The *Rosegarden Codicil:* *Rehearsing Music in Nineteen-Tone Equal Temperament*

* **Graham Hair** (Centre for Music Technology and Department of Music, The University of Glasgow)
* **Ingrid Pearson** (Graduate School, Royal College of Music)
* **Amanda Morrison** (Freelance Soprano: *Exaudi, Synergyvocals, BBC Singers, Steve Reich and Musicians, Scottish Voices* et al)
* **Nicholas Bailey** (Director, Centre for Music Technology, Department of Electronic and Electrical Engineering, The University of Glasgow)
* **Douglas McGilvray** (Centre for Music Technology, Department of Electronic and Electrical Engineering, The University of Glasgow)
* **Richard Parncutt** (Institut für Musikwissenschaft, University of Graz, Austria)

[The authors would like to acknowledge the financial assistance of the Arts and Humanities Research Council of the United Kingdom, which made possible the research presented in this paper.]

The subject of this paper is microtonal tuning and a project which has developed several technologically-assisted methods of rehearsing music based on one particular microtonal scale.

We are using the term 'microtonal tuning' here to mean simply any tuning with more than twelve notes per octave, though we realise that the term is sometimes defined more broadly. For example, the definition could be so broad as to include all tunings other than the one based on the ubiquitous Western tempered scale with twelve equally-spaced steps per octave (hereinafter abbreviated as 12–ET). Scales with seven or eleven equally-spaced steps per octave, for example, could reasonably be considered microtonal. On the other hand, it would seem somewhat eccentric to call the various historical varieties of just and mean-tone tuning with twelve tones or fewer microtonal just because they are different from 12–ET. So we have framed our definition as indicated above merely with our present practical purpose in mind.

The composer's wish to explore the expressive possibilities of such alternative tuning systems presents new challenges to performers who, even at the highest level, may have limited experience with microtonality. Through the extension of a powerful piece of existing open-source software, the *Rosegarden* sequencer, a rehearsal aid has been developed which is highly accessible to expert musicians and sufficiently flexible to handle many different microtonal scales. We call it the *Rosegarden Codicil* : a mildly tongue-in-cheek adoption of the legal term for an addition or supplement that explains, modifies, or revokes a will or part of one (from the Latin *codicillus,* 'little book'). The software *modifies* standard MIDI-based sequencer code, so as to overcome the tuning limitations of most such programmes, imposed by their adoption of MIDI as an internal structural template (which means that they are able only to deal with standard 12–ET notation, and that badly). Using this software, we attempt here to *explain* how singers might approach microtonal scales in a way which links up with their existing musical skills and to *revoke* the idea that microtonality is accessible only to specialists or freaks.

Live preliminary demonstrations of the new software were presented by the clarinettist in our team, Ingrid

Pearson, accompanied by a synthesised microtonal keyboard part, played by the composer, Graham Hair, at the *Microfest* conference in London in October 2005 and by the soprano in our team, Amanda Morrison, at a colloquium in the University of Glasgow in February, 2006, both with software support from our Engineers, Nick Bailey and Doug McGilvray. Further demonstrations were given at a colloquium in the University of Edinburgh in November, 2006 and at a workshop of the Digital Music Research Network at Queen Mary College, London in December 2006. These presentations were followed by an intensive workshop at the Royal College of Music in London in February 2007, and will be followed further in due course by performances of new microtonal works whose preparation has been facilitated by this new rehearsal tool. The sixth member of our team is systematic musicologist Richard Parncutt of Graz University, who has provided input on the perception issues raised by this project, and suggested ways of taking the ideas explored further.

The aforementioned *Microfest* conference also provided resounding confirmation of the expressive possibilities of microtonal singing – if indeed one were needed – in the shape of a fine performance of a number of selections from *Seventeen Lyrics by Li Po*, a work composed between 1931 and 1933 by the 'father' of modern microtonality, the American composer, Harry Partch, given by the German mezzo-soprano Alfrun Schmid accompanied by the Dutch violist Elisabeth Smalt on her Adapted Viola, an instrument whose design is based on the instrument of that name invented by Partch himself. The concert also included a performance of *Psalm 137, By the rivers of Babylon* (1931), in which Schmid and Smalt were joined by keyboardist Bob Gilmore.

The following quote from critic Meirion Bowen, writing nearly 40 years ago about these Partch settings of texts from the Chinese and the Hebrew, suggests some aesthetic and expressive reasons for his interest in microtonal tuning. Although Partch deploys microtonal accompanying instruments, it is above all the voice, the vocal line and the declamation of the text which is the centre of attention.

> In his musical treatment of the *Seventeen Lyrics by Li-Po* (1931-33), Partch brought new life to the words in ways that do not unhinge them from their ancient origin. So, too, with his setting of *Psalm 137, By the rivers of Babylon* (1931): the wailing voice-intonation lends an age-old, timeless quality to the lament, enhanced by comparable sounds from the instrumental accompaniment of chromelodeon, kithara, and adapted viola. (Bowen 1968)

These are by no means the only possible aesthetic and technical reasons for adopting microtonal tuning, but certainly something of the same *general* intentions drove the compositional aspect this project: a search for the technical means of enabling texts from remote times, places, languages or modes of utterance to find expression through the musical language of the twenty-first century and contemporary vocal performance practice. The *specific* intentions were however quite different from the restraint and austerity of Partch: a search for a musical language to reflect something of the ecstatic Islamic mysticism of various couplets from the celebrated fourteenth-century Turkish sufi poet Yunus Emre in the English translations of Süha Faiz (Emre 1992: 63). The first of these couplets provided the text for Hair's song *Ash* – for soprano, clarinet and harmonium – which is the 'test piece' in the research process documented here.

The *Rosegarden Codicil* is a project of the Network for Interdisciplinary Studies in Science, Technology and Music ('*n*-ISM')[1].

1 Network for Interdisciplinary Studies in Science, Technology and Music ('n-ISM') home page. http:// www.n-ism.org

1 19–ET: Rationale and Context

Joel Mandelbaum's doctoral thesis *Multiple Division of the Octave and the Tonal Resources of 19-Tone Temperament*, presented to Indiana University in 1961 (see Mandelbaum 1961), has still not been superseded as a well-considered investigation of the pros and cons of historical and current microtonal practice, even though there have been several subsequent studies (for example: Blackwood 1985 and 1991) which are more sophisticated mathematically and methodologically. The following quote from Mandelbaum's introduction remains largely true forty-five years later: 'In spite of the tremendous changes in musical taste and style which have taken place in recent years, very few musicians appear to see any general departure from the limit of twelve tones per octave as likely in the foreseeable future.' (Mandelbaum 1961: 1). However, there is an extensive minority tradition, with antecedents dating back decades, so one cannot regard microtonality as radically new in any literal sense. In fact, the existence of such a minority tradition was an important driving-force behind the present project, inasmuch as it arises out of the conviction that novel musical structures should not simply be arbitrary, but should embody plausible extensions of constraints which are currently culturally accepted. Nevertheless, the diversity of contemporary microtonal practice and the absence of the widespread adoption of anything which could be called a common practice militates against the likelihood of a change any time soon in the general situation Mandelbaum observed back in 1961.

Despite this, the voices of the dissenting minority tradition continue to make themselves heard. Mandelbaum suggests that '... the advocates of multiple division, however disparate their views, share a feeling of discontent with the prevailing system, 12-tone equal temperament' (Mandelbaum 1961: 1). His dissertation, following on from the studies of Bosanquet (1875) and Yasser (1932) devotes considerable space to the mathematics of intonation theory, documenting many interesting intonational properties of the scale with nineteen equal steps to the octave, in particular (hereinafter abbreviated as '19–ET'). This places his dissertation within a tradition of music theory which reaches back many centuries, of course. Nevertheless, his chapter on 'An Acoustic Comparison of 19-Tone Systems' is forced to conclude that '...it is difficult to claim for 19-tone temperament anything better than the compromise status which is traditionally claimed for 12-tone temperament...' (Mandelbaum 1961: 336).

