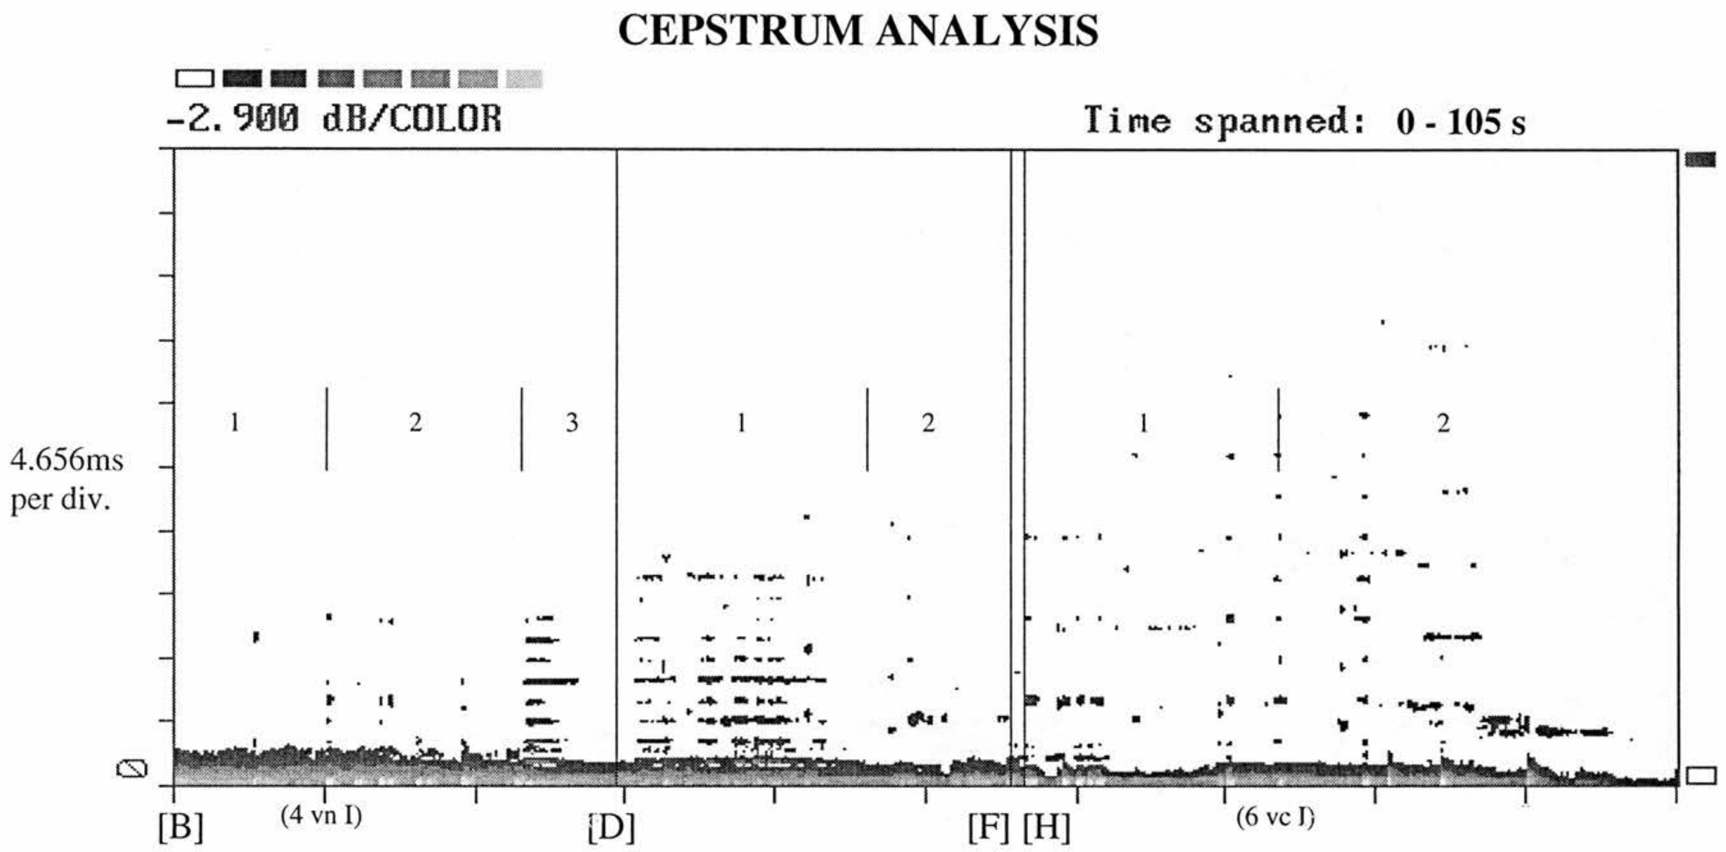

Figure 3.25 B-textures - recording 1

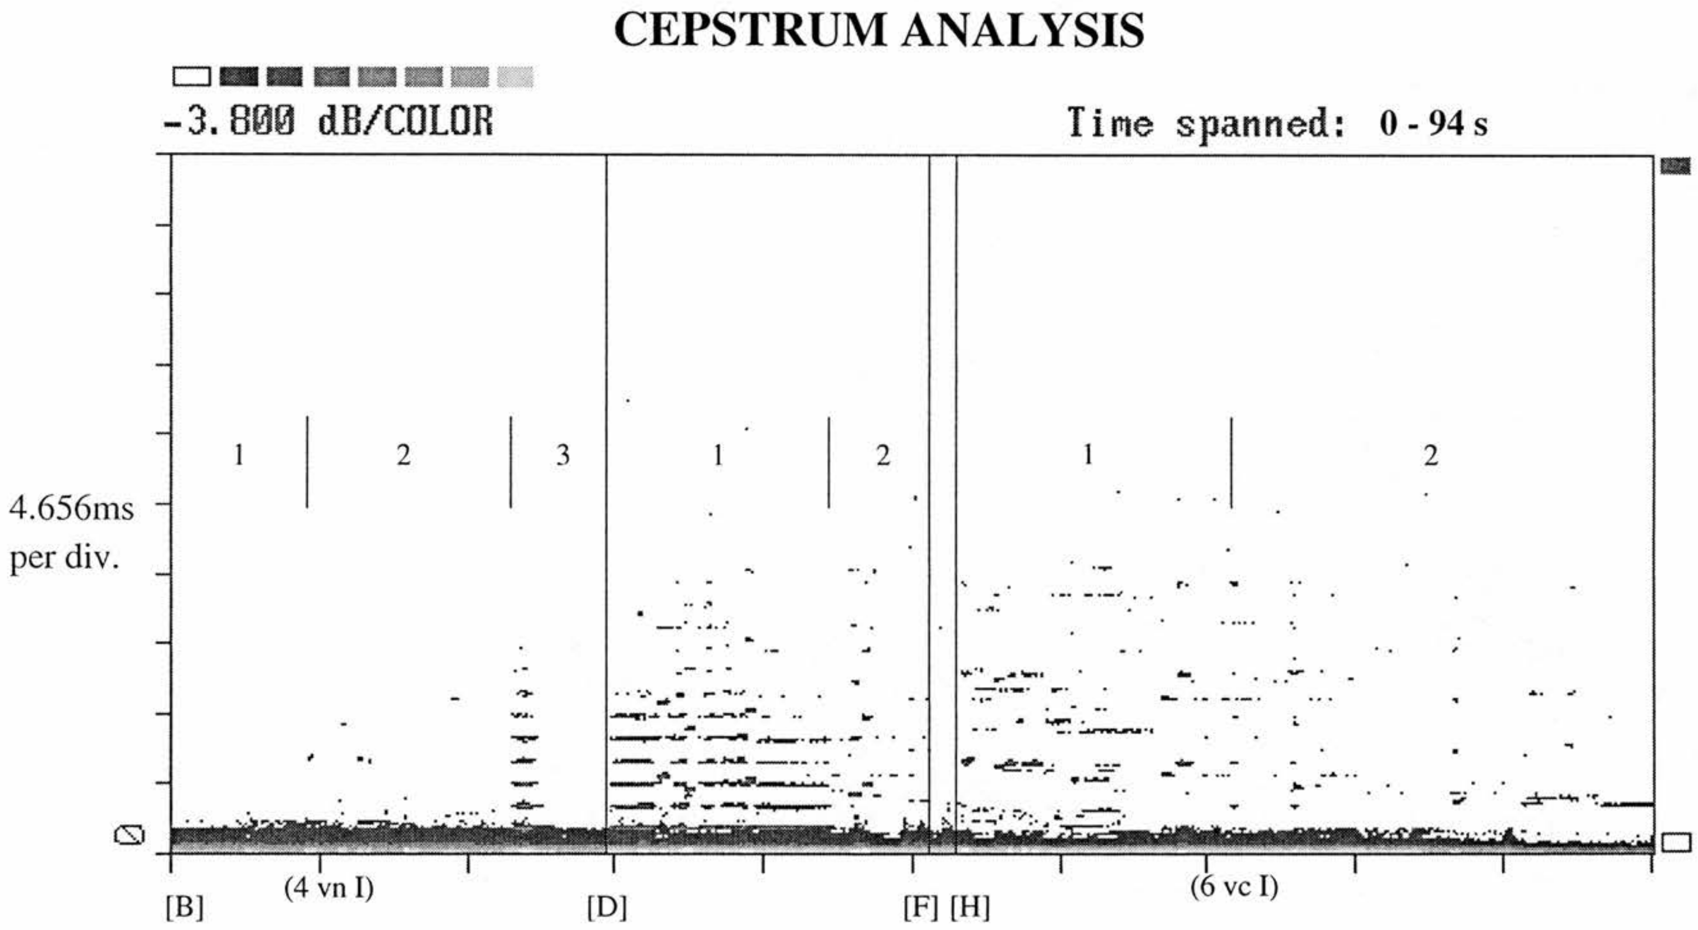

Figure 3.26 B-textures - recording 2

Figure 3.27
Timbral weight: recording 1

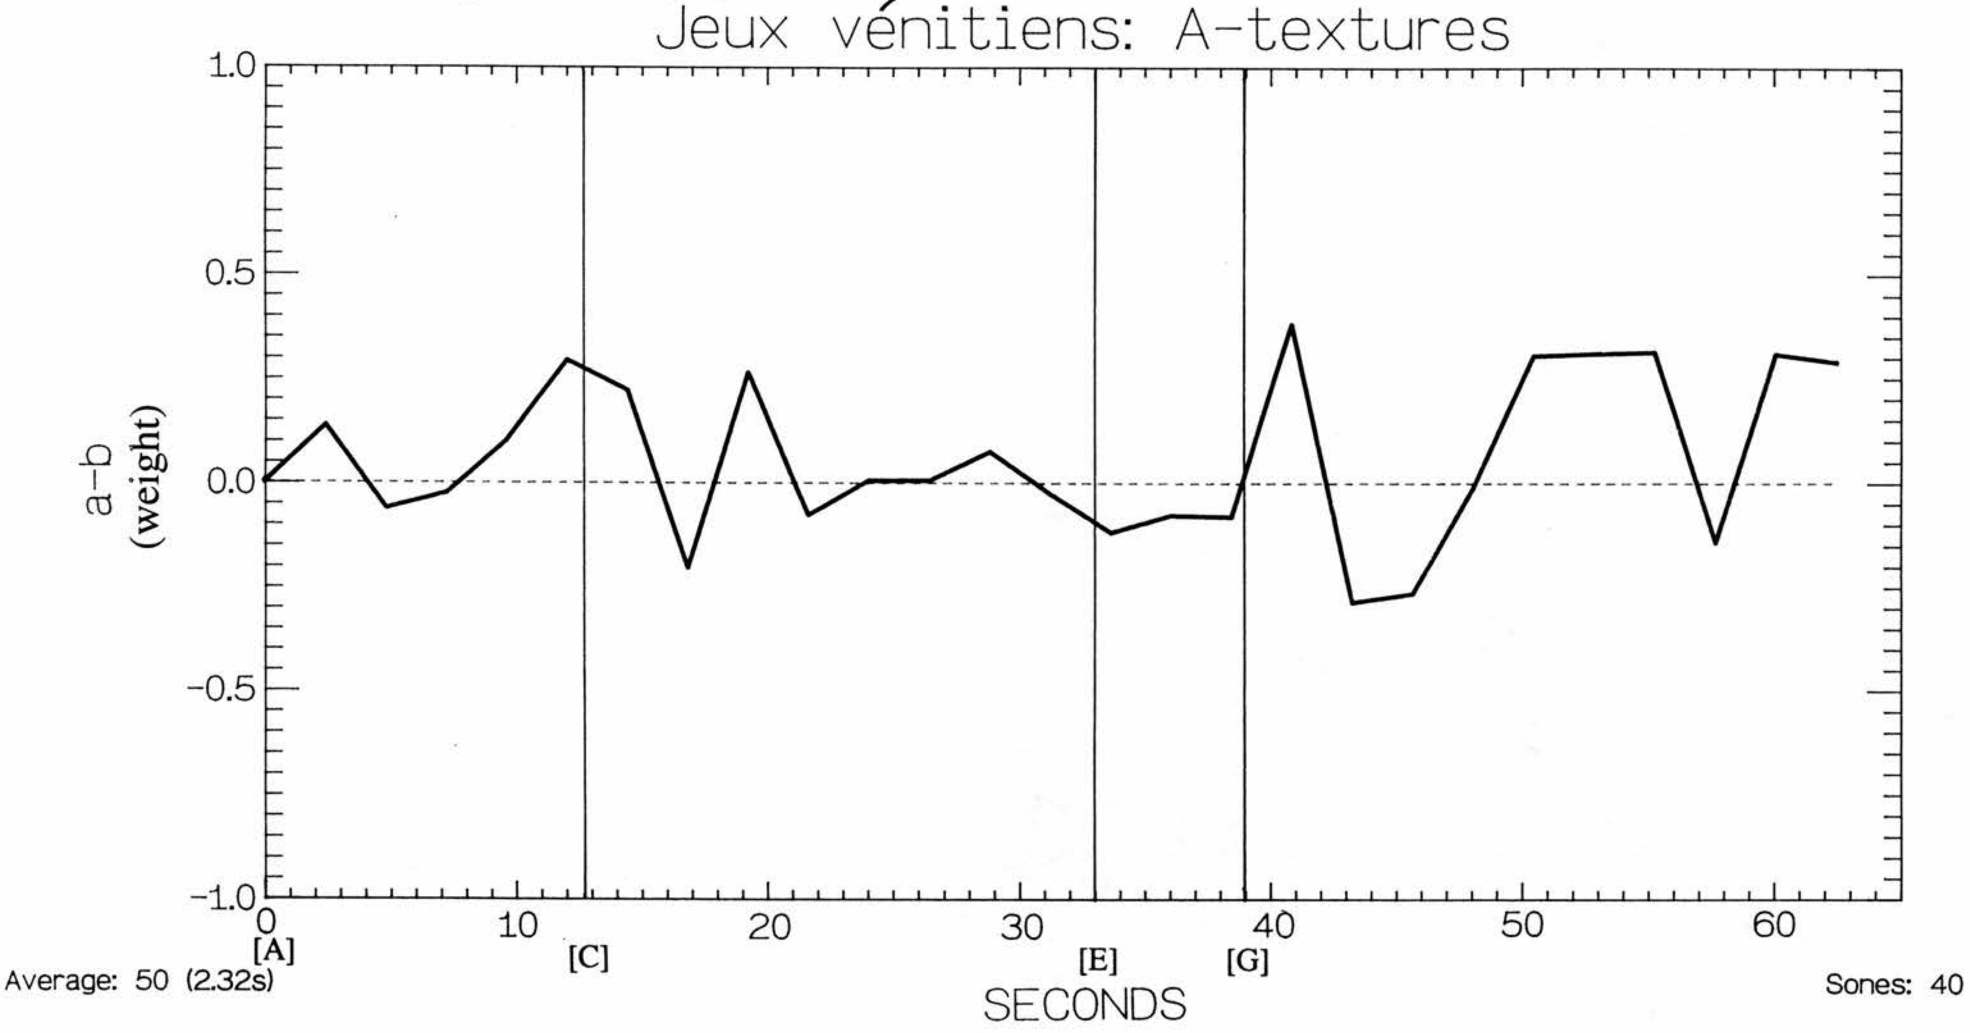

Figure 3.28
Timbral weight: recording 2

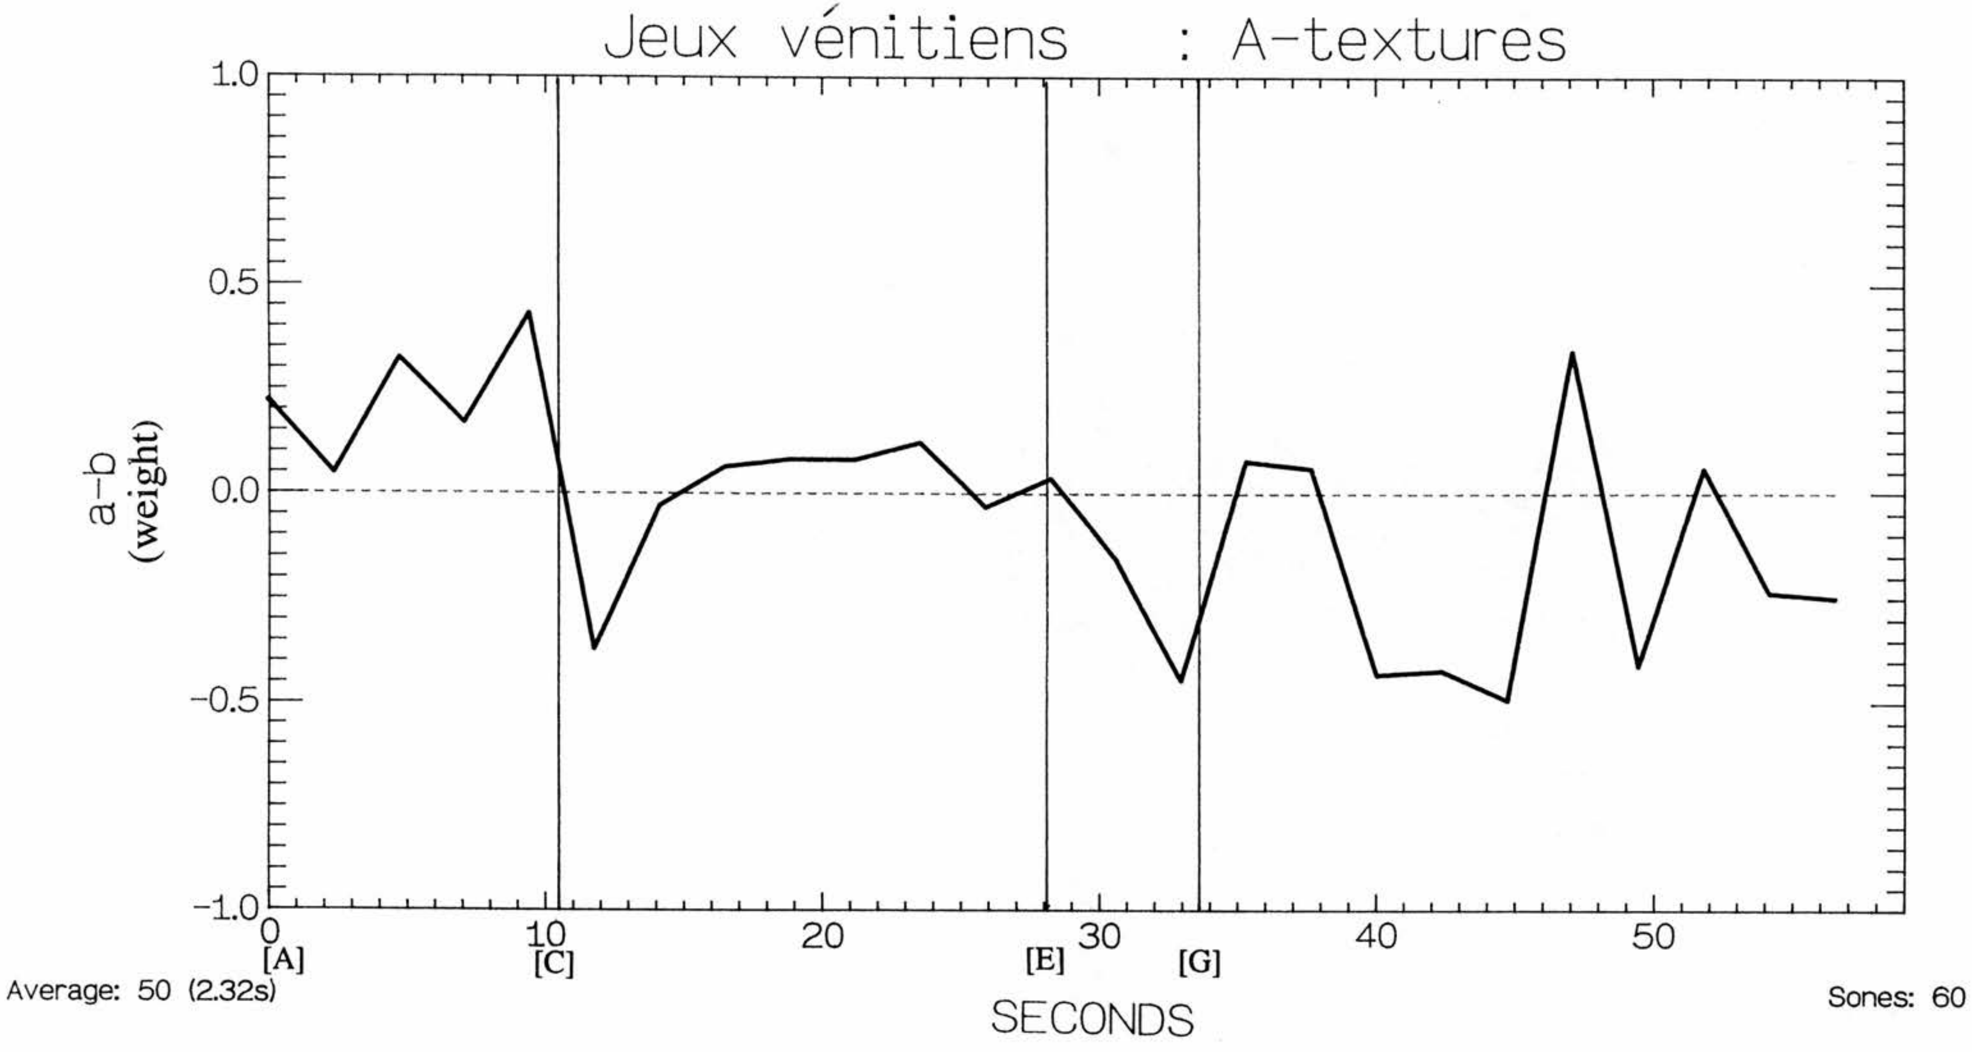

Figure 3.29
Timbral pitch: recording 1

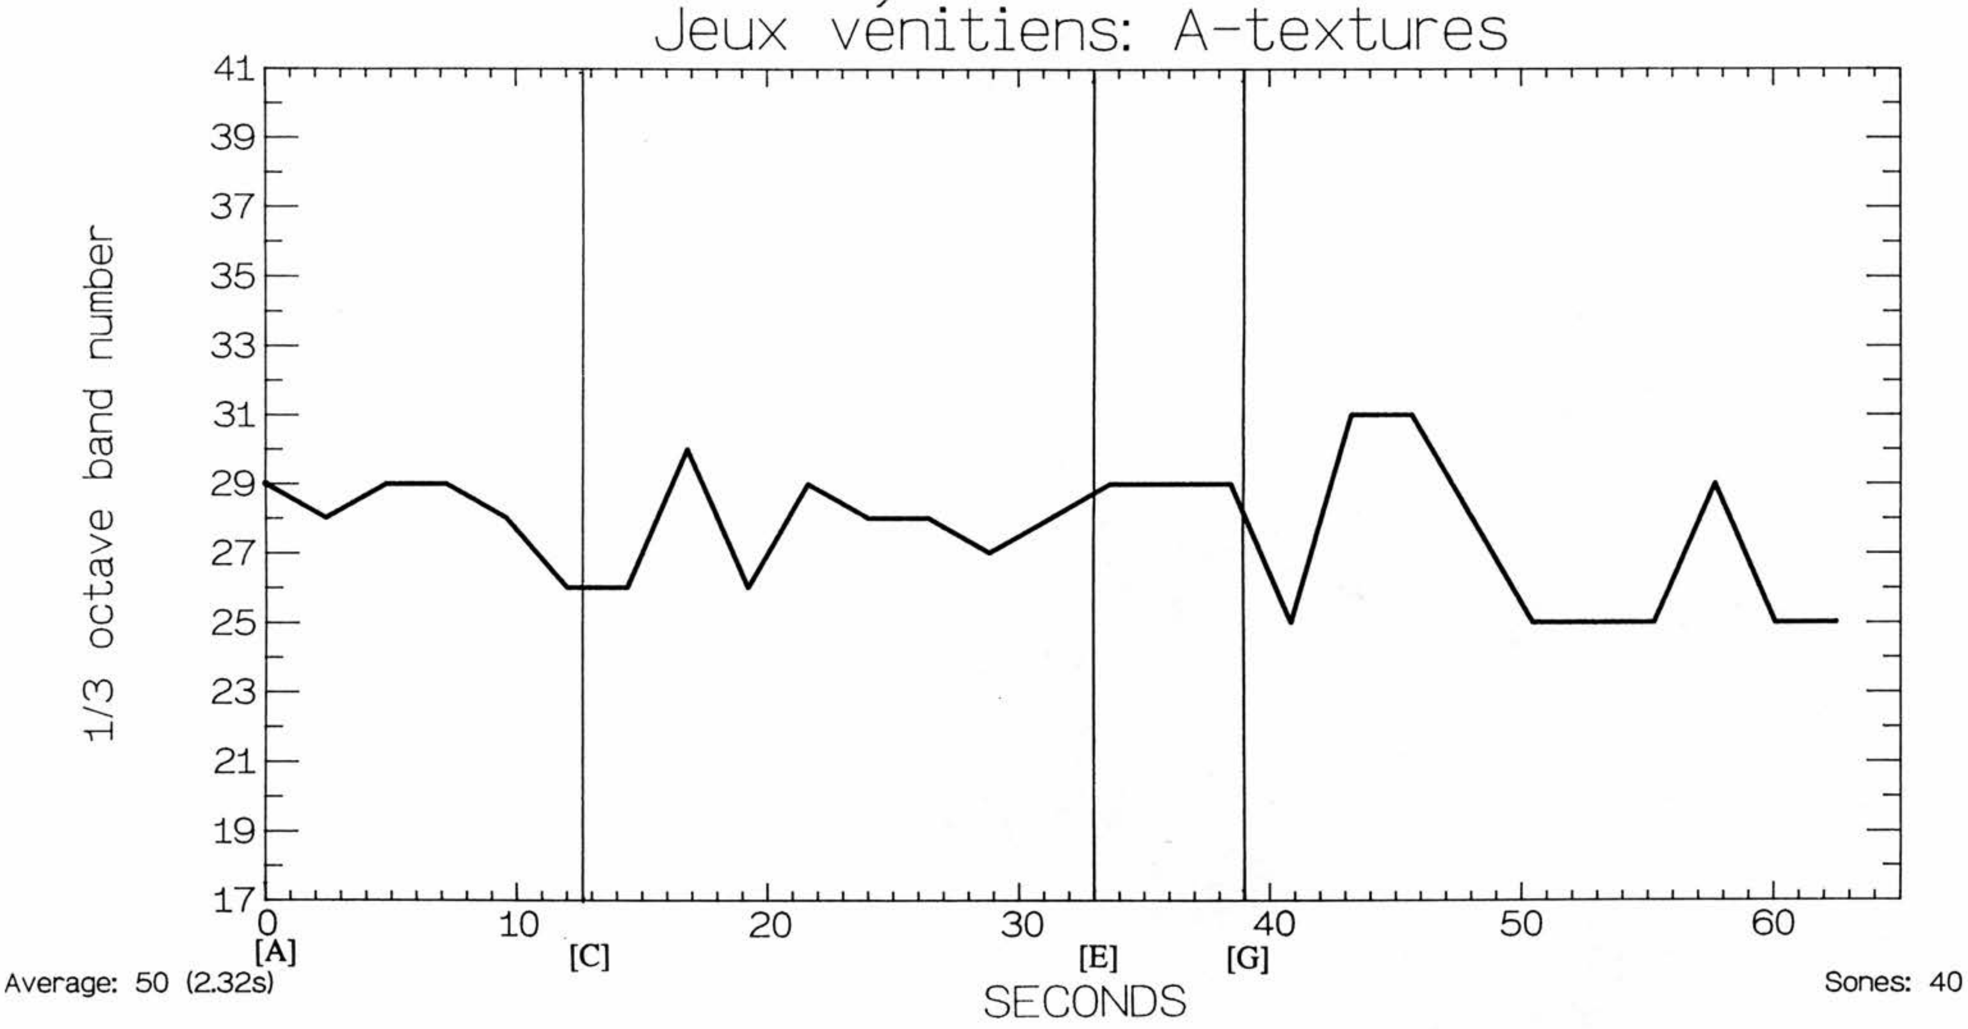

Figure 3.30
Timbral pitch: recording 2

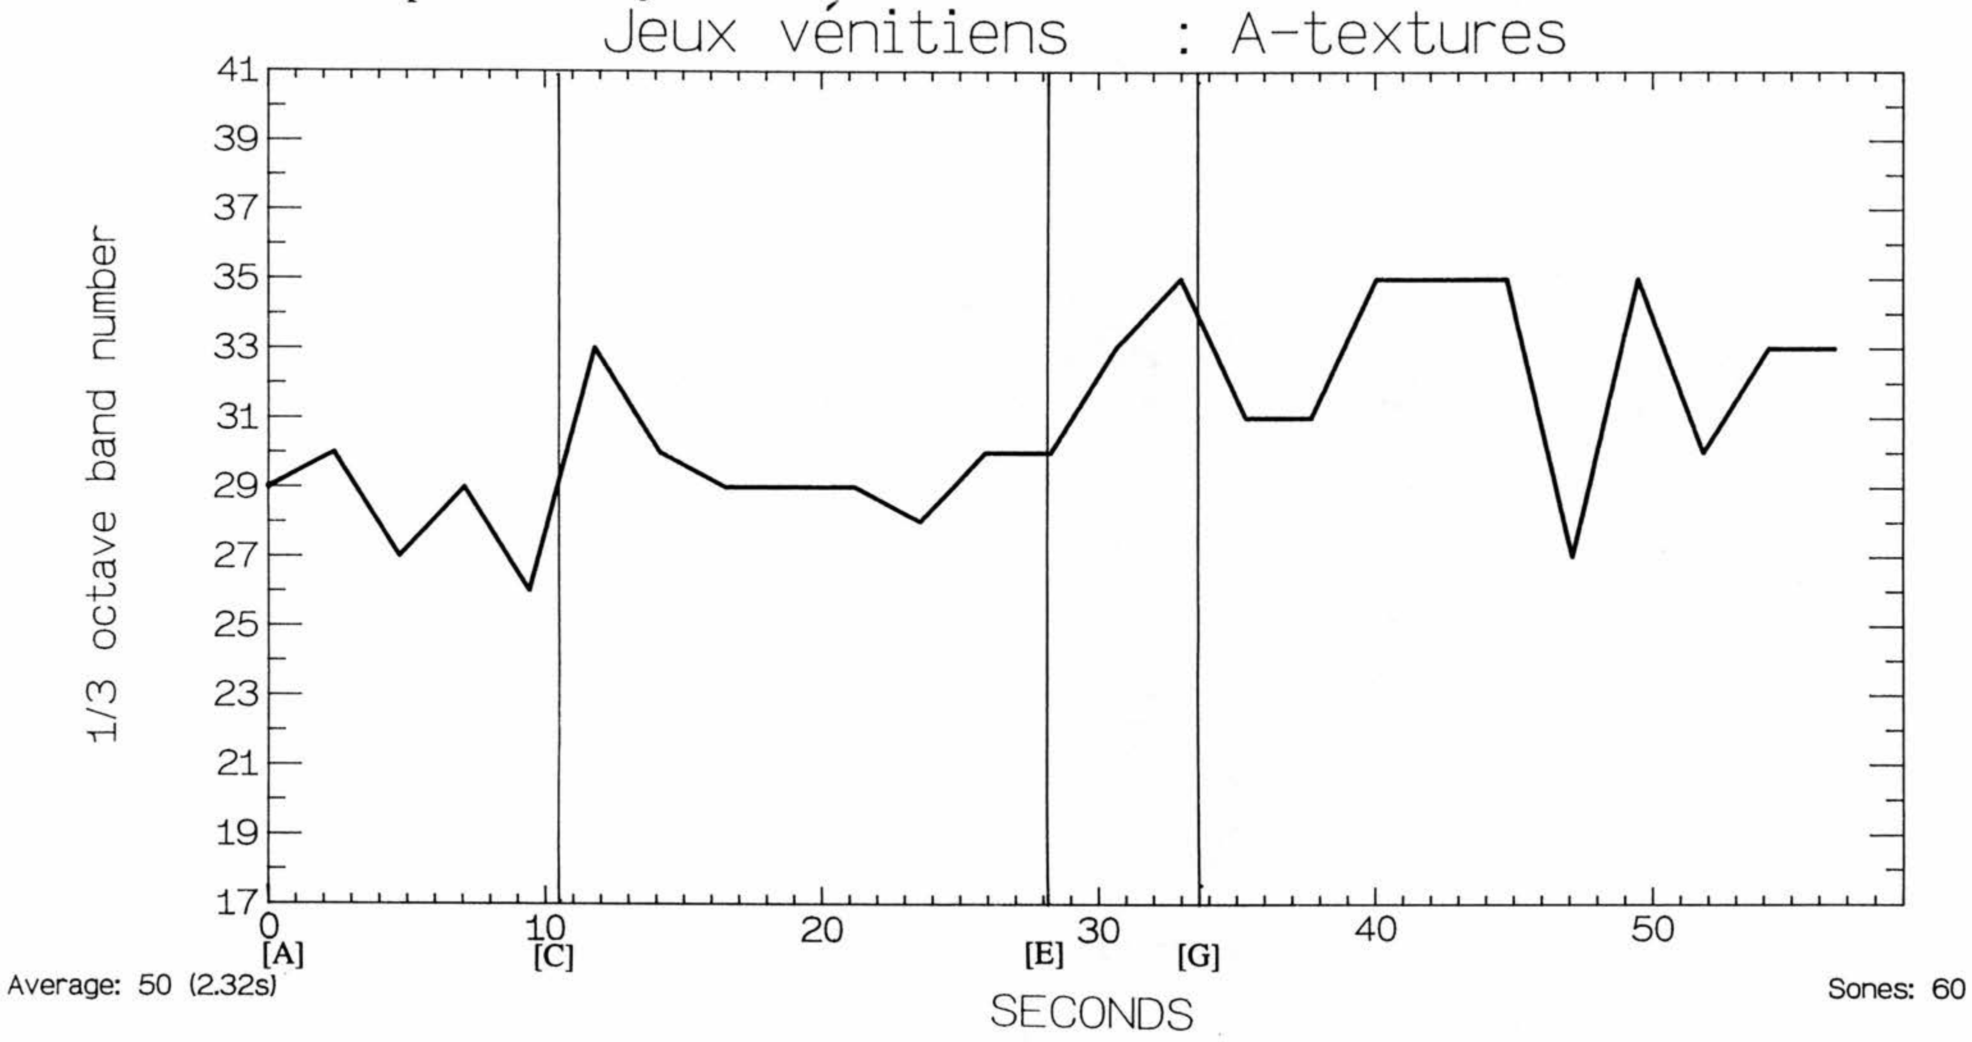

Figure 3.31
Timbral width: recording 1

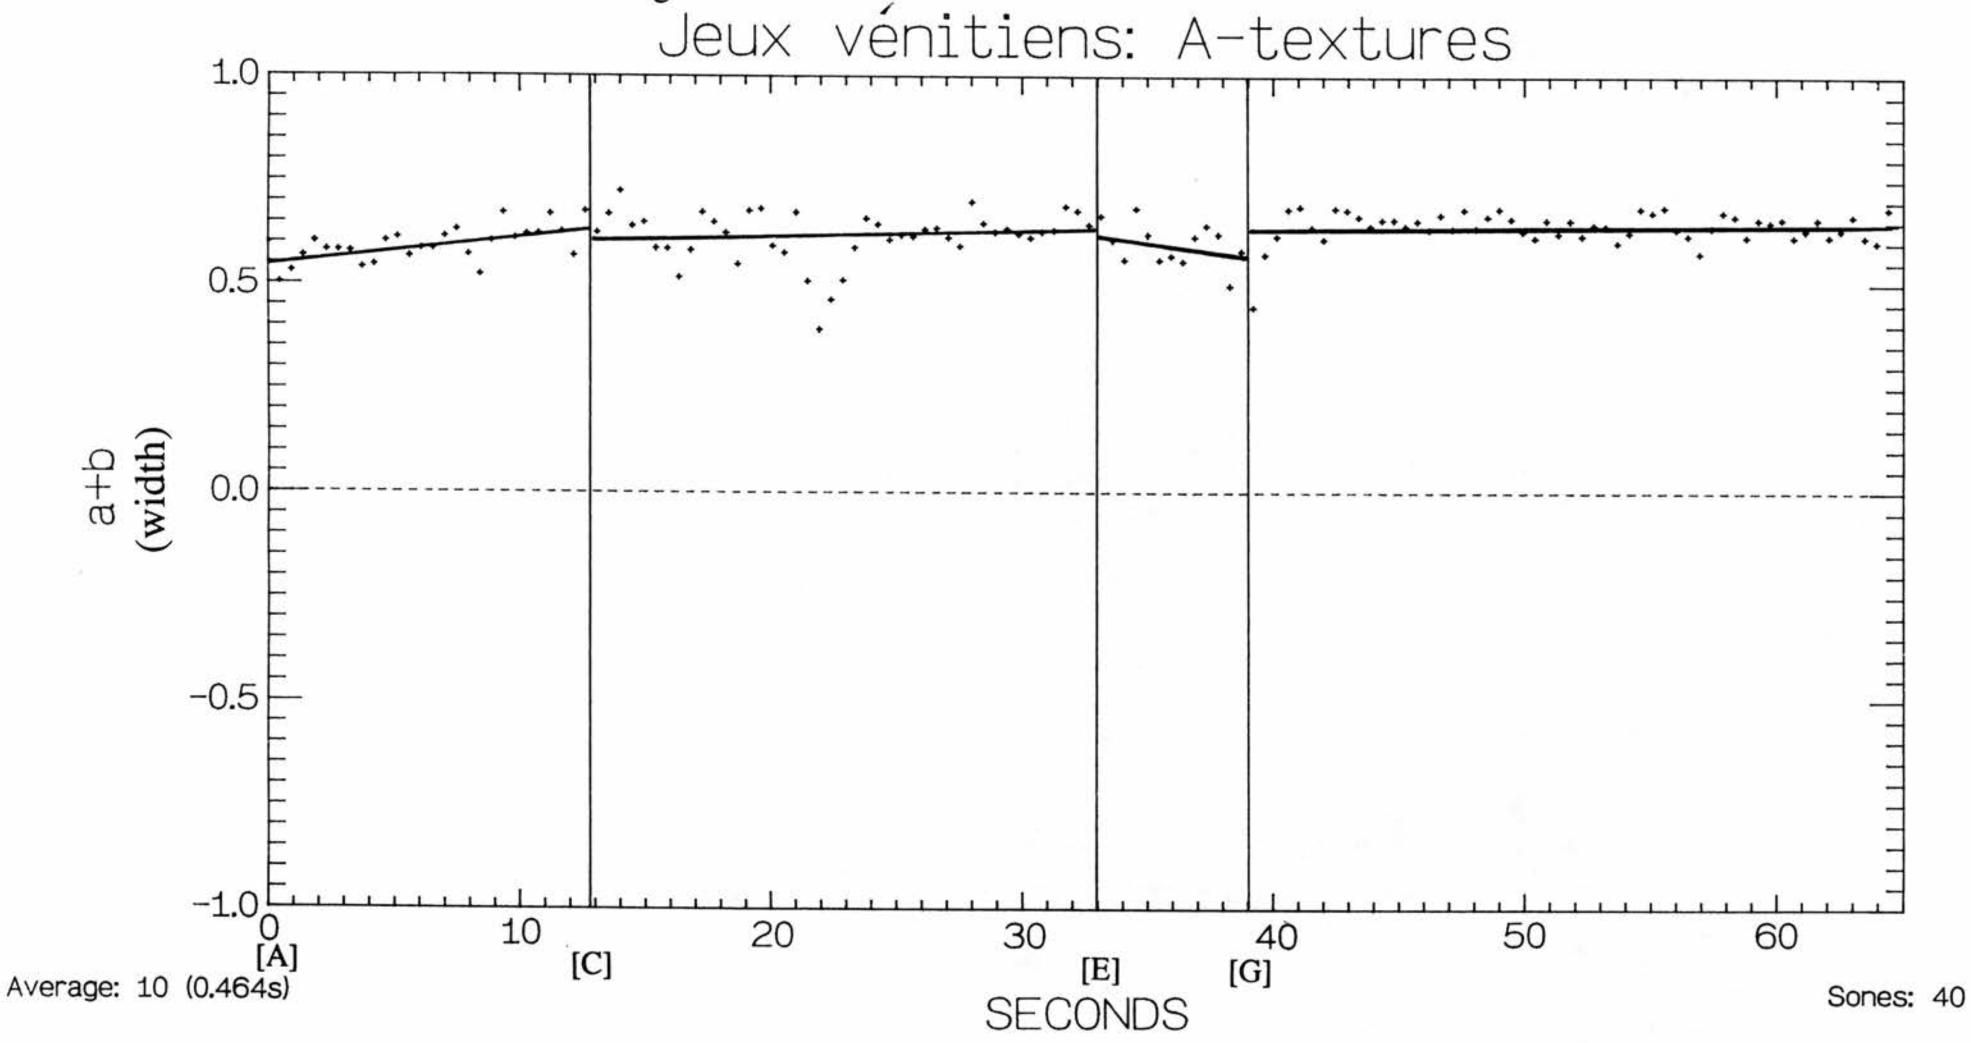

Figure 3.32
Timbral width: recording 2

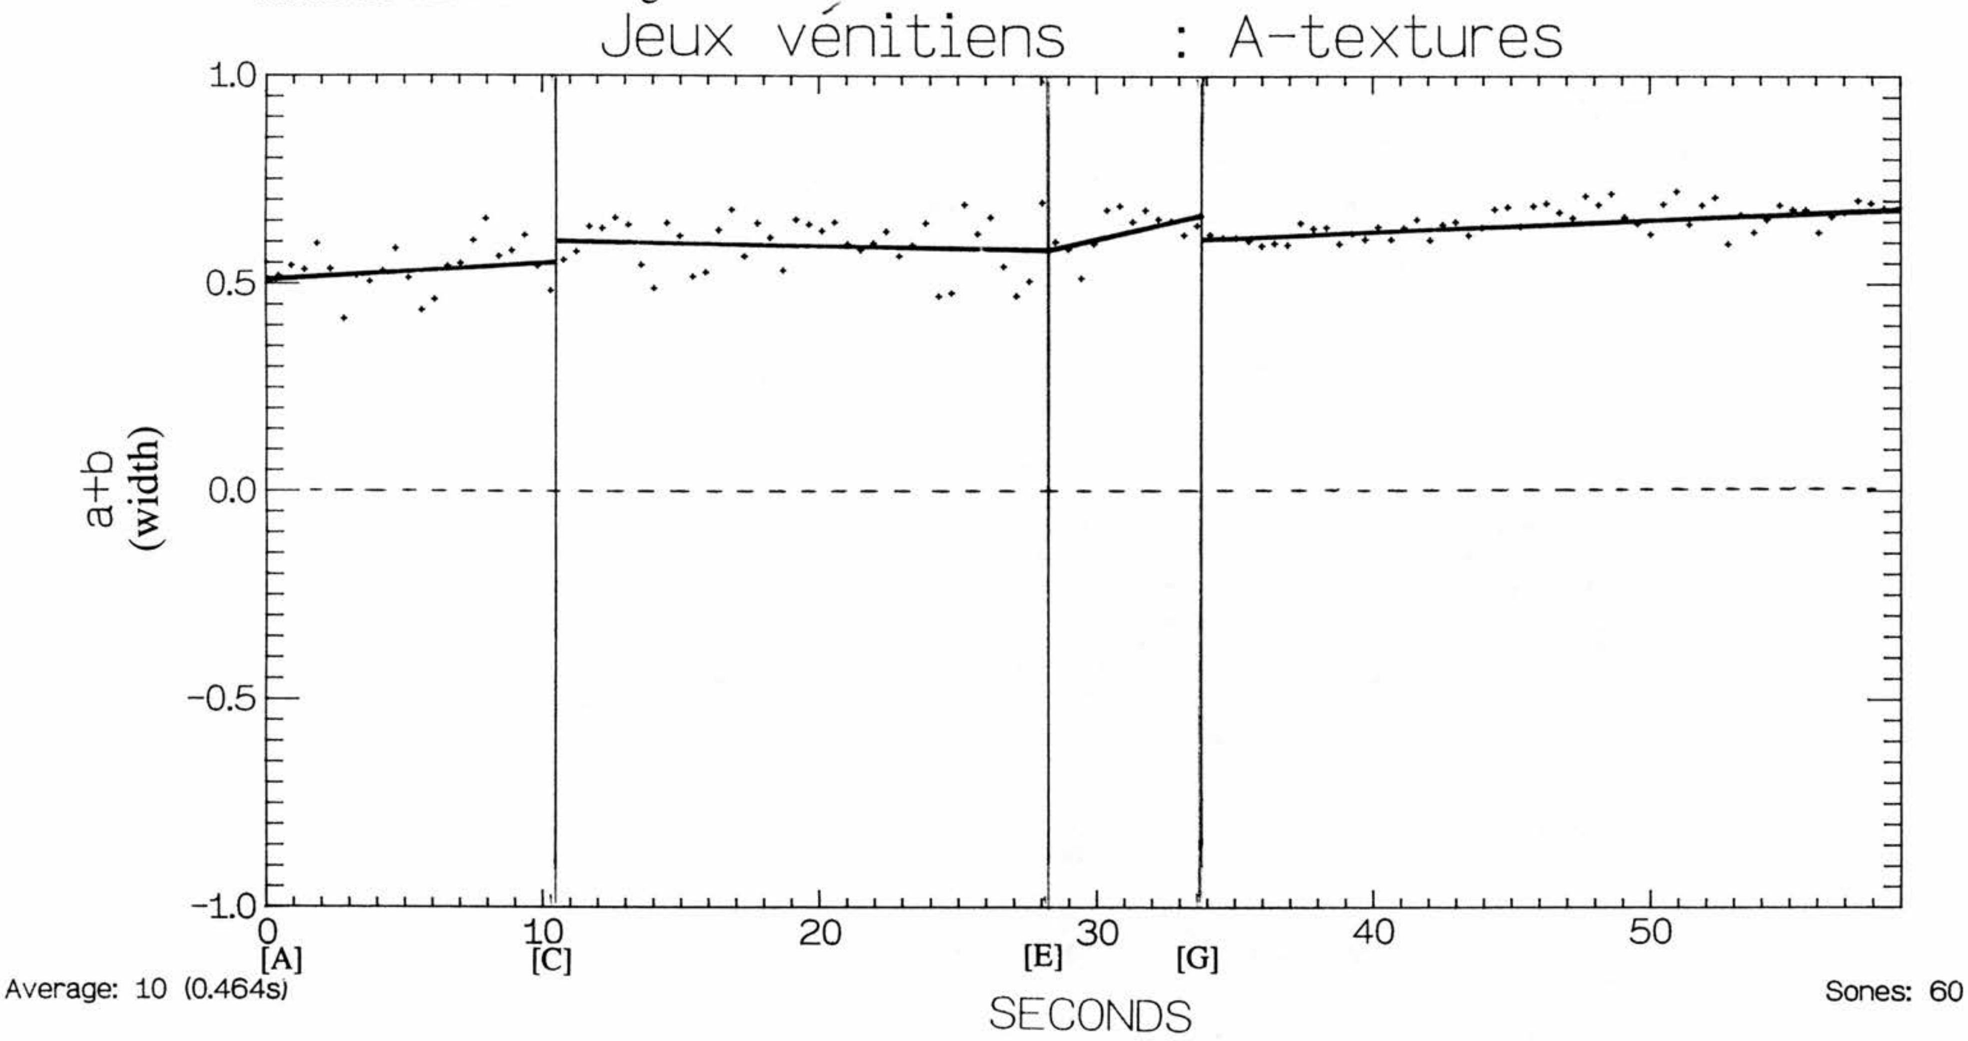