Hence, what has been driving the current project has been not so much discontent with 12–ET as the possibility of drawing upon a rich, alternative source of expressive possibilities, particularly in vocal music. Mandelbaum's chapter concludes by reporting four general points made by the American musician Leigh Gerdine (1917–2002). Gerdine (see McLaren 2006) held, amongst other appointments, that of president of Webster College in St Louis for twenty years. Broadly speaking, our approach emerges from a viewpoint similar to that put forward by Gerdine, so we adopted 19–ET as our tuning system for the purposes of this project. Mandelbaum wrote: 'Leigh Gerdine, in writing to Yasser, lists the following as assets of 19-tone temperament: it includes within its framework the 'good' music of the past; it allows for indefinite extensions; it affirms the validity of the natural harmonic system; it makes possible a future in vocal as well as instrumental music.' (Mandelbaum 1961: 336). We shall occasionally refer to these four 'Gerdine properties' in what follows.

The last of these 'Gerdine properties' is perhaps the most relevant to this project. Most of Hair's pieces are composed for voices, so the especially interesting compositional feature of the *Rosegarden Codicil* project is the development of a rehearsal tool which might be used by singers: not necessarily those with a special interest in modern, experimental or microtonal music, but simply by those with a regular Conservatory training and background and a mind open to the possibilitity that performing microtonal music may, with appropriate attitudes and rehearsal techniques, be a practical proposition.

Amongst the various challenges faced by such performers is that, unlike the position of (some) composers of microtonal music, most singers will probably be unable to specialise in microtonality, and devote their career to it. They will need to spend most of their time performing the standard repertoire, whether opera or early music, lieder or whatever. Secondly, singers cannot be equipped with frets or keyboards, which will mechanically control their intonation, but will have to use their built-in biological equipment – aural, oral and vocal tract – and their conceptual capacities to produce microtonal melodies.

Our composer comes from a background which includes both a Conservatory training in performance (piano) and University-based work in composition and research, and he works most of the time with singers and instrumentalists who usually perform standard repertoire. He has thus been engaged for a long time with 12–ET intonation and continues to be so engaged, most of the time. His work in microtonality is, on the other hand, rather recent, so he approached microtonality from what he considered might be a useful direction for someone with a lot of general musical experience but rather less experience specifically in microtonality, viz by imagining himself a singer and engaging with microtonality through the specific tasks – conceptual, perceptual and physical – which singers require to perform. The question is: what is required to conceptualise microtonally (and in particular, to conceptualise music using the 19–ET scale) in such a way as to be able to perform microtonal musical works without mechanical aids such as frets, keyboards or other mechanically-assisted intonation controllers, but bringing to the problem long practical experience in performance and an otherwise relatively well-developed capacity for musical conceptualisation in other contexts?

Hair has in recent years been involved compositionally with what is sometimes called New Tonality, in which (to generalise with brutish approximation), tonal centres, diatonicism, consonance, and textural continuity play a role alongside acentricity, chromaticism, dissonance, and textural discontinuity. Harmony is – like atonal harmony – *contextual* (or referential), not functional. His approach to the 19–ET scale came out of this background, conceptualising the 19–ET scale as a superset of the diatonic scale (i.e. comprised of the seven-note diatonic scale plus twelve chromatic tones) rather than expanding the chromatic scale from twelve chromatic tones to nineteen. This approach is really inspired by the first of Leigh Gerdine's properties (the relationship of the 19–ET's conceptual framework to that of existing repertoires of the past generally agreed to be 'good'), rather than by his third (that 19-ET 'affirms the validity of the natural harmonic system'), since 19–ET's intonational compromise with the just diatonic scale is only a little less approximate than that of 12-ET. In this context, 19–ET is just one among the many diverse features to be integrated under the umbrella of New Tonality: different from those enumerated above only in degree, not in kind.

Consequently, the near-to-just tuning of the intervals of a third – major and minor – in 19–ET were a lesser or perhaps even negligible factor. Reinforcement of this pragmatic attitude comes from Murray Barbour's classic text *Tuning and Temperament* (1951): 'To modern ears, accustomed to the sharp major thirds of equal temperament, the thirds of 379 cents, 1/3-comma flat, would sound insipid in the extreme'. Indeed, the literature seems to go further; it tends to suggest that most musicians, if tested objectively, tend to prefer, in both perception and performance, Pythagorean thirds to just ones, and tend to prefer equal temperament over both (see Sundberg 1982, Duke 1985, Rakowski 1990, Loosen 1993, Karrick 1998). But even if this is wrong, or if such answers stem merely from cultural factors or are a residue of habit, Hair decided that in order to develop microtonal skills, the best way would be to proceed not from ideas about just intonation, but from ideas about the scale, viz that it is a practical proposition from a singer's viewpoint to think of the 19–ET scale as a superset of the seven-note diatonic scale, despite the small adjustments to the intonation of that ultra-familar musical construct which will be required. More recently than Mandelbaum, Joseph Straus (1990) has demonstrated that even 12–ET music which is (avowedly) totally and continuously chromatic (such as middle-period Bártók and twelve-tone Schoenberg) often manifests residues of diatonic thinking, such that the chromatic twelve might nevertheless best be conceptualised in terms of a superset of the diatonic seven, and so Hair decided to apply this way of thinking about the chromatic collection to the practicalities of 19–ET performance, ie thinking of the nineteen tones in terms of adding twelve extra-diatonic tones to the the diatonic seven rather than just five, as in 12–ET.

In this context then, Pearson, is an honorary singer (using mainly aural, oral and embouchure capacities but with help from special, non-conventional fingerings). The intention is to develop the rehearsal exercises which we demonstrate here into some real music, viz a set of 19–ET duets for both a real singer (Morrison), and an honorary one, with – eventually – an accompaniment for 19–ET 'continuo' (viz, for the moment, pending the building of an 'acoustic' 19-ET chamber organ: a USB *MIDI* keyboard [2]).

This '7+12' view of 19–ET is essentially based merely on a 'hunch': that this train of thought is a practically viable proposition for singers. On the other hand, it must be admitted that Joseph Yasser (author of the most detailed and extensive study of 19–ET apart from Mandelbaum's) assumes a '12+7' model, and advances some sophisticated arguments from historical and music-theoretical perspectives in support of it. He makes an initial observation that the 7-note diatonic scale is best considered as an expansion from 5 principal tones, ie the pentatonic scale (eg C D F G A), by the addition of 2 auxiliary ones (eg E and B). This '5+2' model of the diatonic scale sees in it the basis of a concept of 'evolving tonality' in which an evolutionary process progressively absorbs the auxiliary tones into the collection of primary tones. In this view 5+2 evolves into 7+5 into 12+7. This is not the place to take issue with Yasser's view, however, despite the fact that our pragmatic approach does appear at first glance implicitly to question it on perceptual grounds.

2 Developing a New Technology

In a seminar presented at Glasgow University in 2003, Paul Davies, one of the foremost authorities on programming for musical applications, concluded that

- writing good applications is hard
- writing good graphical applications is harder
- writing good graphical audio applications is harder still
- writing good graphical audio & *MIDI* applications is an exercise in masochism
- many details to get right

2 *MIDI* (Musical Instrument Digital Interface) is an obsolescent *de facto* standard for the interconnection of musical controllers (primarily keyboards) and synthesizers. Since the adoption of the now-ubiquitous Universal Serial Bus (USB) on most computers, the hardware part of the MIDI standard is increasingly uncommon, the protocol for instrumental control being transported instead between USB-equipped devices. Hence the oxymoron USB-*MIDI* has been adopted.

* everything you ignore will eventually bite you
* fixing things offers many chances to release early and often[3]

The *Rosegarden Codicil* is certainly a demonstration of this, having been declared finished on at least three occasions! Not including the many lines of code on which an application of this complexity relies, the source code of the entire *Rosegarden* program runs to some 300,000 lines, making its construction hard by dint of sheer volume of work. Graphical applications are harder to write than old-fashioned command-line-based ones because the design extends beyond the required functionality to include ergonomic and aesthetic considerations, and because the components of a graphical application operate asynchronously, outside the influence of the application itself, as the user selects various check-boxes, menu items, buttons and so forth. Since audio applications intended for real-time use impose strict constraints on system latency, the asynchronous nature of the user interface demands have the potential to disrupt the timely nature of the delivery of audio samples. Adding the requirements of *MIDI* to the system considerably increases timing difficulties: *MIDI* data arrives in unpredictable bundles, all requiring very accurate timing, at a rate which can vary dramatically, whereas audio streams are continuous bandwidth demanding regular and timely servicing. The huge number of potential interactions between the *MIDI*, audio, and user-interface subsystems makes it extremely difficult to test such systems formally, so involvement of users at the earliest possible stages to detect potential problems is vital.

Davis also pointed to the different incentives involved in software development in different development models:

* commercial: emphasis on easily discovered, cheap solutions driven by marketing deadlines and margins. But: stable sales offer chances to refine and expand product capabilities.
* academic: emphasis on deep explorations of the problem space, generally insufficient focus given to refining user interfaces. But: lack of marketing considerations allows totally new ideas to be investigated.
* philanthropic: whatever you want!

The sheer size of *Rosegarden*'s source code base all but rules out the commercial model for application development; the sloccount (Lines of Code) program estimates the cost as exceeding US$5.5m as of November 2006. The open source development model means that all of this programming effort is available to the *Codicil*'s developers at essentially no cost; such open access to a major piece of software, even without the direct access this paper's authors have had to the *Rosegarden* development team, would be highly unlikely to be forthcoming from a commercial software supplier. Most such suppliers usually cite copyright security as reasons for not allowing developers access to their source code, although exactly how making available source code facilitates copying of programs isn't at all obvious.

Since this microtonal project was funded at a level considerably less than US$5m, it is clear that collaborators in the musical domain would have to learn to love the open-source development and execution environments

3 Centre for Music Technology, Glasgow University. http://markov.music.gla.ac.uk/Papers/PDTalk/conclusions.html

which make this work possible – in this case, the *KDE*[4] desktop running on the *GNU/Linux*[5] operating system. Without a flexibility amounting essentially to little more than fear of the unknown, there would have been little point in attempting to develop such sophisticated tools. Fortunately, the Engineering team found excellent collaborators in Pearson and Hair, who have proved to be quite comfortable with the technology.

Collaborative software development has implications for the software development team as well as the composer and performers involved, and the feature-set required differed somewhat from that incorporated in the initial implementation. The *Rosegarden* development team had made the design decision that there should be no alteration to the program which makes it difficult to perform the program's main function: the delivery of *MIDI* data to a synthesiser in a timely fashion. Unfortunately, *MIDI* is a very poor protocol for microtonal work: it does not even properly represent enharmonically equivalent notes in 12-ET, being essentially a numbering scheme for the keys of a piano keyboard. While there is a *MIDI* tuning standard which permits the notional piano to be retuned arbitrarily prior to synthesis, it is poorly supported especially by commercial synthesisers. The alternative method of retuning a note – use of the pitch bend facility – is unsuitable, as it operates only on an entire *MIDI* channel rather than a single note-event.

A first attempt at reconciling the microtonal requirements of this project with the limitations inherent in a program based on *MIDI* for its music representation involved using *Rosegarden*'s note-attribute datastructures, which store additional information internally on a per-note basis. We observed that the small second interval in 19-ET equates to 63 cents, and as such, is actually closer to a 12-ET quarter-tone (50 cents) than to a 12-ET semitone (100 cents). Using this methodology, the 19-ET hyperchromatic scale can be represented by a subset of a 24-ET scale notated using semi-sharp/sharp/triple-semi-sharp notation.

Because of the different harmonic structure of the various instruments which might be presented to the pitchtracker, a choice of several different algorithms is presented to the user when the program runs. It is a matter of experimentation for the performer as to which algorithm they find most reliable, and this can also be influenced by microphone positioning and so on. A further technical challenge is imposed by the unusual requirement that the application use only the most recent audio samples in generating its pitch estimate. For most audio recording programs, a higher priority is placed on not missing any audio samples than on reducing latency at all costs. However, if the pitchtracker's performance is to be even the equal of expert musicians, only the most recent samples acquired through the computer's audio interface can be considered. Code was written to persuade the computer's audio subsystem to behave differently from its normal strategy of avoiding data loss in favour of the most timely possible data delivery.

---

4 *KDE*: a portable desktop environment for small computer systems ('Kommon [sic] Desktop Environment'). The letter 'K' distinguishes this community-developed project from *Sun Microsystems*'™ CDE.

5 *GNU/Linux* is a community-developed operating system of great complexity and reliability, consisting of a Kernel ('Linux'), developed by Linus Torvald and others, and a large number of support programs developed by the Free Software Foundation (http://www.fsf.org). On this foundation has been built many very substantial applications, most of which are free. Free in this sense means both that the source code is available to developers and that the applications are available at zero monetary cost. Such applications include *KDE* (qv) and *Rosegarden*. GNU is modelled on the mature *Unix*™ operating system's toolkit. It is an acronym for 'GNU's Not Unix'.

This technically elegant solution turns out to have very little musical utility. The problem arises from the assumption that a musical score is a pitch-time graph, which it clearly is not. Especially where the music is diatonic, a system which makes no sense enharmonically greatly hinders the use of the program as a real-time rehearsal tool. It was therefore necessary to devise a dual-layer system for mapping musical notation to expected frequency. This permits notes which are normally (in 12-ET) regarded as enharmonically equivalent to be mapped to different expected frequencies in 19-ET notation.. A text file is supplied which maps the notational spelling to the expected frequency, and it is this system which we are currently distributing.

As well as the availability of the *Rosegarden* sequencer and its source code as an excellently written and stable starting point, the successful conclusion of this project depended on many of the aspects of the 'extreme programming' approach adopted by the software development team. Particularly:

* flexibility must be exercised over the application's specification. The specification is often unclear either to the musicians or engineers at the outset, so the development cycles where programs are seen through to release before user-testing commences are inappropriate.
* even the most seemingly obvious preconceptions as to the semantics of the data being handled must be regarded with suspicion.
* expert users must be on hand almost continuously to prevent waste of valuable development effort. In short, interdisciplinary effort must involve the participating disciplines!

An important, yet much overlooked ingredient in the successful application of technology in performance and composition comes from a full understanding of the purpose to which the new technology is to be put. In practice, however, because it is also important to respond to artists' requirements on a sufficiently short timescale, it is usually best to adopt an incremental approach to software development: we are, so to speak, 'building the camera while shooting the film' (Bailey 2001 and Berend 1997). Imaginative repurposing of existing technology is impossible without full access to source code, and while this is apparently obvious, it passes by many researchers who continue to deploy closed, proprietary solutions, because "everybody uses it", essentially short-circuiting the creative process in the technological arena. Fortunately for the provision of a microtonal rehearsal tool, an excellent candidate technology in the form of the *Rosegarden* sequencer[6] was available in open-source.