Figure 3.33

Timbral pitch: recording 1

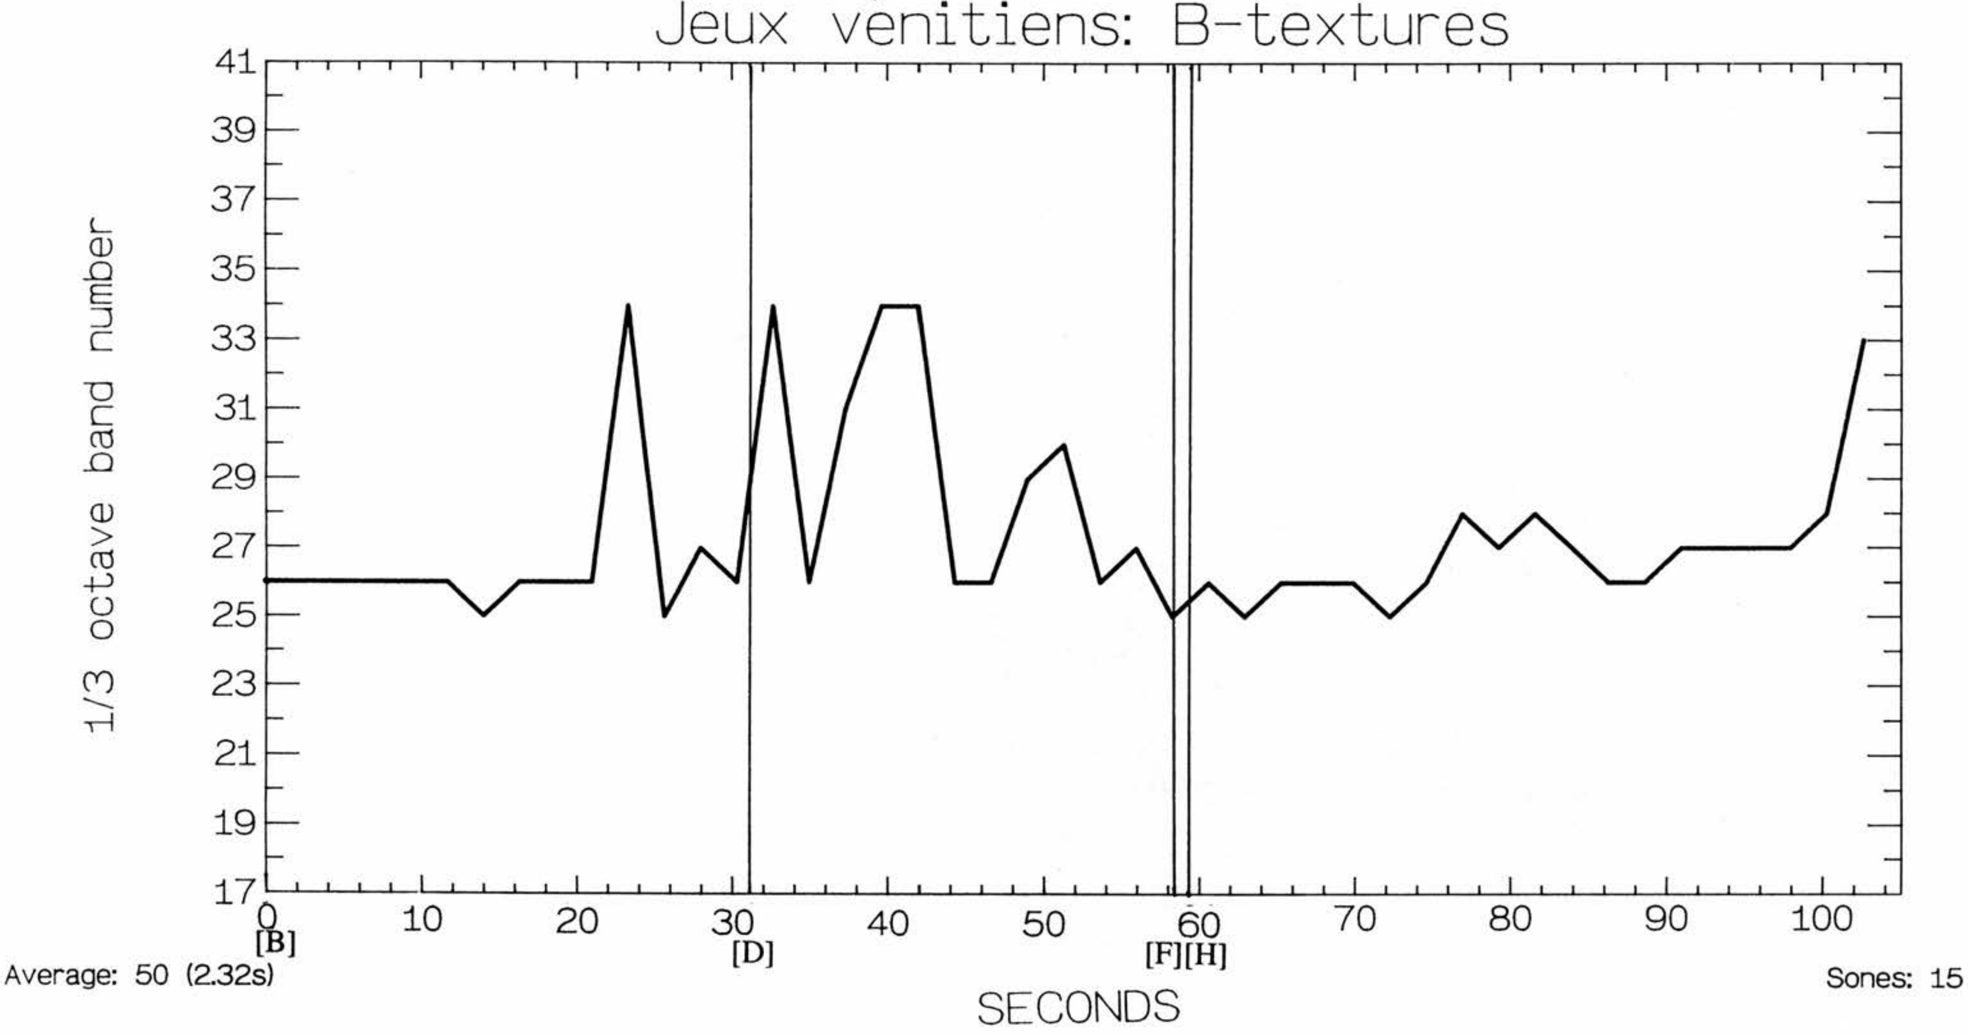

Figure 3.34

Timbral pitch: recording 2

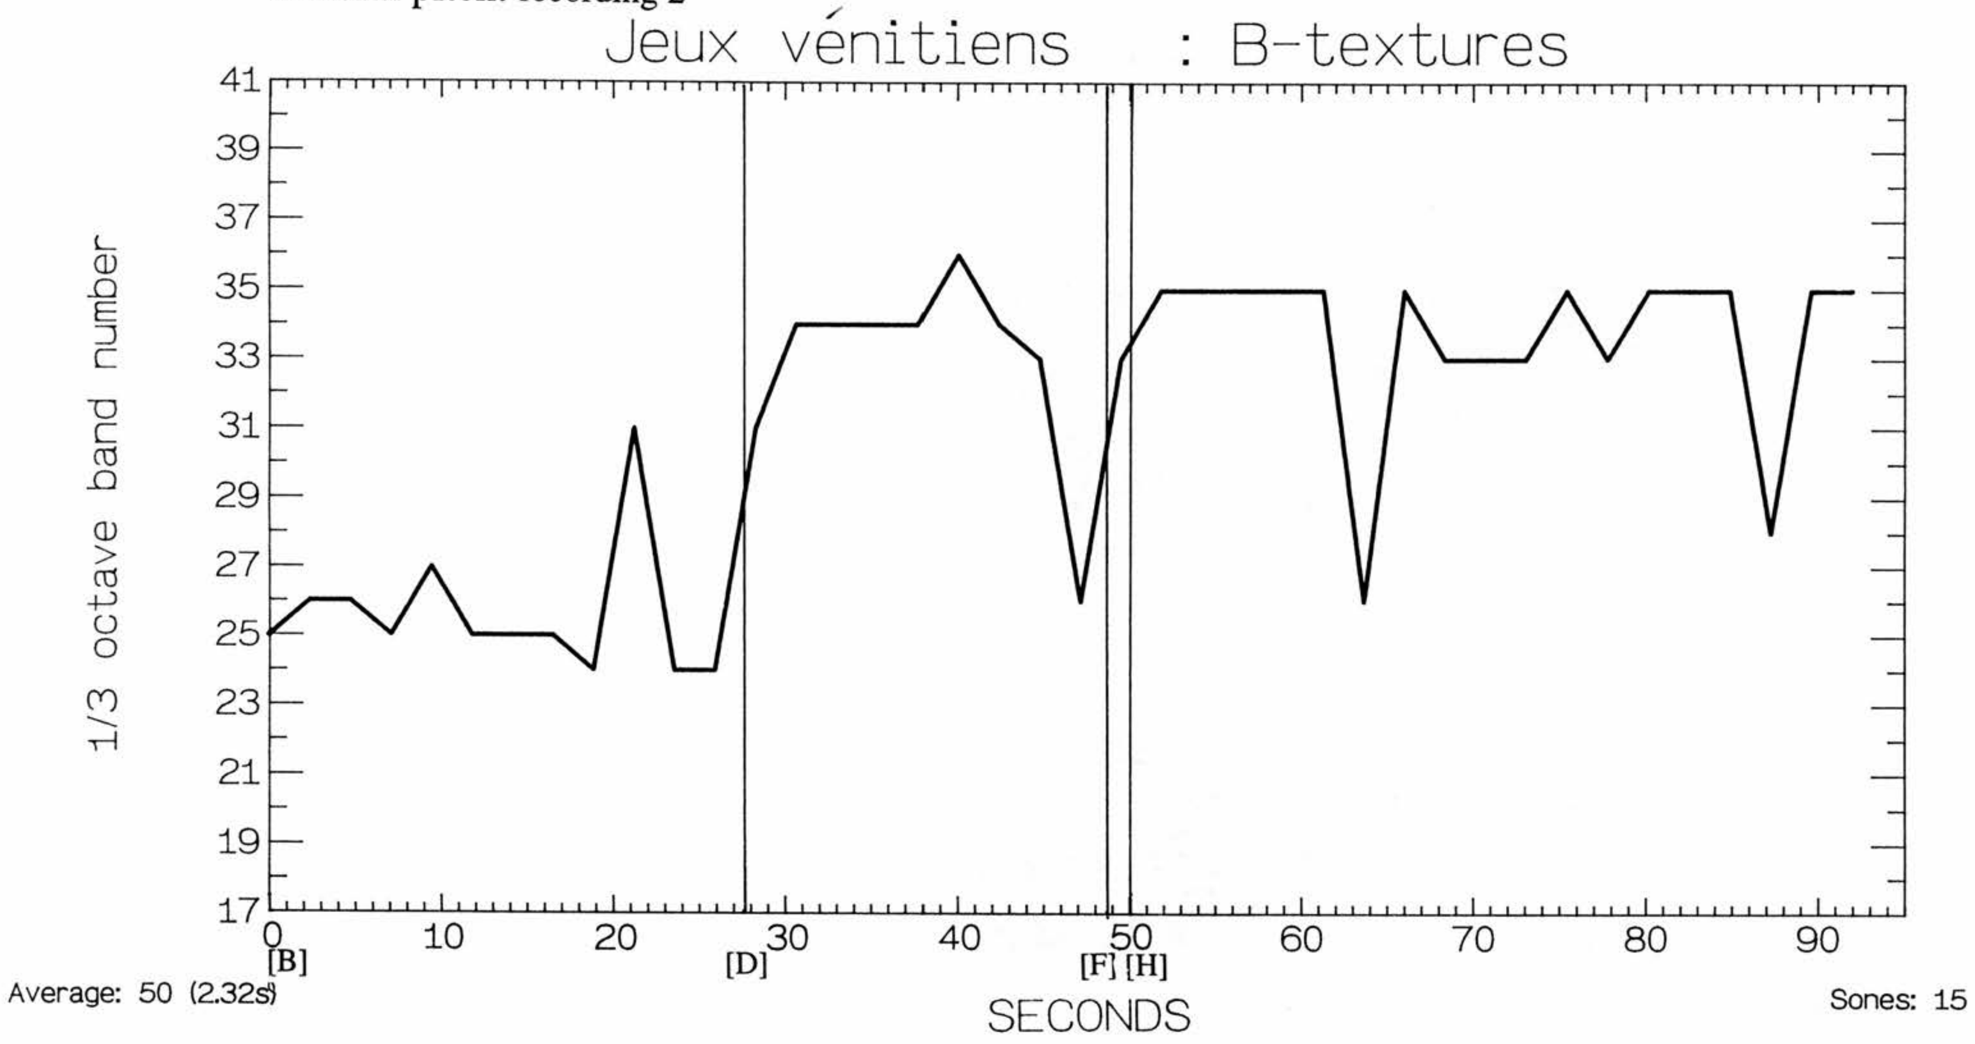

Figure 3.35
Timbral weight: recording 1

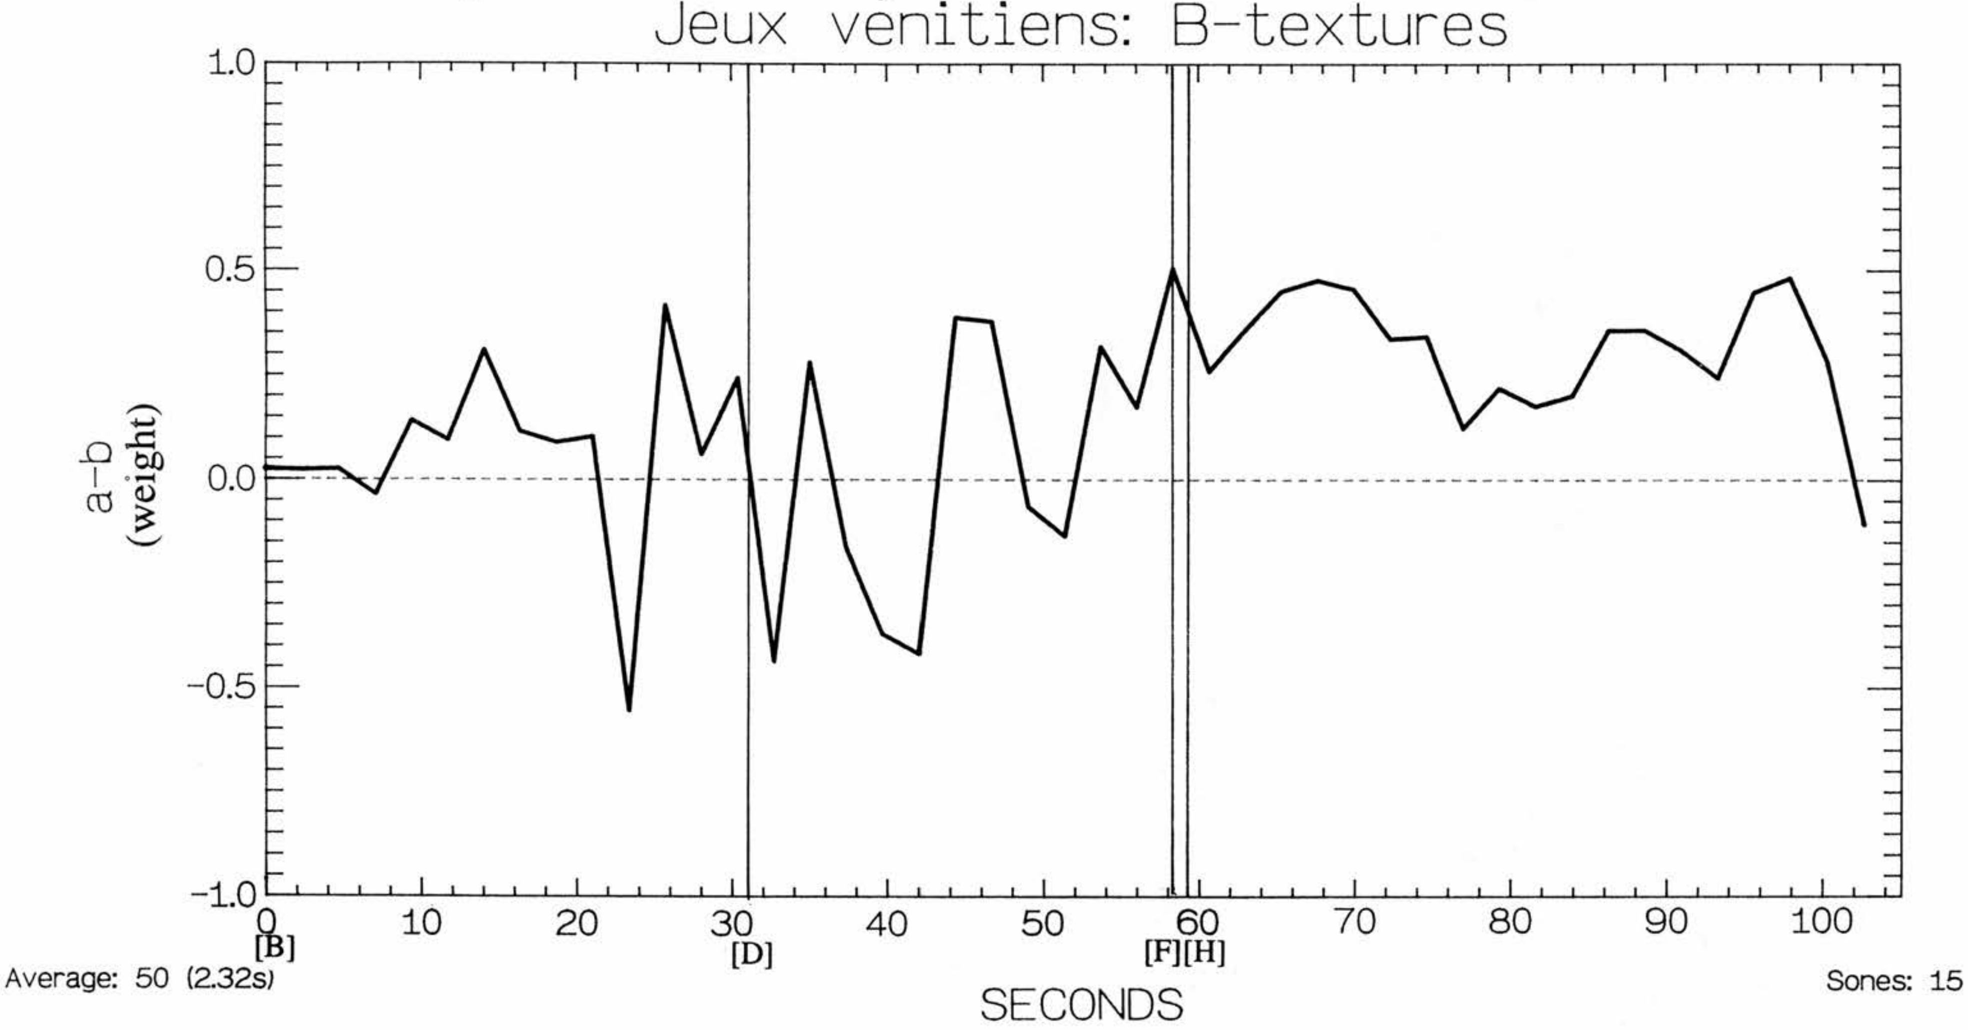

Figure 3.36
Timbral weight: recording 2

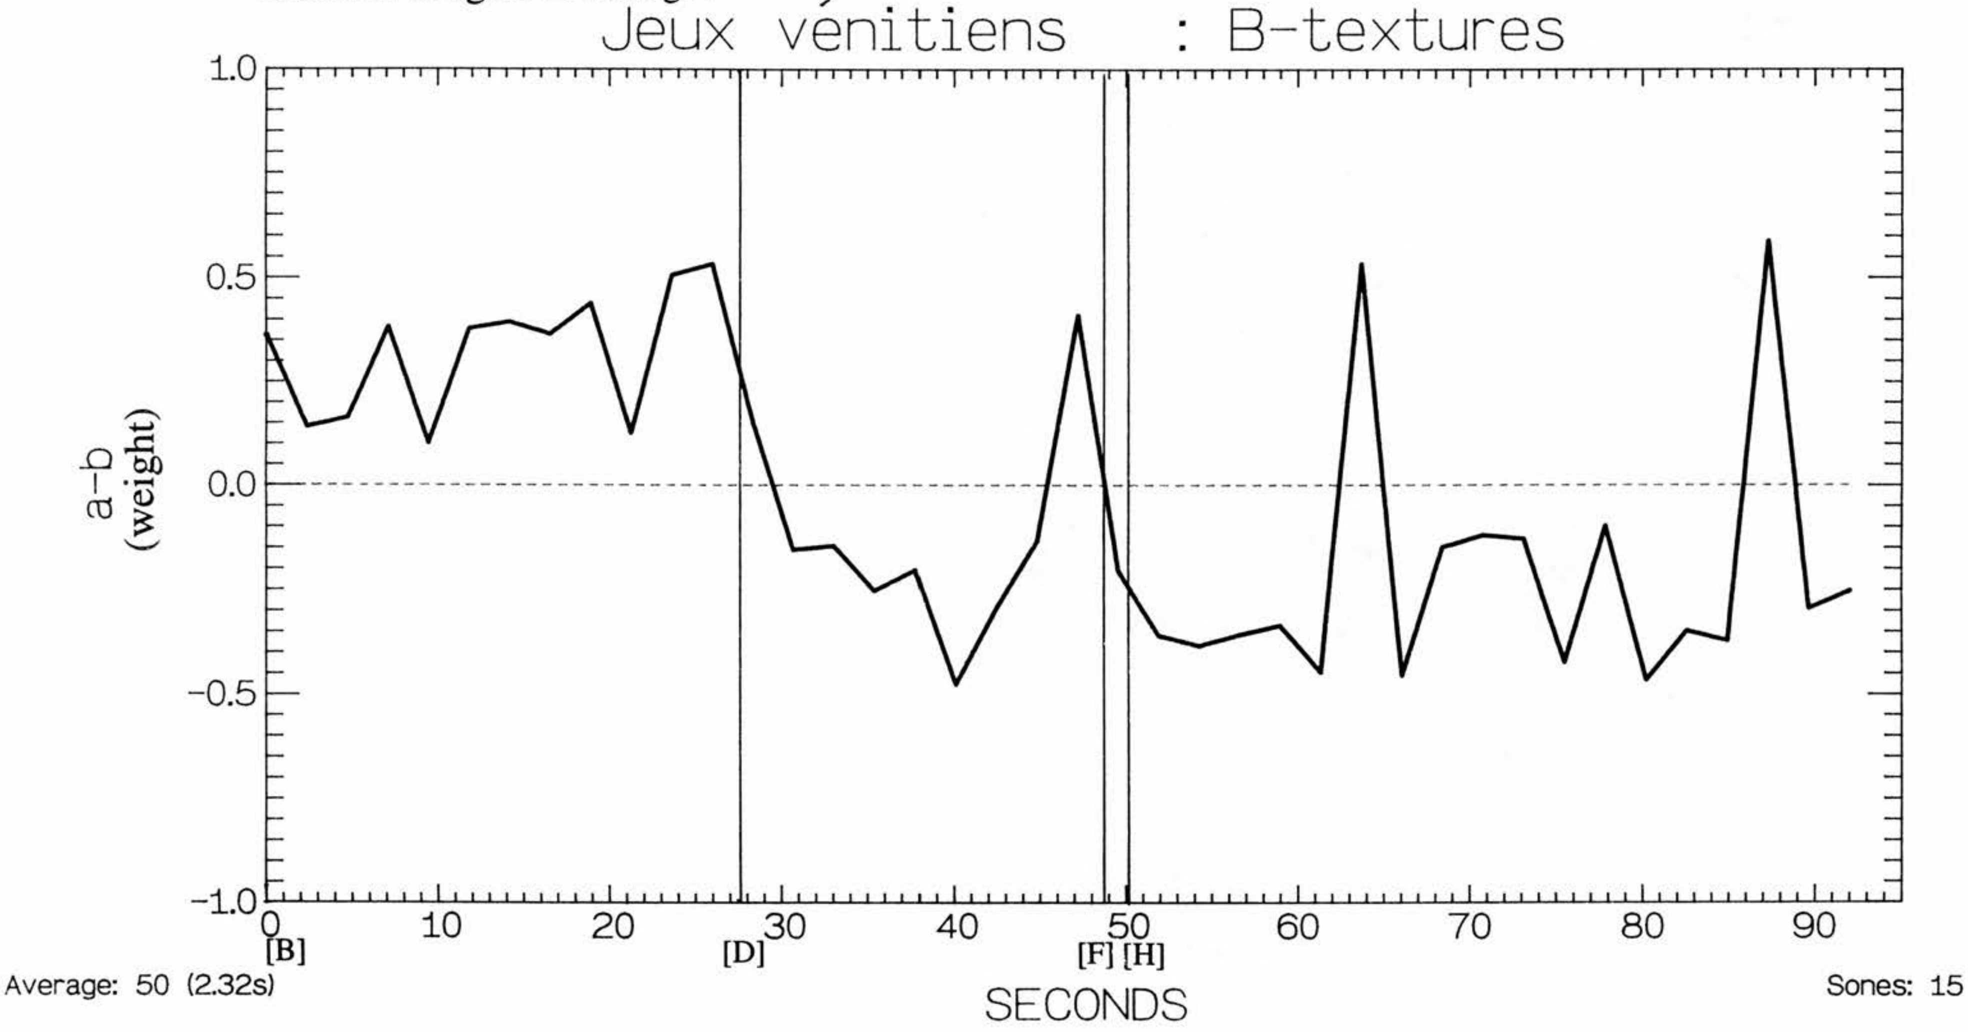

Figure 3.37

Timbral width: recording 1

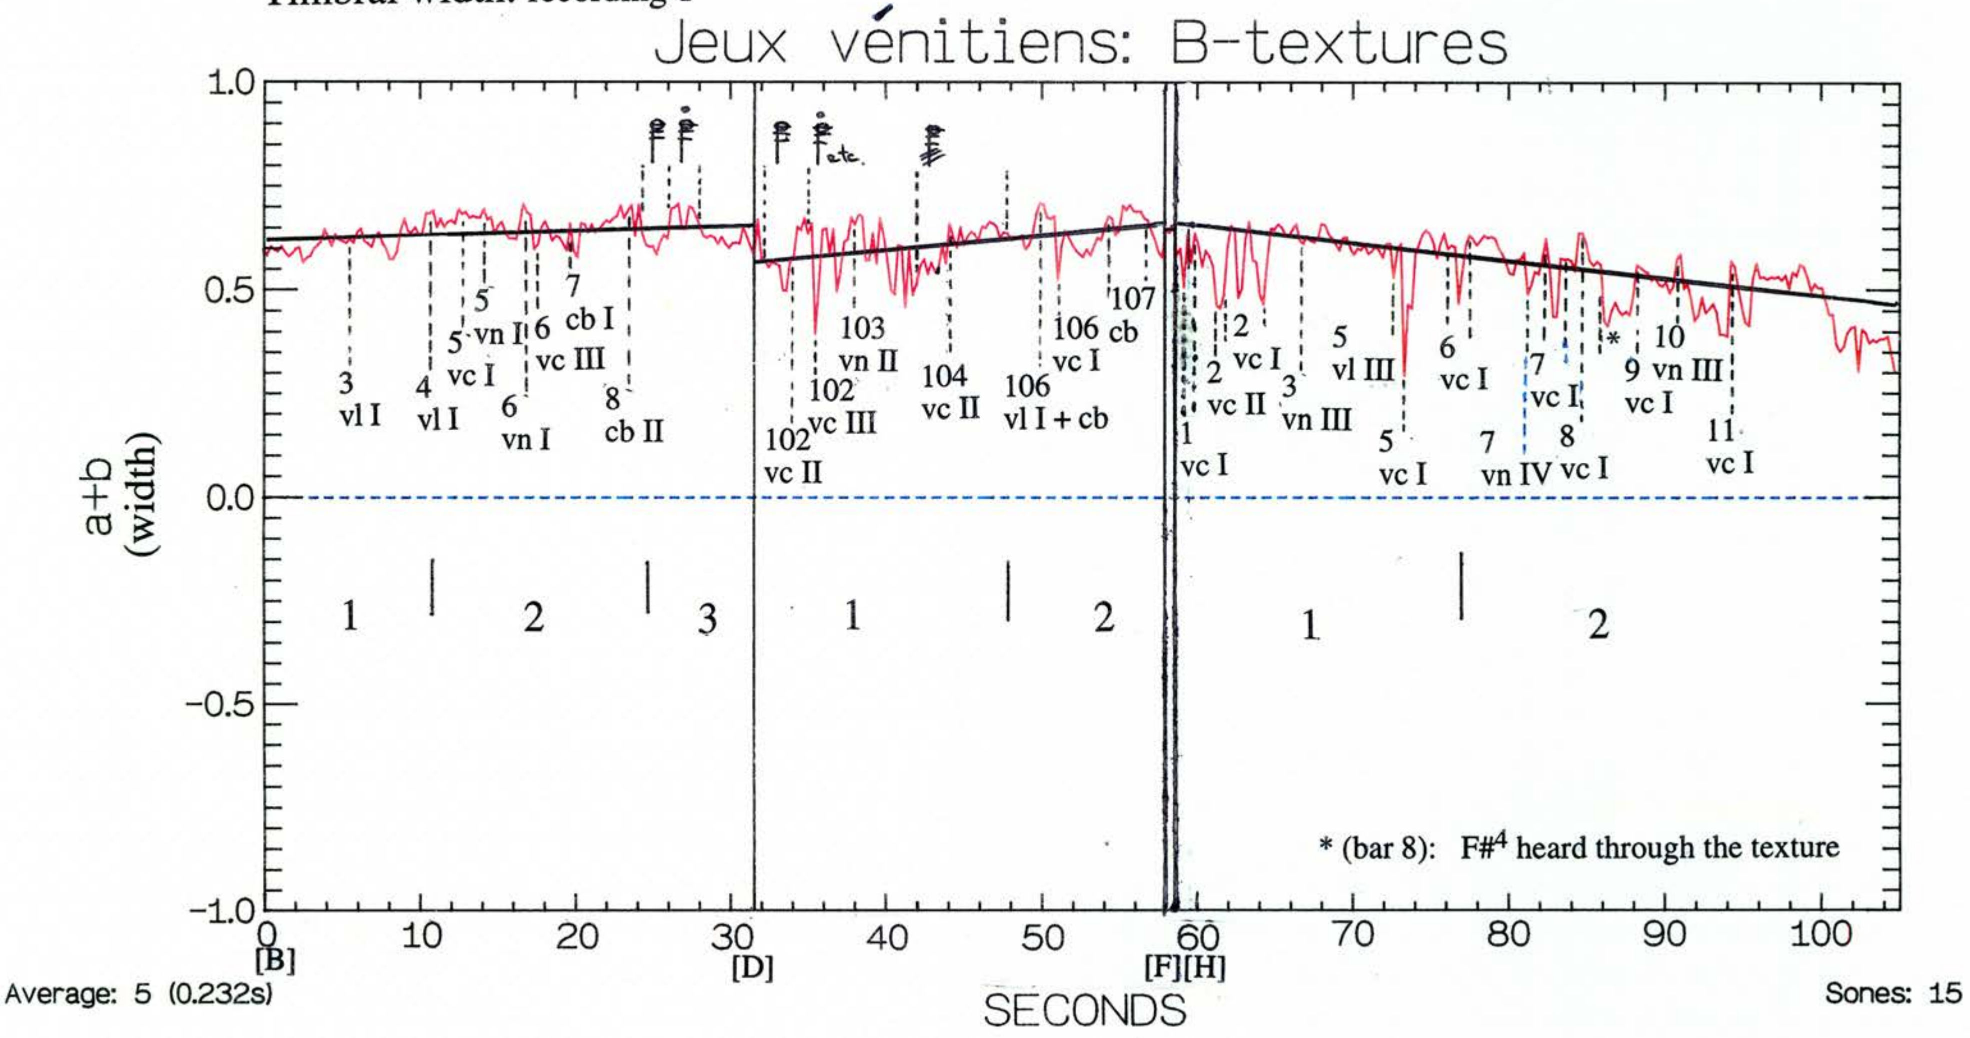

Figure 3.38

Timbral width: recording 2

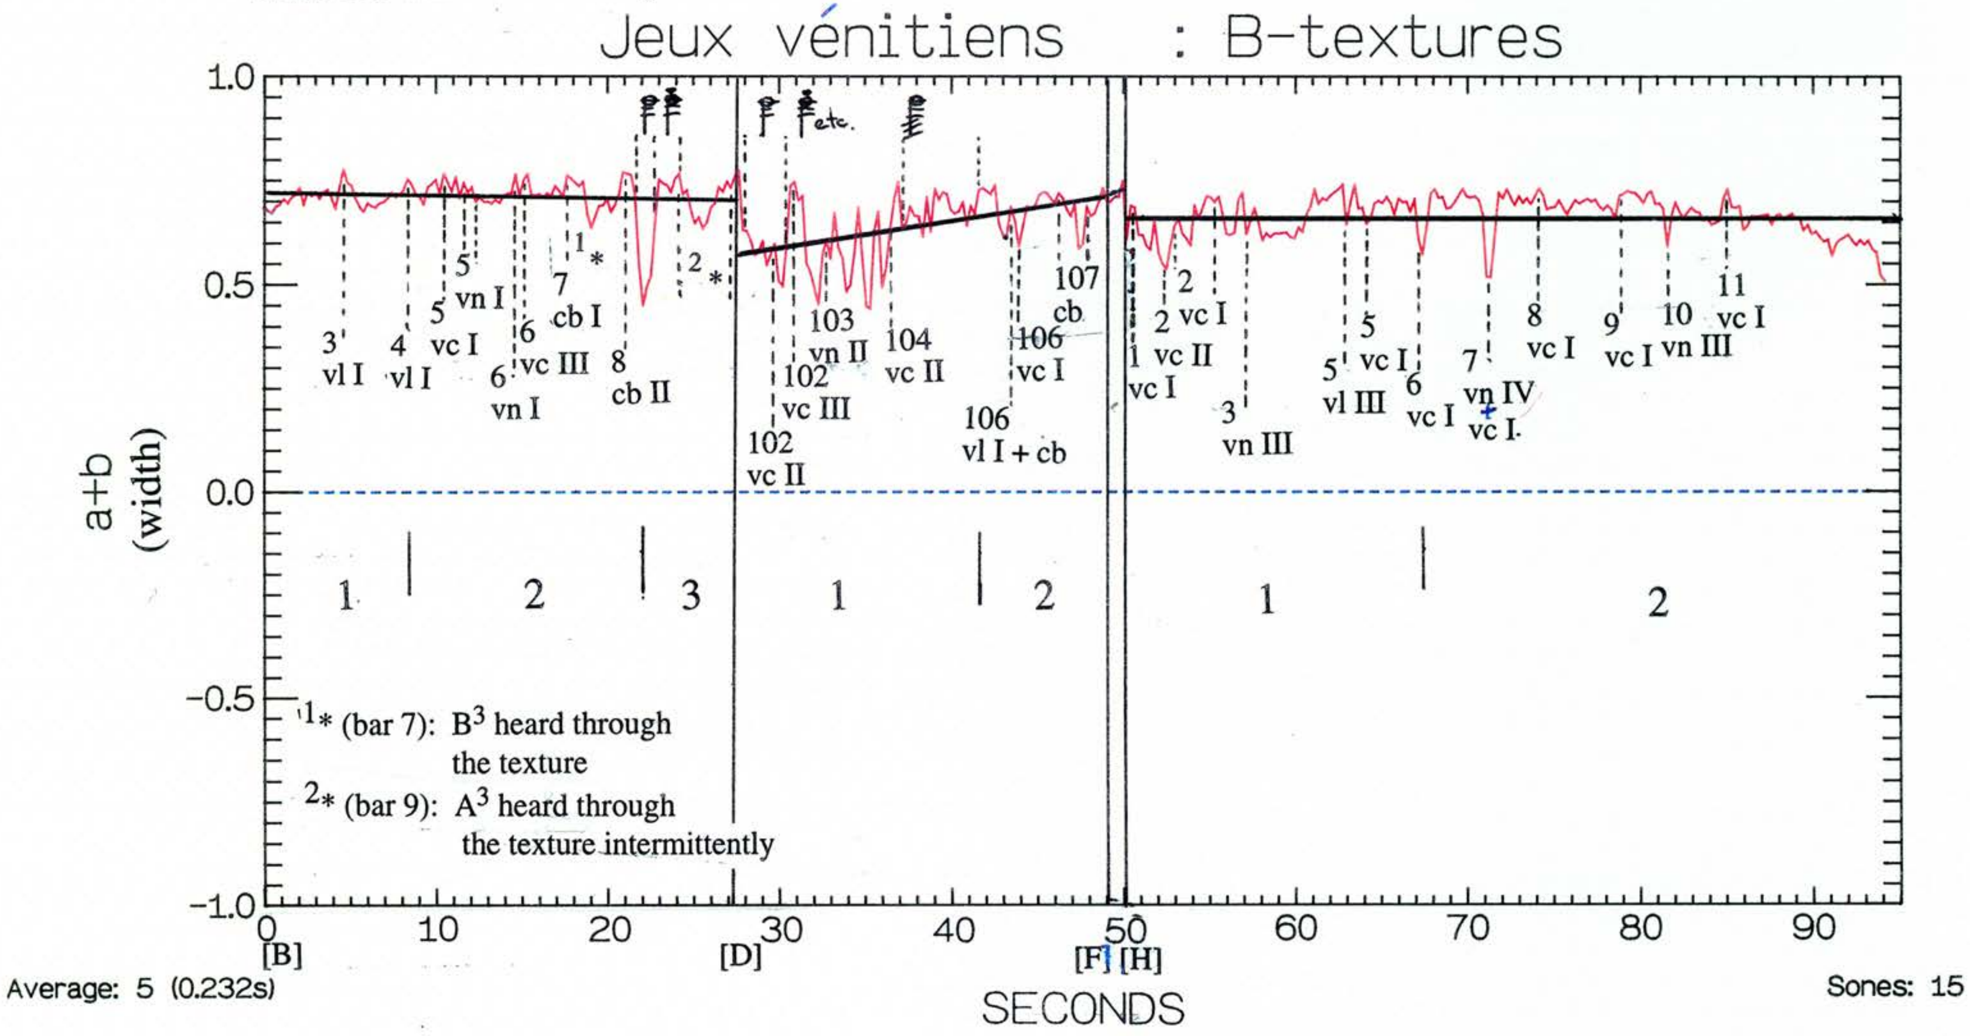

Figure 3.39

Combined Timbral Graphs: recording 1

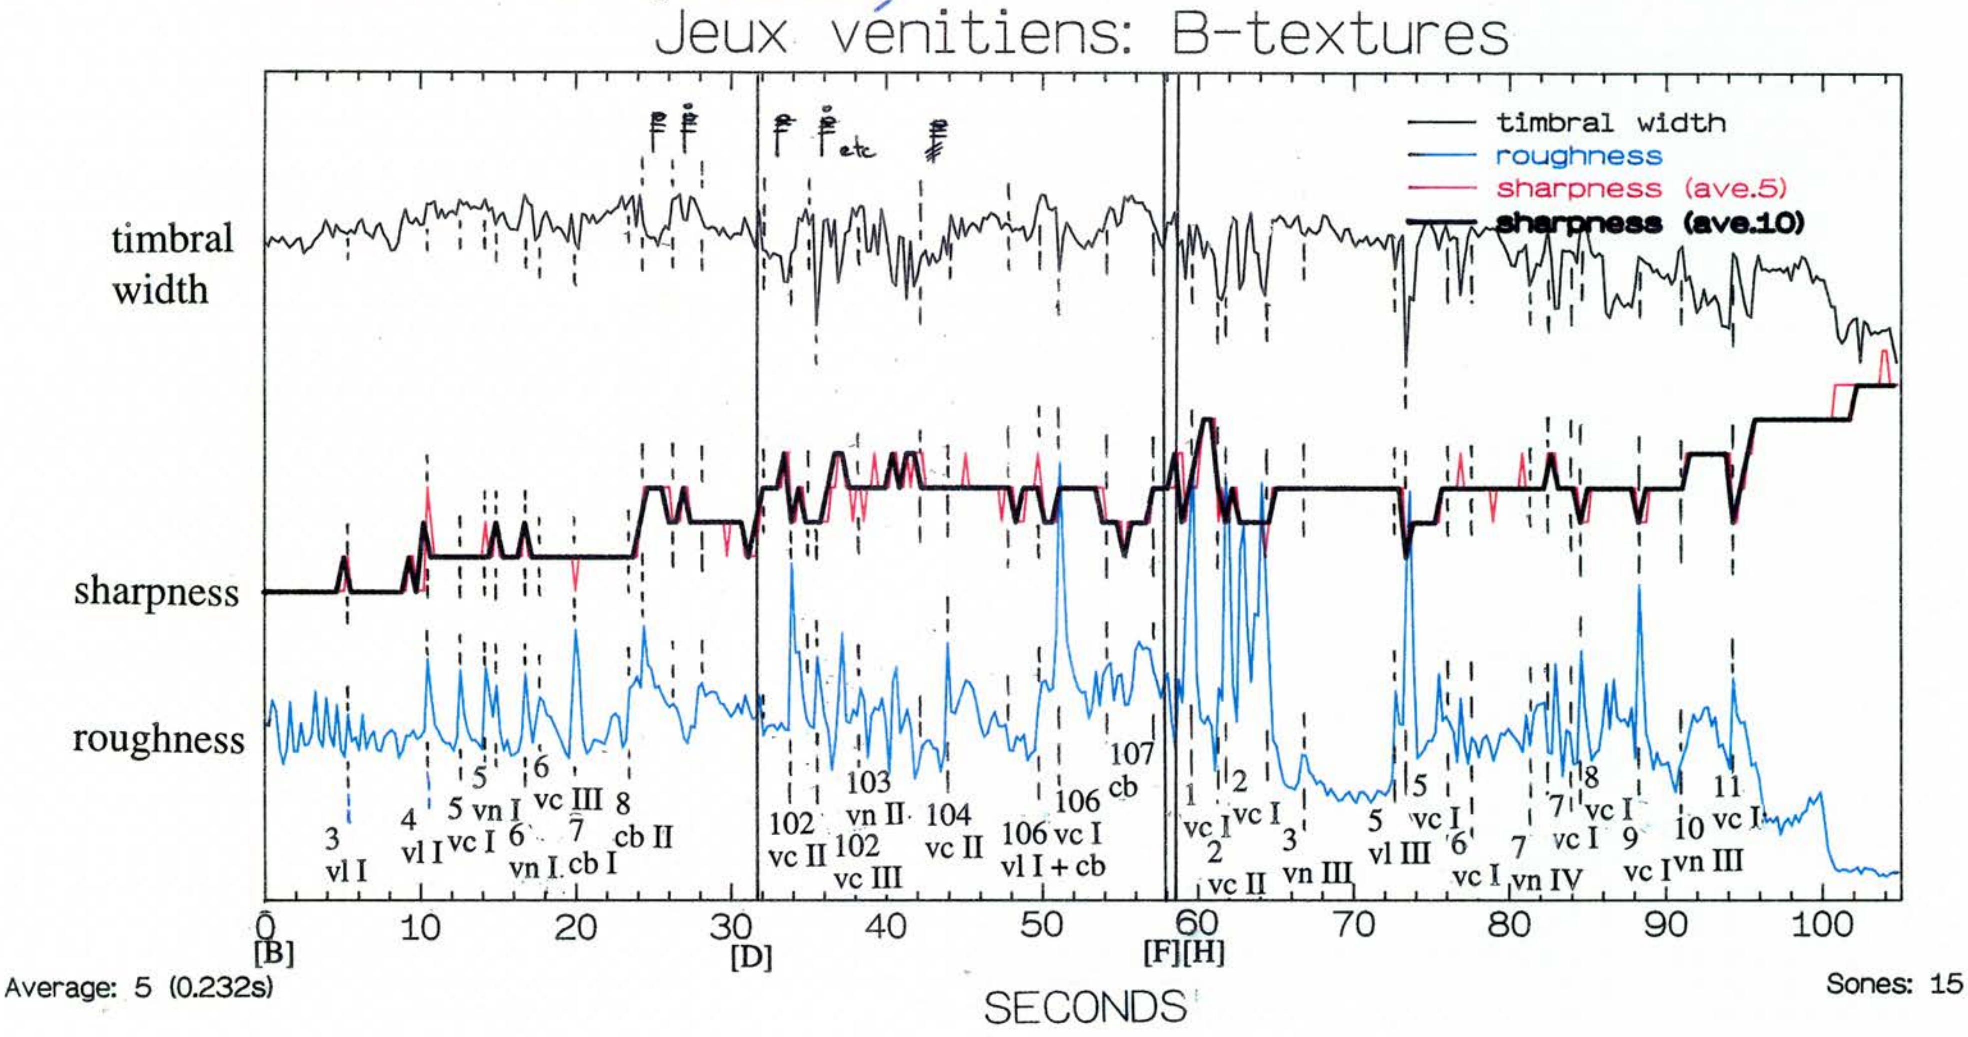

Figure 3.40
Combined Timbral Graphs: recording 2