Rosegarden is a sophisticated sequencer with high quality music typesetting available as a bonus. It is written for the *GNU/Linux* operating system and is therefore portable across a large number of hardware platforms. It comprises some quarter of a million lines of source code. The main developers are active and, more importantly, available, which is rarely, if ever, the case with proprietary software. The possibility of communicating with the developer directly is particularly important when changes are being considered to such an extensive and complex application.

The next question to consider is the notation of 19–ET music. A basic principle here is that one should not move too far away from practices to which performers are already accustomed. This might suggest that one should adopt the notation universally used for 12–ET music and simply add seven new symbols. Consider as a model some recent Peters Edition scores of the kind sometimes described as 'The New Complexity' (involving composers such as Brian Ferneyhough and James Dillon), which are often based on the quarter-tone scale.

6 *Rosegarden* sequencer. http://www.rosegardenmusic.com

Four new notatation symbols are introduced to cope with such music: symbols for half-sharp, one-and-a-half sharp, half-flat and one-and-a-half-flat. Perhaps one could adapt this for 19–ET? Might one conceiveably simply map the pitches of the 19–ET onto the nearest quarter-tone? In such a circumstance, the interval of one nineteenth of an octave (63 cents) would be represented, for the most part, by the interval of one twenty-fourth of an octave (50 cents), an approximation which is not too bad – more accurate than (for example) the representation of a just major third in 12–ET. Perhaps the amount of approximation might not be readily noticed? Such a system of notation would be simple to implement as an extension of a sequencer program, especially one which could already cope with 'half-sharp' notation. The result is given is Example 1.

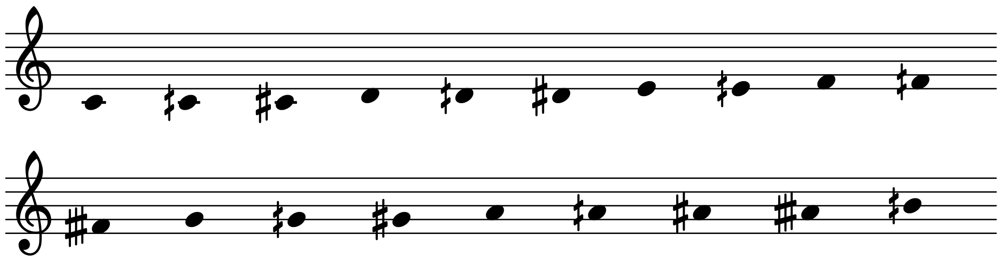

Figure 1: A 'Hyper-chromatic' with 19 Degrees and Tone-Centre on c′, using 'nearest' notation

This system has an obvious short-coming – namely that, whereas 14 of the scale's 19 intervals are of 50 cents, five (scale-steps 3-4, 6-7, 10-11, 14-15 and 18-19) will need to be 100 cents – but that was not the only thing which made us decide against this solution. For the psychological reasons mentioned earlier, we decided to approach 19–ET as a superset of the diatonic seven, rather than of the chromatic twelve, and this called for an alternative approach to notation as well.

In our current version, the 19–ET system is instead notated using only traditional flat and sharp signs, with the division of the scale into extra degrees being accomplished by assigning different pitch values to what would normally be enharmonic equivalents in 12–ET. In addition to practicality, such notation also has historical precedents. For example, it is compatible with the notation used in highly chromatic music of the 17th century, for which, it is thought, one of the various versions of extended meantone tuning could have been used, and for which, in certain areas, notably Flanders and Italy, split-key keyboards with more than twelve keys per octave were sometimes built (Stembridge 1992, 1993 and 1994). Mandelbaum also mentions one or two more modern keyboards which were designed to play in 19–ET, beginning with Peter Samuel Munck's harmonium, built c1850 and now in the Stockholm *Musikmuseet.* The notation, using this system, of a 19–ET 'hyper-chromatic' scale with nineteen degrees, based on c′, is shown in Figure 2.

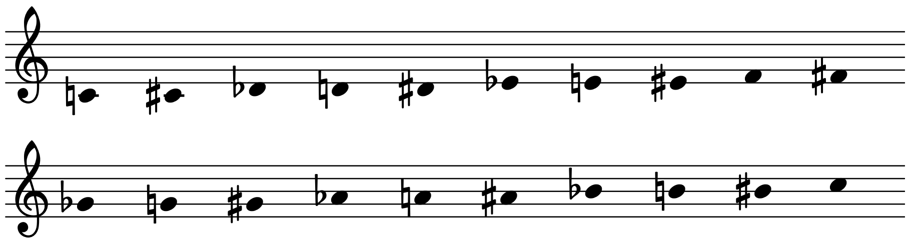

Figure 2: A ‘Hyper-chromatic’ Scale with 19 Degrees and Tone-Centre on c′, using the adopted notation.

The lesser tones between E and F and B and C are divided only once. Note the term ‘lesser tone’, which, in this context, is clearly more appropriate nomenclature than ‘semitone’, since its size is two-thirds (not half) of a (‘greater’) 19–ET tone (126 cents, compared to 189 cents). Note also that the enharmonic equivalences are not the same as in 12–ET. For example: C♯ is equivalent to D♭♭, and D♭ is equivalent to C𝄪. E♯ is the note *between* E and F, and is equivalent to F♭.

One possible shortcoming of this second alternative – that the interval between a note and its sharped or flatted form, which performers are accustomed to conceptualise as 100 cents, is actually nearer to half that (63 cents) – suggests that the use of half-sharps and half-flats might be more precise. This observation prompts consideration of a third alternative system, which encapsulates the best qualities of both of the others (see example 3).

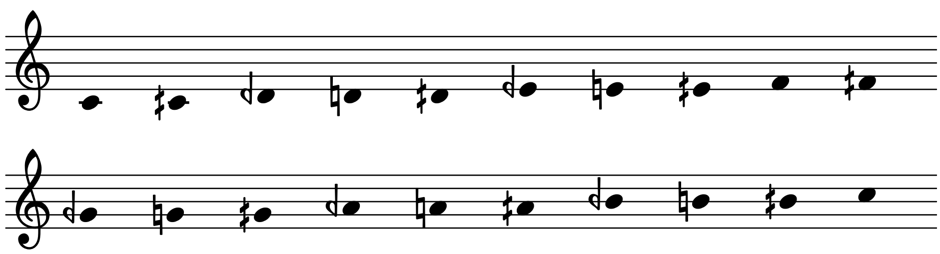

Figure 3: The same ‘Hyper-chromatic’ Scale, using ‘half-flat/half-sharp’ notation.

But in the end, we decided to use the second of our three alternatives, essentially because it involves no unfamiliar notation symbols at all.

A few of the many properties of the 19–ET system which may be exploited by the composer are:

* Nineteen tones to the octave offers the possibility of modulation to eighteen tonal centres other than the tonic. As mentioned earlier, Hair’s previous pieces, composed in 12–ET, deal with tonal centers, diatonicism, consonance, and textural continuity, alongside acentricity, chromaticism, dissonance, and textural discontinuity. 19–ET offers extended use of all these features, thus embodying the second of Gerdine’s

Ash

Text by Yunus Emre (Turkish, 14th century), trans Süha Faiz

Graham Hair: *Microtonal Songs: Ash*

Figure 4 (beginning): Score of *Ash* by Graham Hair (bars 1–8) with 'Hot Notes' for the soprano and clarinettist shown as additional staves

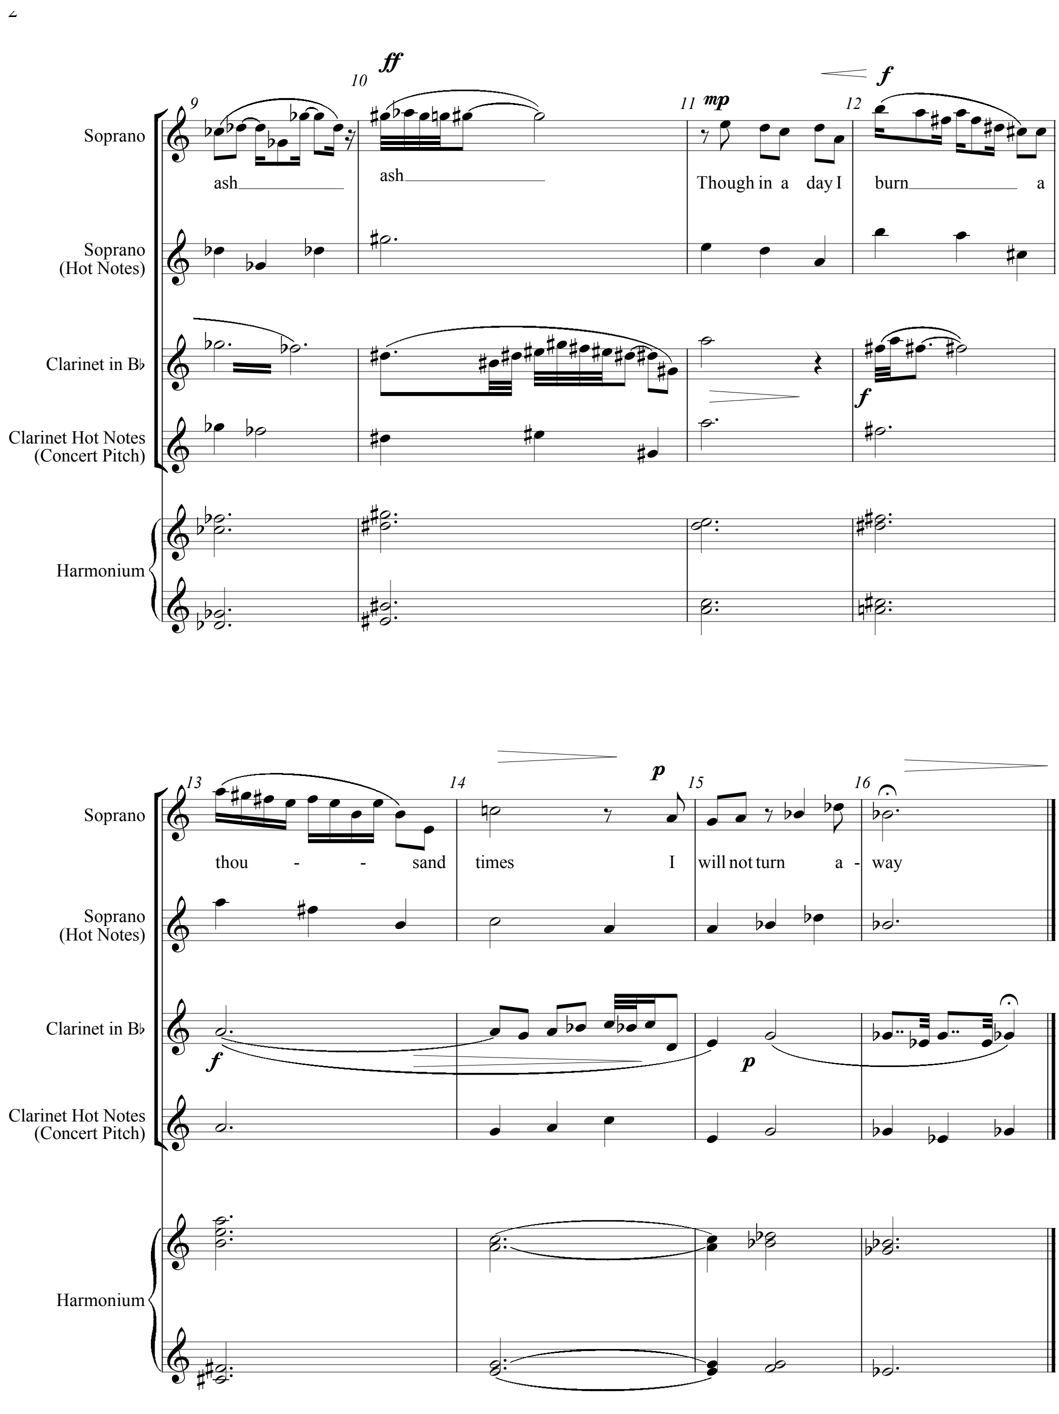

Graham Hair: *Microtonal Songs: Ash*

Figure 4 (conclusion): Score of *Ash* by Graham Hair (bars 9–16) with 'Hot Notes' for the soprano and clarinettist shown as additional staves 'properties'.

* Because nineteen is a prime number, that most familiar of features of 12–ET, the cycle of fifths, is parallelled, in 19–ET, by cycles of all of the other intervals in the tuning system, with considerable consequences for musical structure (see again Gerdine's second 'property').
* As well as simply offering additional modulation possibilities, quite a number of novel melodic progressions are made available. For example, the divisions of the scale permit melodic motion through an interval of a perfect fourth in two equal steps, or through a major third in three equal steps.
* Notation and performance are both connected to existing practice: there are no additional symbols to learn, and the additional tones are contextualised around a version of the familiar diatonic scale.

Three rehearsal devices based on *Rosegarden* were implemented for this project in the first instance:

* the capacity to play 19–ET *singalong* recordings from a computer;
* a 19–ET *scordatura* keyboard which enables the singer to check her intonation of a melody drawn from the 19–ET scale on a USB *MIDI* keyboard designed to play 12–ET melodies; and
* a pitch-tracker (the *Rosegarden Codicil*), which 'listens' to what the singer is singing, and provides feedback on how she's doing, either in real time or as a *post hoc* comparison of a recorded audio track with the score.

Our term 'scordatura keyboard', is a neologism adopted (and extended) here from its more traditional usage, in which (according to *Wikipedia*[7]) a *scordatura* (literally Italian for 'mistuning') is an alternative tuning used for the open strings of a string instrument: an extended technique used to allow the playing of otherwise impossible melodies, harmonies, figures, chords, or other note combinations.

These three devices provide progressively less spoon-feeding for the singer. The *singalong* recordings provide support for every note, the keyboard a facility to check intonation when desired, while the pitch-tracker provides just a few points of reference, ie a few strategic notes which are placed amongst unfamiliar pitch configurations, to which they act as a guide. The concept of *hot* and *cold* notes – which, loosely speaking, permits additional metadata to be represented in a performance system concerning whether a note is melodically primary or subordinate – provides the means by which this is done: *hot* notes are the important points of reference for the performer in rehearsal, while *cold* ones add melodic embellishments grouped around the *hot* ones. The score of the piece presented in Figure 4 includes the *hot* notes; as can be observed, the actual vocal and clarinet melodies elaborate these with a considerable amount of ornamentation.

Despite the fact that we started out on this project believing that these three devices had common sense on their side, we must nevertheless confess that, as a result of workshopping *Ash* with their help, we changed our approach in one or two respects, as described in section 4, below.

3 Modifying Rosegarden

3.1 Incorporating a Pitch Tracker

Like most *MIDI* sequencers, *Rosegarden*'s internal representation is highly dependent on the *MIDI* standard.

7 *Wikipedia* home page. http://www.wikipedia.com

MIDI pitch is encoded as an integer in the range 0–127 representing performance pitch, with no information regarding the enharmonic spelling of a note. To support notation within *Rosegarden*, *MIDI* pitch is augmented by an optional accidental, which allows the user to specify how a note should be notated on a conventional staff.