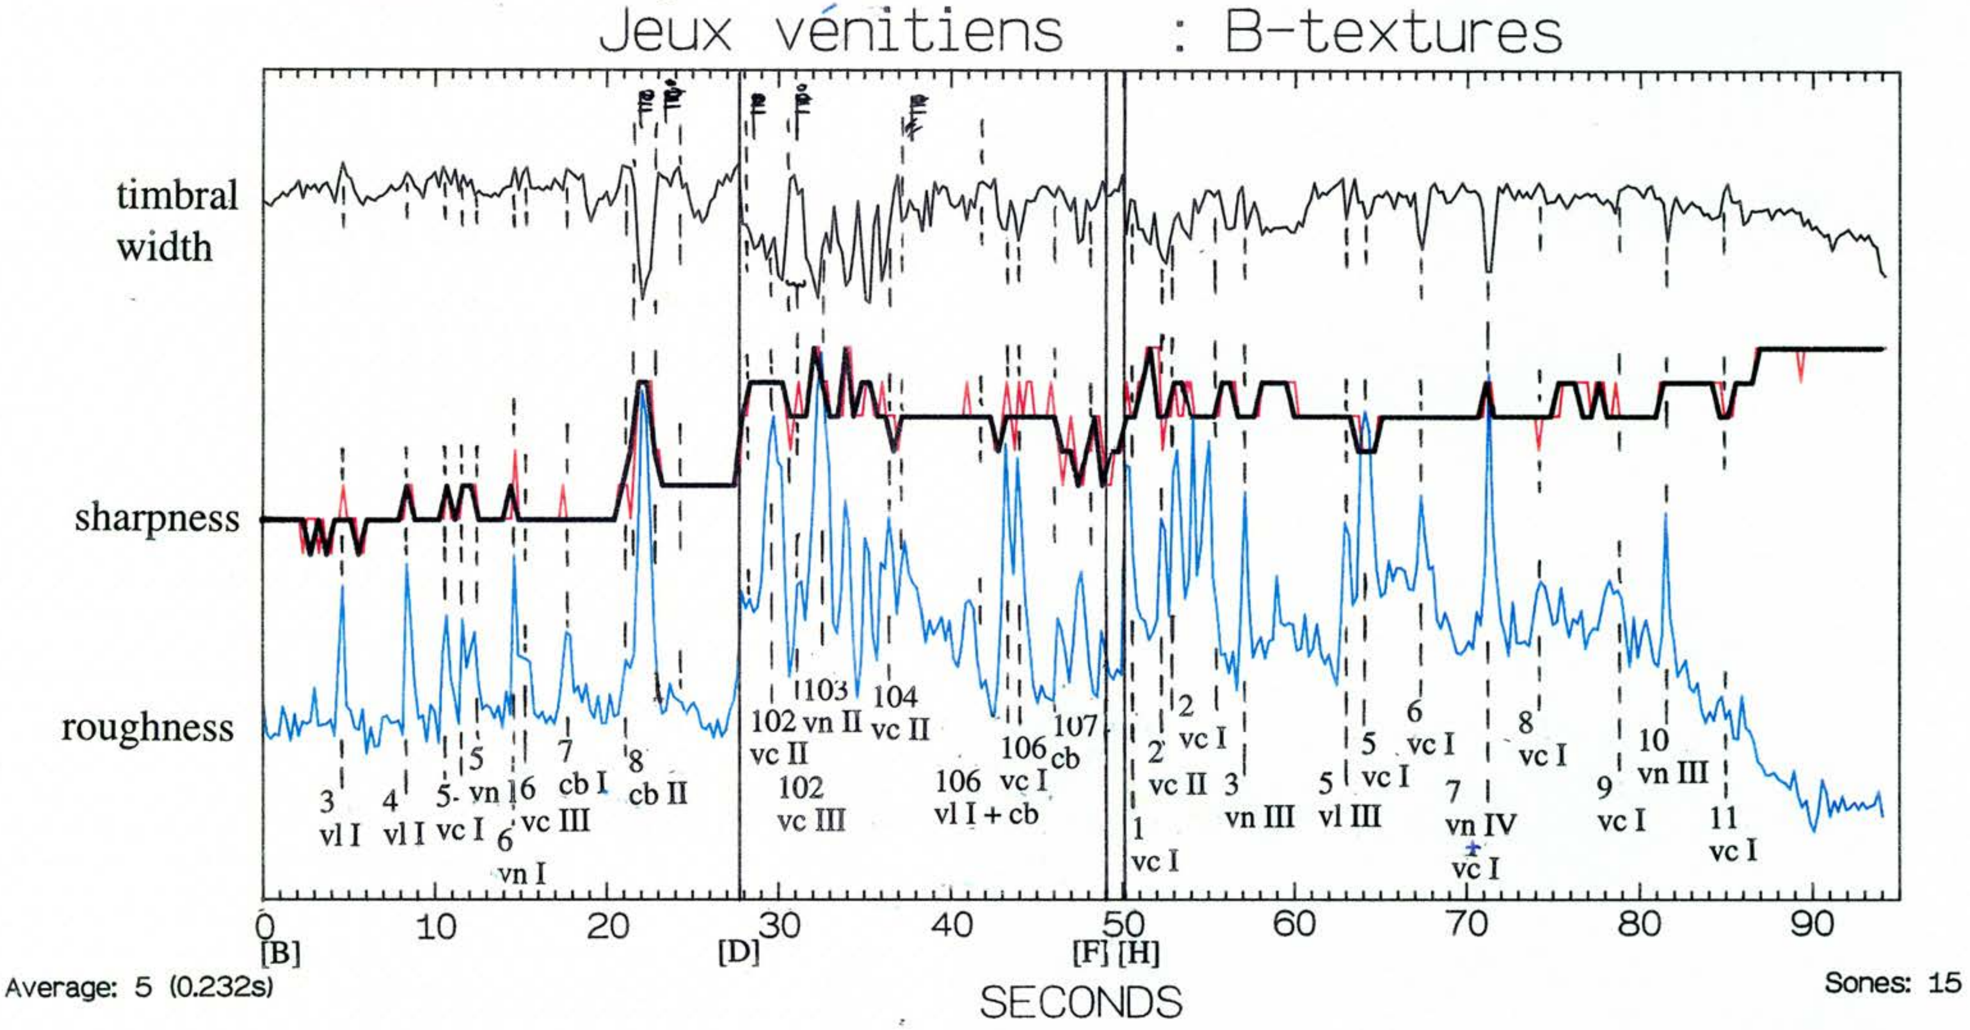

Figure 3.41

Combined Timbral Graphs: recording 1

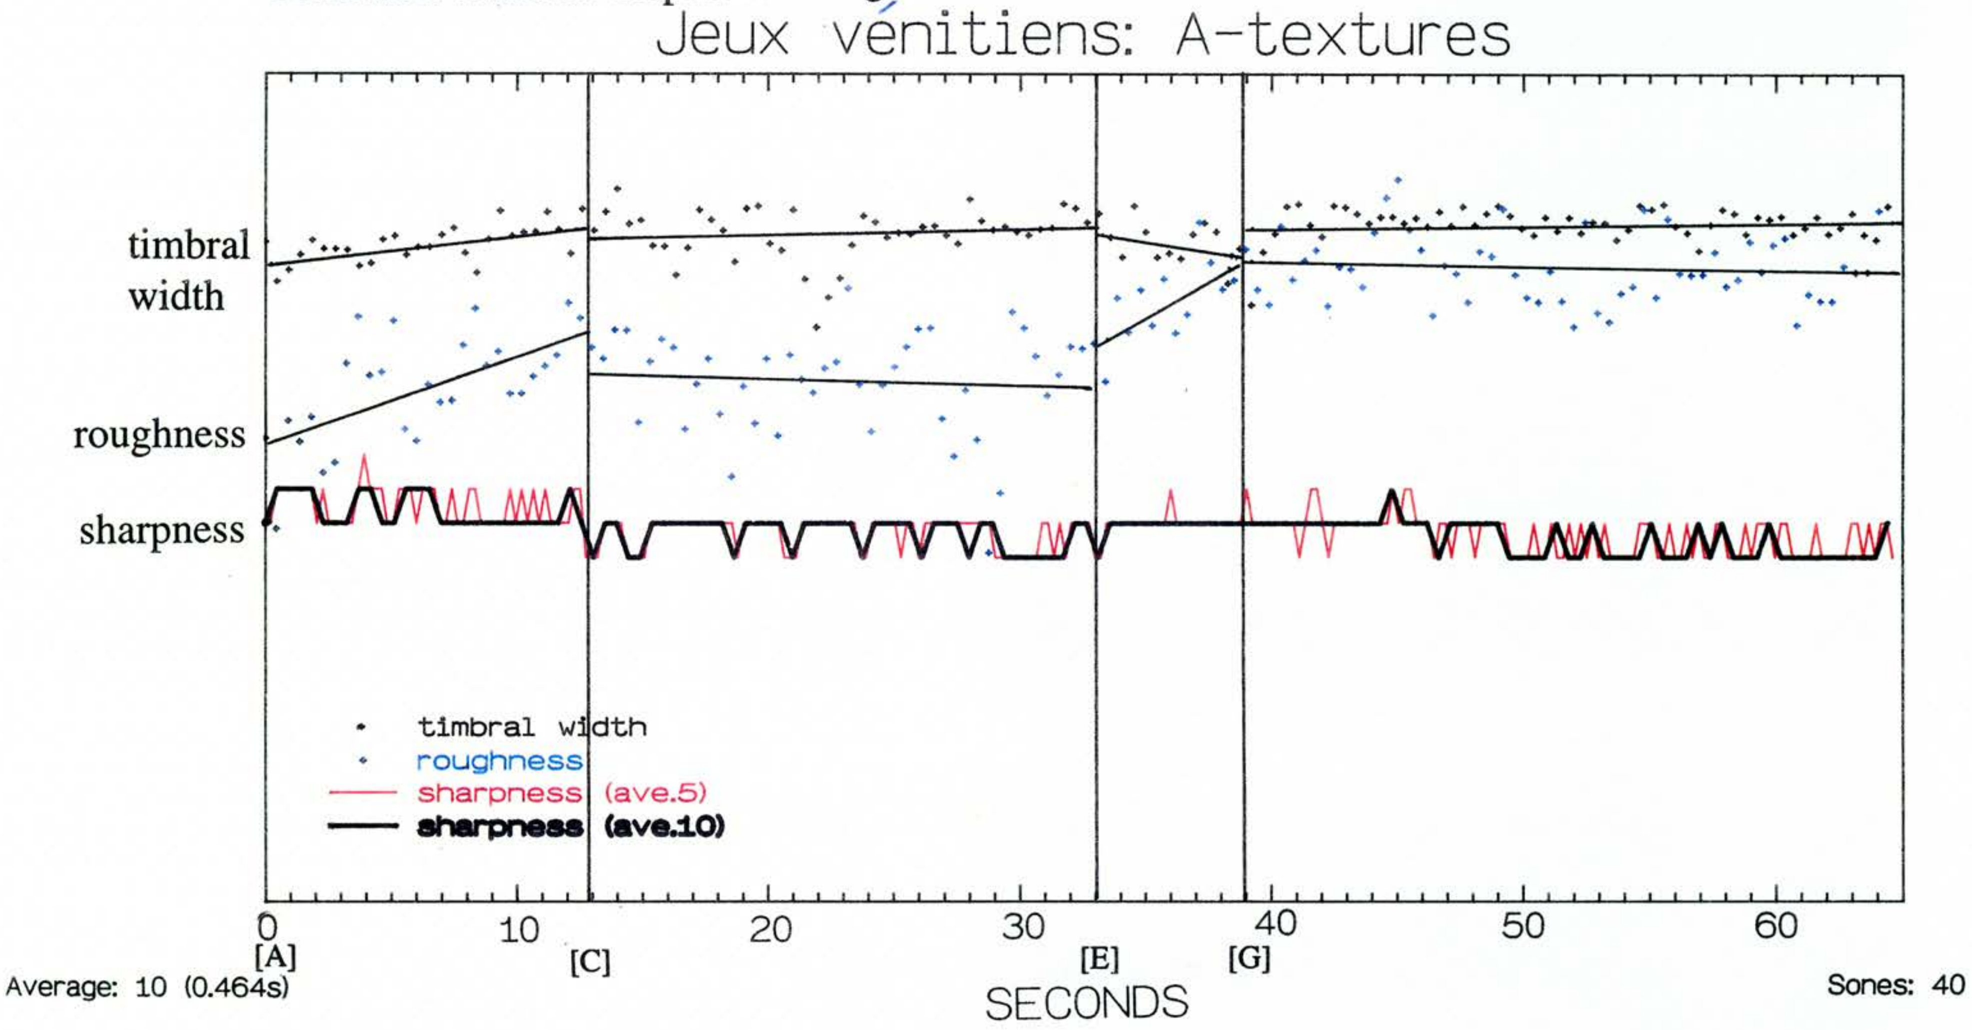

Figure 3.42

Combined Timbral Graphs: recording 2

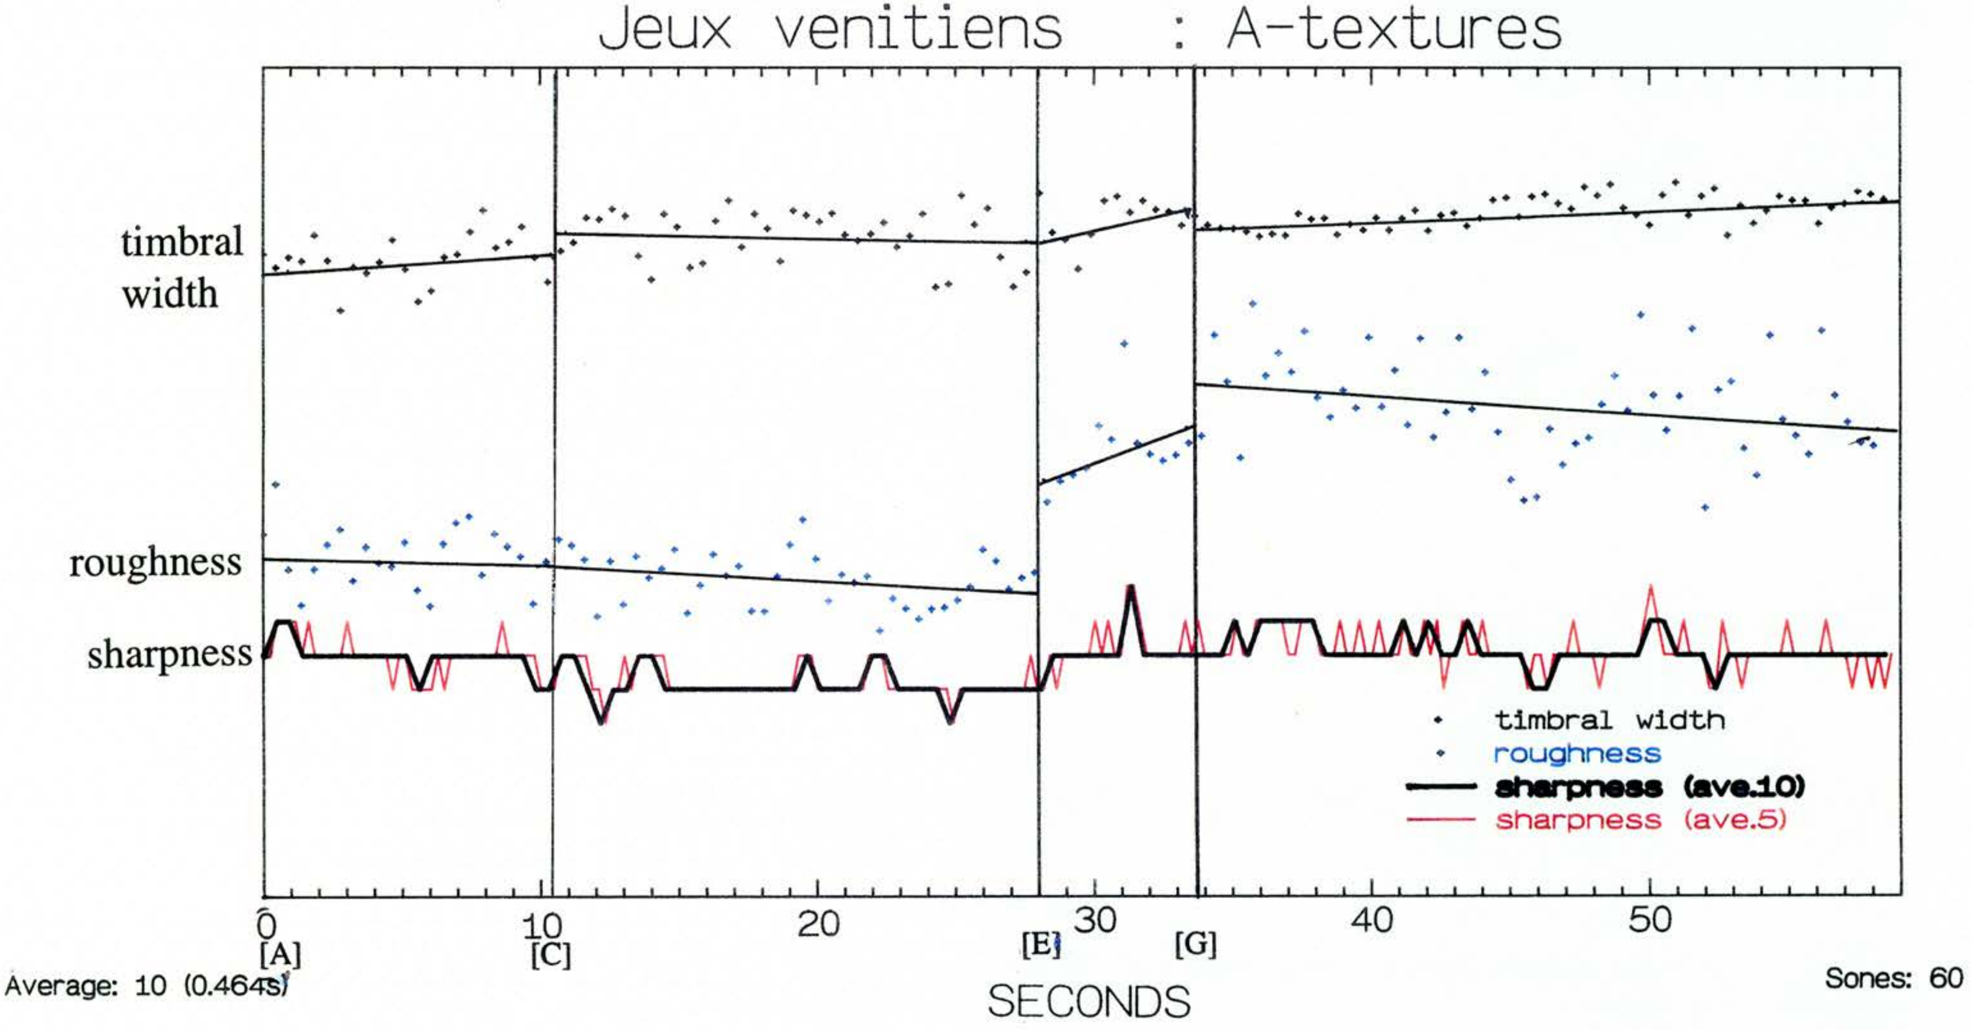

Figure 3.43

Combined Timbral Graphs: recording 1

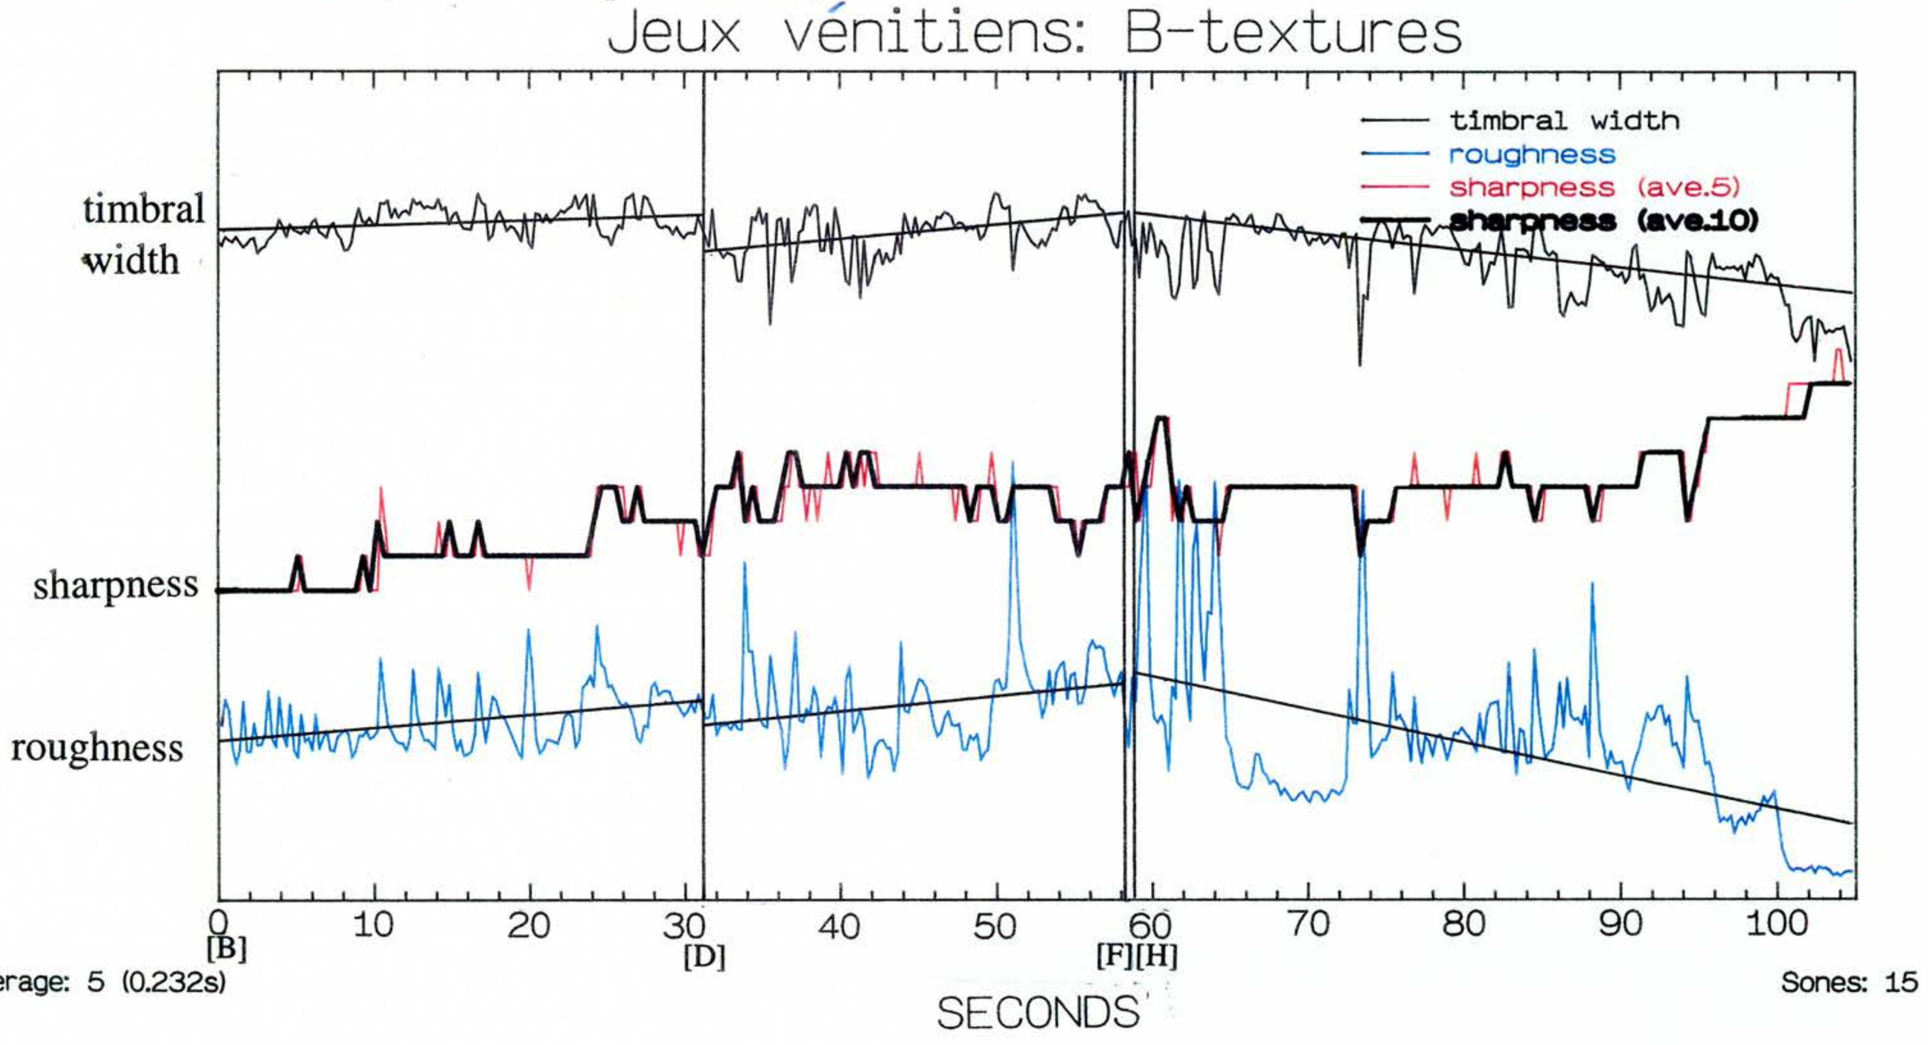

Figure 3.44

Combined Timbral Graphs: recording 2

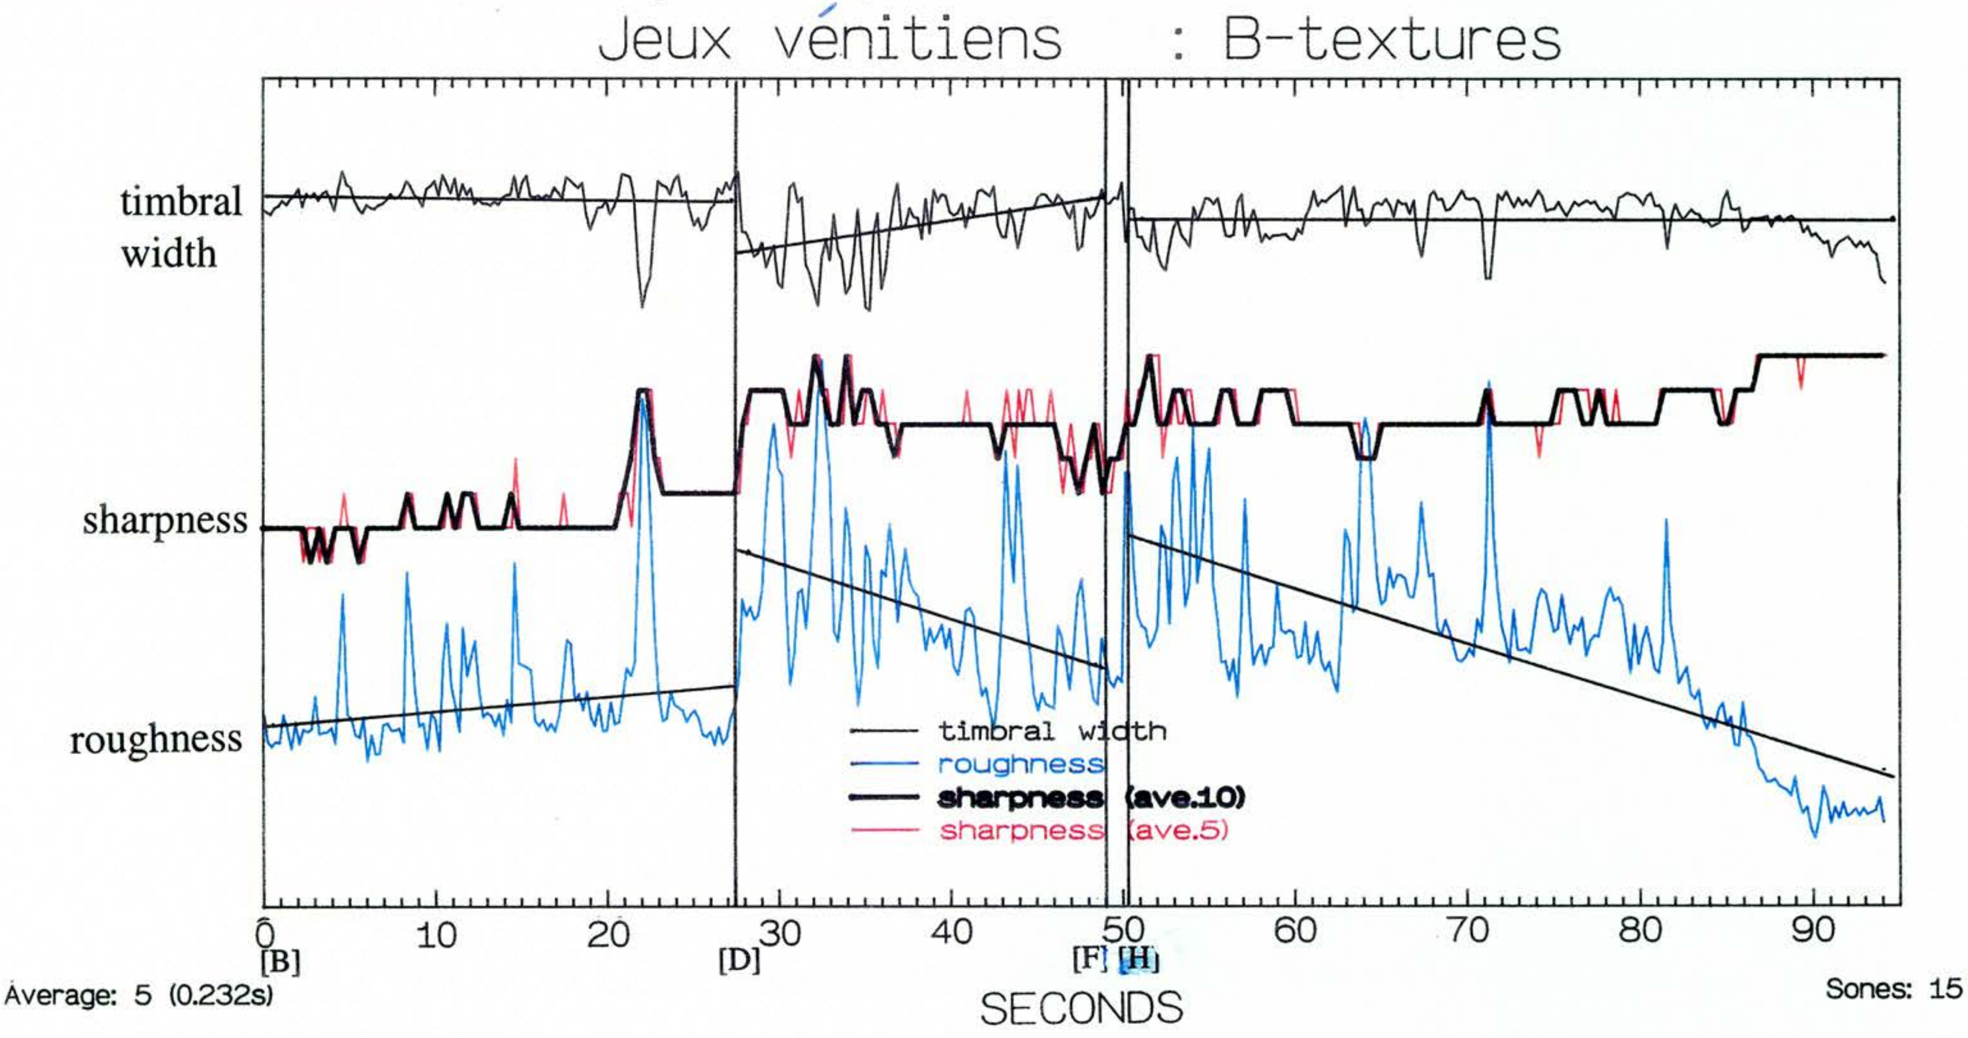

Figures for

chapter 4

Analysis of *Atmosphères*

Figure 4.1 (part 1) Level 4

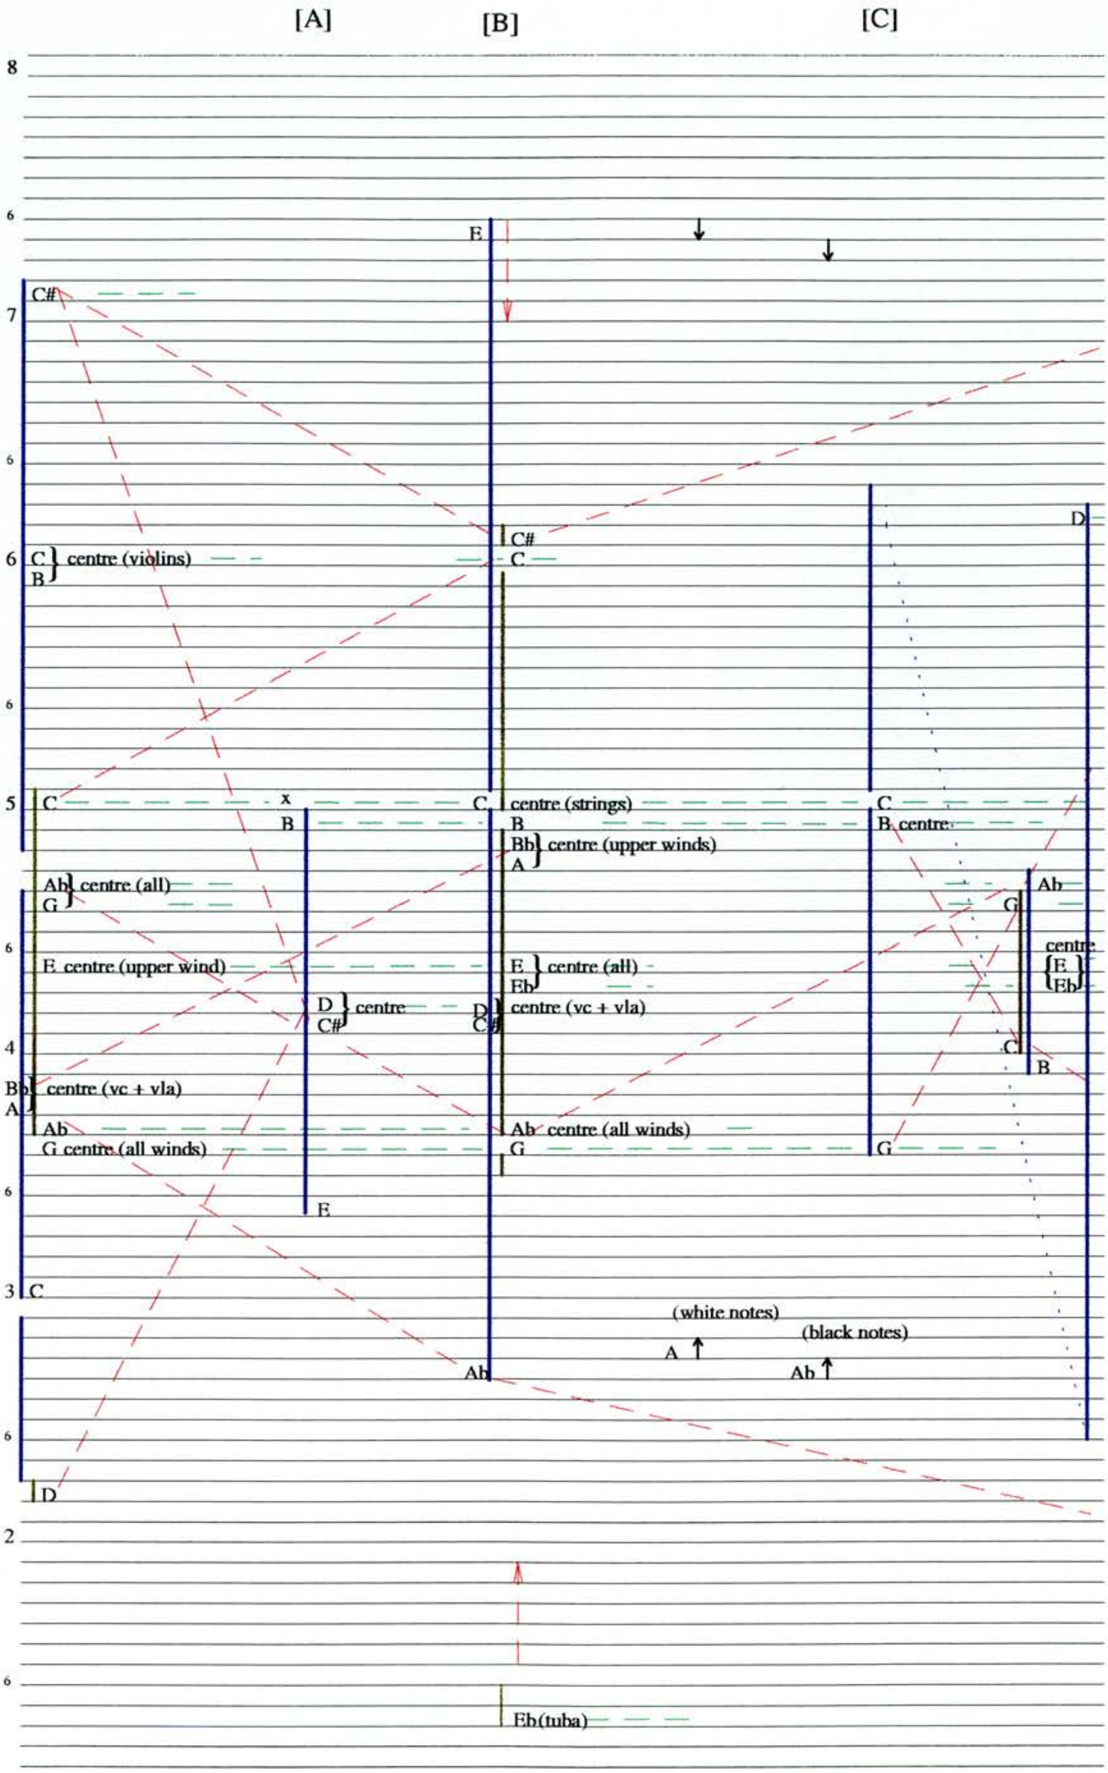

Figure 4.1 (part 2) Level 4

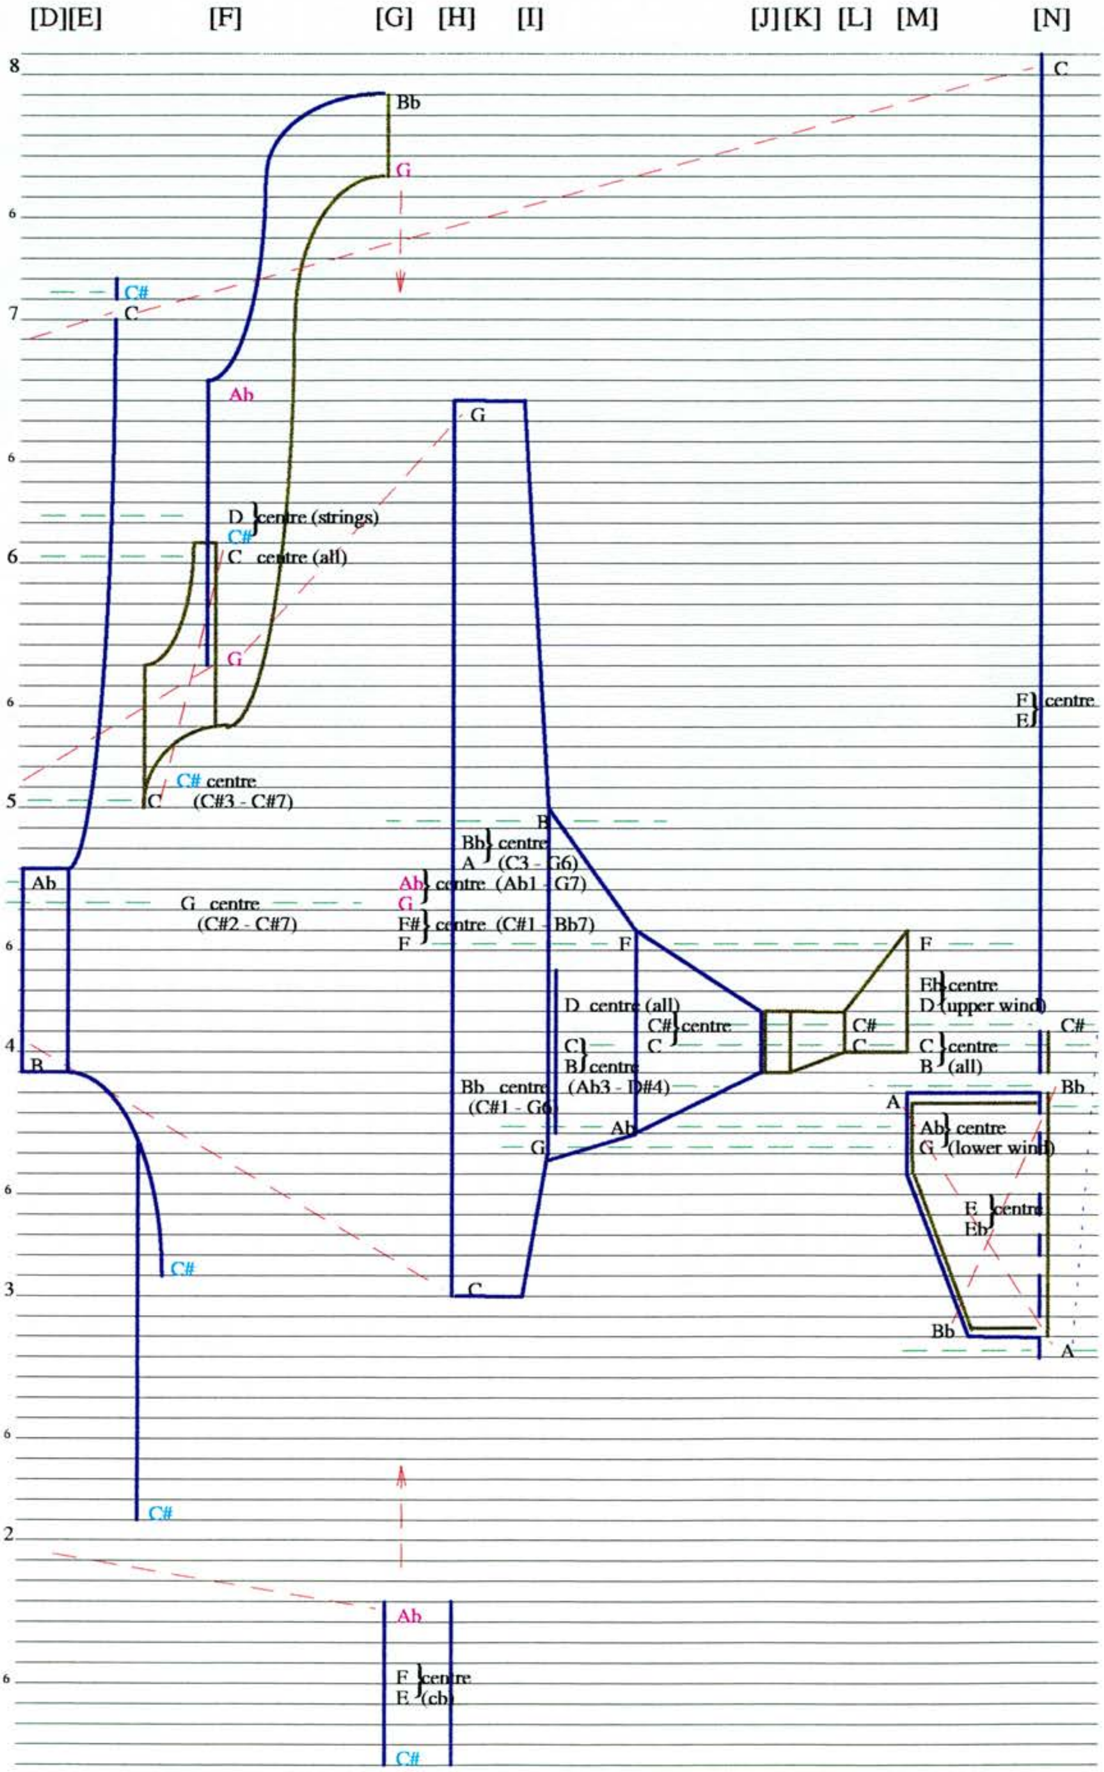

Figure 4.1 (part 3) Level 4

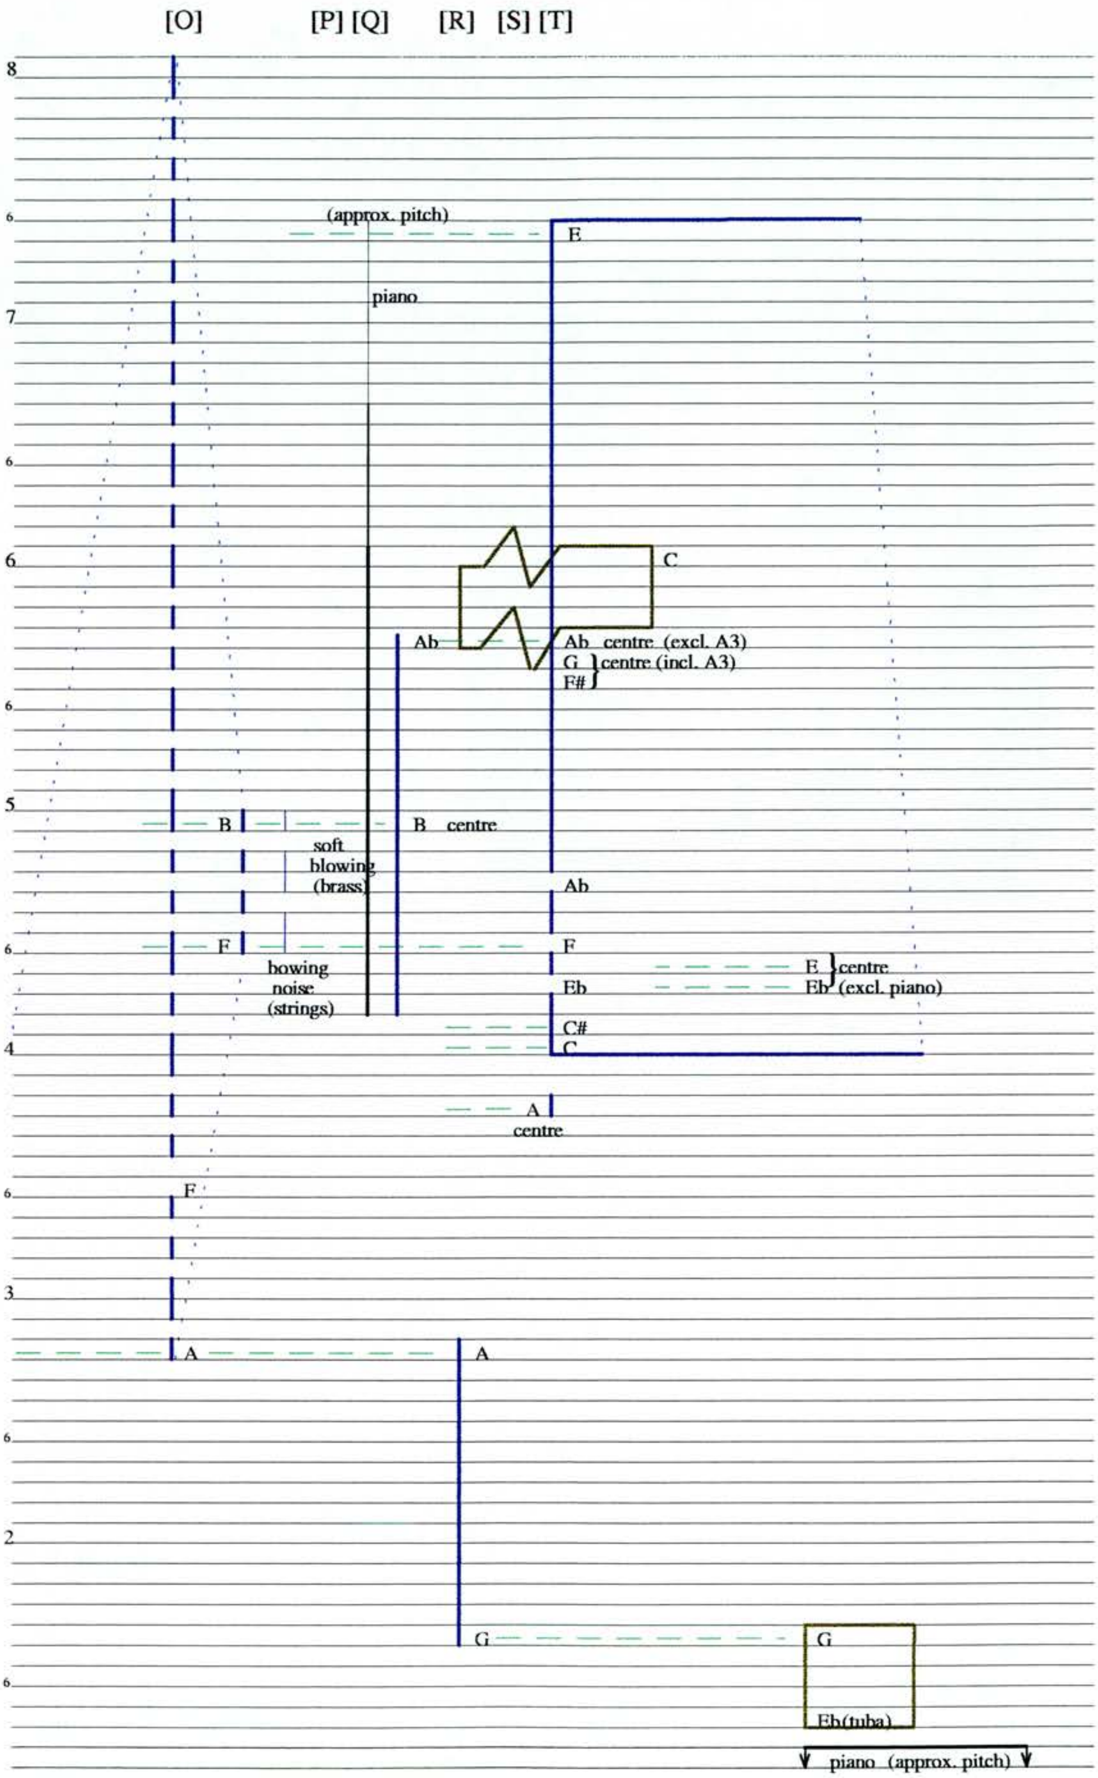