Real-time performance is a *sine qua non* for a sequencer program, so one of the basic design principles set out by the *Rosegarden* development team is that information should be represented in a way that makes the music near trivial for the sequencer to play. Consequently, any changes had to be made within the existing representation, especially if they were to be adopted by the development team into the mainstream release of the *Rosegarden* package.

Within this framework, it was indeed possible to extend *Rosegarden* to accommodate microtonal accidentals, and modify the graphical classes to display these notes accordingly on the staff. During playback, the *MIDI* pitch and accidental of the current note are translated into an enharmonic spelling. The tuning system indicates which pitch this spelling represents and calculates the target frequency. The concept of a tuning system has been added to facilitate the translation of the standardised enharmonic representation into the target frequency. Commands to change tuning systems can be embedded within the score, so that the flexibility of microtonal representation is sufficient to accommodate any future extension.

3.2 Data Capture

The process by which the performer maintains pitch is analagous to a closed control loop. The musician adjusts the performed frequency according to instantaneous feedback. In a normal situation the musician's auditory system provides this feedback with perceptually no latency. In the rehearsal situation, the musician is relying on the feedback provided by the pitch tracker since she is not trained to recognise 19–ET intervals. Unlike a recording situation where it is essential to capture every sample of data, this is a hard realtime application dependent on the currency of the data it processes. A data capture thread copies samples into a ringbuffer which holds sufficient to provide a complete frame for the pitch tracker. Each new sample is written over the oldest sample, so when the pitch tracker requires data, the capture process can immediately provide a window stretching backwards from the current instant. To achieve this, the pitch tracker uses the *Jack* Audio Connection Kit (*JACK*)[8]. This is a high-performance, low latency audio server for Linux/Macintosh OS X which is also used for audio routing by *Rosegarden*.

3.3 Pitch Tracking Algorithms

The pitch-tracking class uses a choice of algorithms to provide a high-resolution estimate of the fundamental frequency of harmonic complex tones. It comprises two stages: a low-resolution search for the fundamental frequency followed by a high resolution localisation of the fundamental component.

Alternative methods have been implemented in the first stage. The first, autocorrelation, is one of the most widely used methods of fundamental frequency estimation. It provides a measure of the periodicity of the signal by comparing the signal with a time-shifted copy of itself. It is capable of detecting the fundamental frequency even if that particular frequency component is missing from the signal. The alternative Harmonic

8 *Jack* audio connection kit. http://www.jackit.sourceforge.net

Product Spectrum or HPS (Schroeder 1968) method searches for a harmonic pattern within the spectrum of the signal. When the frequency axis of the power spectrum is downsampled by an integer factor *n*, the *n*th harmonic aligns with the position of the fundamental. The product of these spectra provides the HPS. Repeated multiplication has the effect of exaggerating the fundamental while at the same time attenuating other partials.

A common error made by pitch trackers (Pardo and Birmingham 2002, Bolton 2005) is to misjudge the fundamental frequency by an octave. In the case where both algorithms are making octave errors, an option exists to make the comparison between target and performed frequency independent of octave differences.

Once an estimation for the fundamental frequency has been obtained through one of the above methods, a high-resolution estimate is obtained using phase vocoding (Dolson 1986) coupled with a phase-unwrapping technique. This uses the phase difference between the same bin in two successive FFT frames to detect the frequency of the partial contained in the bin with greater accuracy, based on the assumption that no more than one partial occupies each analysis bin using Portnoff's algorithm (Portnoff 1976).

Because all of these methods are obtained using the Fourier transform, the data from the first stage of processing can be reused in the second.

3.4 Tuning System

The tuning system must define the intervals in the scale, and how these relate to its graphical notation. Scala[9], an existing tool for experimenting with microtonal tunings, has a file format for defining the intervals of a scale. Each interval is defined by comparison to the tonic either in cents, or as a frequency ratio. A scale is defined by a list of such values in a text file accompanied by some metadata. Although Scala is not open-source software and therefore of very limited use to us, to encourage compatibility, our representation is based on the Scala format. Intervals are defined in the same way, but to service the notational requirements of this application each interval must be followed by a list of all equivalent enharmonic spellings. Variations of a tuning in different keys can be grouped within one system, providing the facility to support tuning modulations in real time.

Internally, each tuning contains a standard C++ hash table which maps a text string representation of each spelling variation to the interval it represents stored in cents. The target frequency is obtained by converting the *Rosegarden* pitch into a string representation of the notated pitch and searching the tuning map for the interval (see Figure 5).

9 *Scala* home page. http://www.xstall.nl/huygensf/scala

```
System
12 tone

Tuning Equal-Tempered
#             cents   spelling
0.0,          A,
100.0,        A2,     B-2
2/12,         B,
3/12,         C,
400.0,        C2,     D-2

endTuning
endSystem
```

Figure 5: An Excerpt from a Tuning File

3.5 Designing the Performance View

As the *Codicil* is intended for real-time use, the user interface required that the score and feedback be displayed simultaneously. Also, the feedback should be clear and obvious so that it can be analysed while reading the score. To achieve this, the pitch view is positioned in a resizeable panel beneath the score view. The feedback is displayed as a moving graph plotting the deviation from the target frequency in cents. Markers traverse the graph when a note on/off occurs to assist in following the score and to mark the transient effects which occur at these points (see Figure 6 below). By scrolling backwards and forwards through the score, it is possible to review pitch accuracy after the practice session.