Figure 4.2

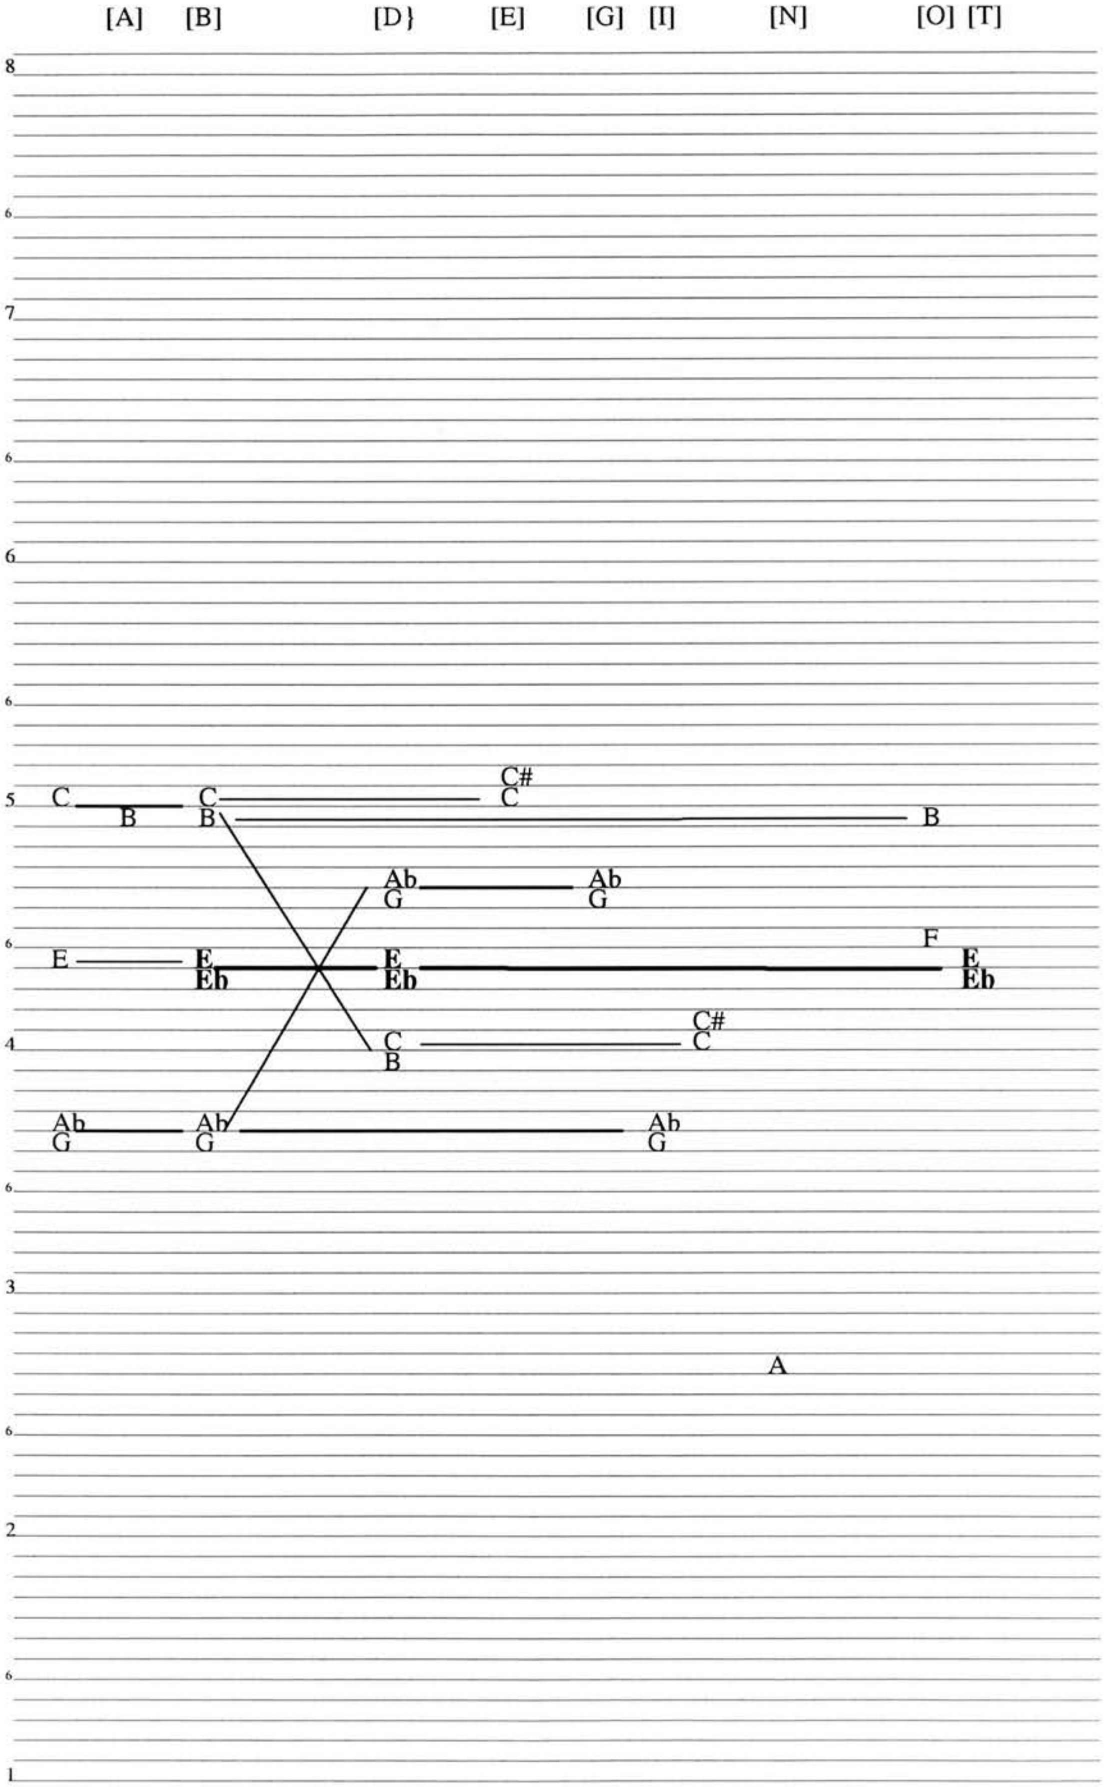

Level 3

Figure 4.3

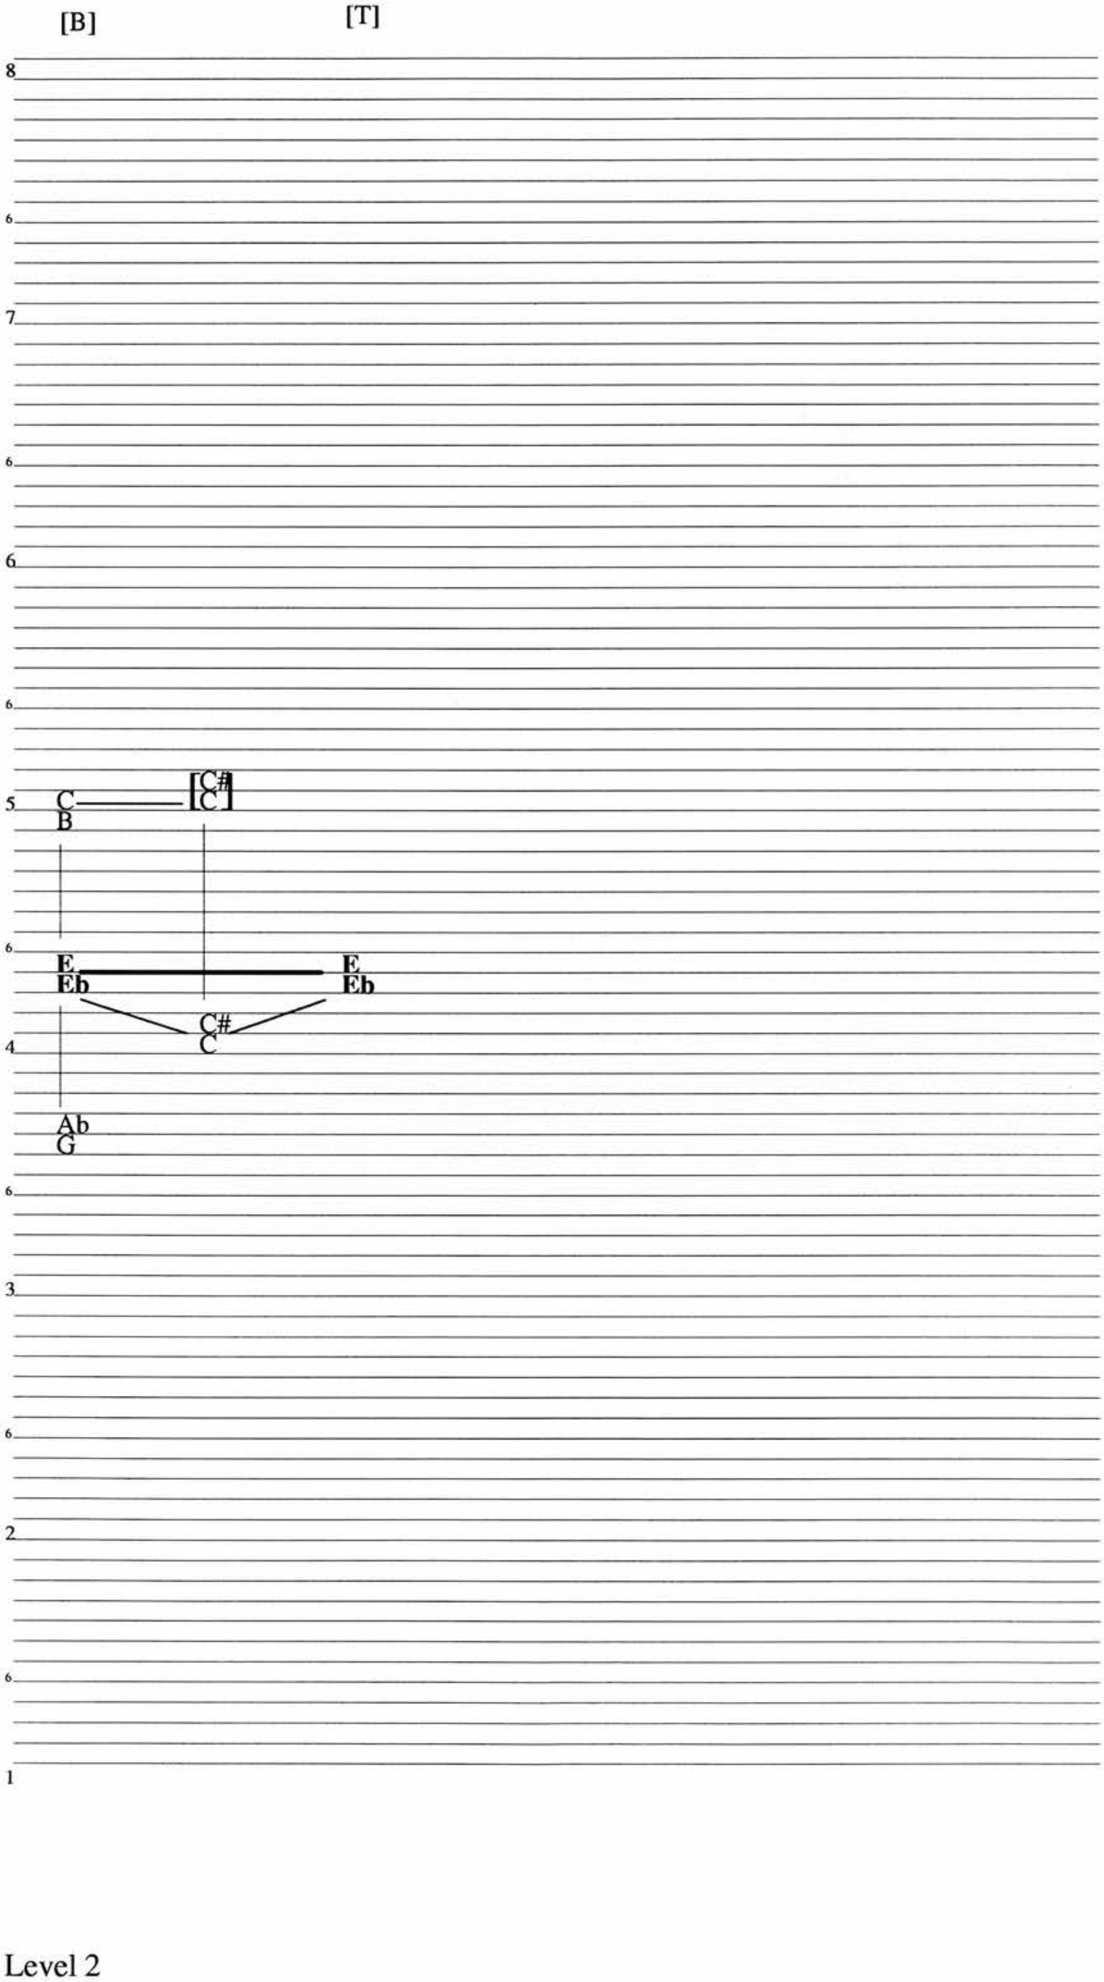

Level 2

Figure 4.4

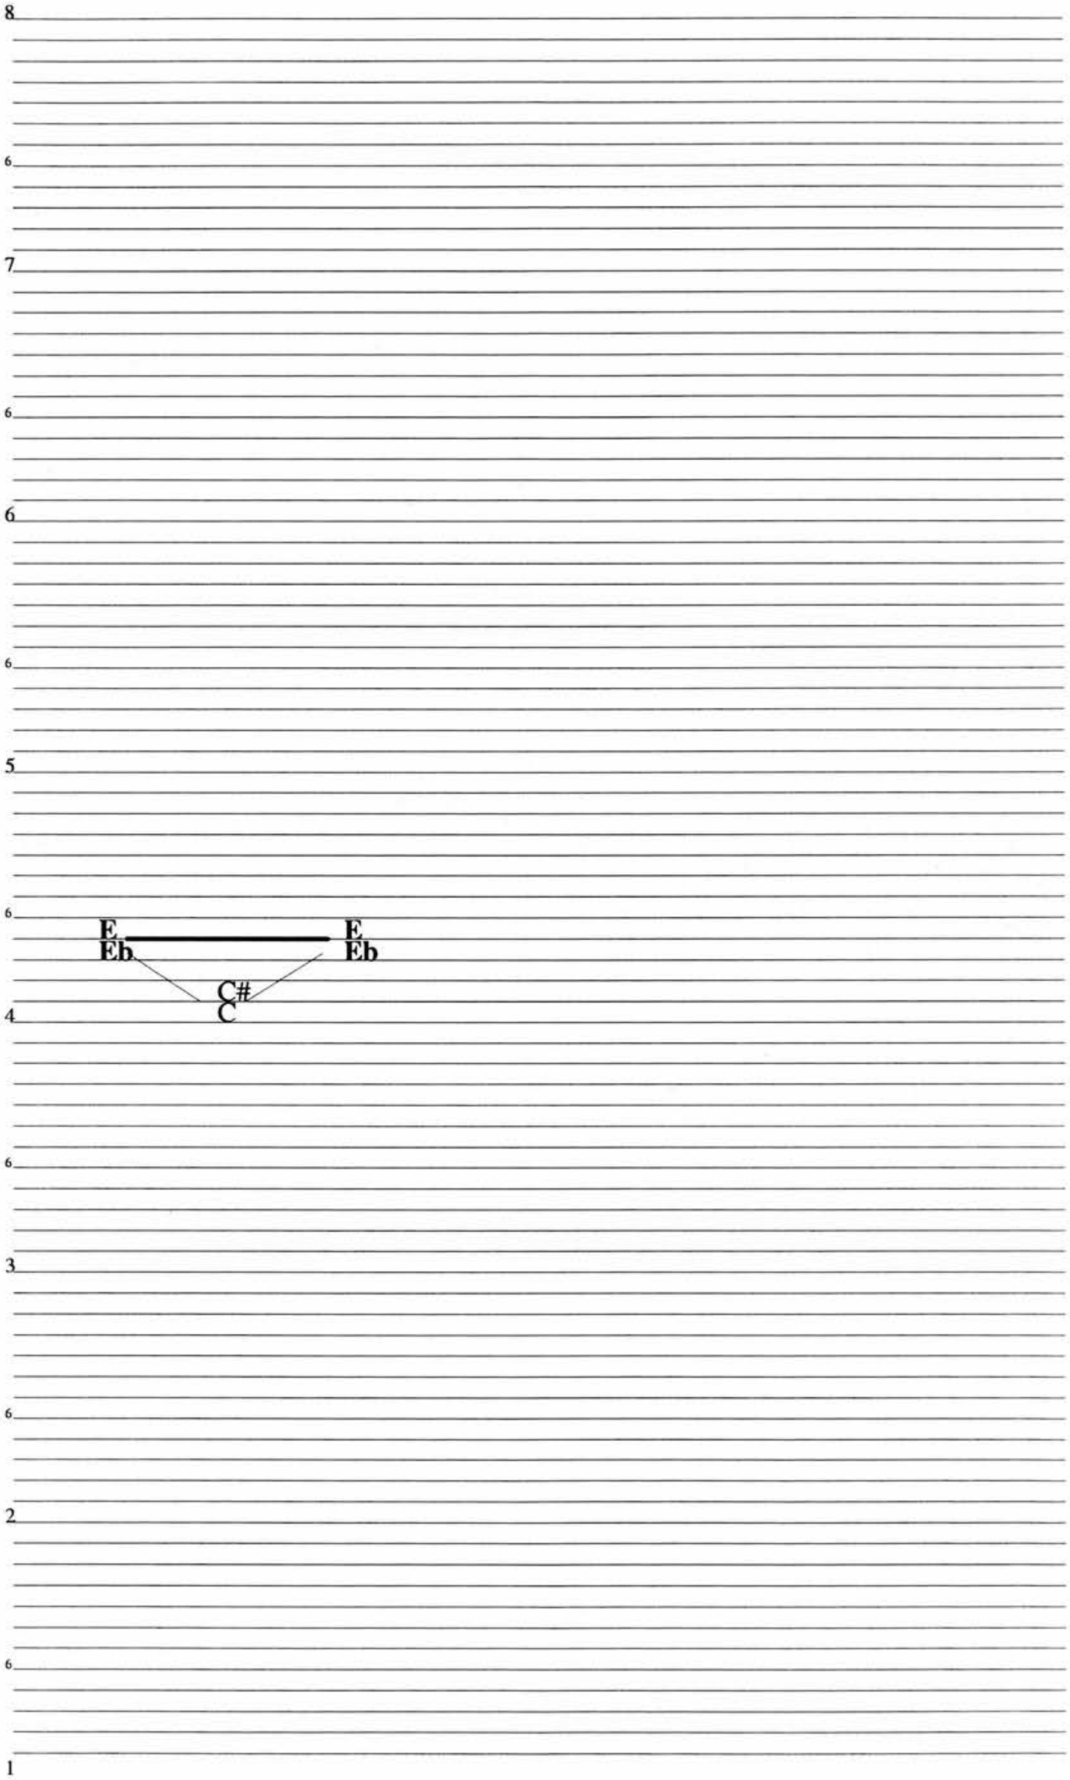

Level 1

Figure 4.5

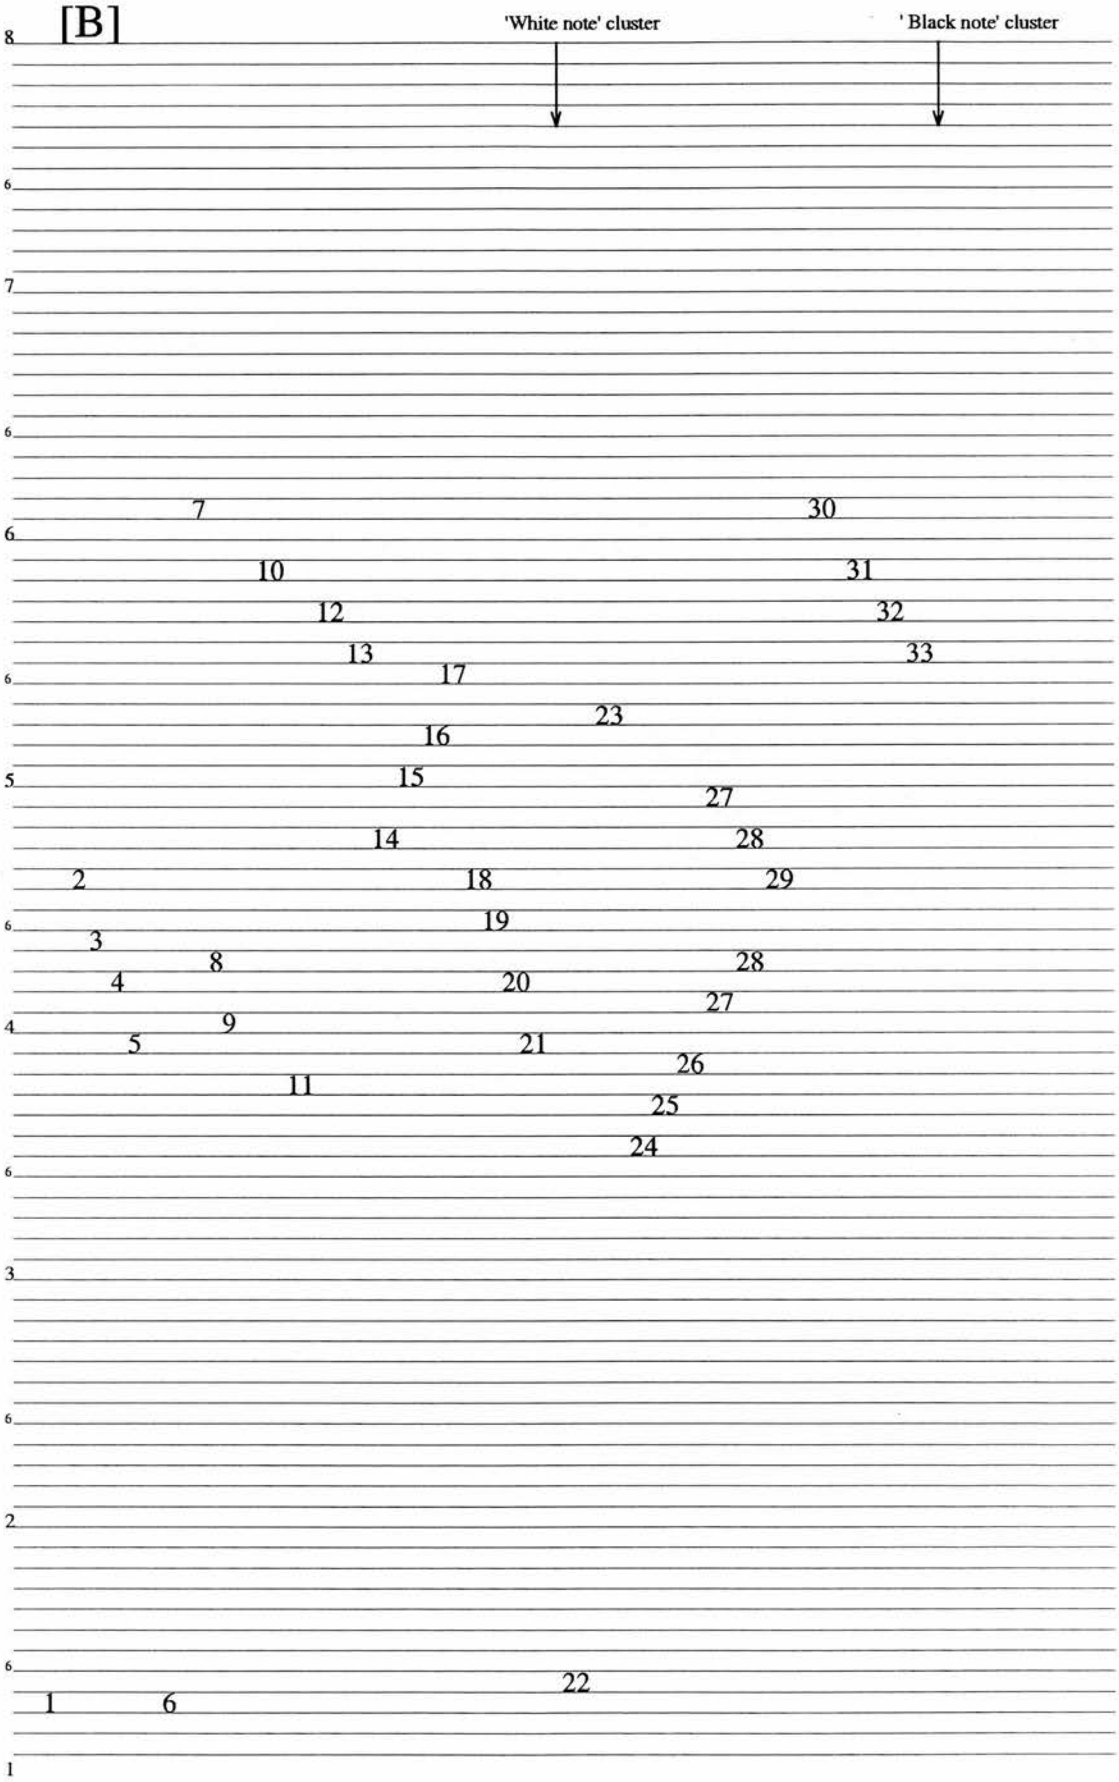

Order of points of re-articulation

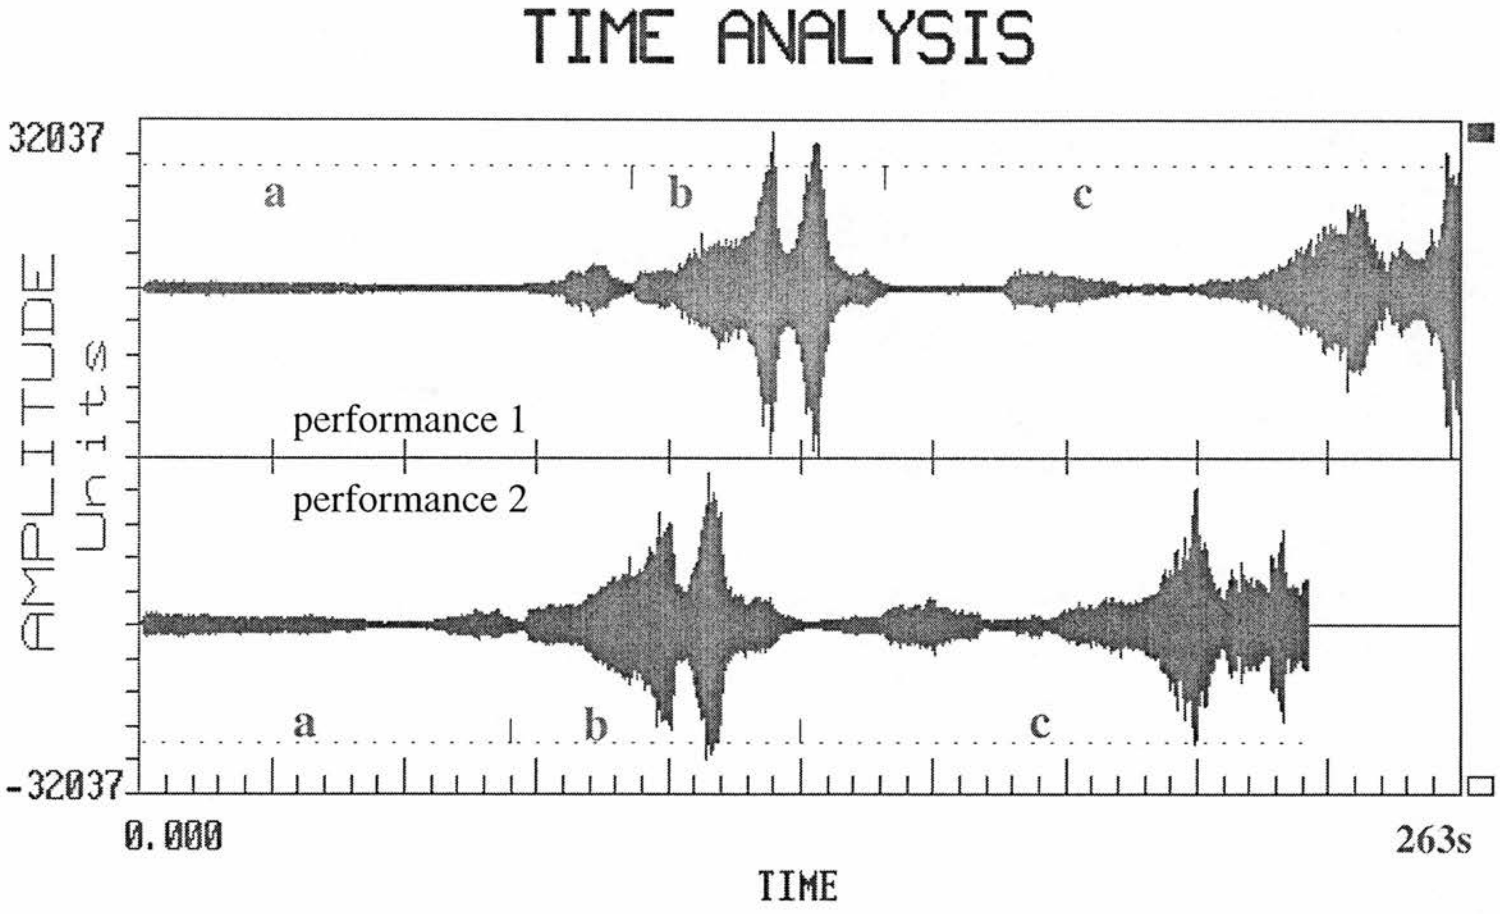

Figure 4.6a Time waveform of first half of performances 1 and 2 of Atmospheres

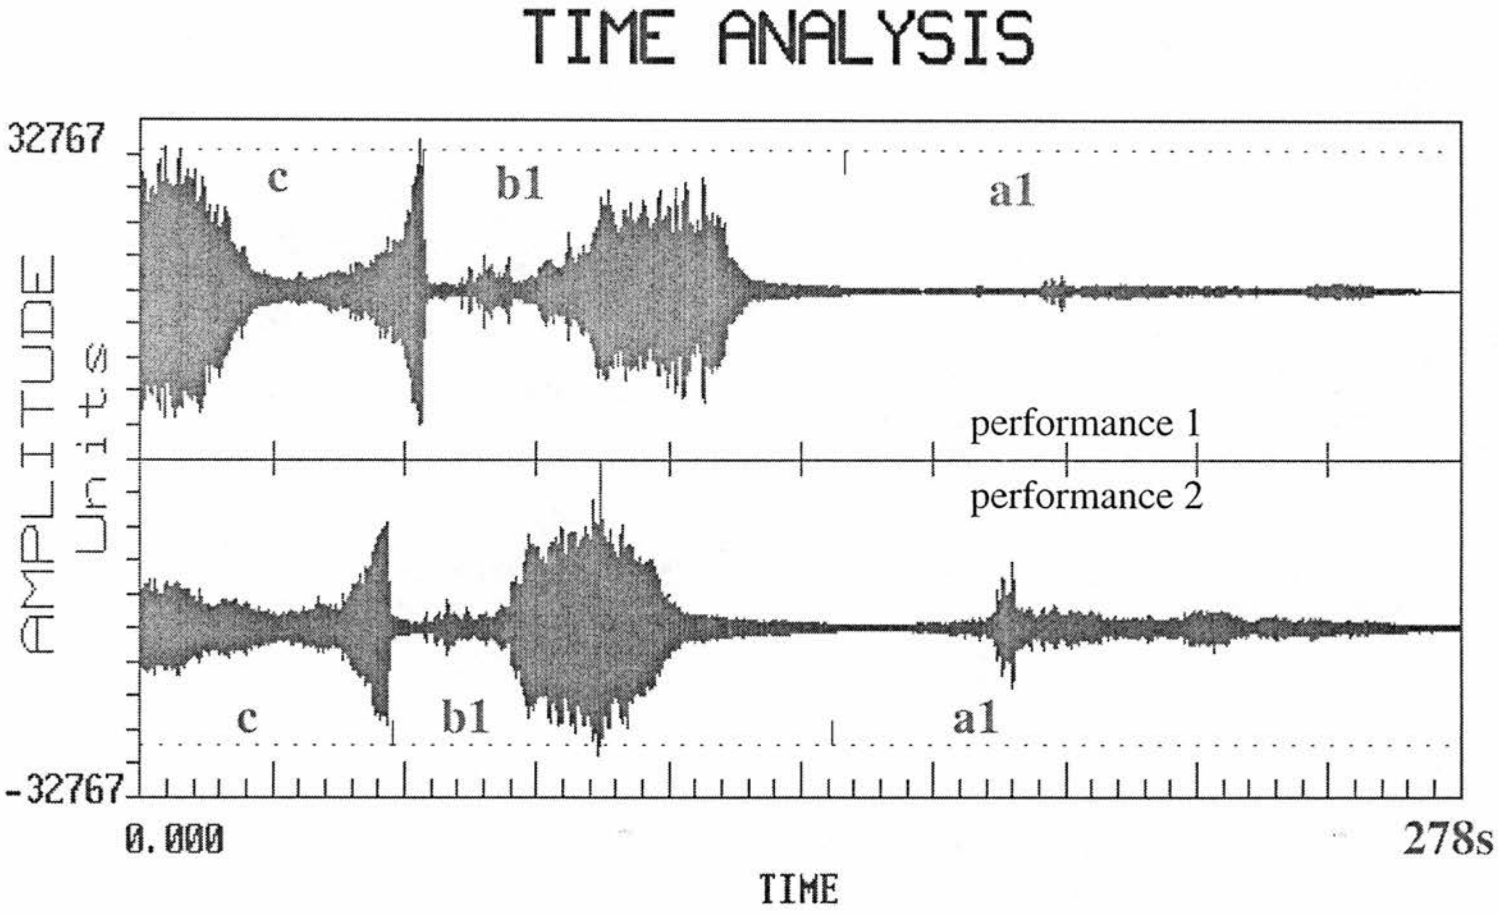

Figure 4.6b Time waveform of second half of performances 1 and 2 of Atmospheres

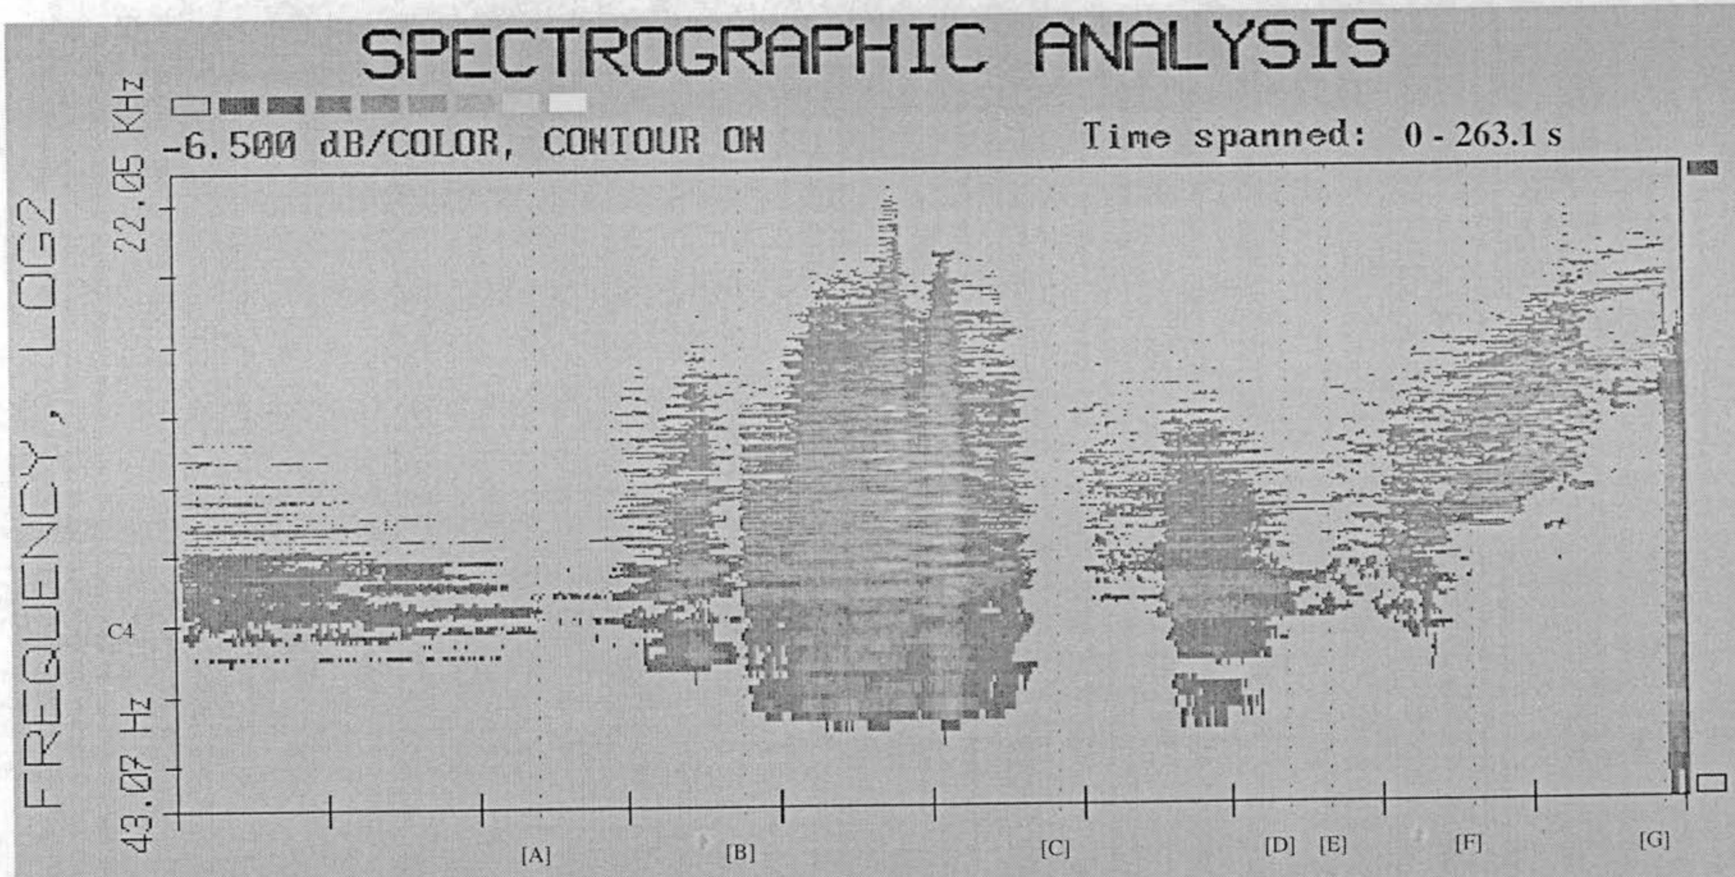

Figure 4.7a Atmospheres: performance 1

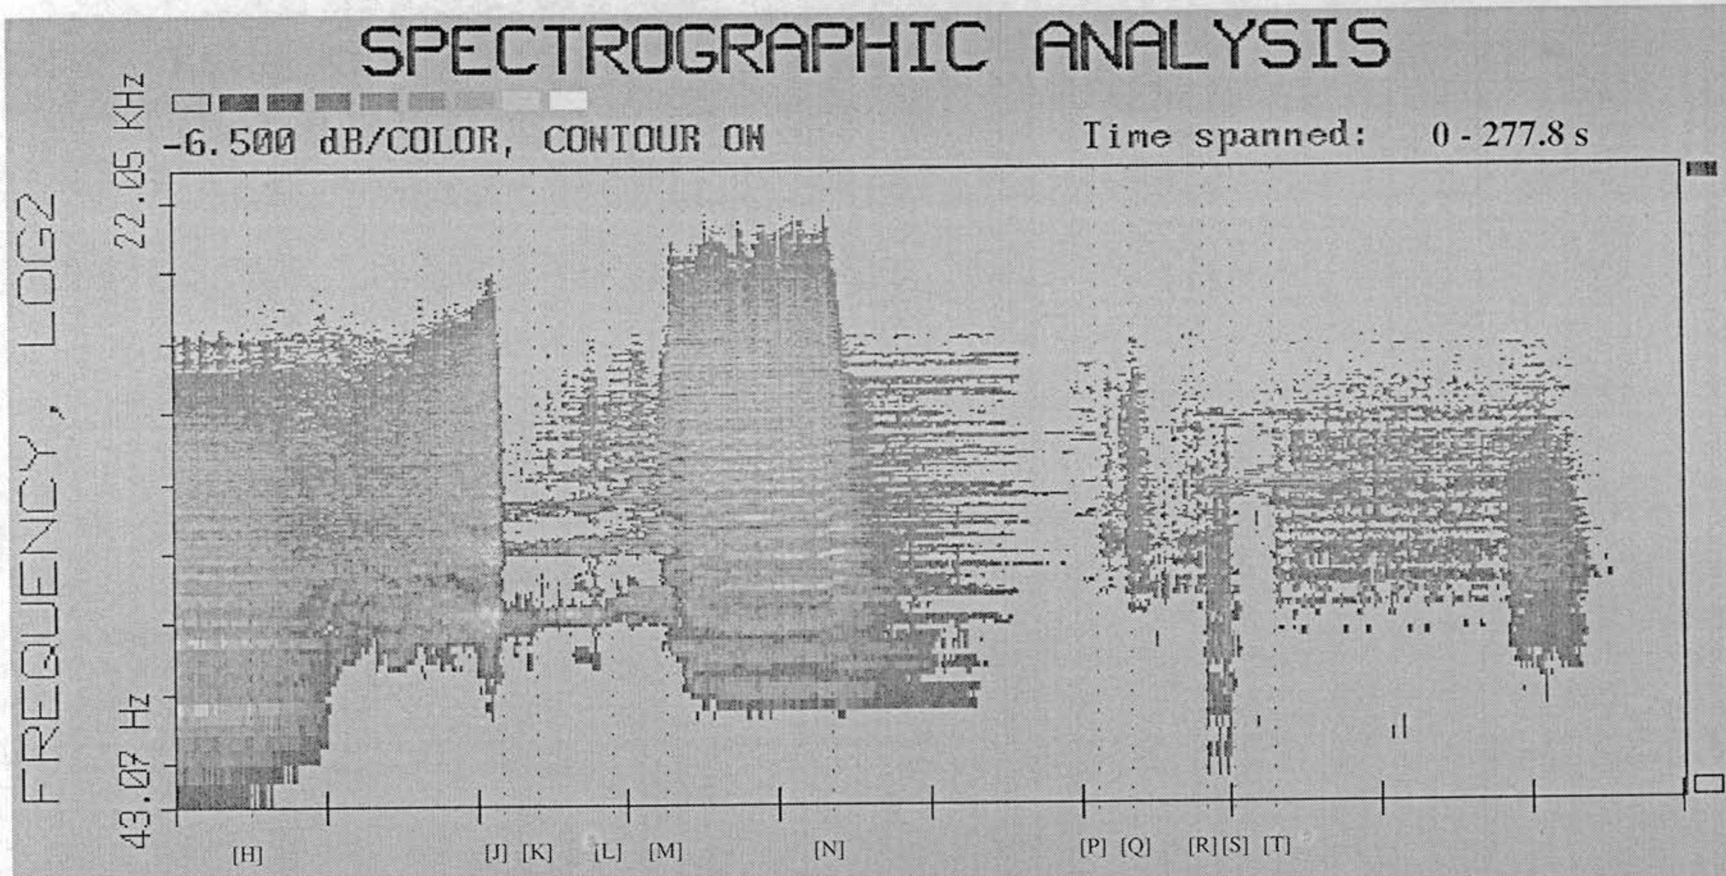

Figure 4.7b Atmospheres: performance 1

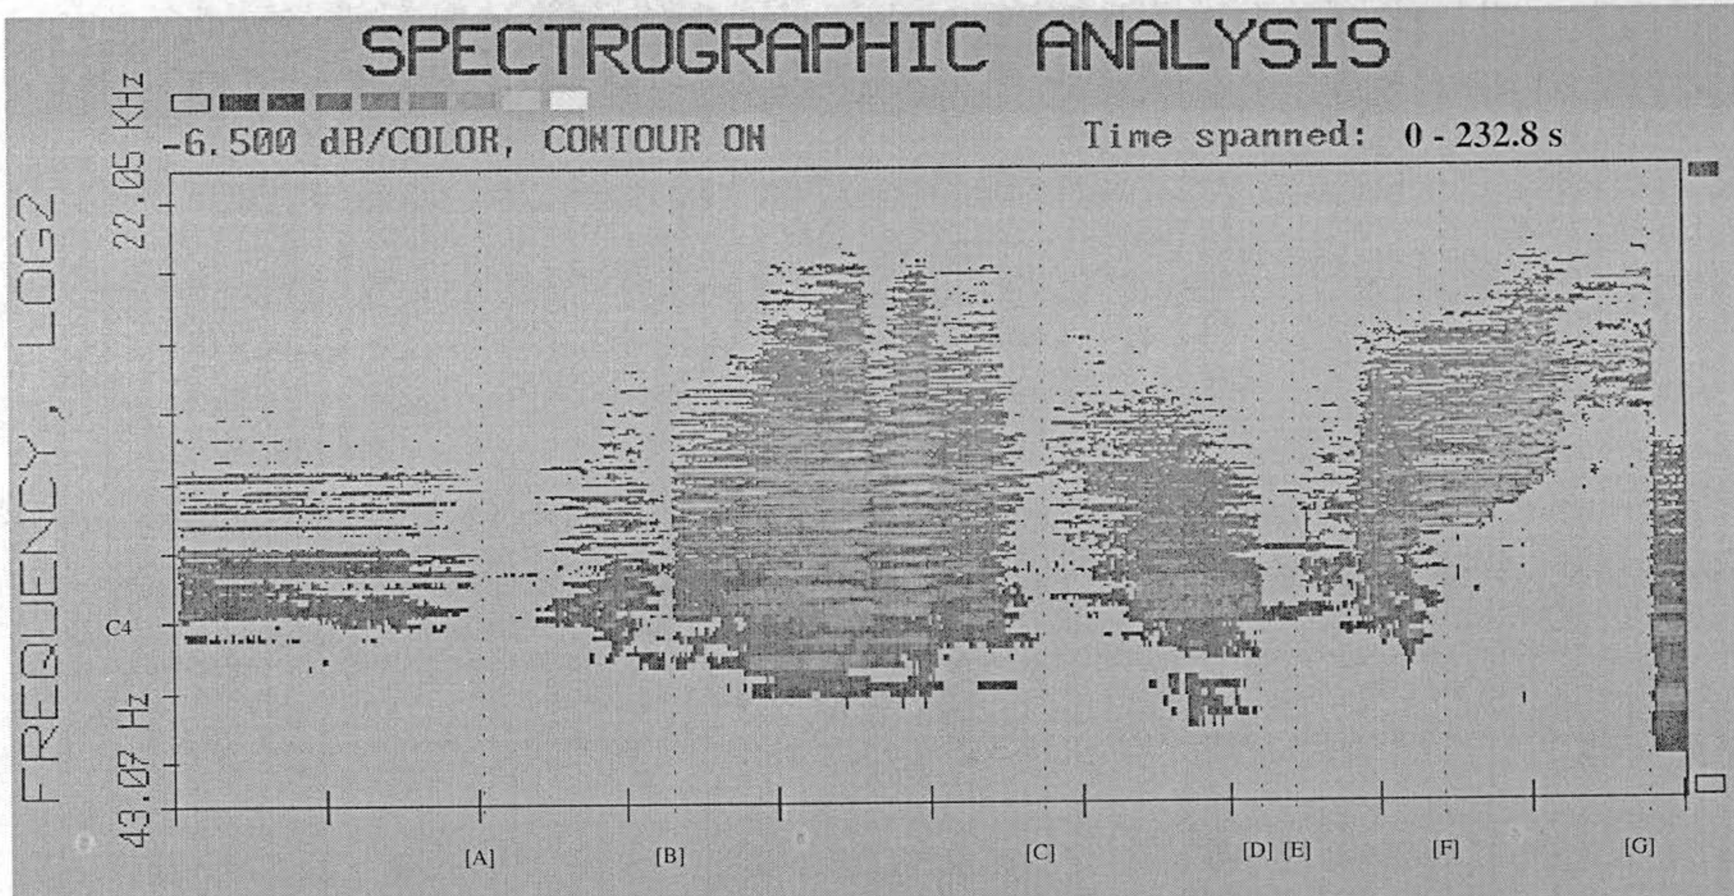

Figure 4.8a Atmospheres: performance 2

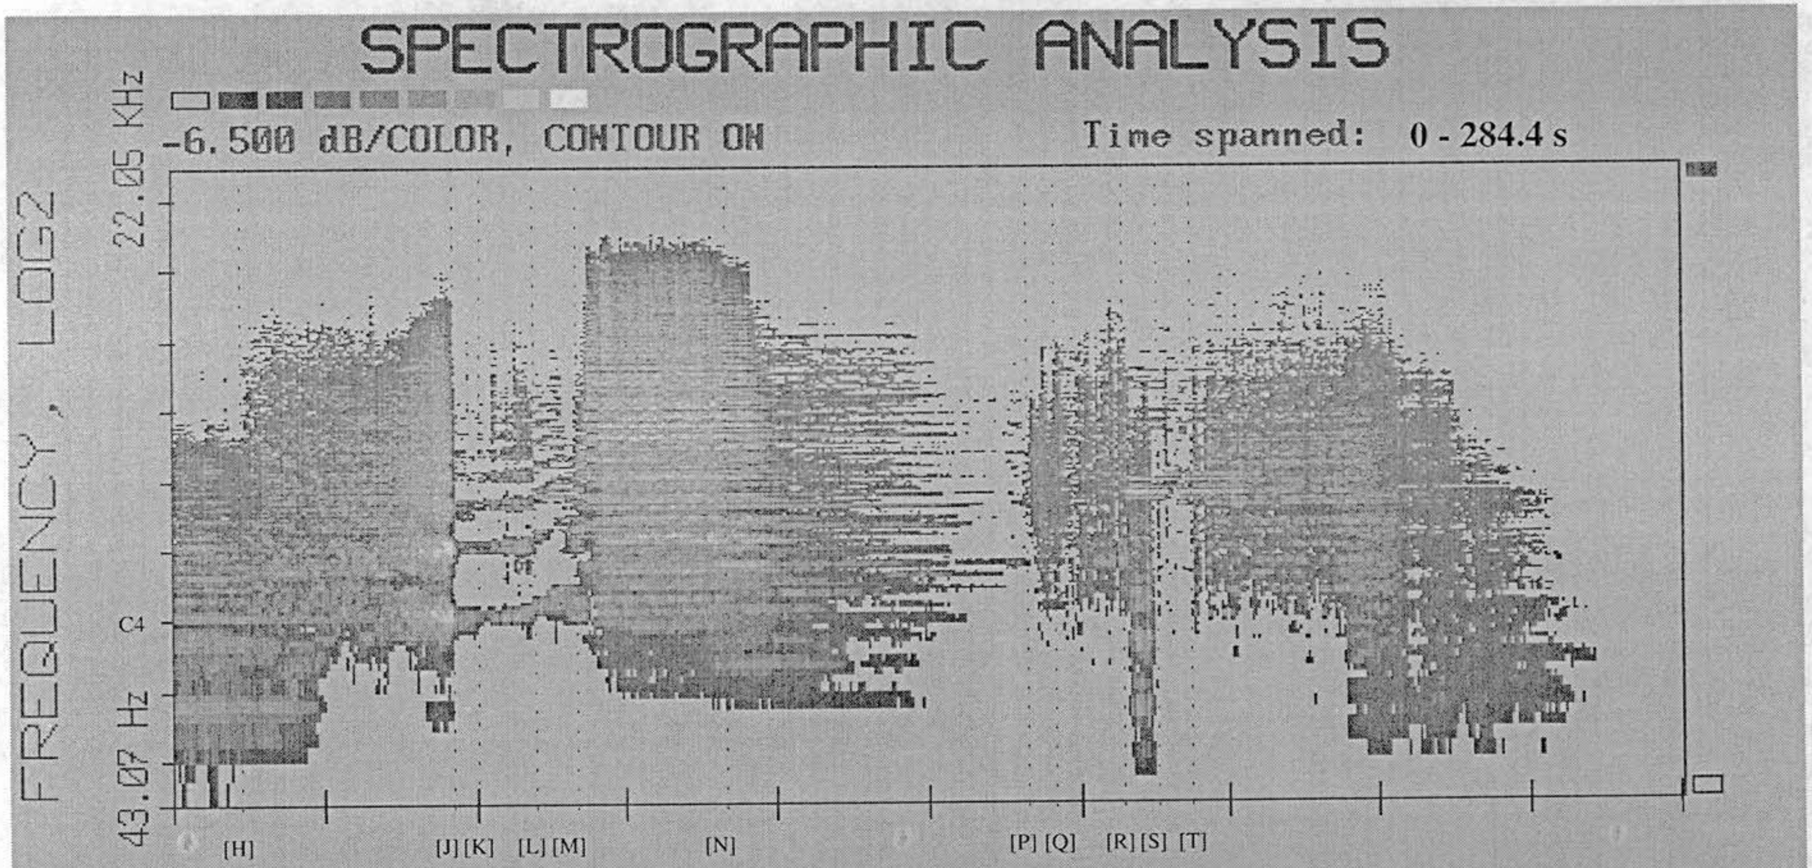

Figure 4.8b Atmospheres: performance 2

Figure 4.9a

Combined Timbral Graphs: performance 1

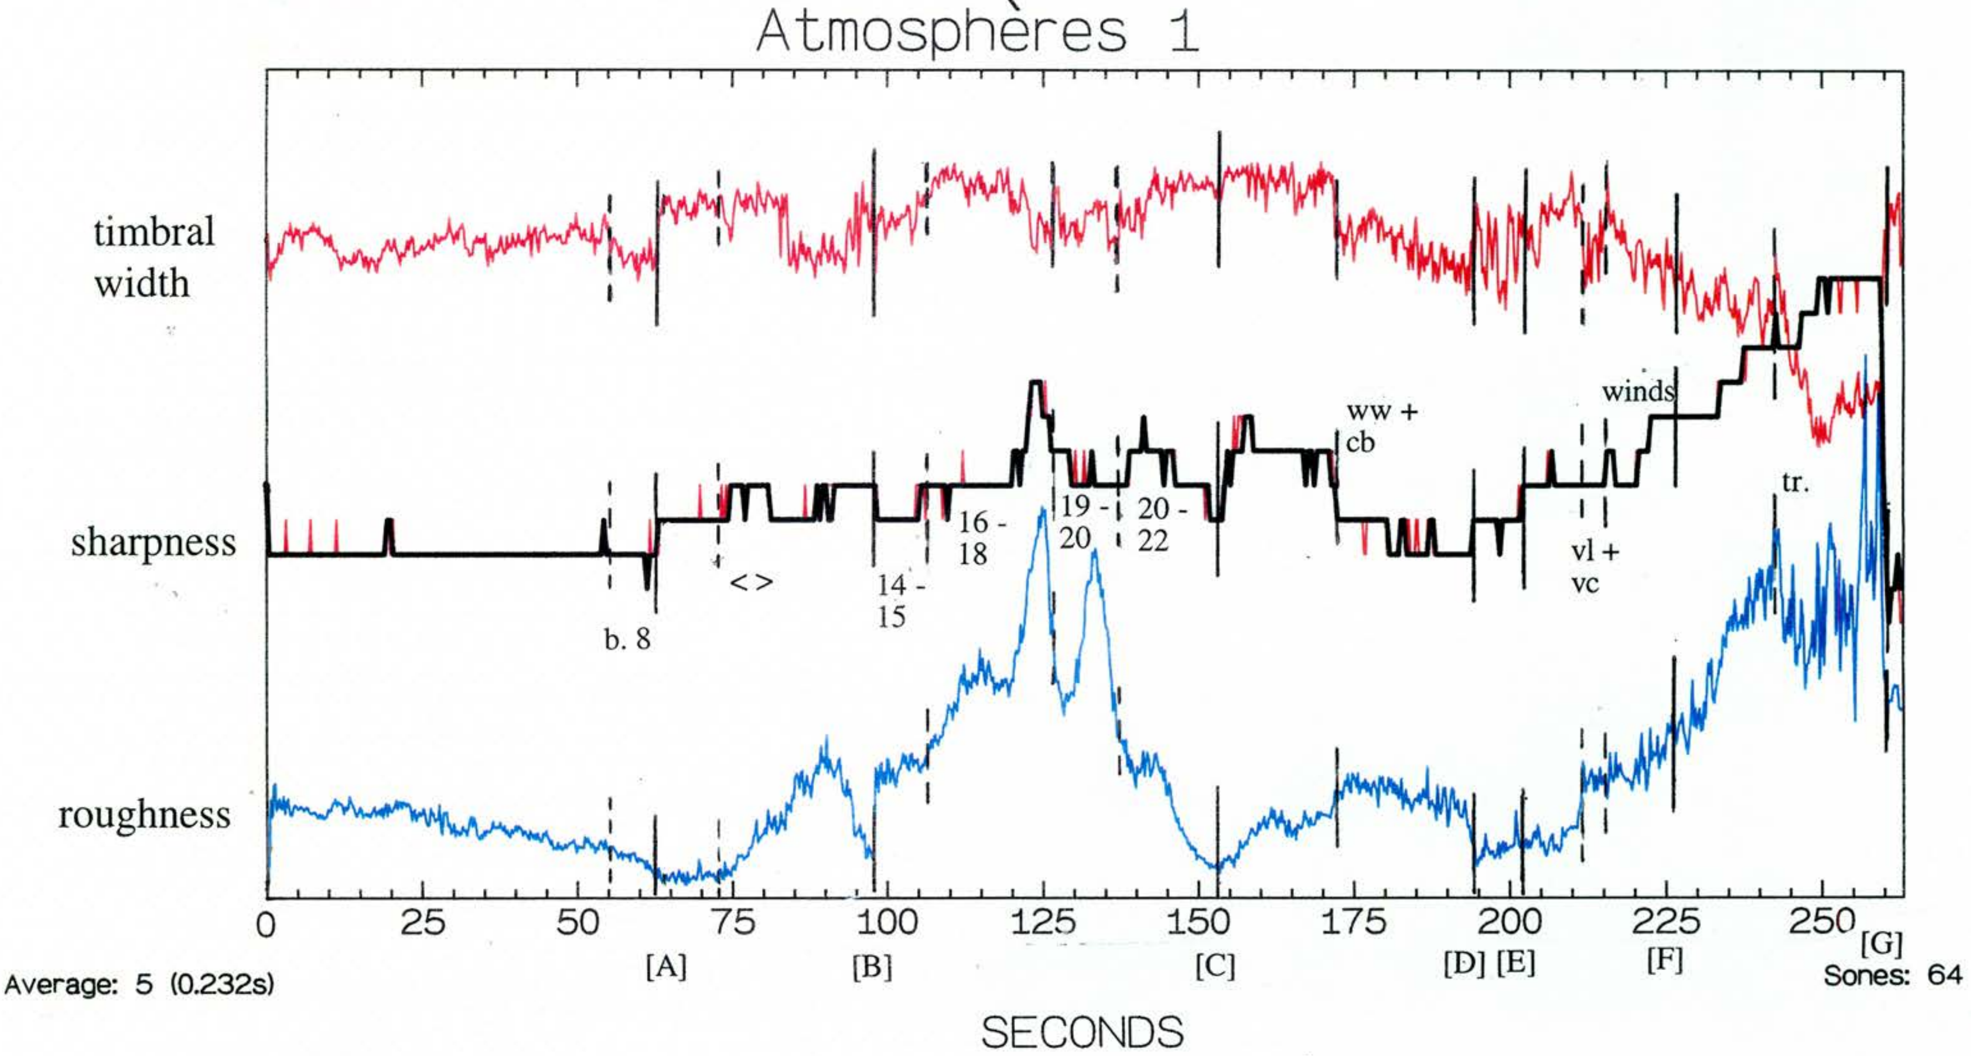

Figure 4.9b

Combined Timbral Graphs: performance 1

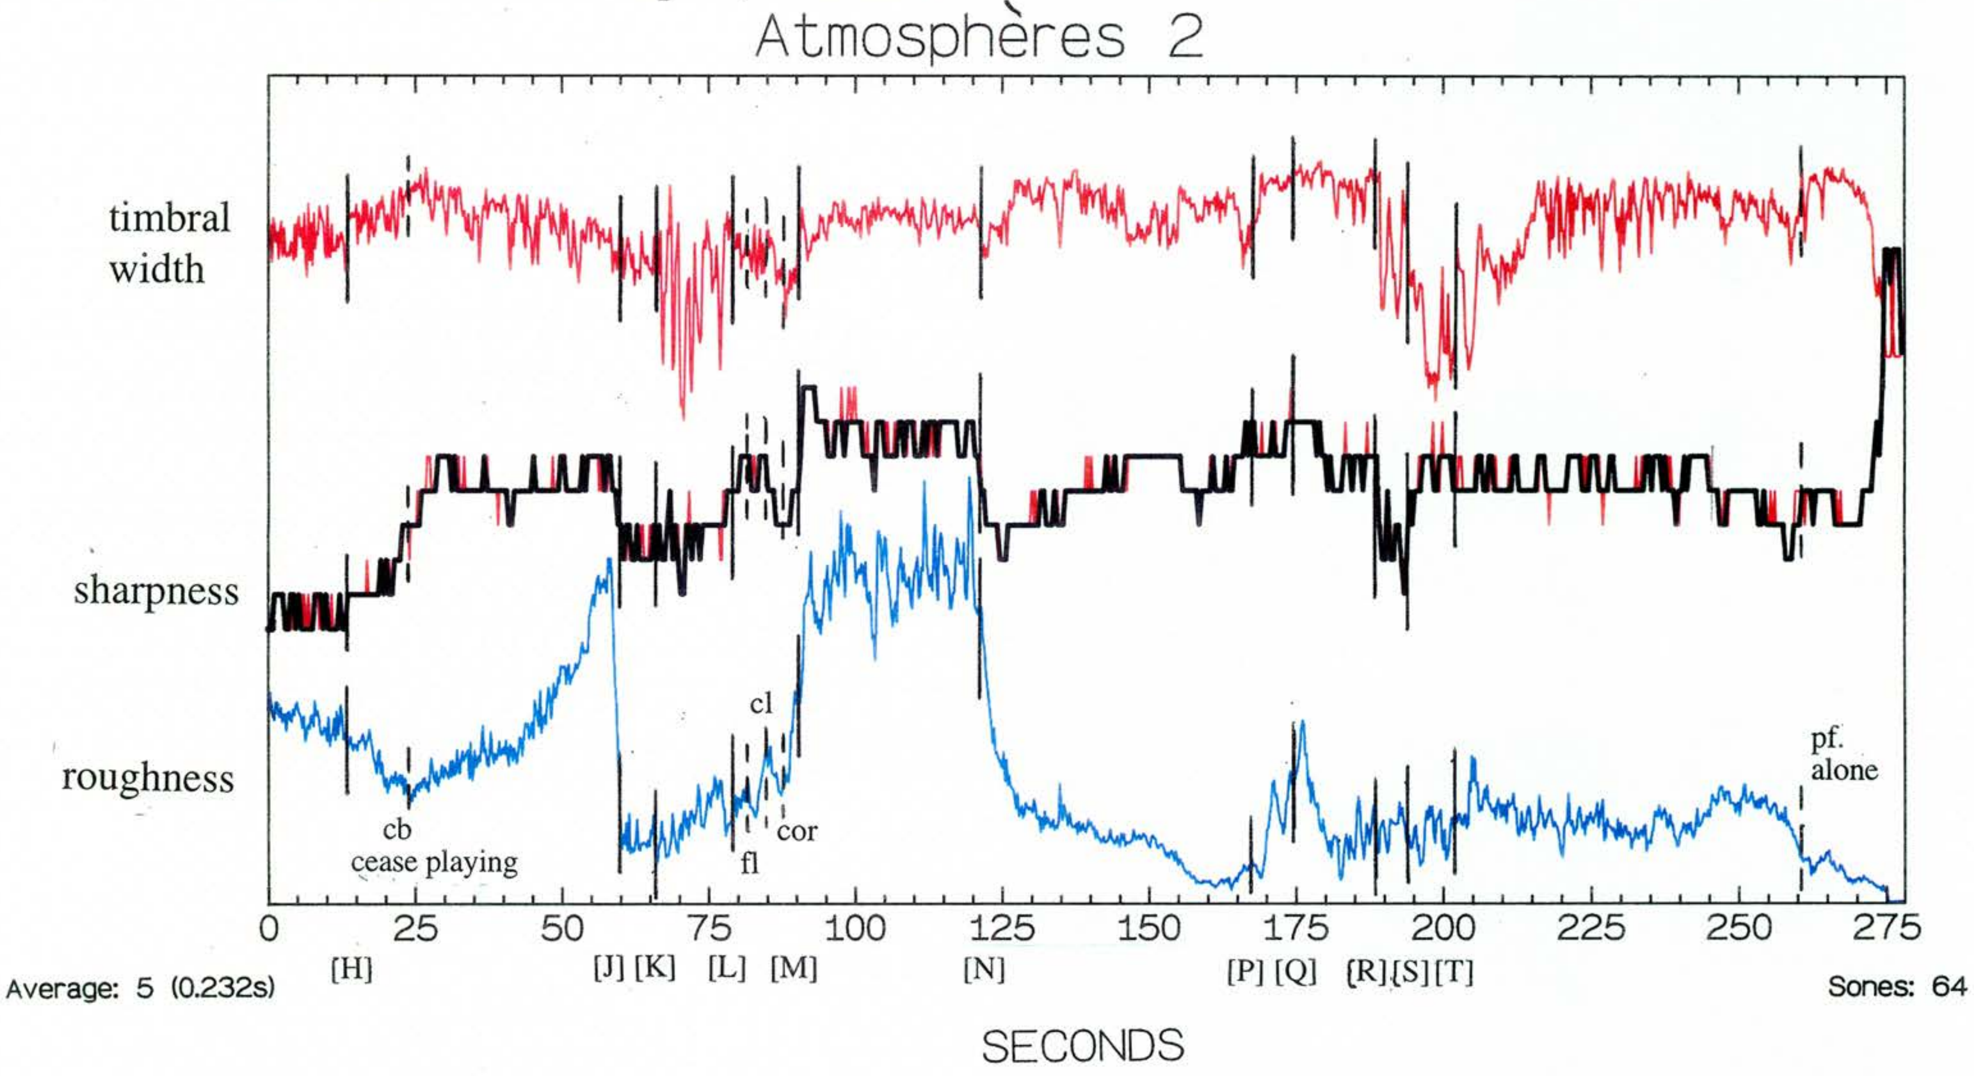

Figure 4.10a

Combined Timbral Graphs: performance 2

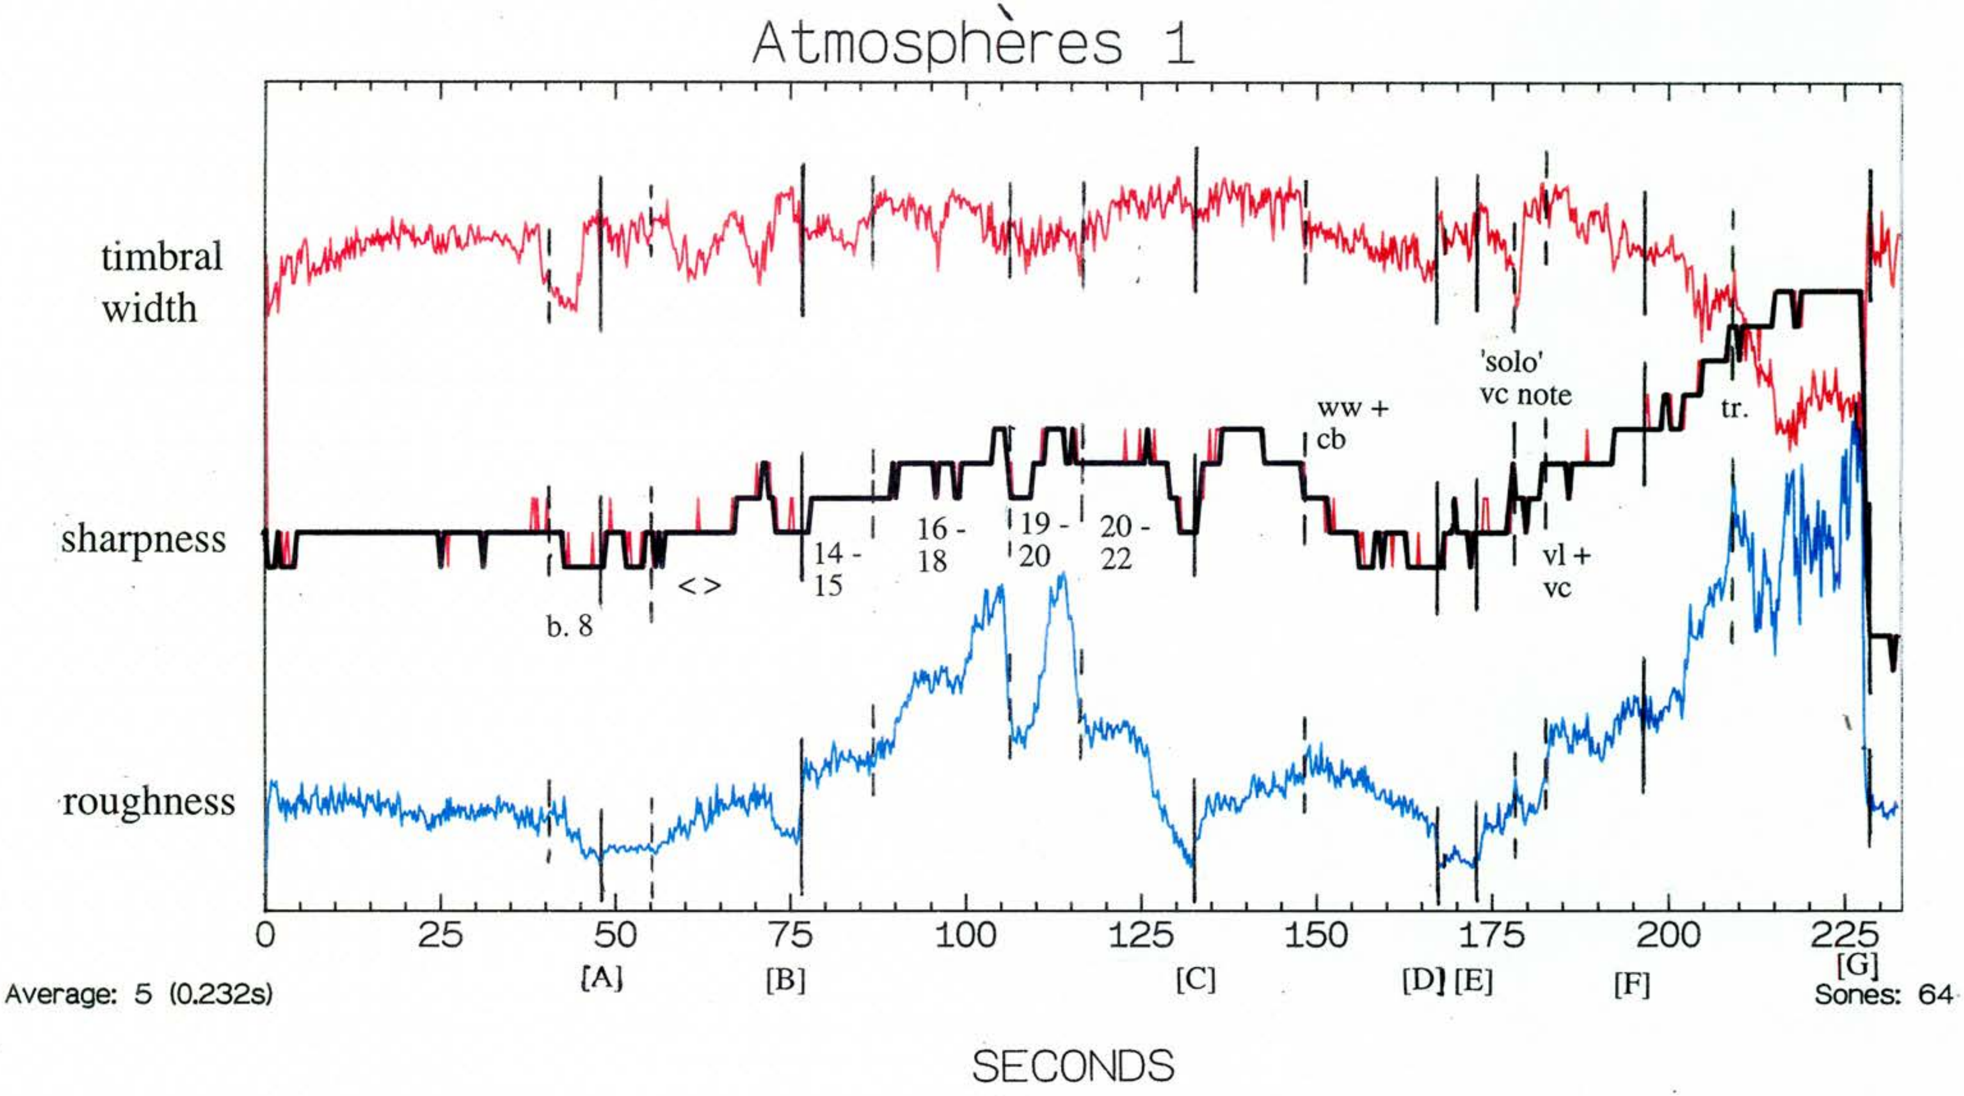

Figure 4.10b
Combined Timbral Graphs: performance 2

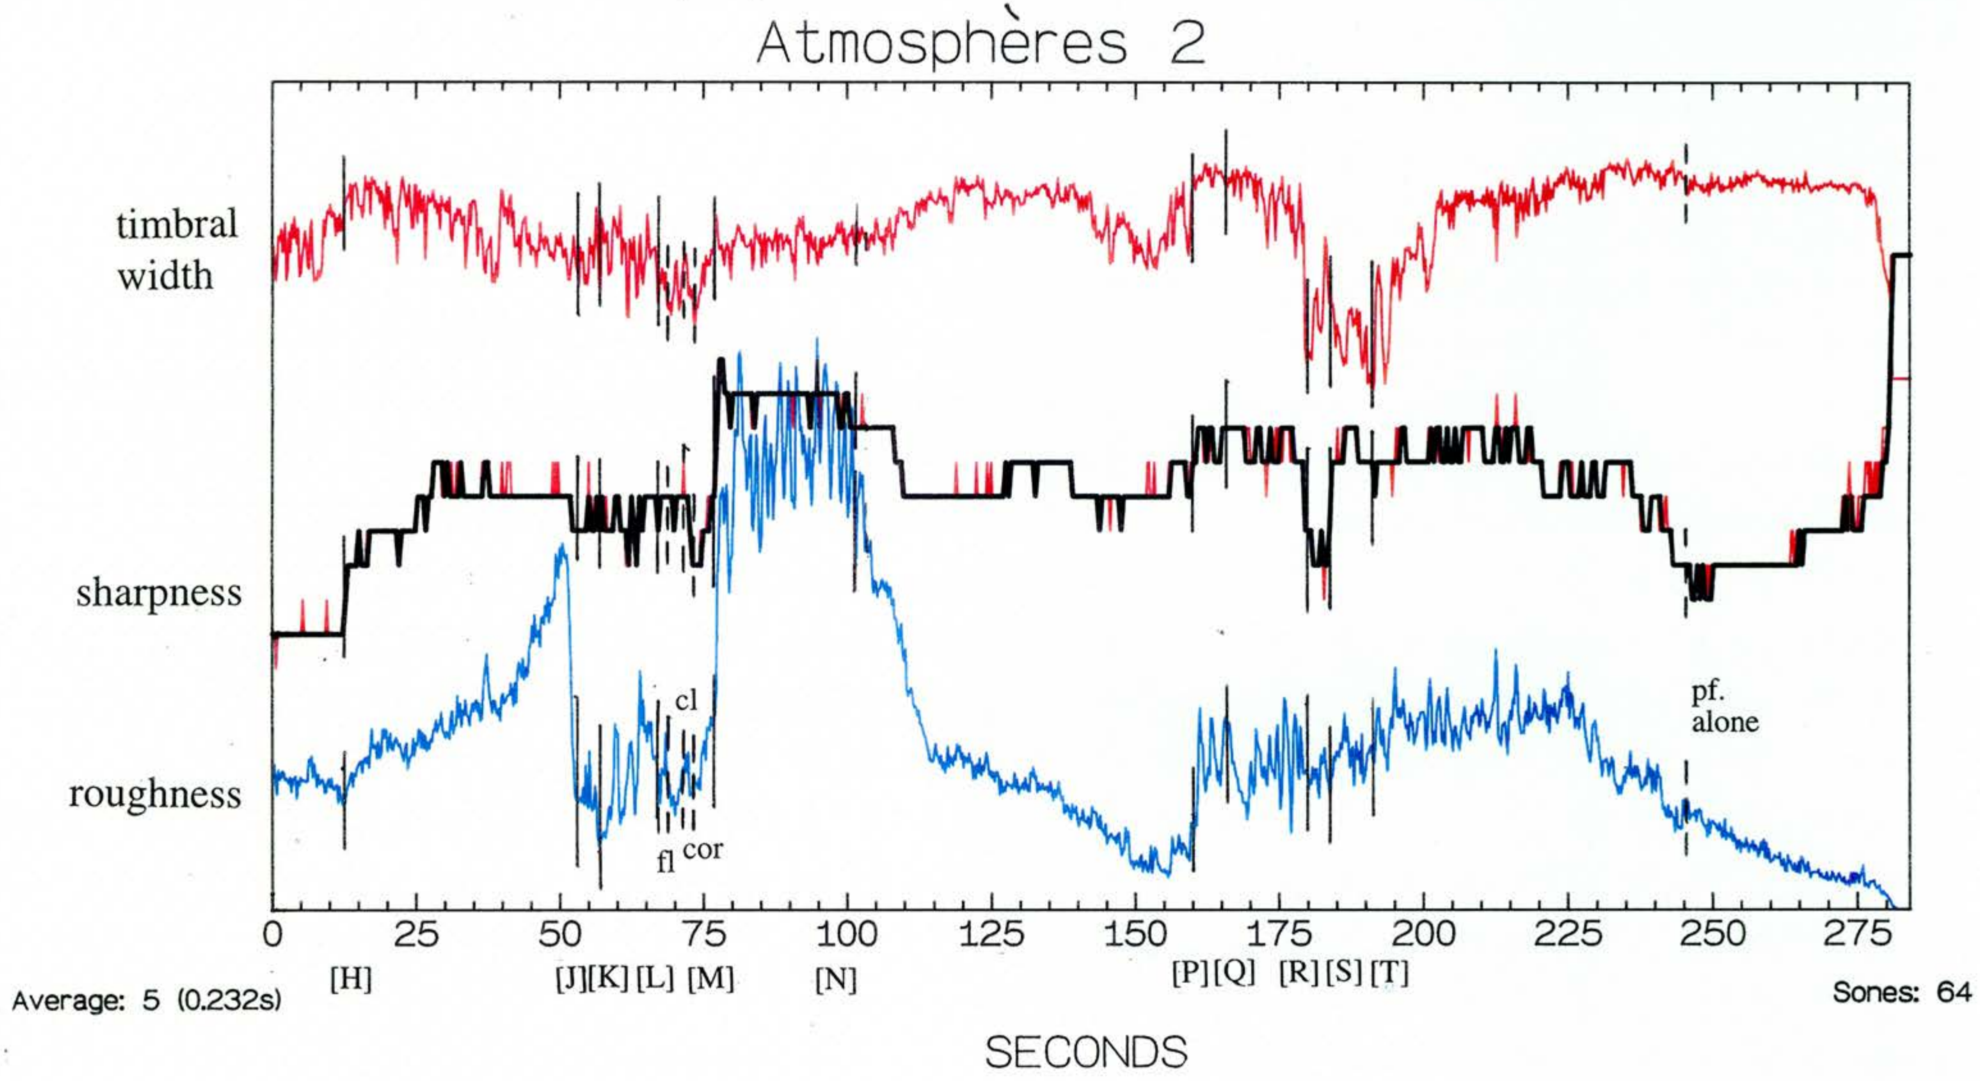

Figure 4.11a

Timbral weight: performance 1