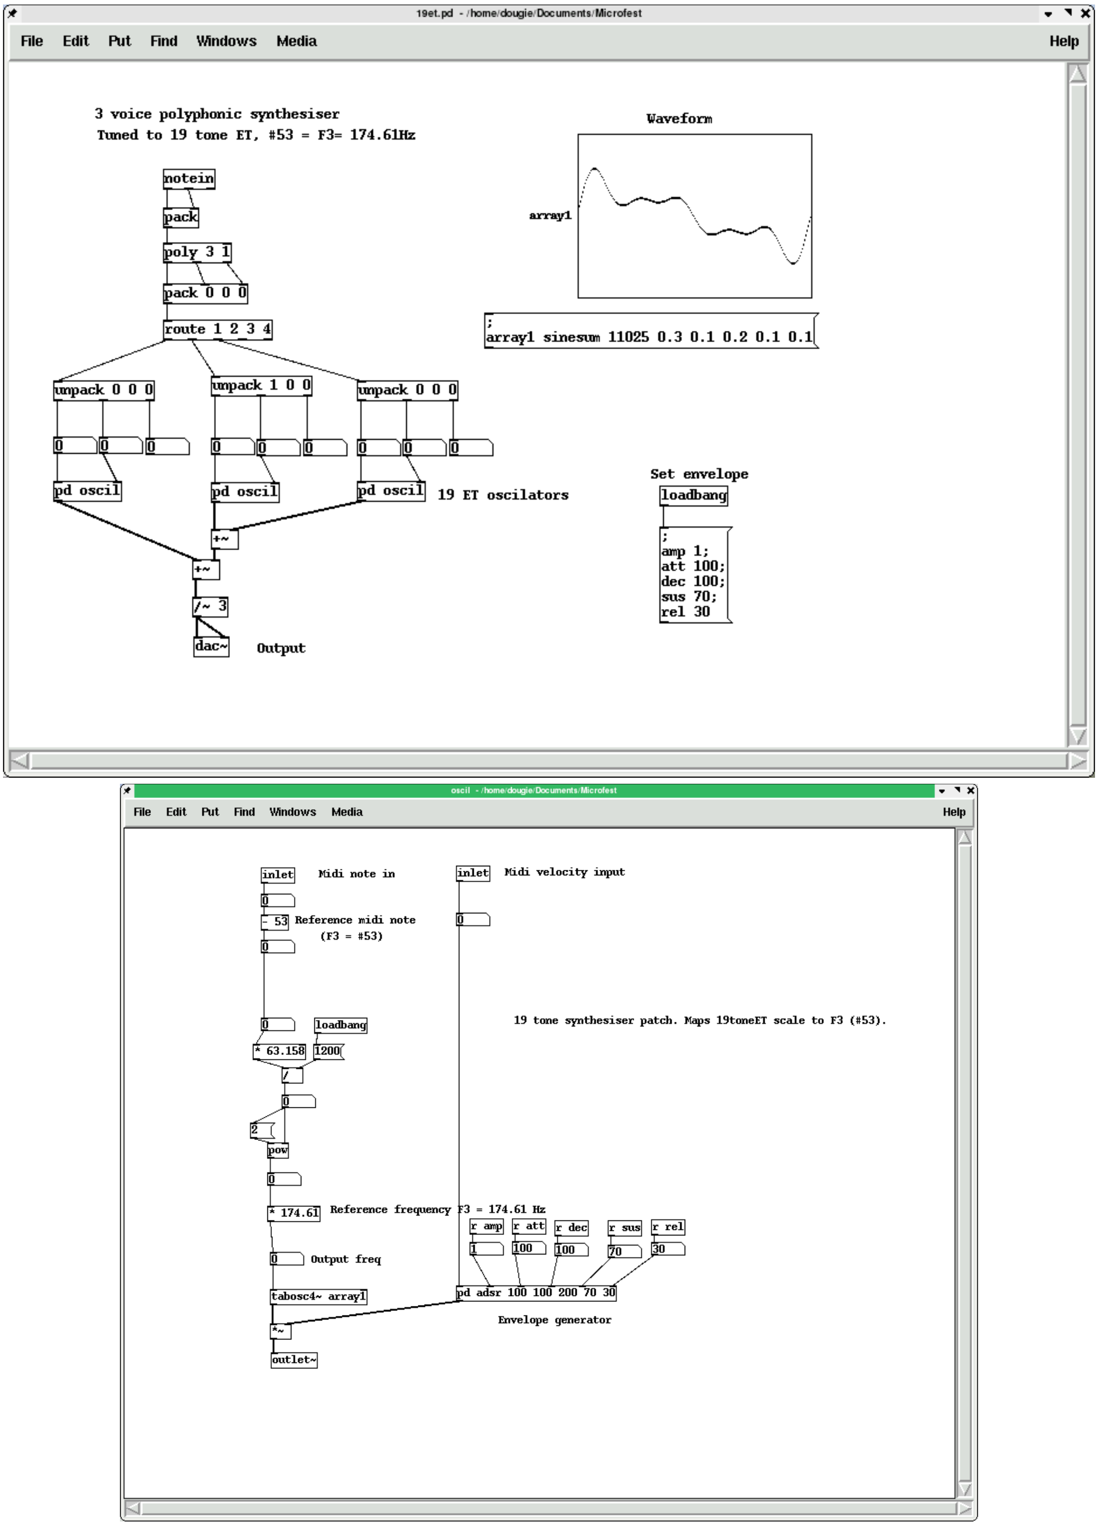

Figure 7: PD patch for a 19–ET Polyphonic MIDI Instrument

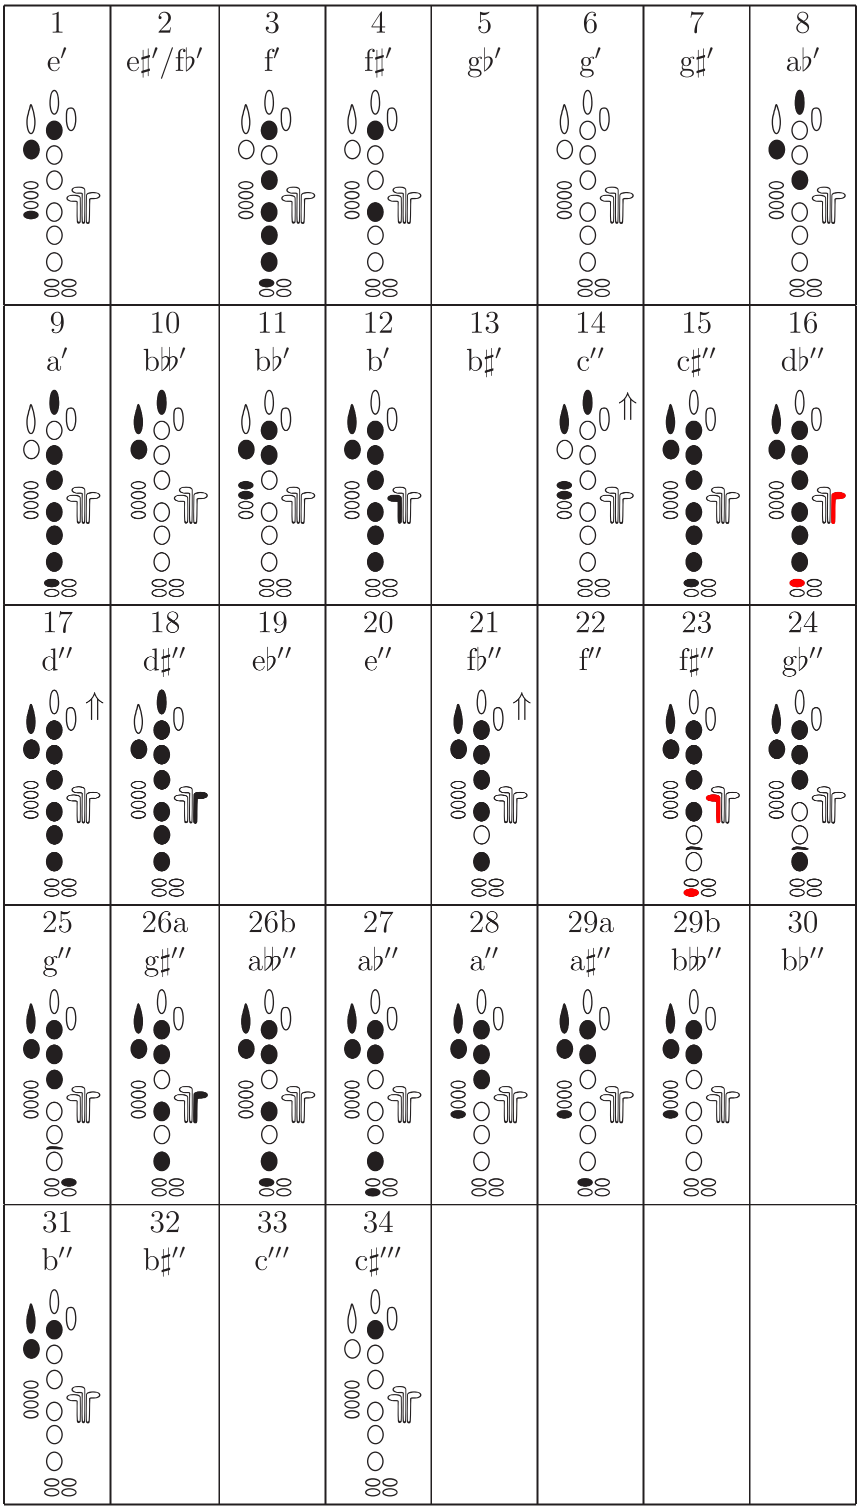

Figure 8: Fingerings for the clarinet gamut

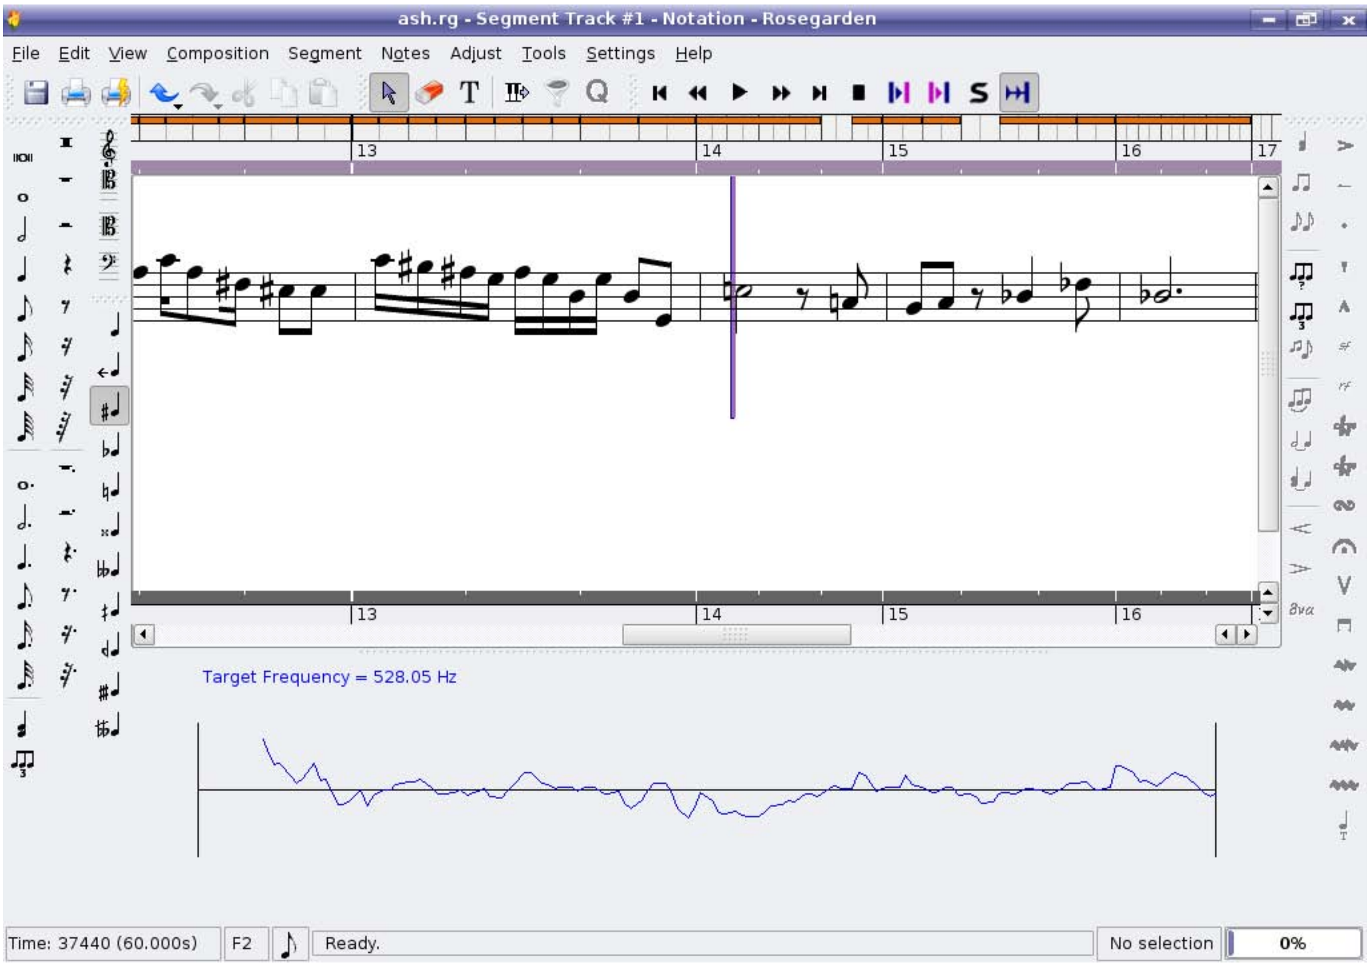

Figure 6: The Rosegarden Real-Time Pitch-Tracker Window

4 Using the Codicil

We conclude with a discussion of some of the issues involved in practising the soprano and clarinet melodies of a short extract of Hair's composition *Ash* with the help of the *Rosegarden Codicil*. This includes some remarks on the keyboard accompaniment, on the pitching problems from the singer's point of view and on the fingering and embouchure problems raised by the clarinet part. The score of *Ash* appears as Figure 4, above, and an audiofile (Hair 2007) of a performance thereof by Morrison, Pearson and Hair is appended.

The 19–ET keyboard part is realised by re-mapping the *MIDI* pitch values from a normal controller. Performers play from *scordatura* notation, so that what they see on the page indicates where to place their fingers rather than the sounds which would normally emerge if the keyboard retained its normal 12–ET tuning. Performing on such an 'ill-tempered' keyboard might be thought to be somewhat confusing, but clearly better than expecting players to work out where the 19–ET note lies on the modified 12–ET instrument from notation which indicates the desired sound. For performance purposes, a patch for a 19–ET polyphonic instrument has been constructed using Miller Puckette's *Pure Data* program or PD[10]. This patch is shown in Figure 7, below. Additionally, it is possible to use the software synthesizer *Fluidsynth*[11] to realise the 19-ET *scordatura* directly by loading a tuning file and an appropriate sound-font. Although this may provide more of a generic solution with an extended range of possible timbres, it also relies on the precise intonational accuracy of the (12–ET) sound fonts, and moreover, since some fonts do not bend well – resulting in timbral inconsistency – care must be exercised using this approach.