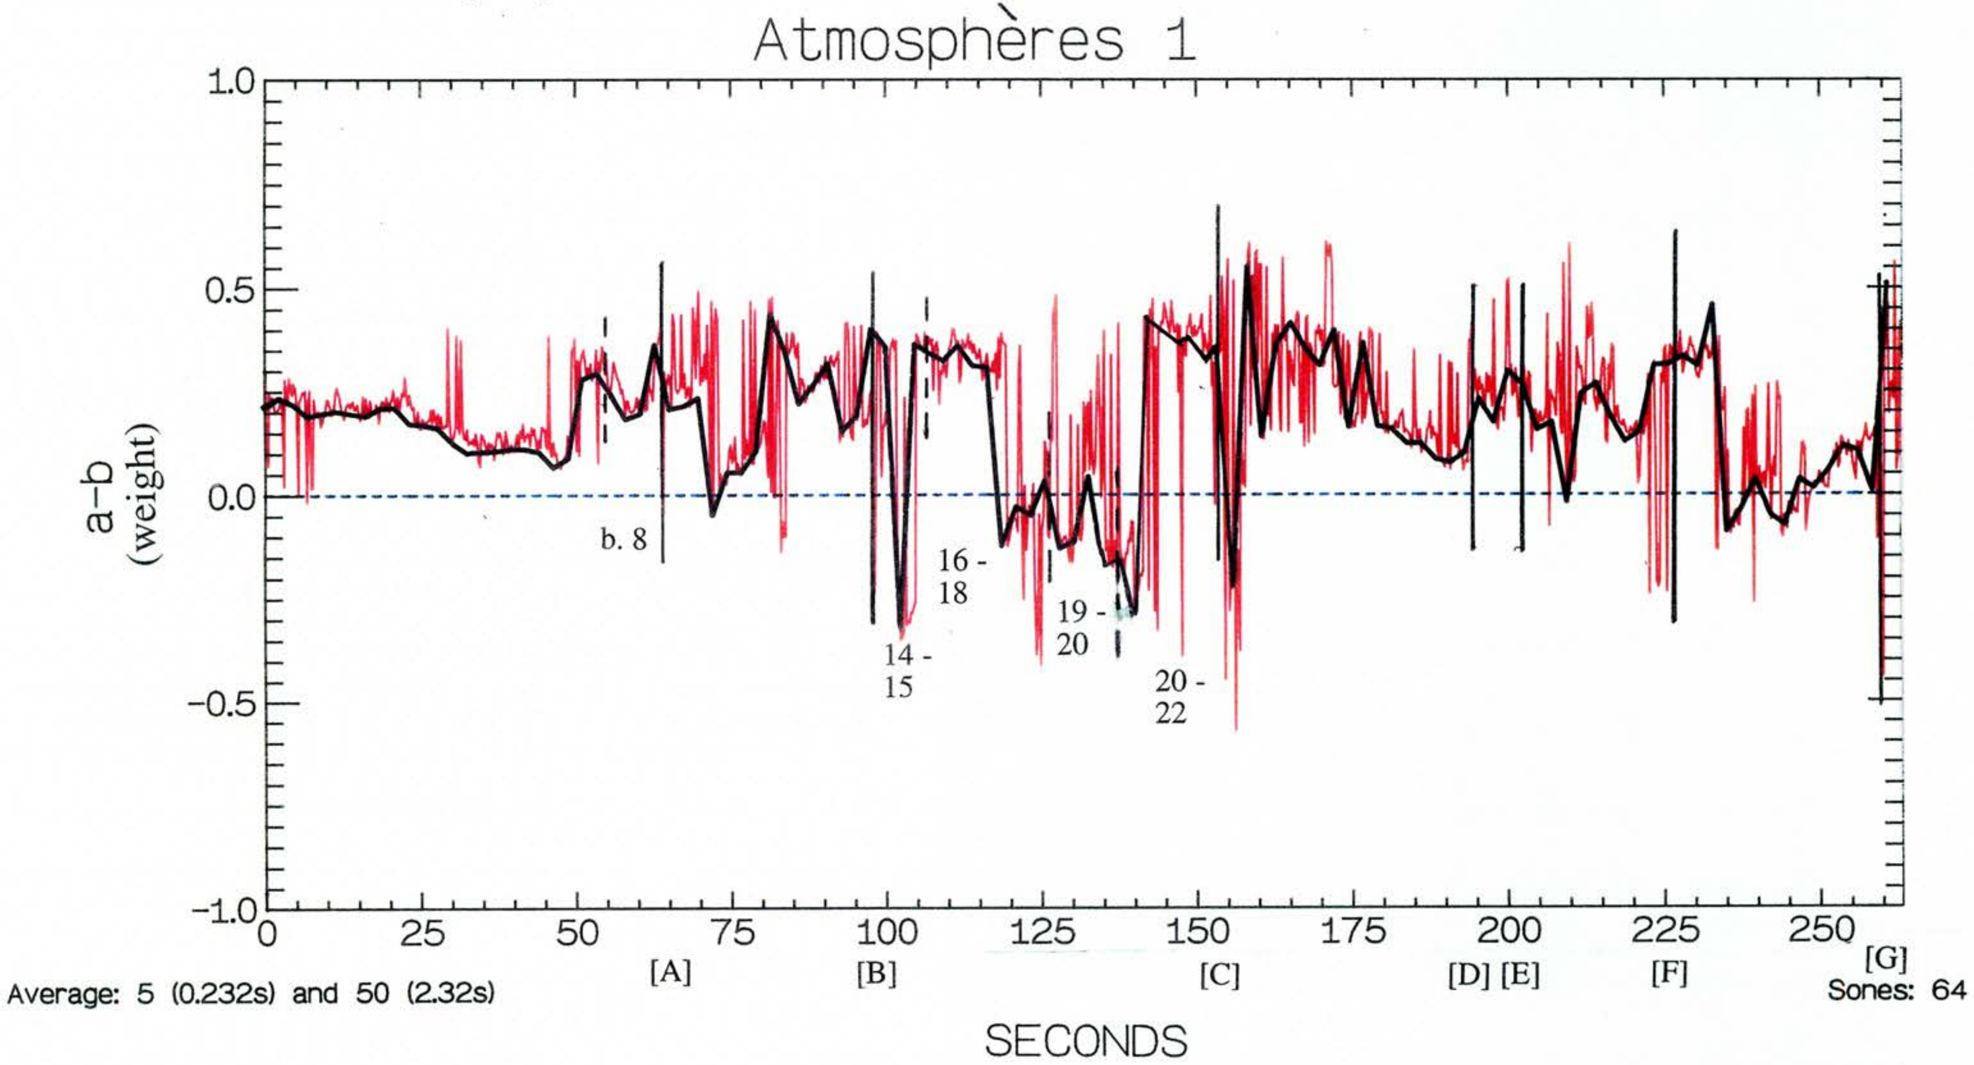

Figure 4.11b

Timbral weight: performance 1

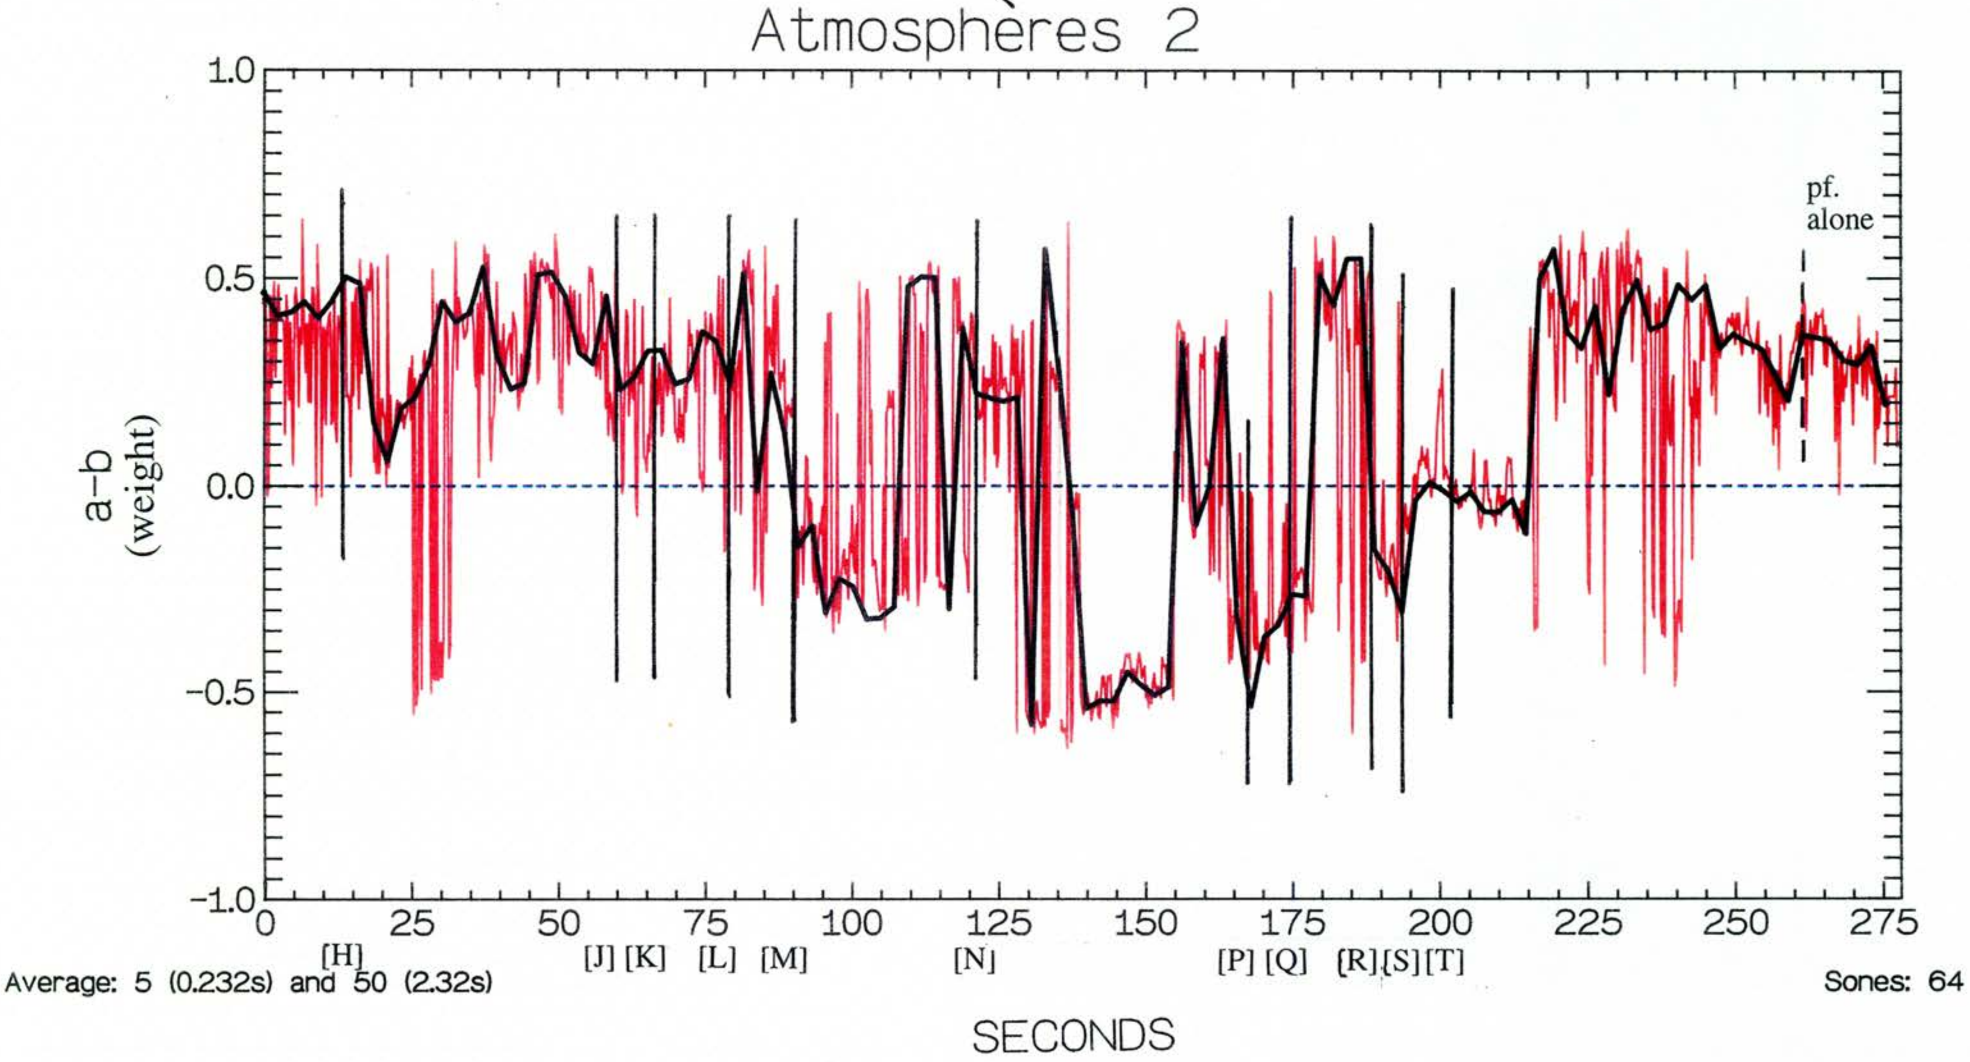

Figure 4.12a

Timbral weight: performance 2

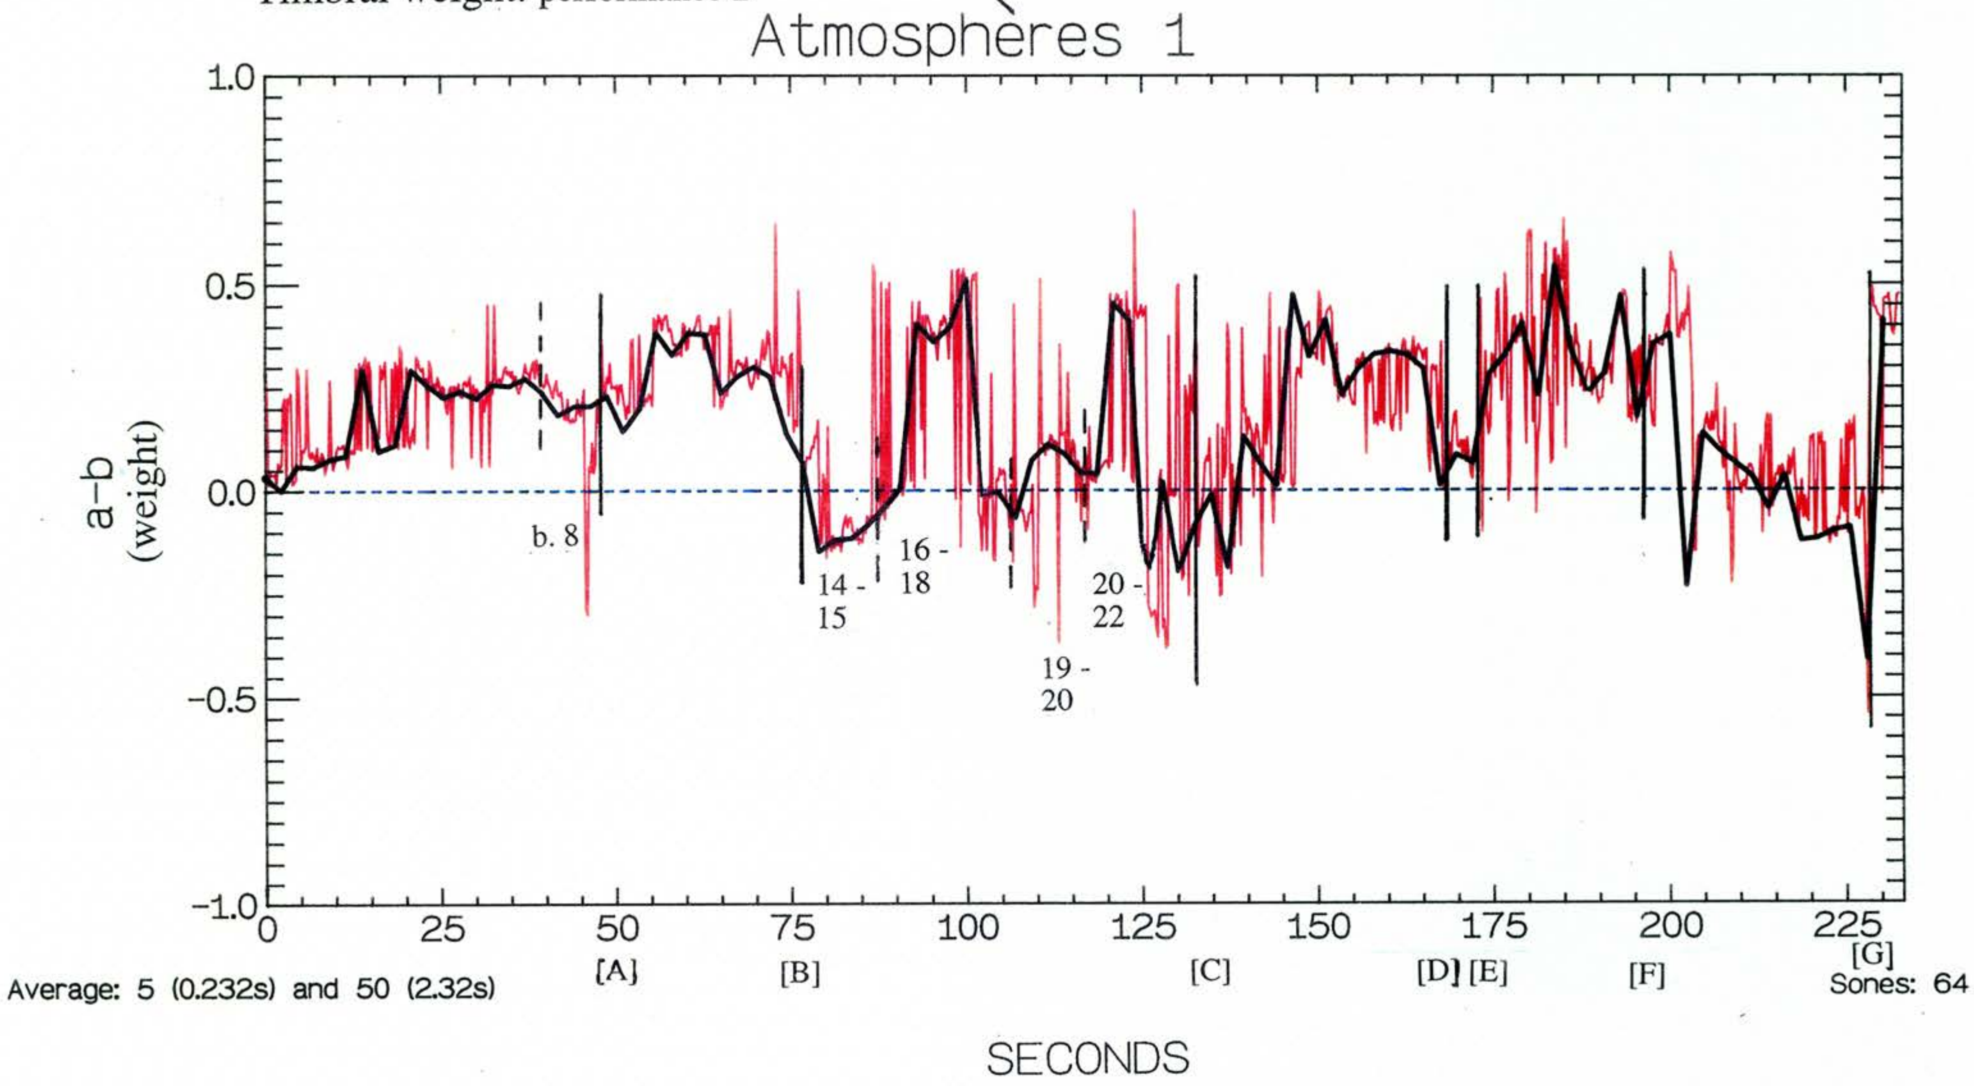

Figure 4.12b
Timbral weight: performance 2

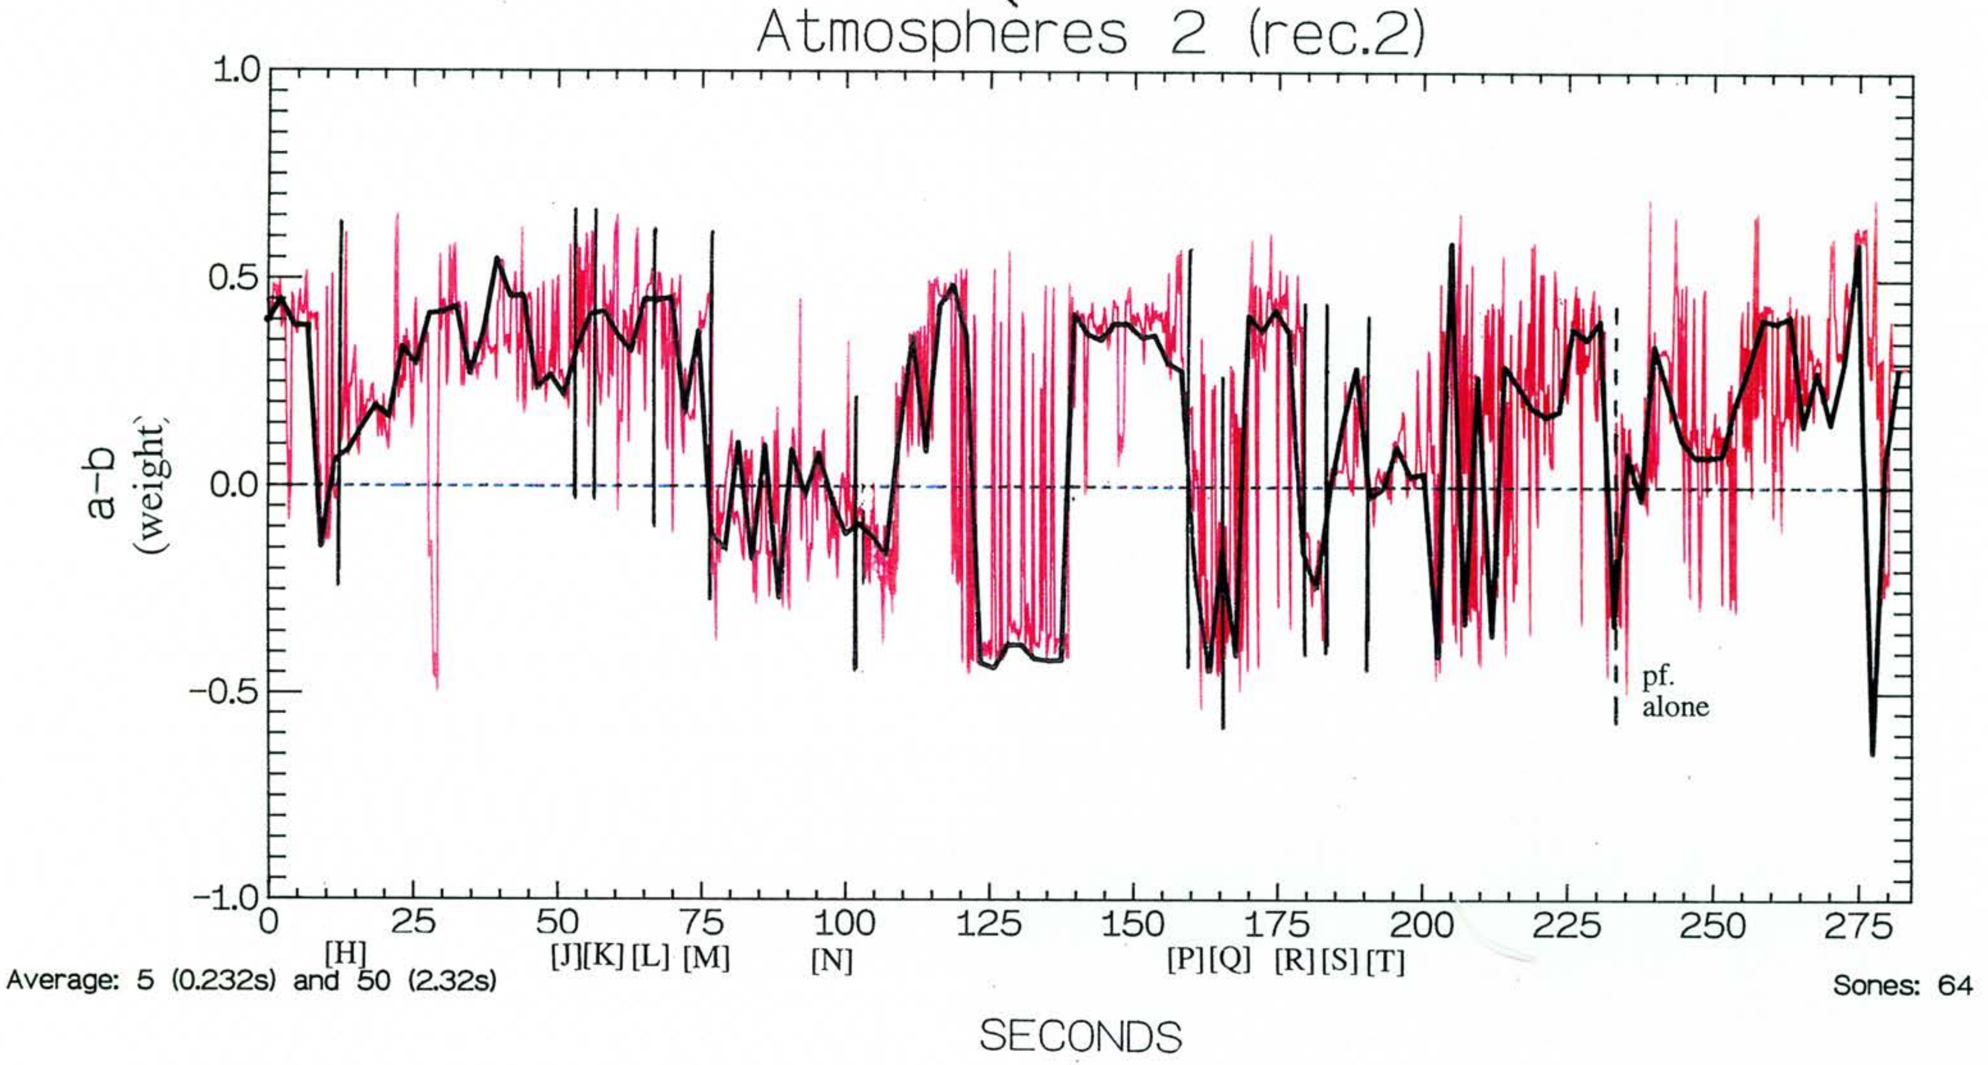

Figure 4.13a
Timbral pitch: performance 1

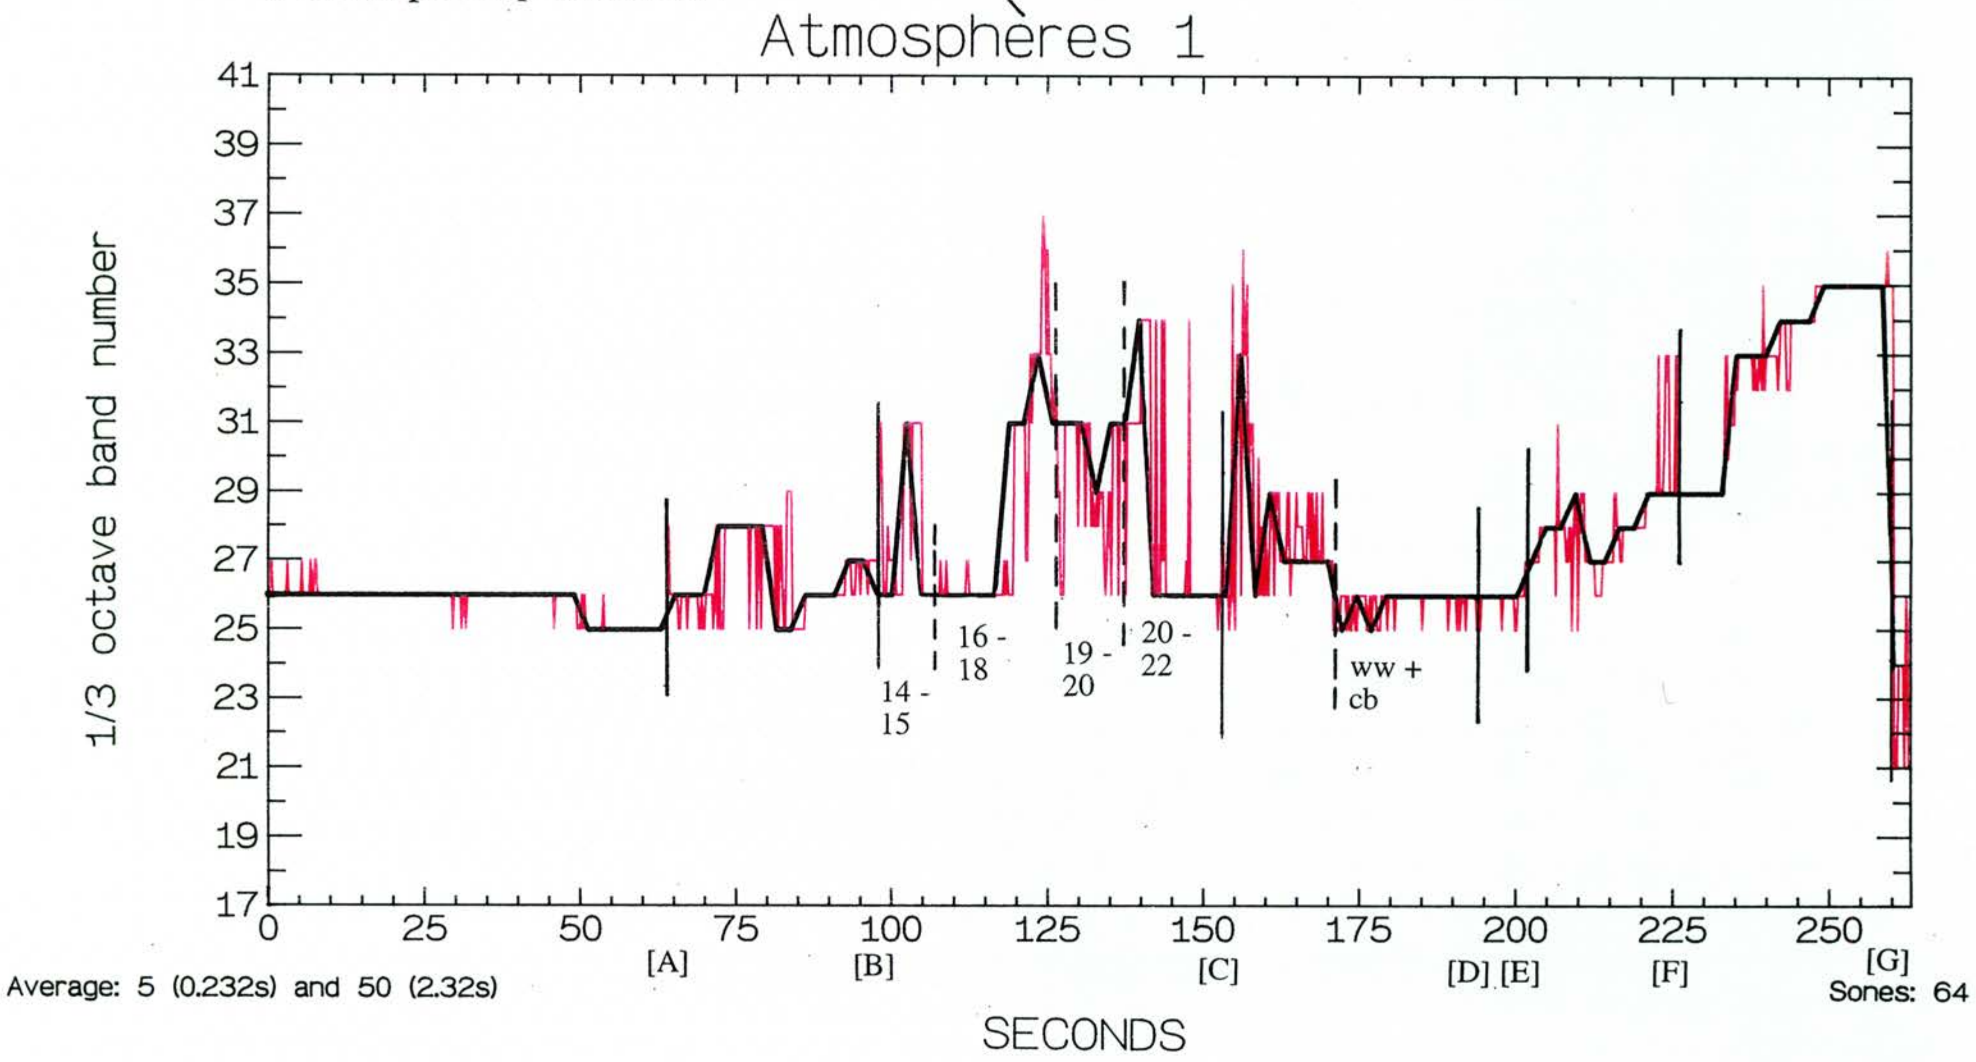

Figure 4.13b
Timbral pitch: performance 1

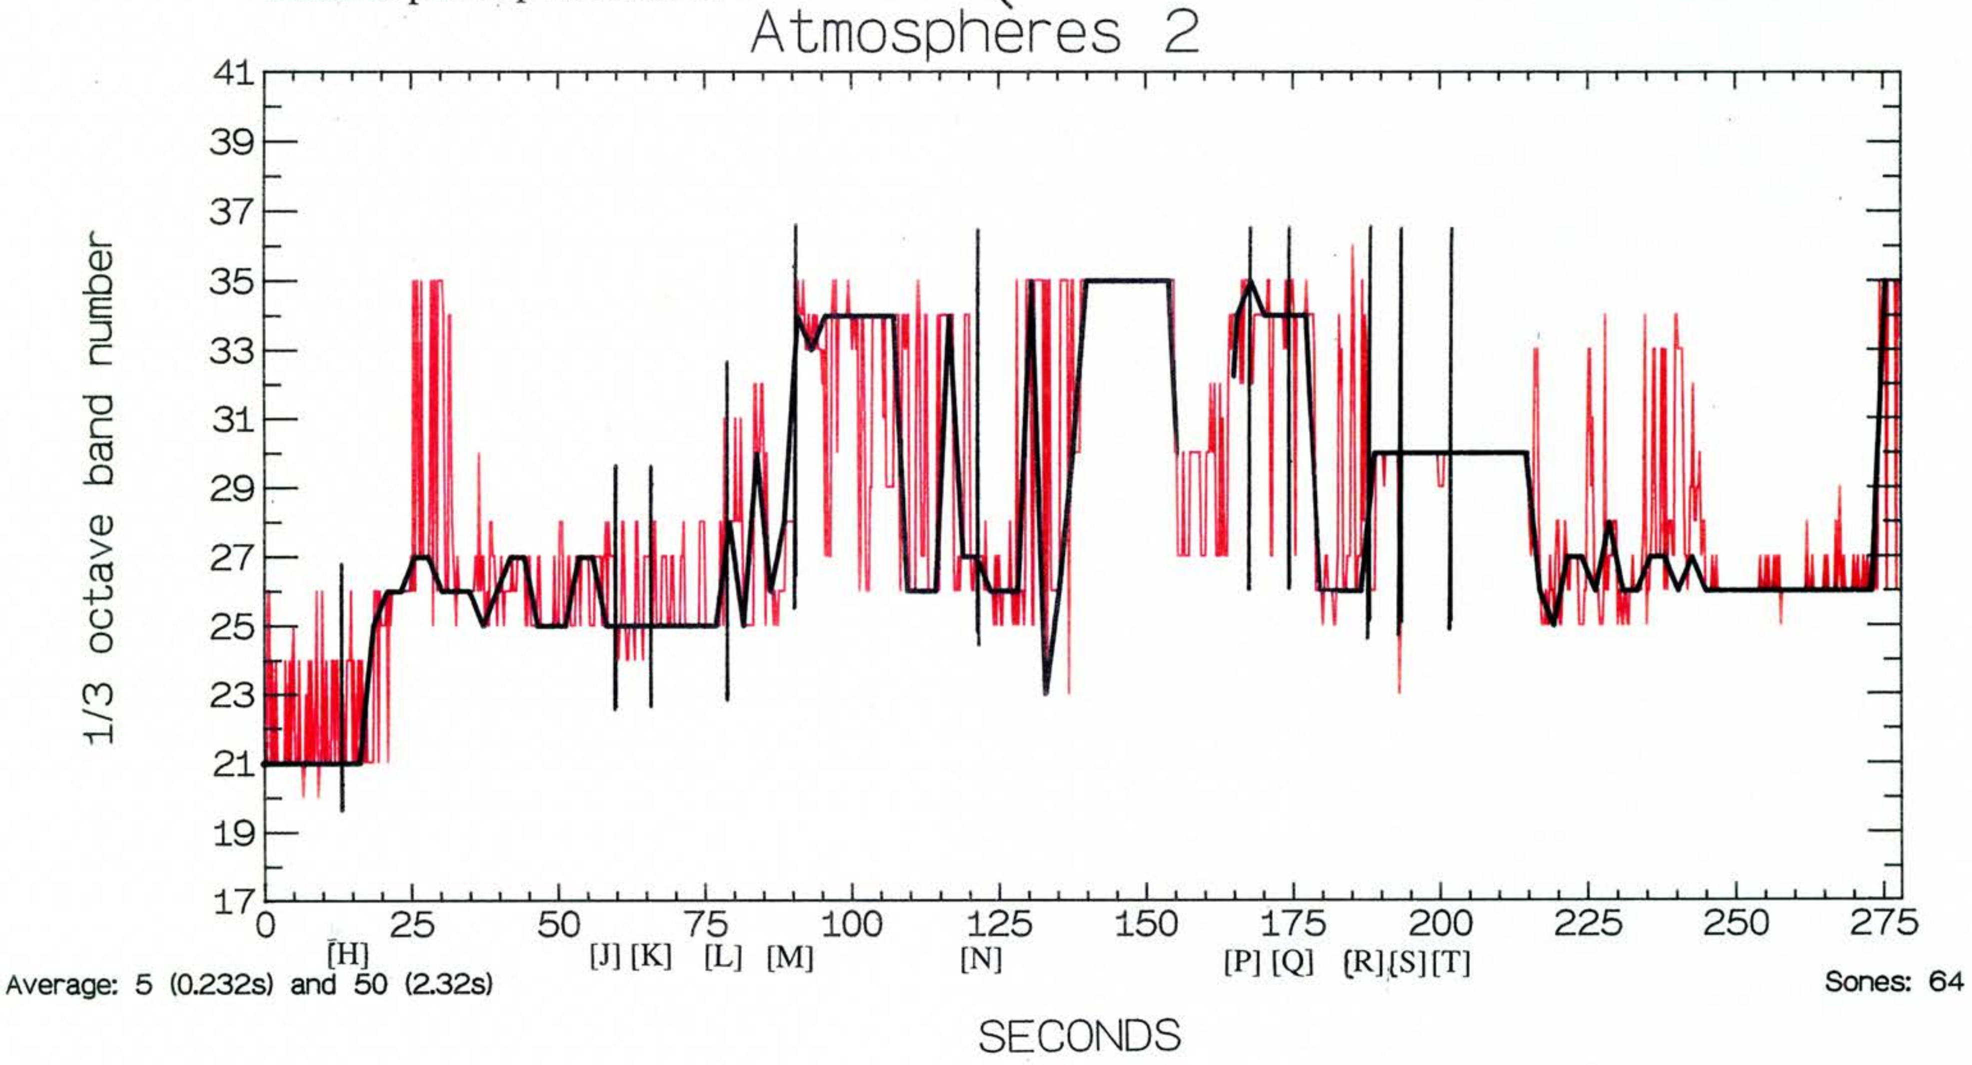

Figure 4.14a

Timbral pitch: performance 2

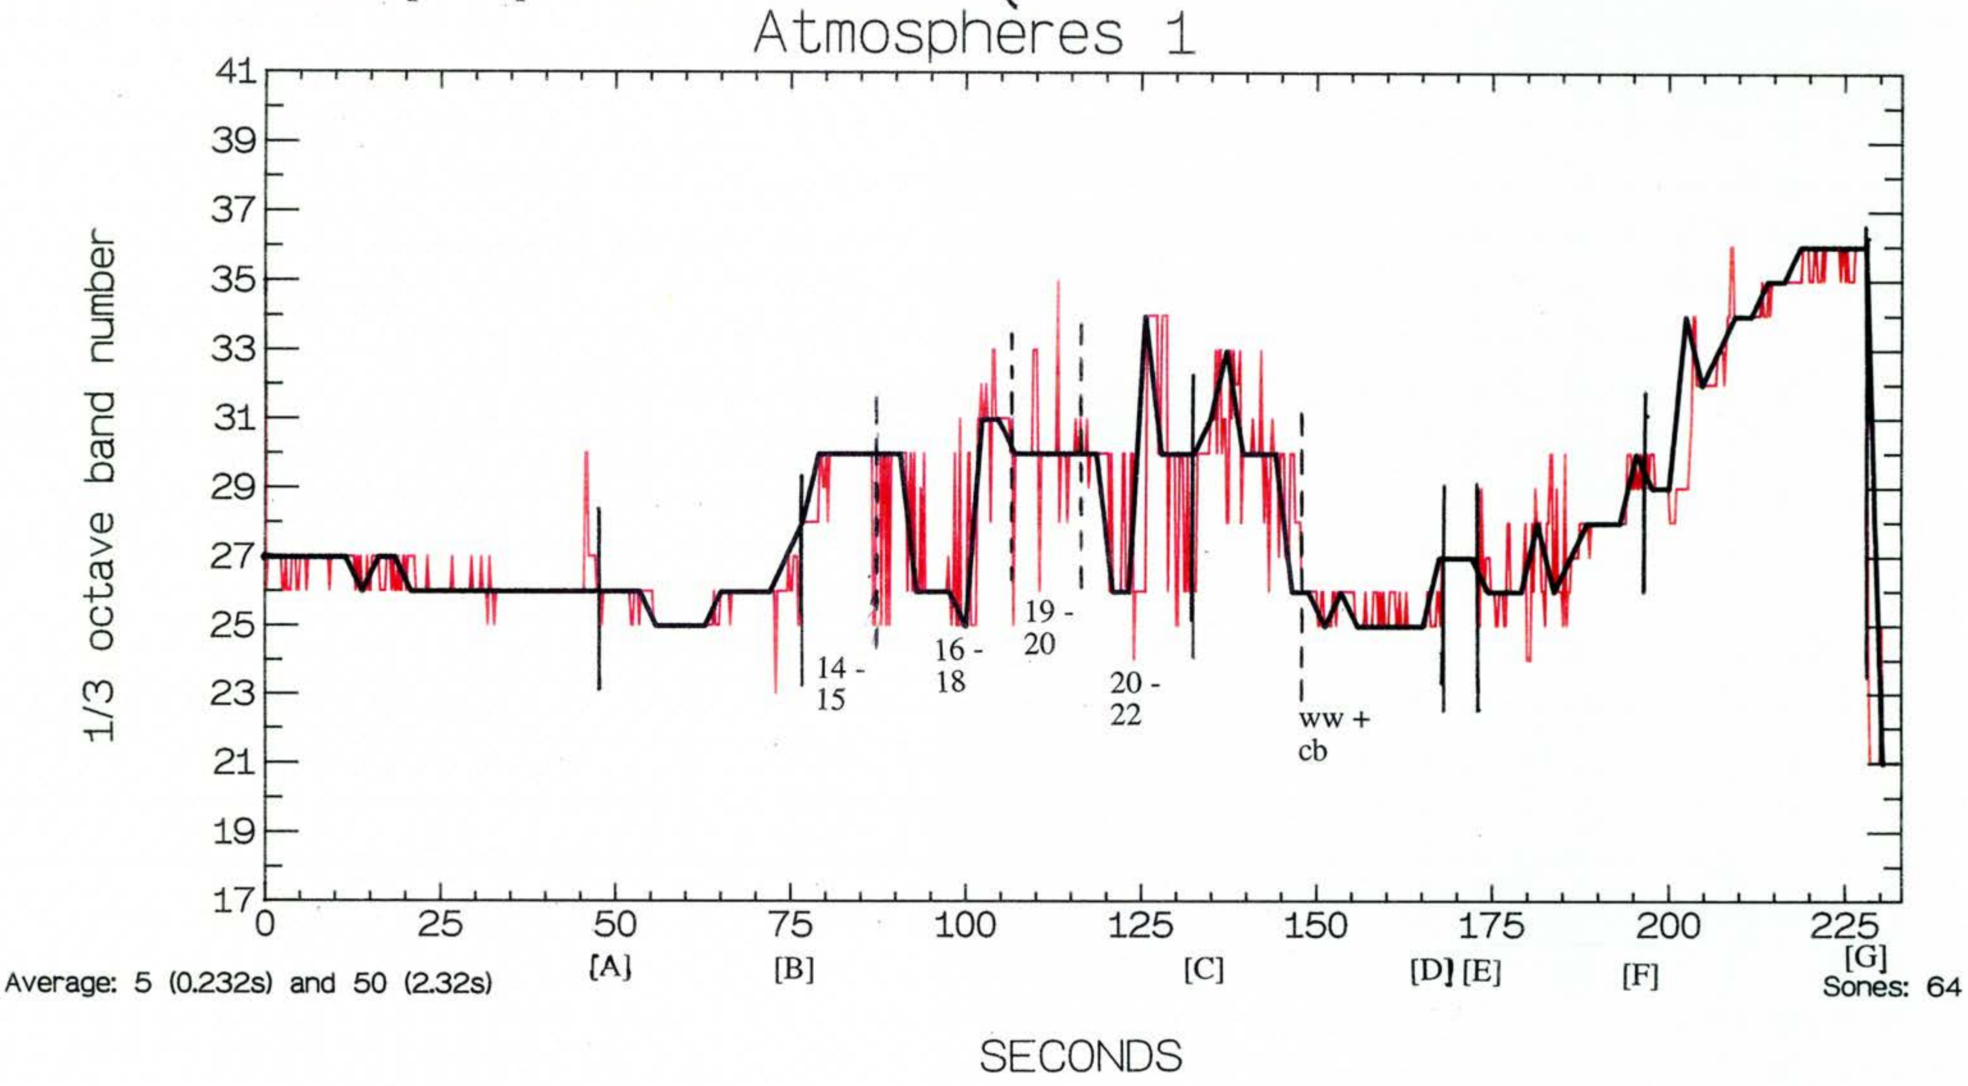

Figure 4.14b
Timbral pitch: performance 2

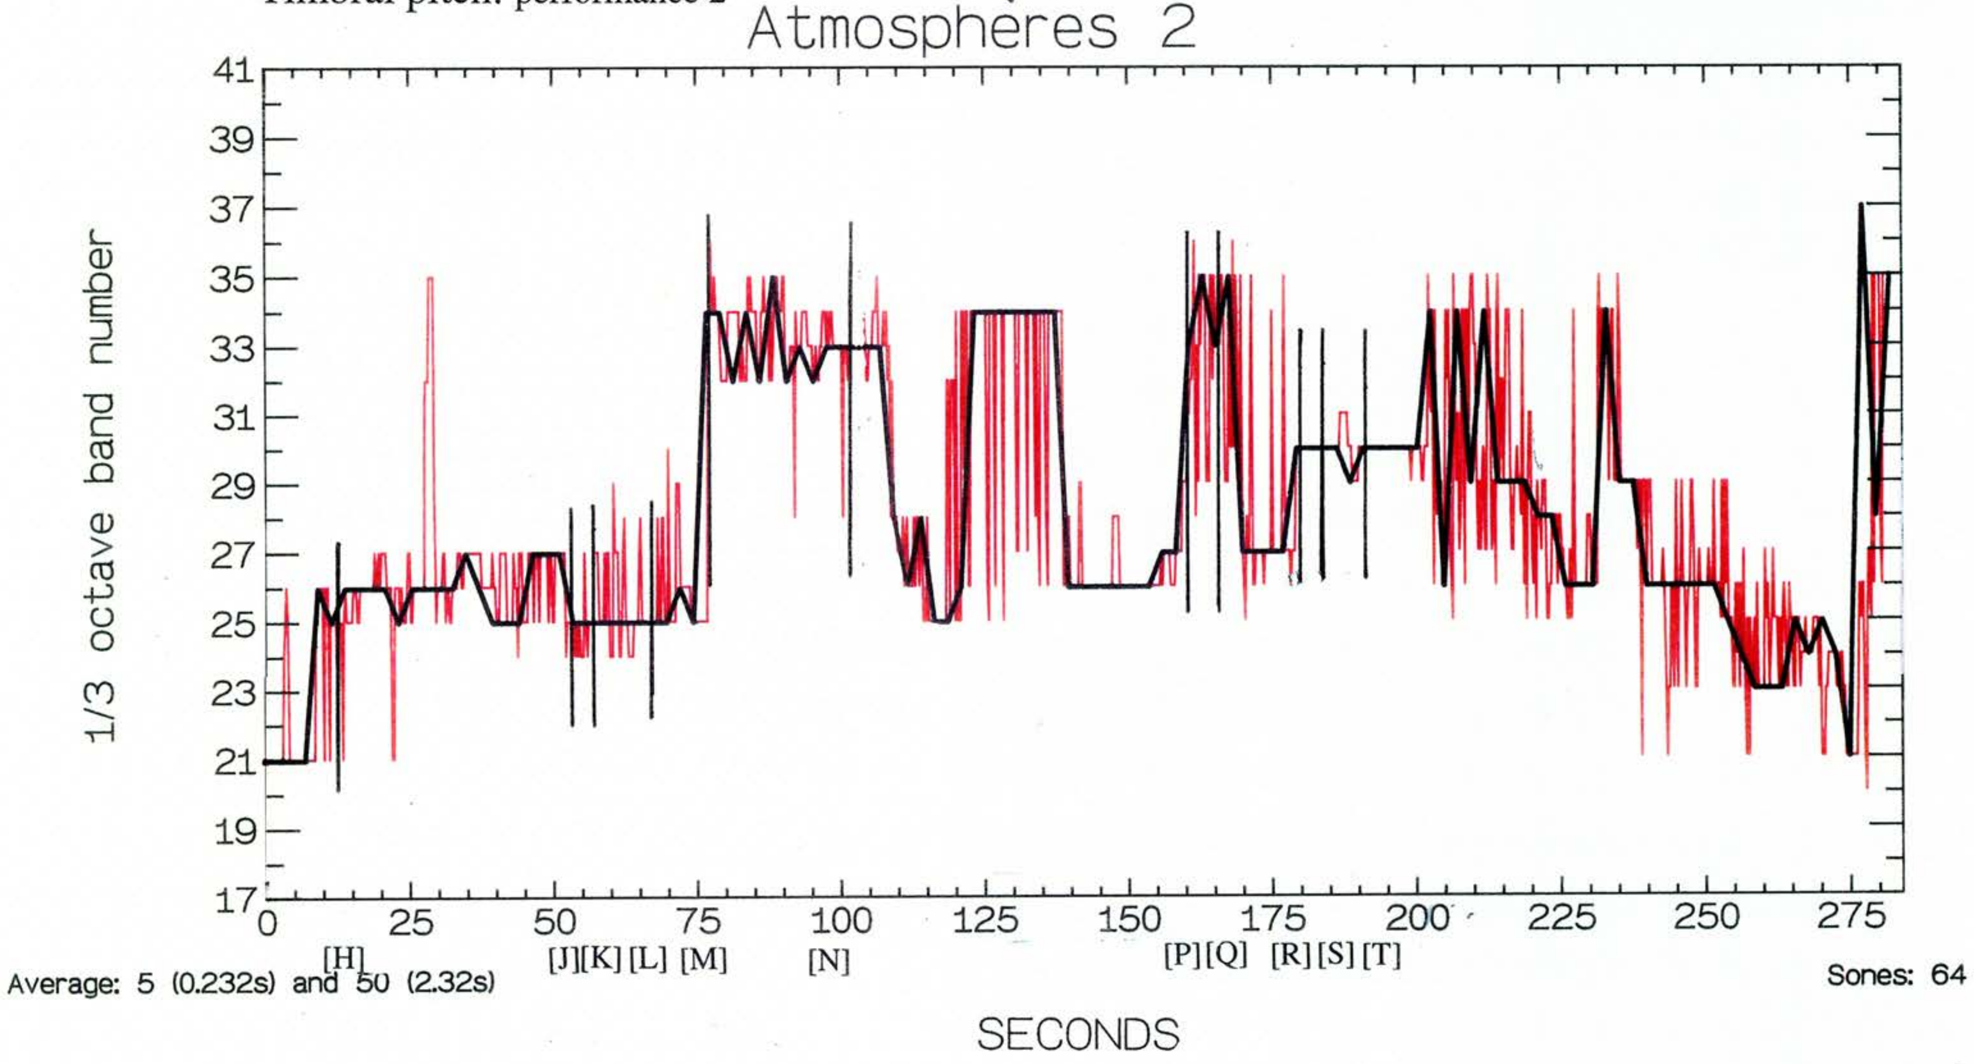

Figure 4.15a
Timbral width: performance 1

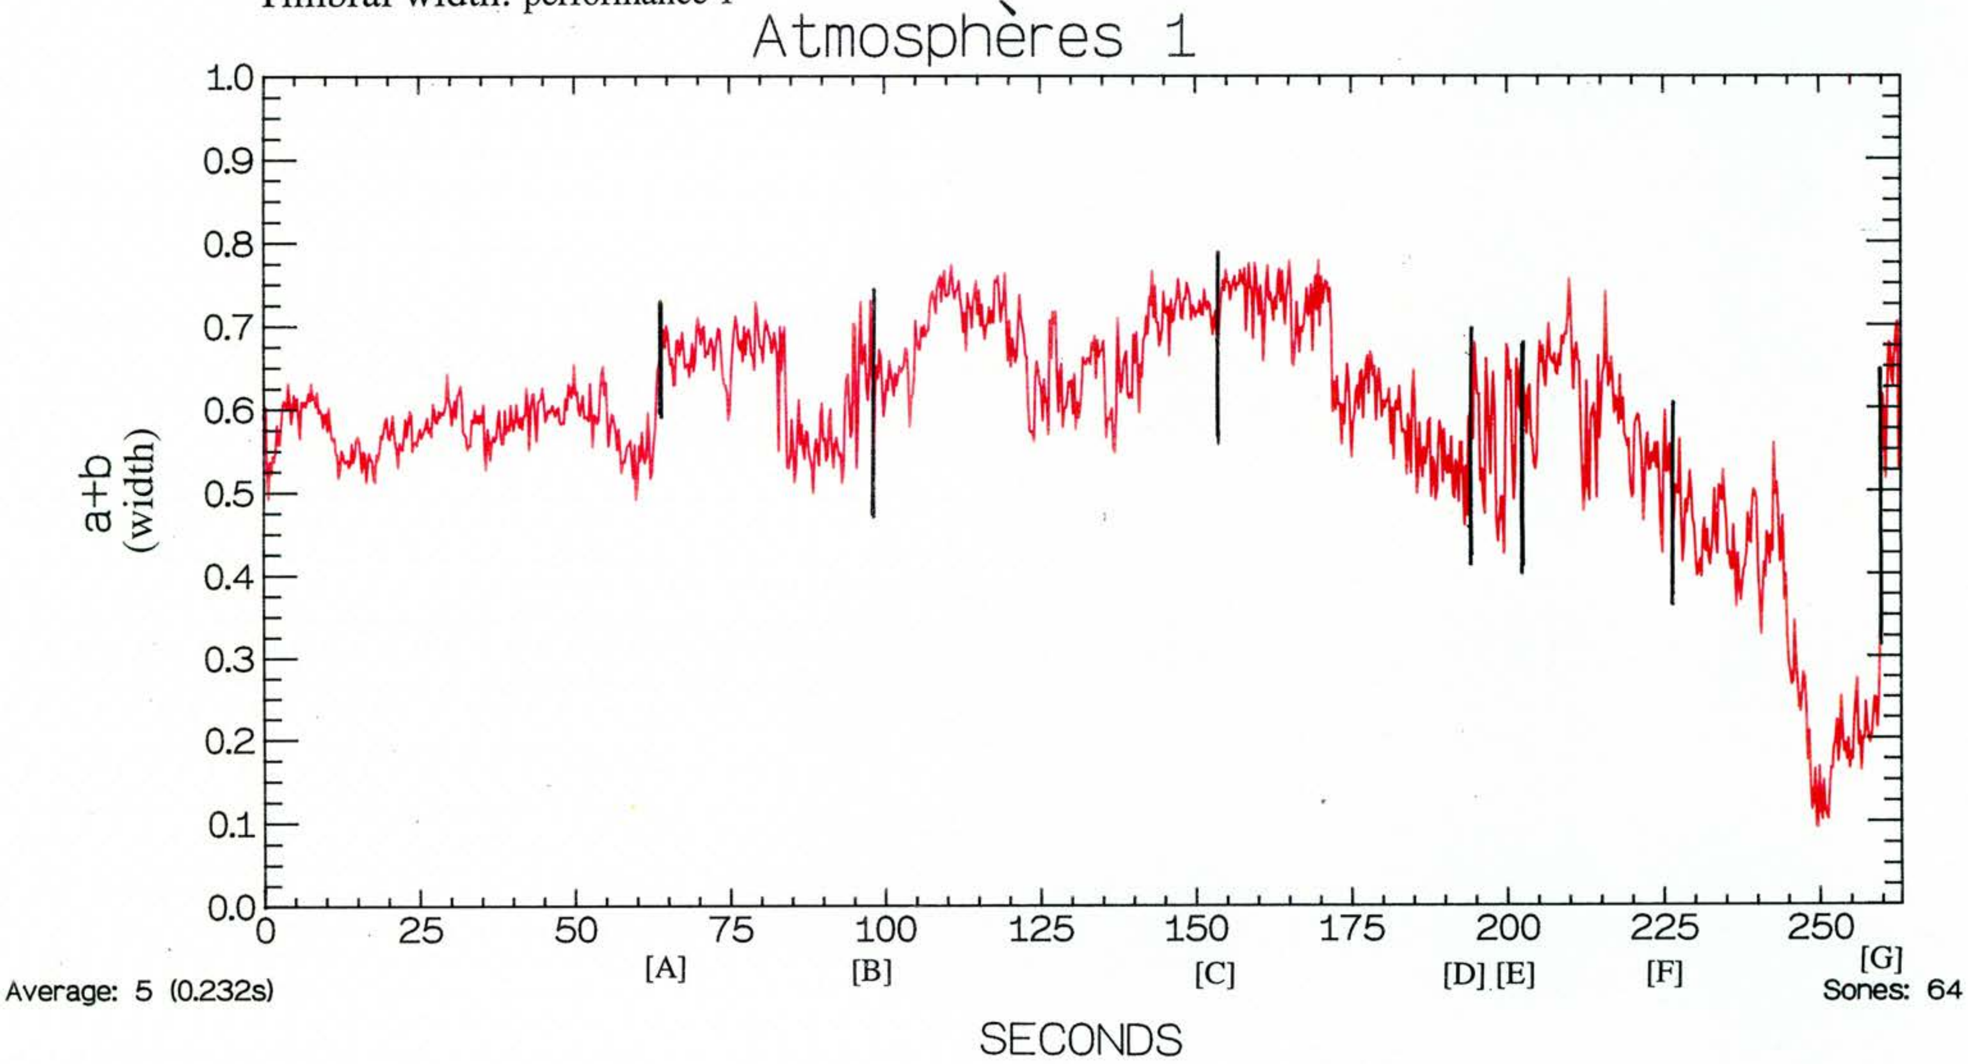

Figure 4.15b

Timbral width: performance 1

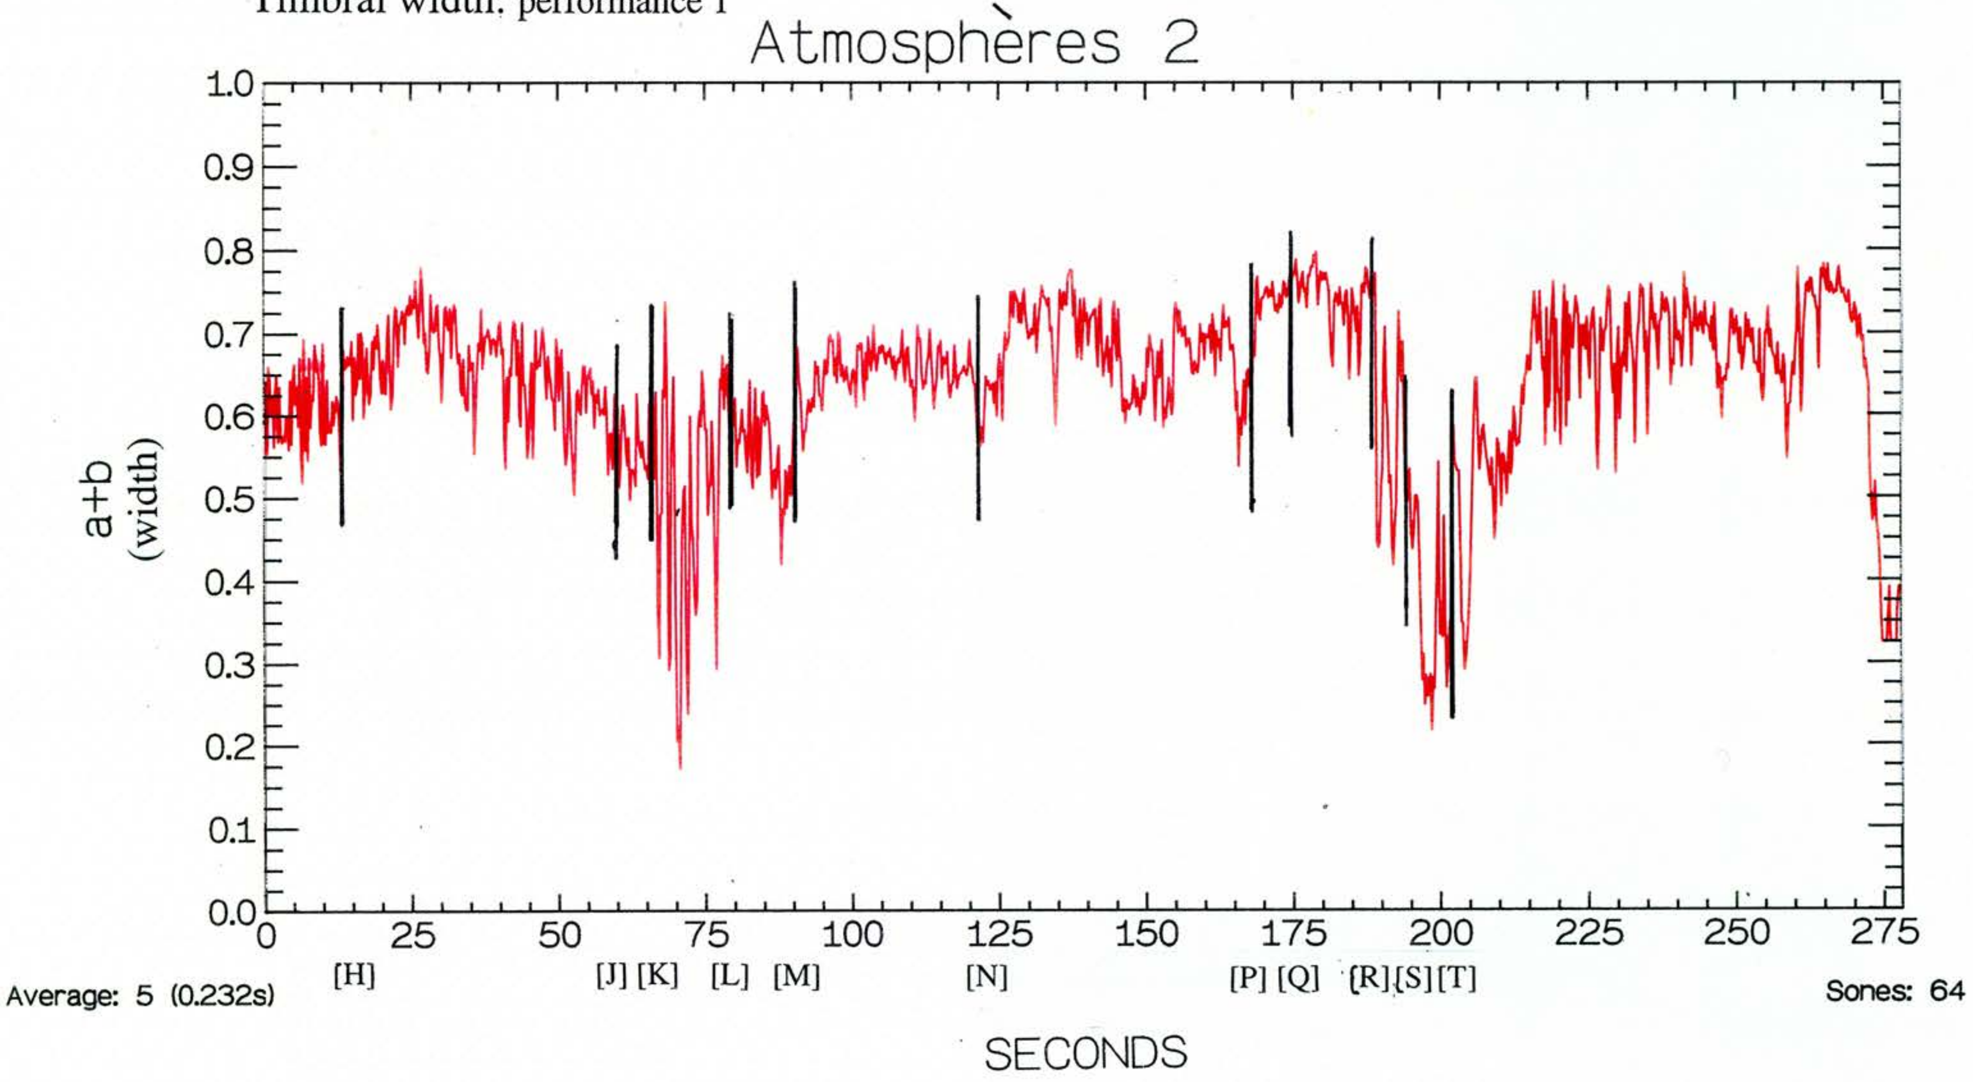

Figure 4.16a

Timbral width: performance 2

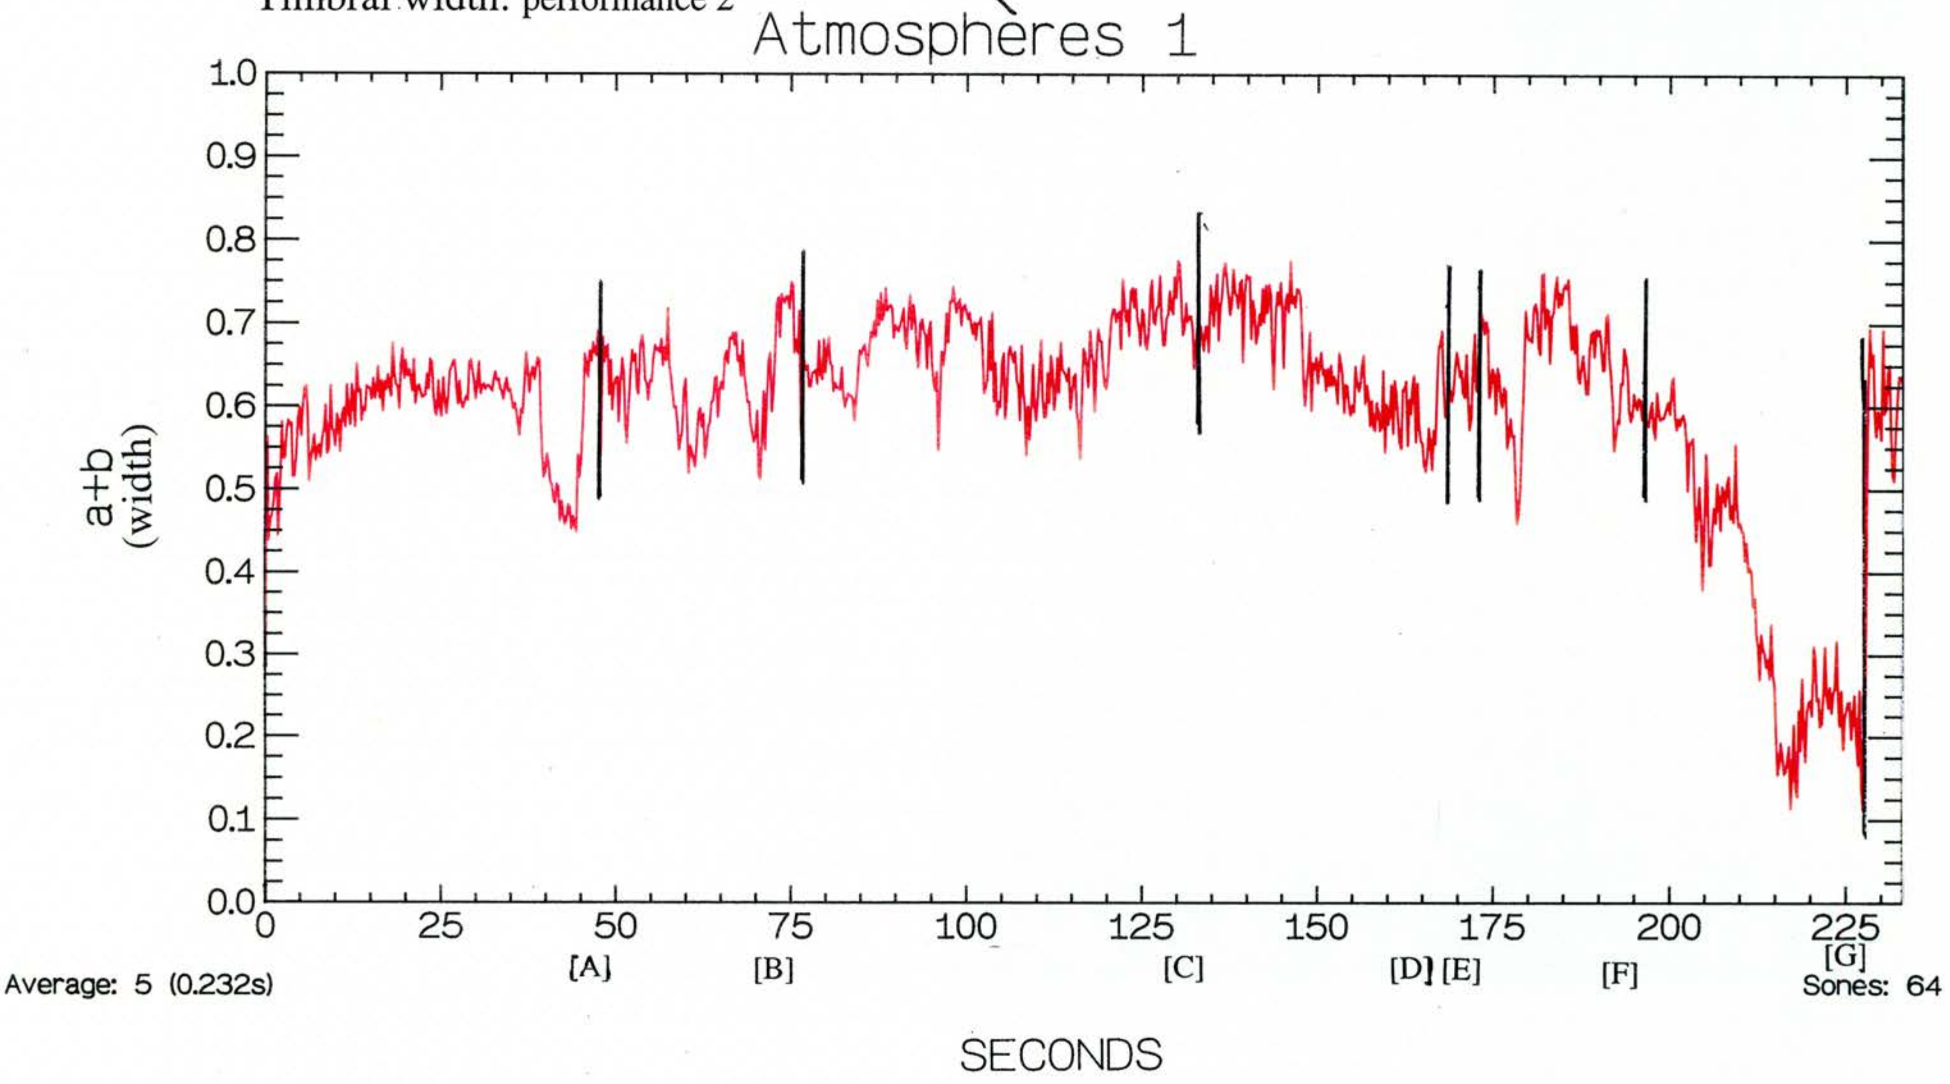

Figure 4.16b

Timbral width: performance 2

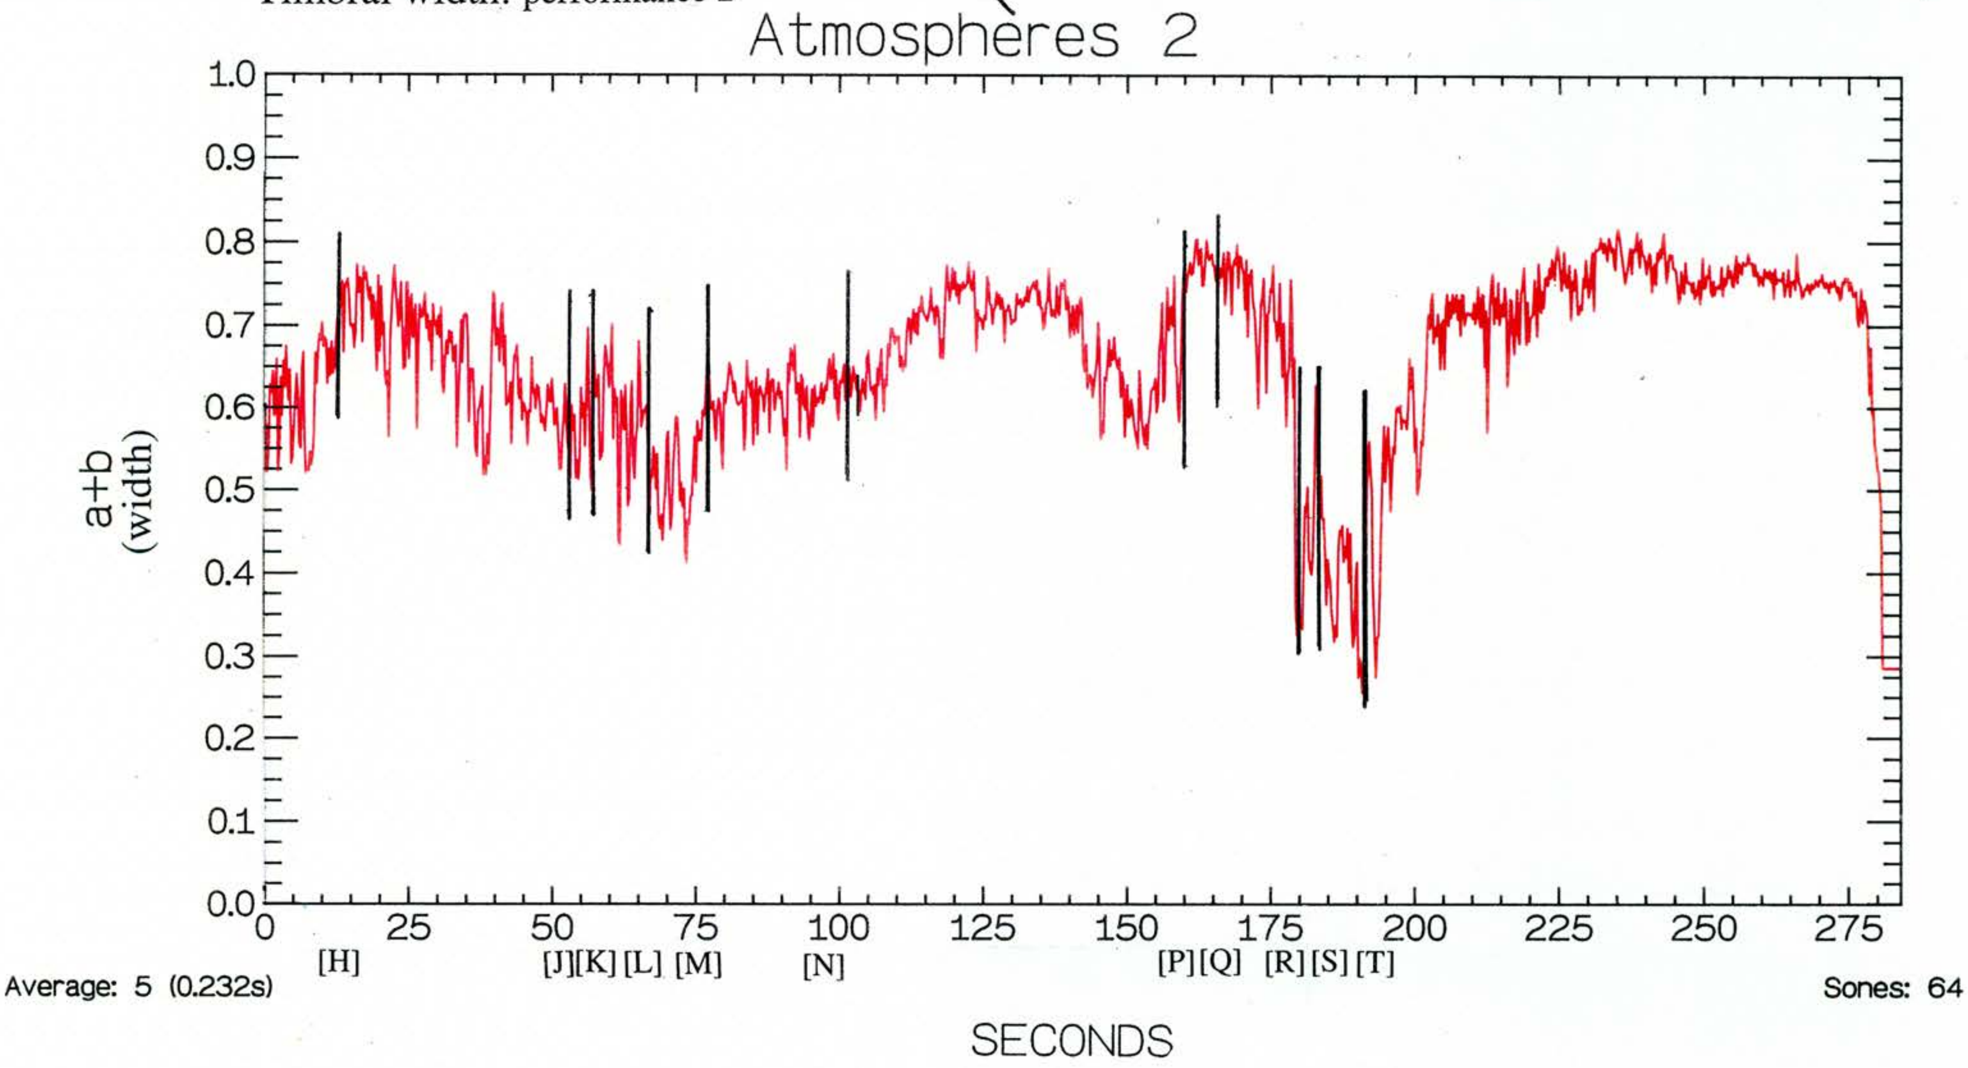

Figure 4.17a
Sharpness: performance 1

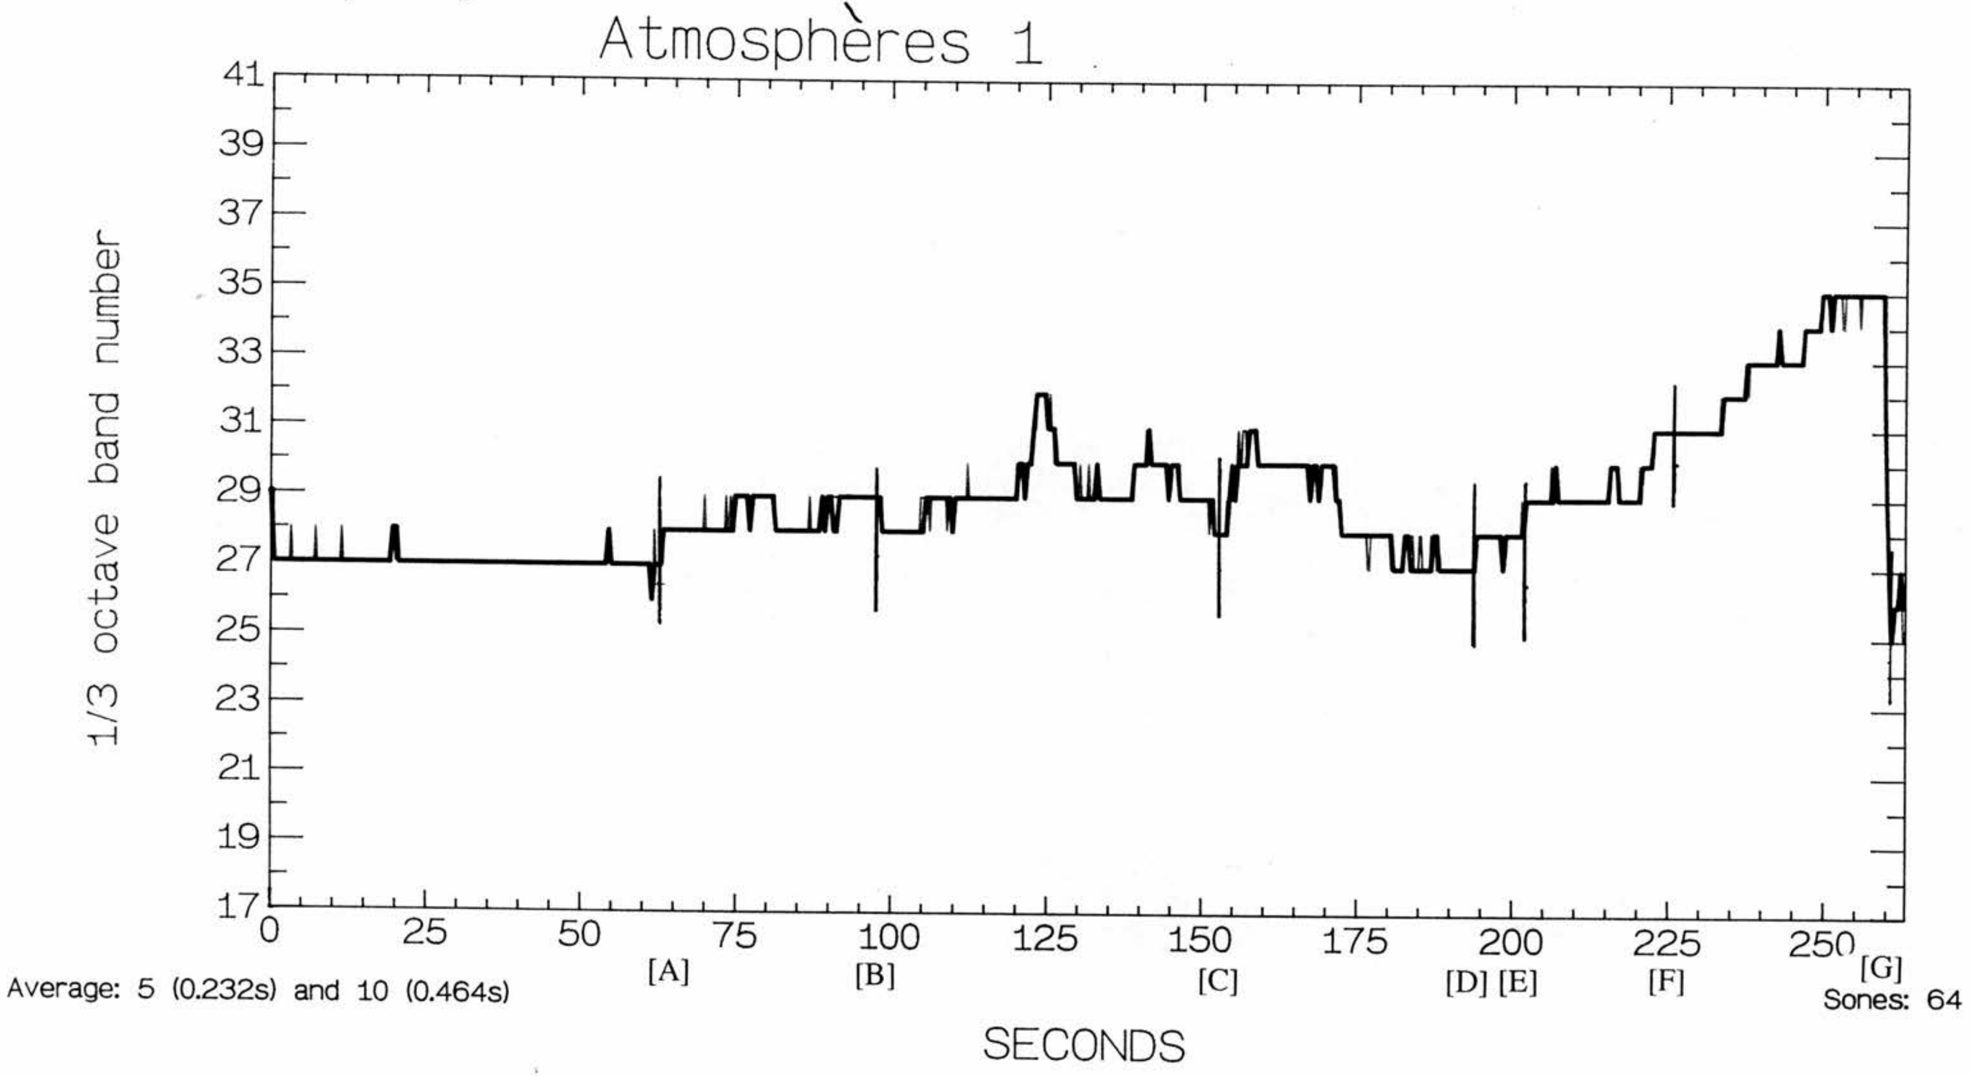

Figure 4.17b
Sharpness: performance 1

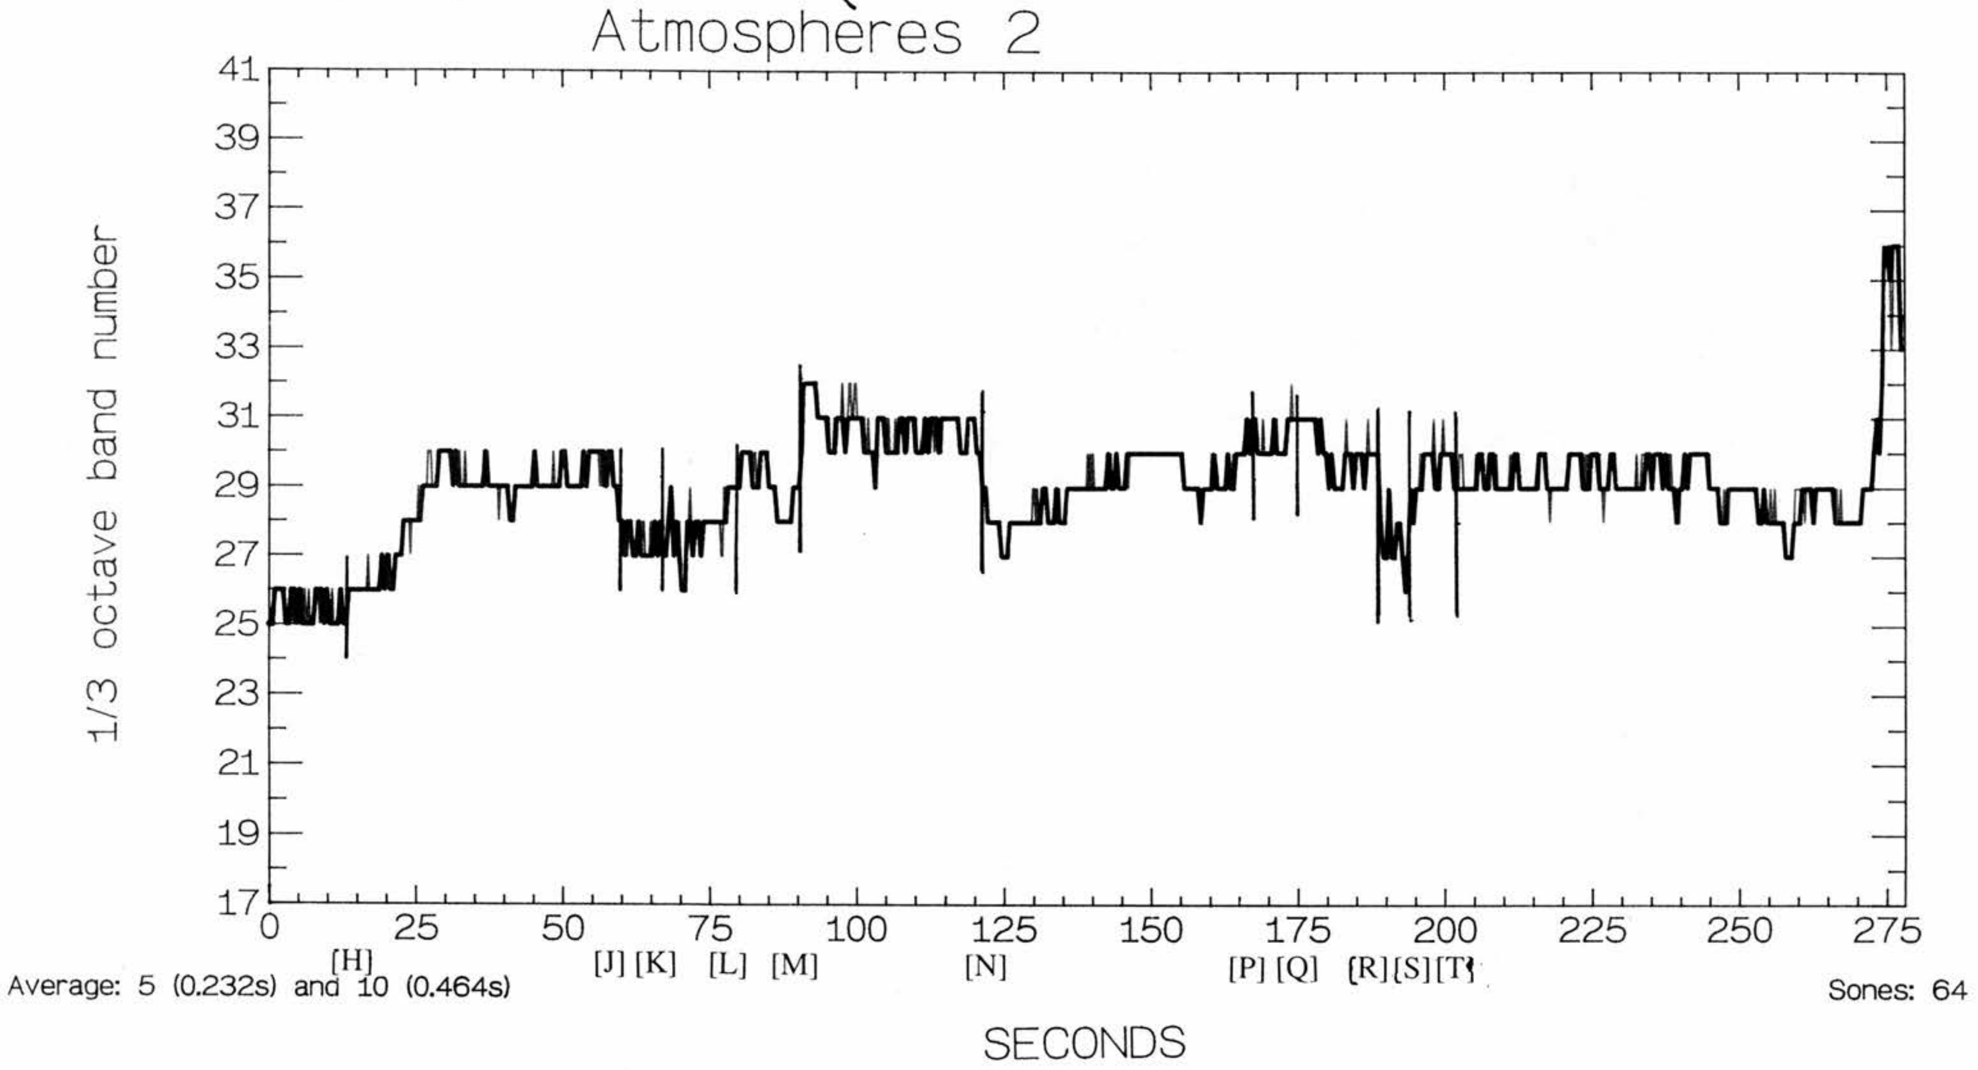

Figure 4.18a
Sharpness: performance 2

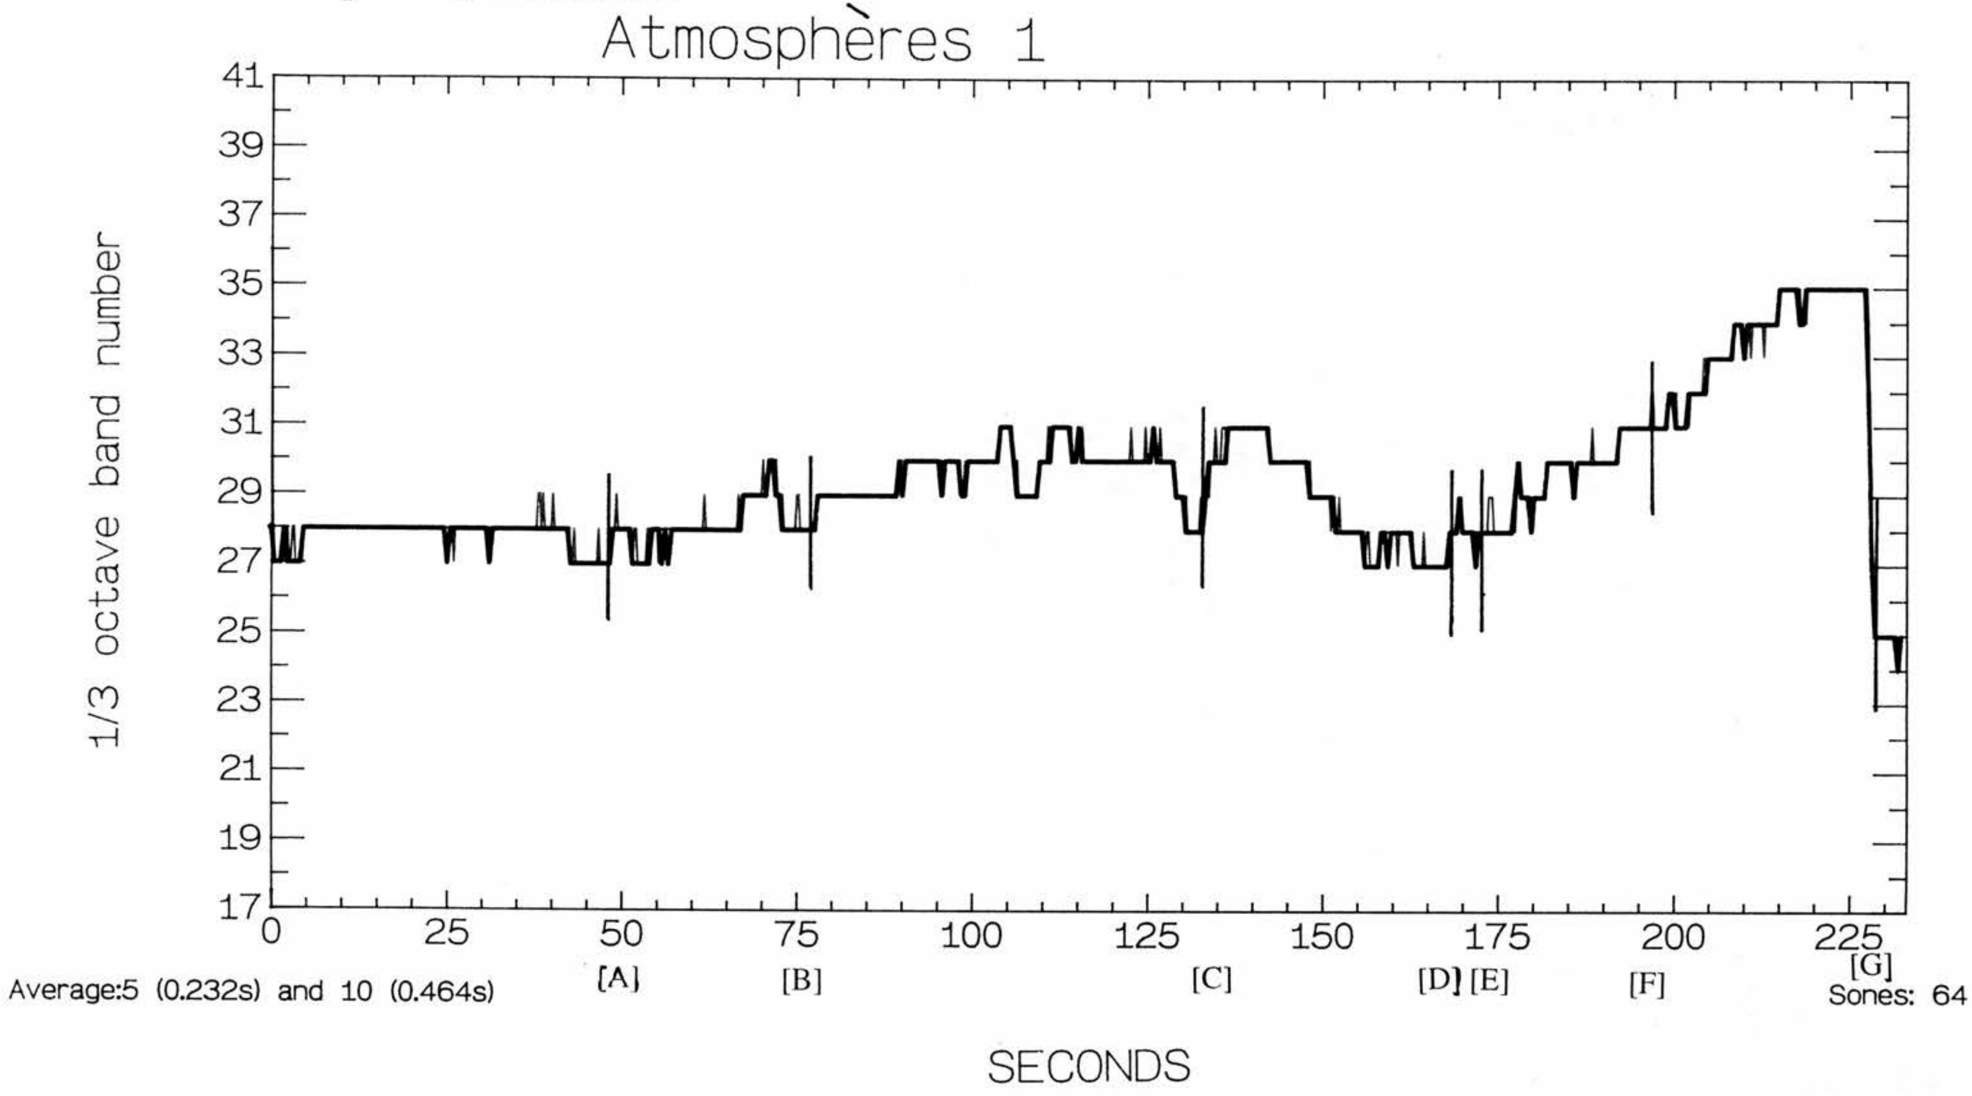

Figure 4.18b
Sharpness: performance 2

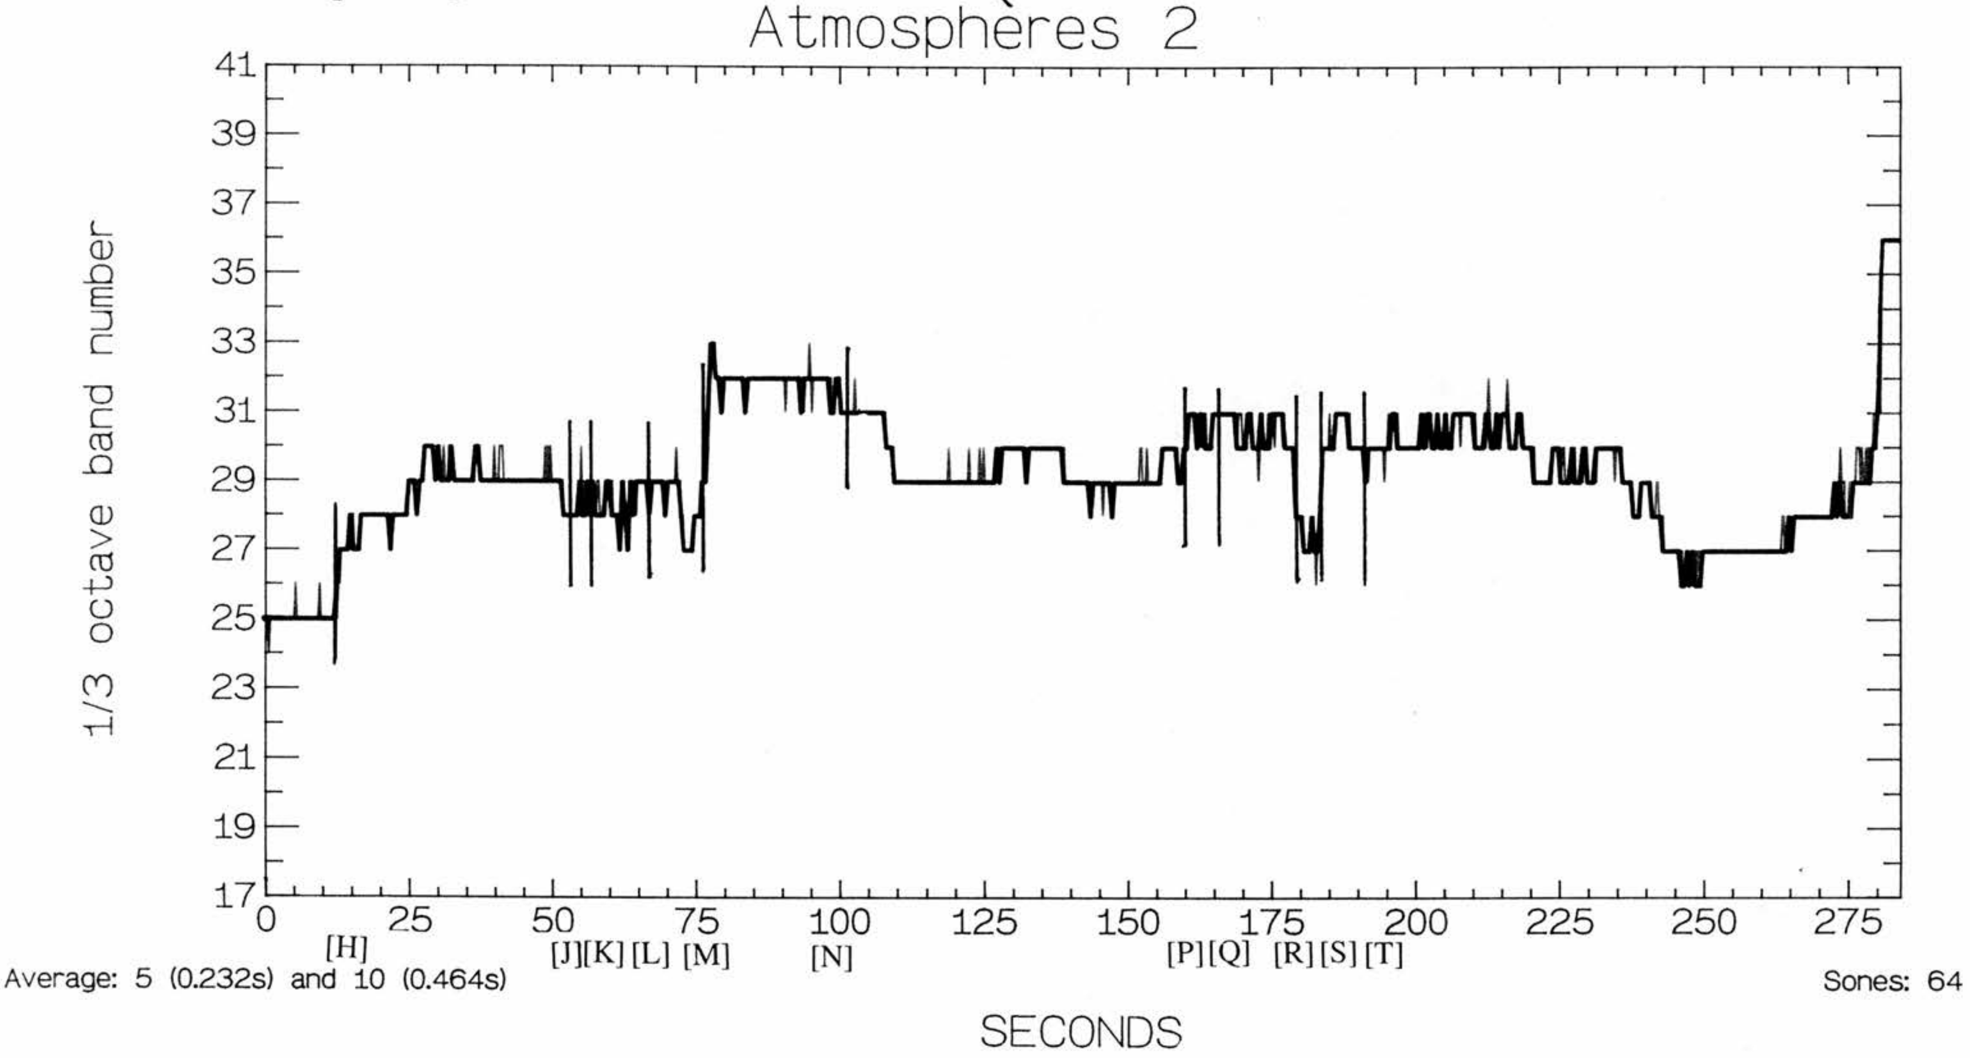

Figure 4.19a

Roughness: performance 1

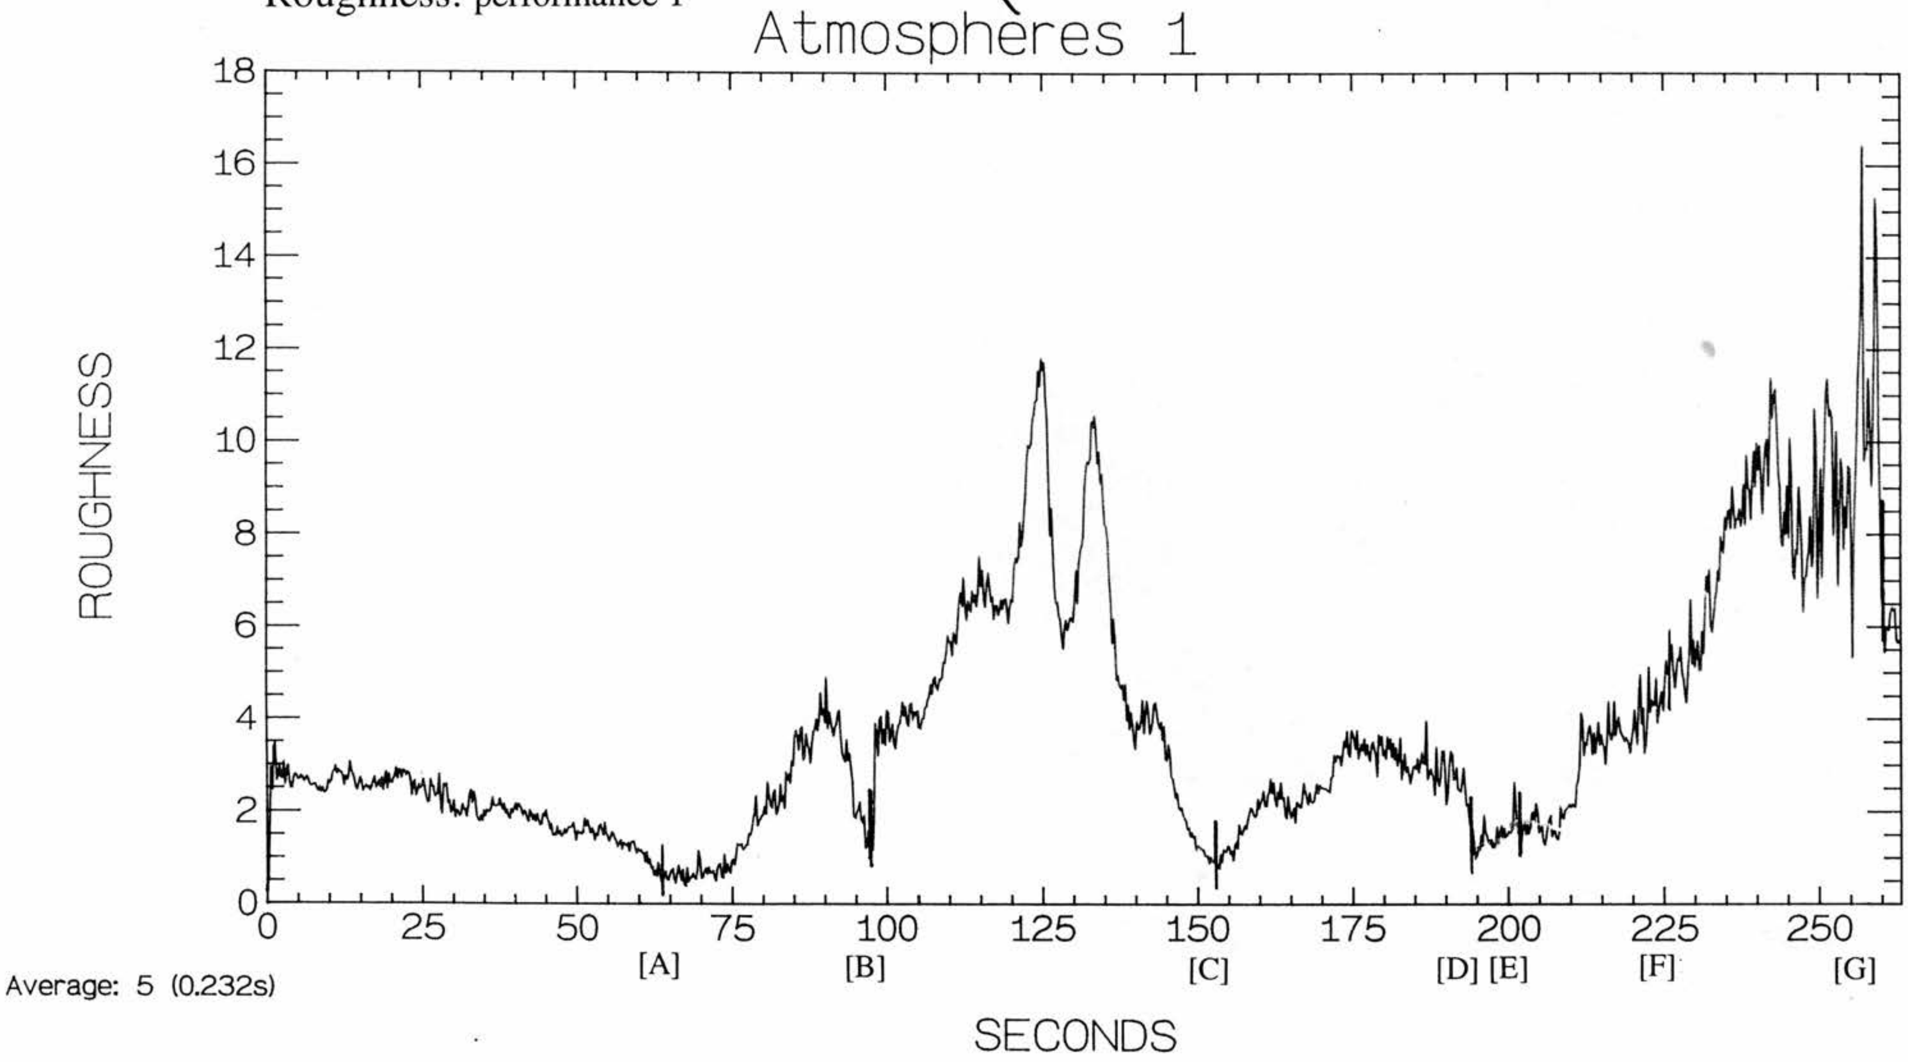

Figure 4.19b
Roughness: performance 1

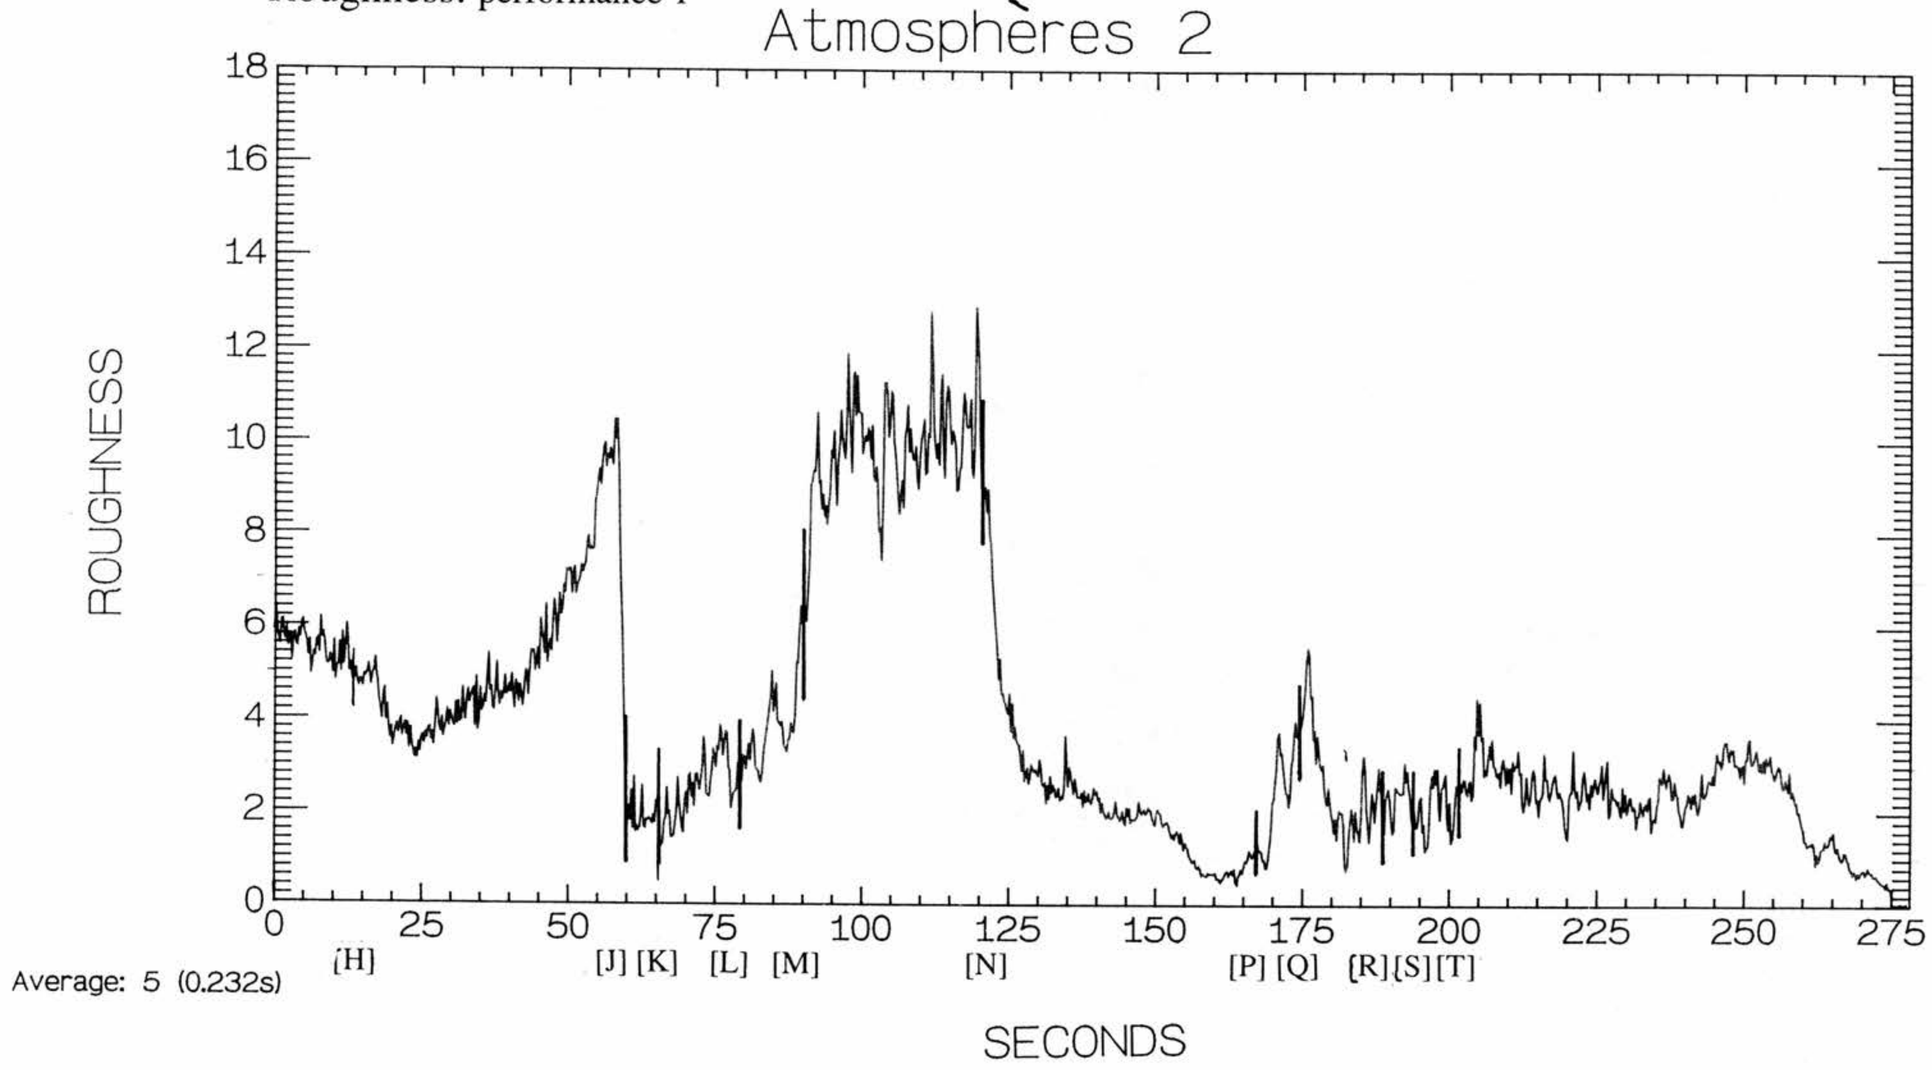

Figure 4.20a
Roughness: performance 2

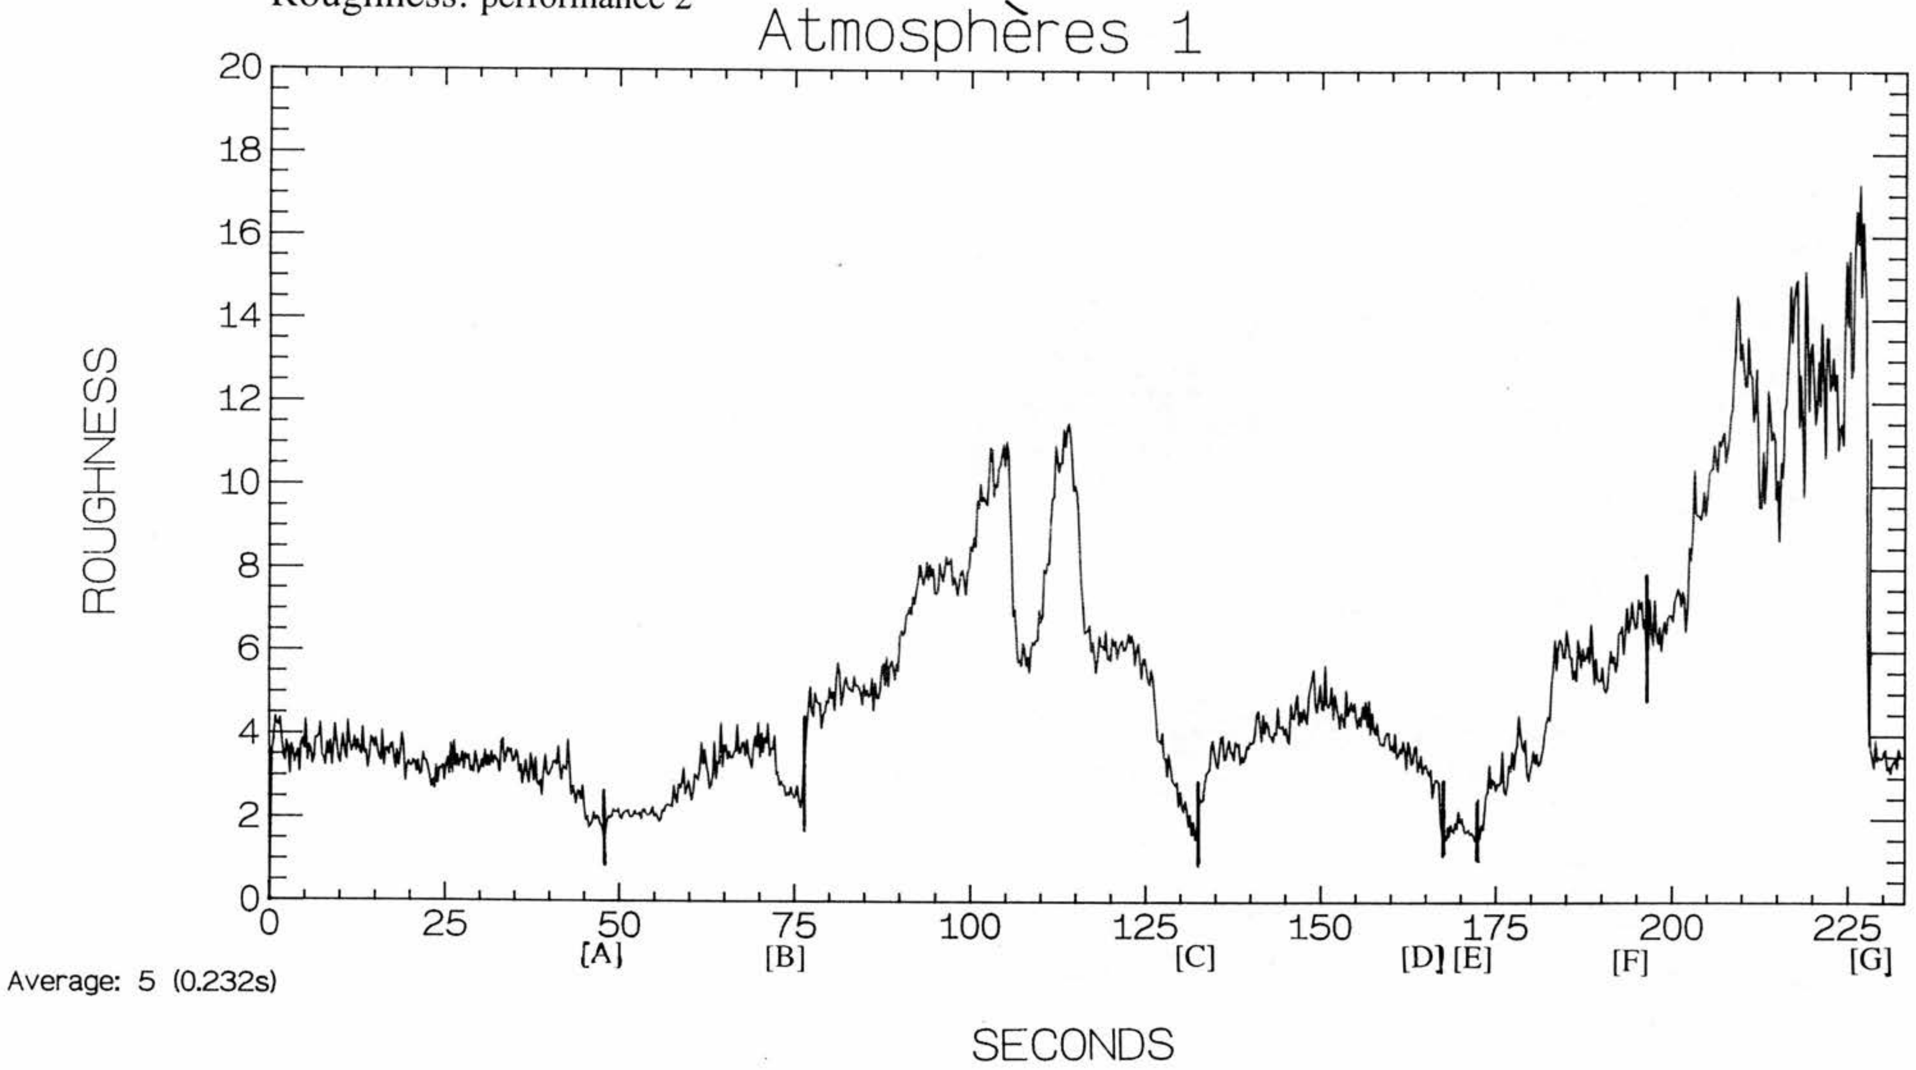

Figure 4.20b
Roughness: performance 2

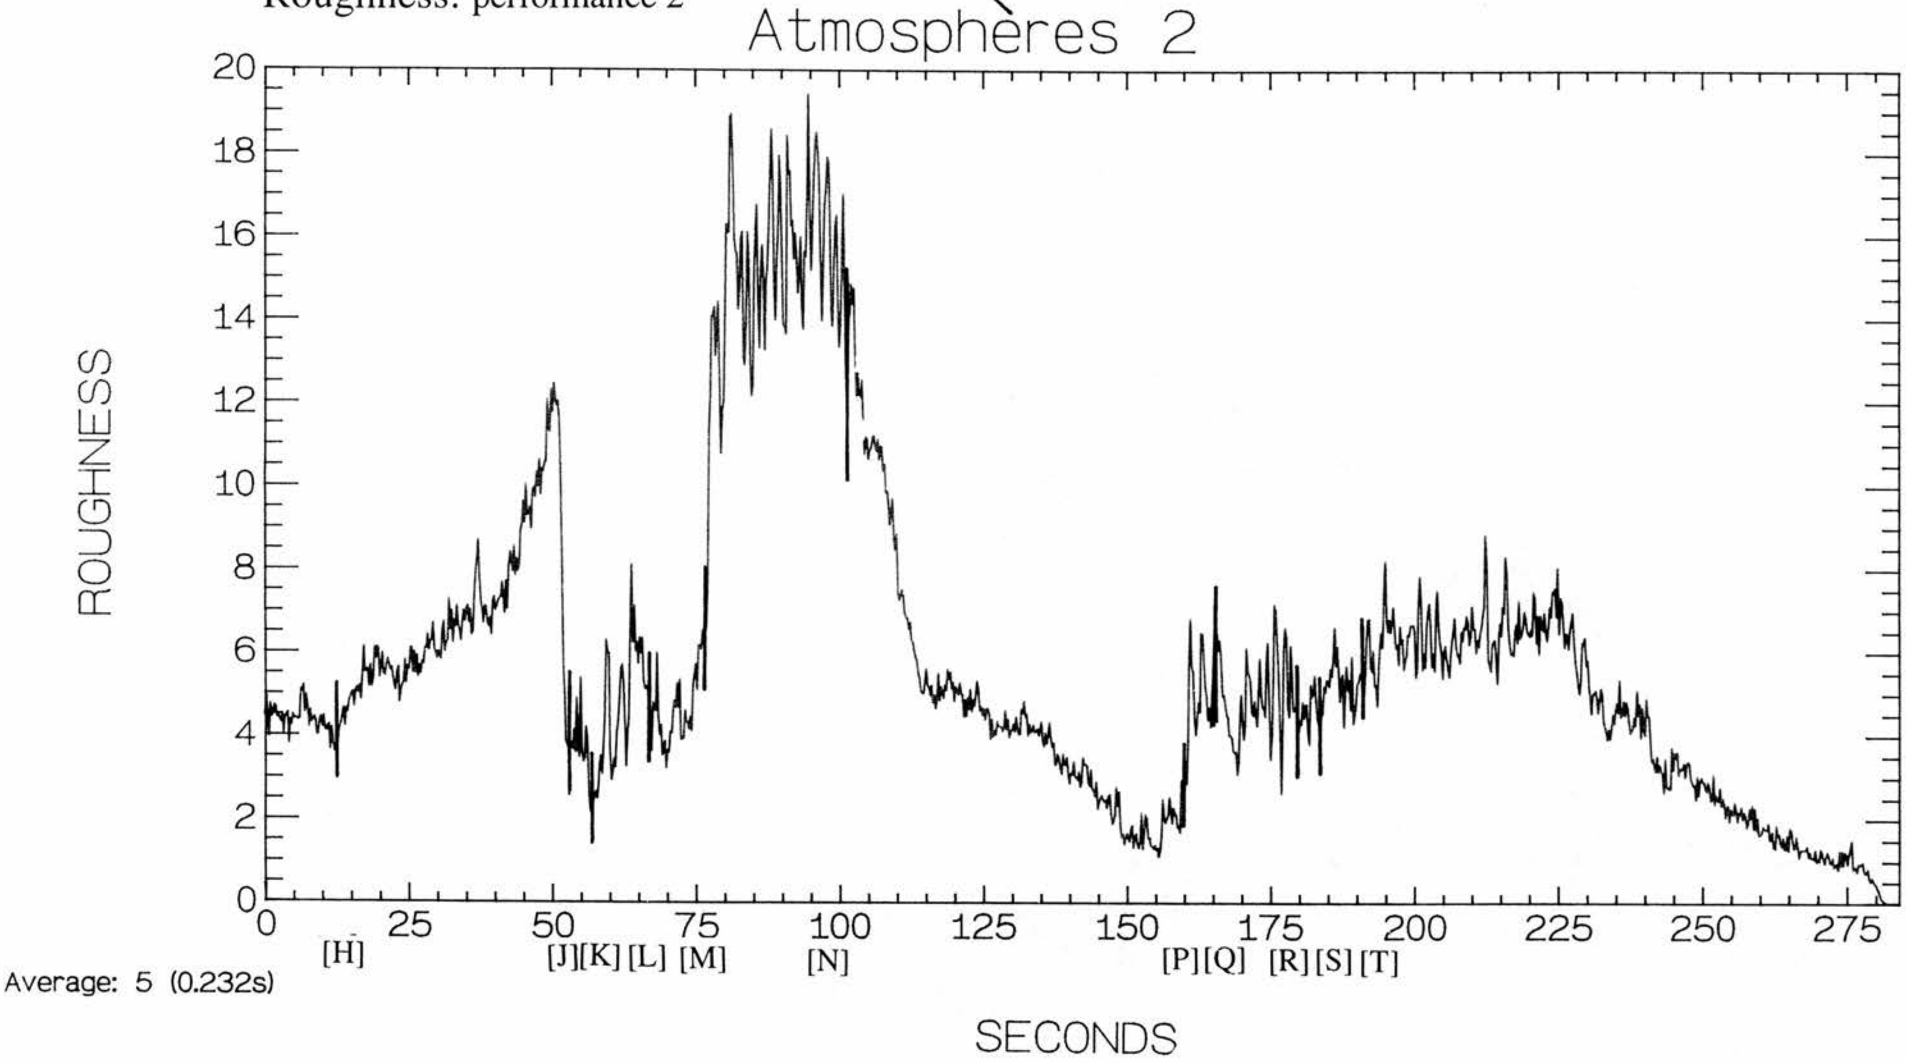

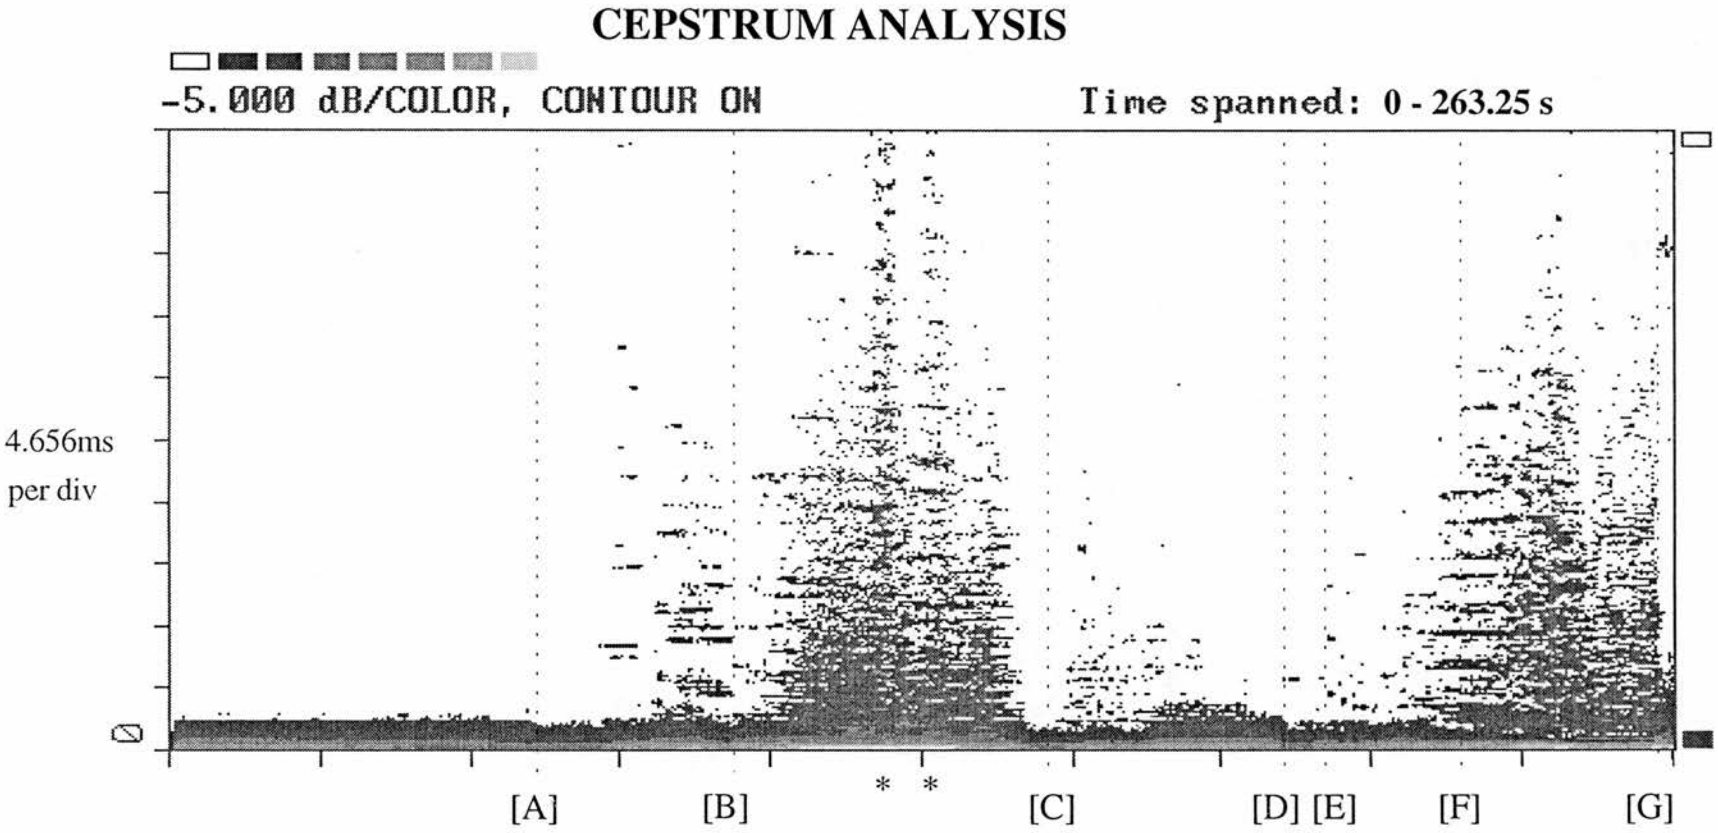

Figure 4.21a

Atmospheres: performance 1

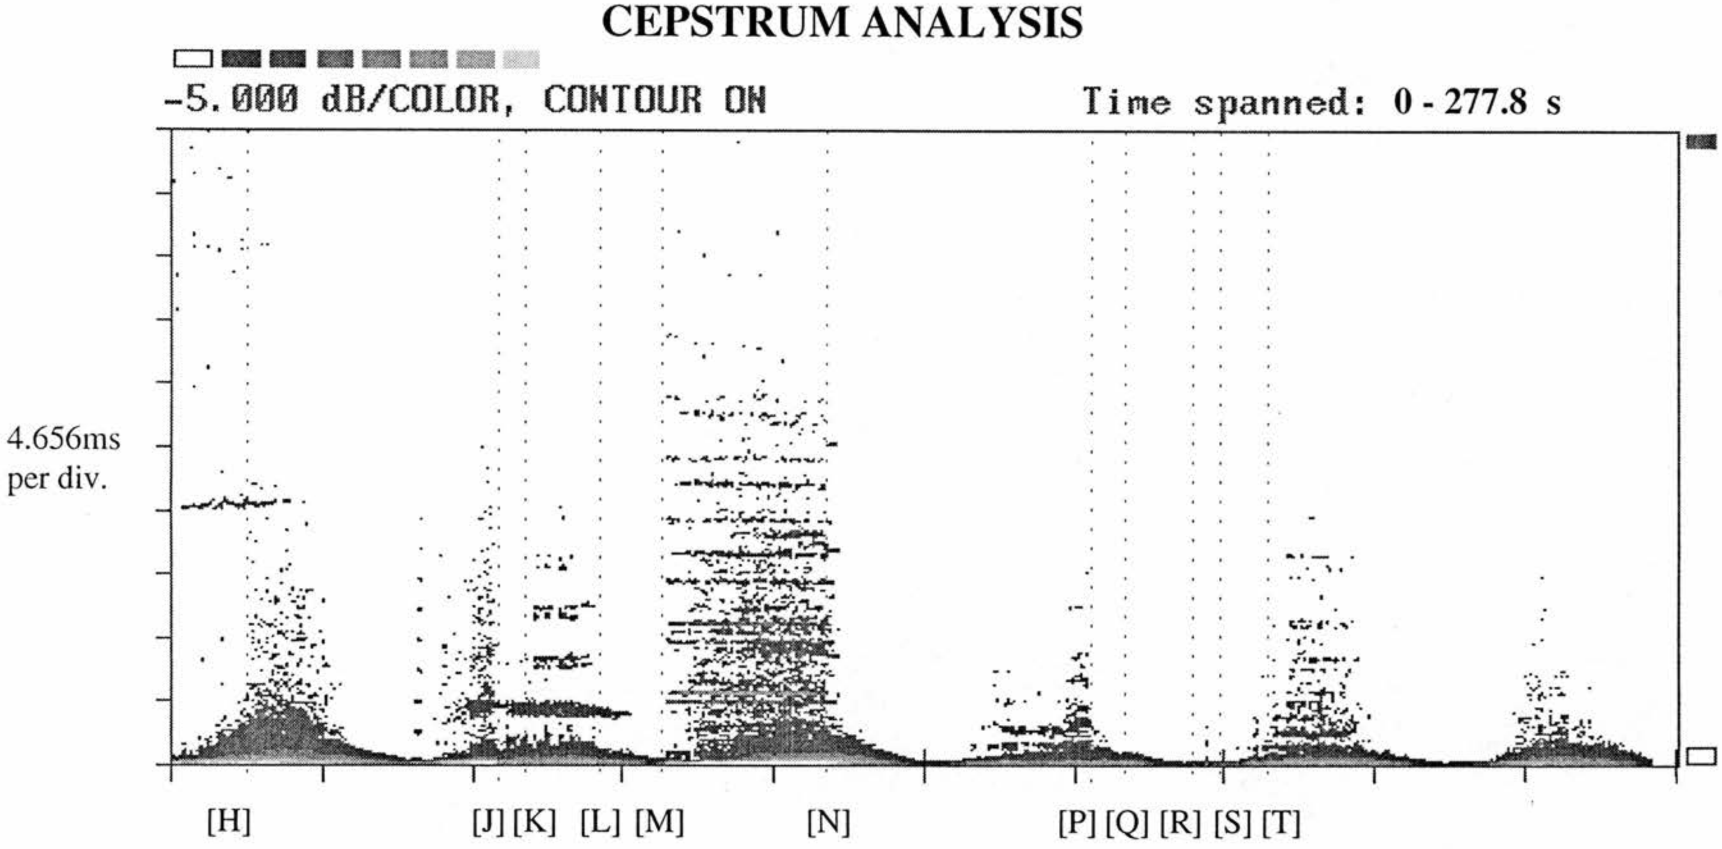

Figure 4.21b Atmospheres: performance 1

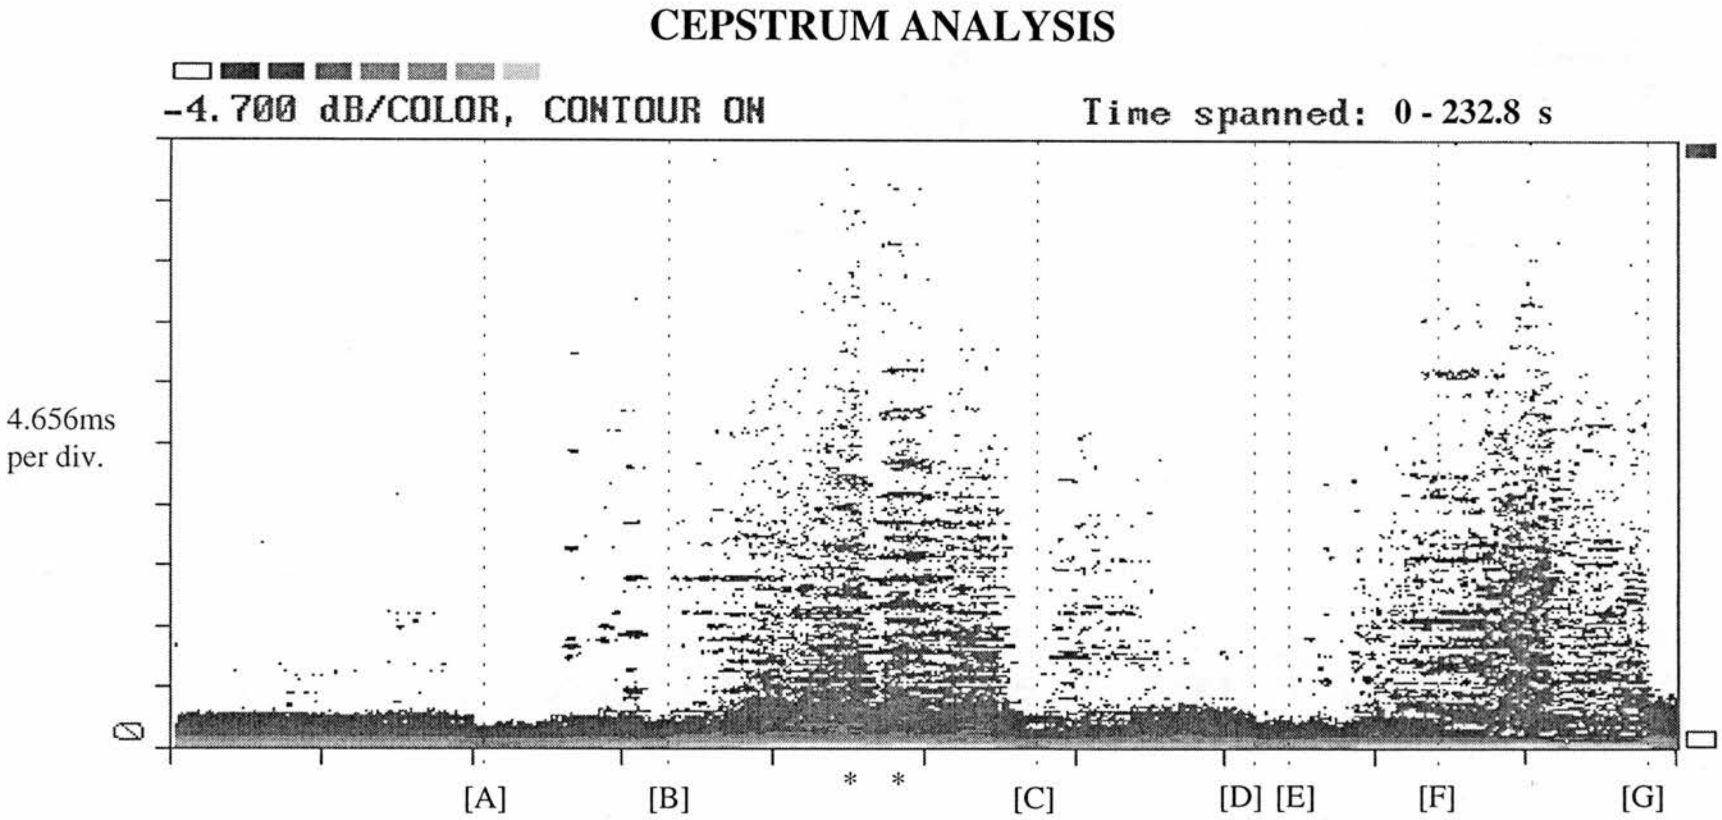

Figure 4.22a Atmospheres: performance 2

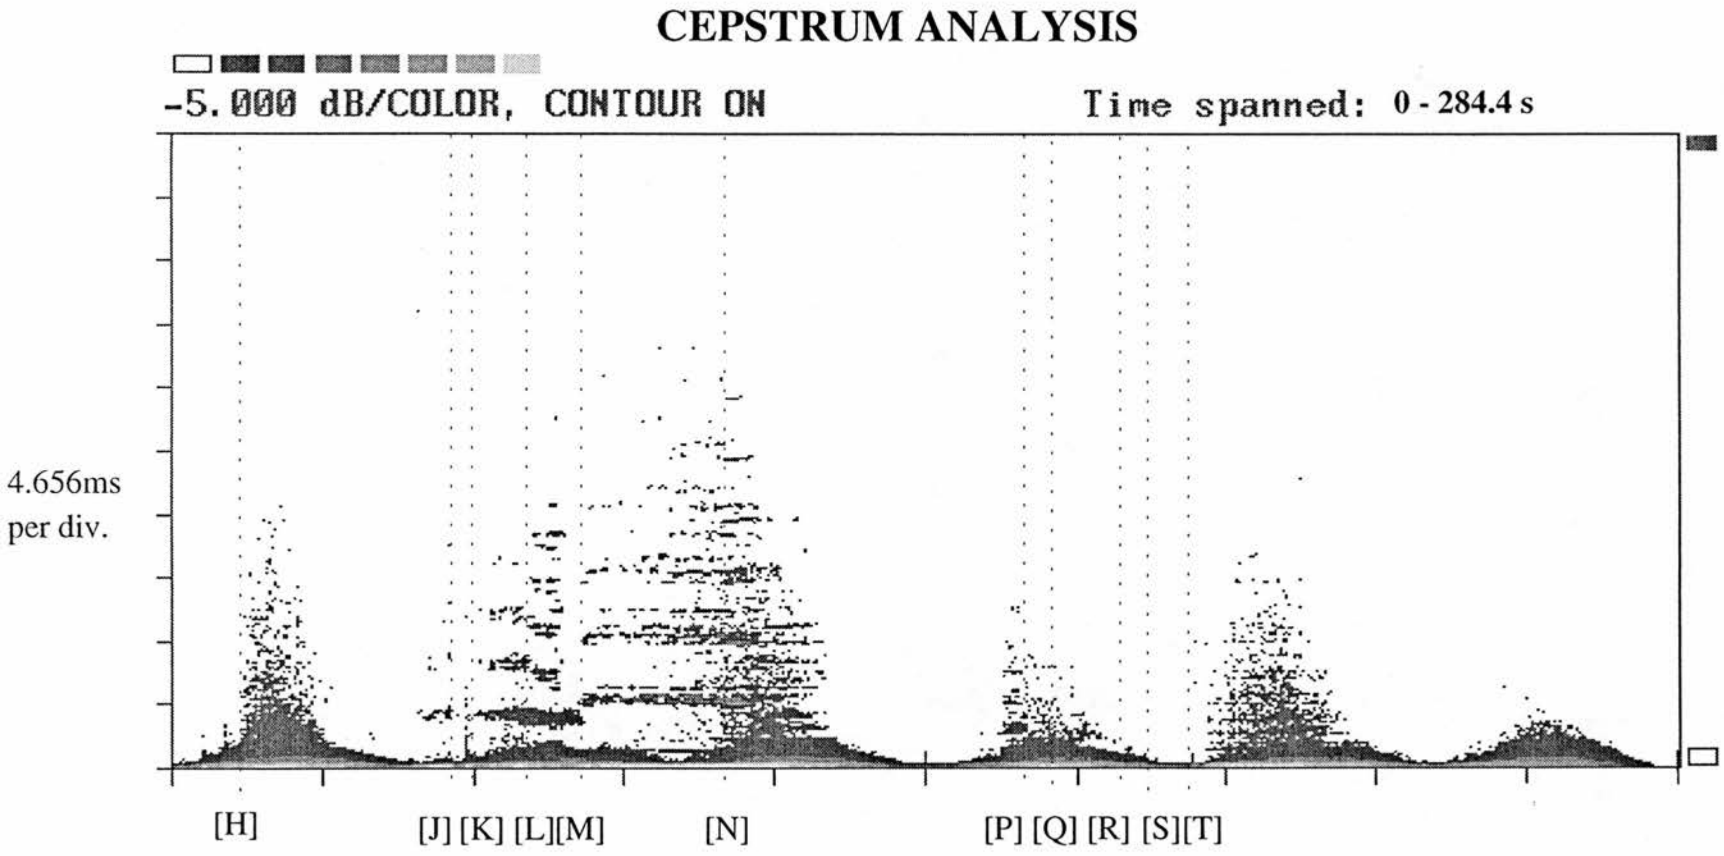

Figure 4.22b Atmospheres: performance 2