Although all of Pearson's prior performance activities have been with the 12–ET system, or with the intonation systems of early music, she brings considerable experience of new music as well as the standard repertoire to problems of microtonal intonation where more than 12 notes per octave are involved. A concentration on historically-informed performance using period instruments has led to familiarity with a wider range of temperaments. None the less, there is still a considerable gap between theory and practice, particularly in orchestral repertoire, where the effect of a particular temperament on the wind section is often somewhat marginal.

Players of wind instruments are continually reminded of the shortcomings of their instruments as regards playing with accurate intonation, whether in 12– or 19–ET. Whilst reed strength, air pressure as well as room temperature and humidity can alter the tuning, most clarinettists mitigate against the instrument's inherent intonation faults *in situ* by adjusting their embouchure. This adjustment is made by drawing upon the player's existing aural perception of the pitch of the note which is then tempered according to the circumstances.

The most significant challenges presented to the clarinettist by the 19–ET scale are the need to establish aural points of reference for each of the individual pitches, and then to devise a chart of reliable fingerings. In order to get these exactly right, we provided Pearson with a tuning file: a 19–ET scale running the whole gamut of notes used in the composition, from (concert pitch) `d'` up to `b''`, an octave and a major sixth higher. Compare the fingering chart (Example 8) with the scalar chart (Example 9). Example 8 gives the fingerings for all the pitches used in *Ash* in ascending scalar order; note that the letter-names used indicate the pitches as transposed for clarinet in B♭. Example 9 gives all the pitches used in *Ash* as they are notated by the composer, at concert pitch, also in scalar order. Of the several enharmonically-equivalent (ie 19–ET enharmonically-equivalent!) pitch names available, Example 9 conforms to those which were actually chosen by the composer. Ten pitches in the gamut are omitted, and two occur in different enharmonically-equivalent (19–ET enharmonically-equivalent!) nomenclature. In workshopping the piece, however, Pearson found it generally useful to have at hand a chart with the fingerings written out in the order in which they occur in the music.

10 *Pure Data* home page. http://www.crca.ucsd.edu/~msp/software.html
11 *Fluidsynth* home page. http://www.fluidsynth.com

Ash

Gamut of clarinet notes used in *Ash*

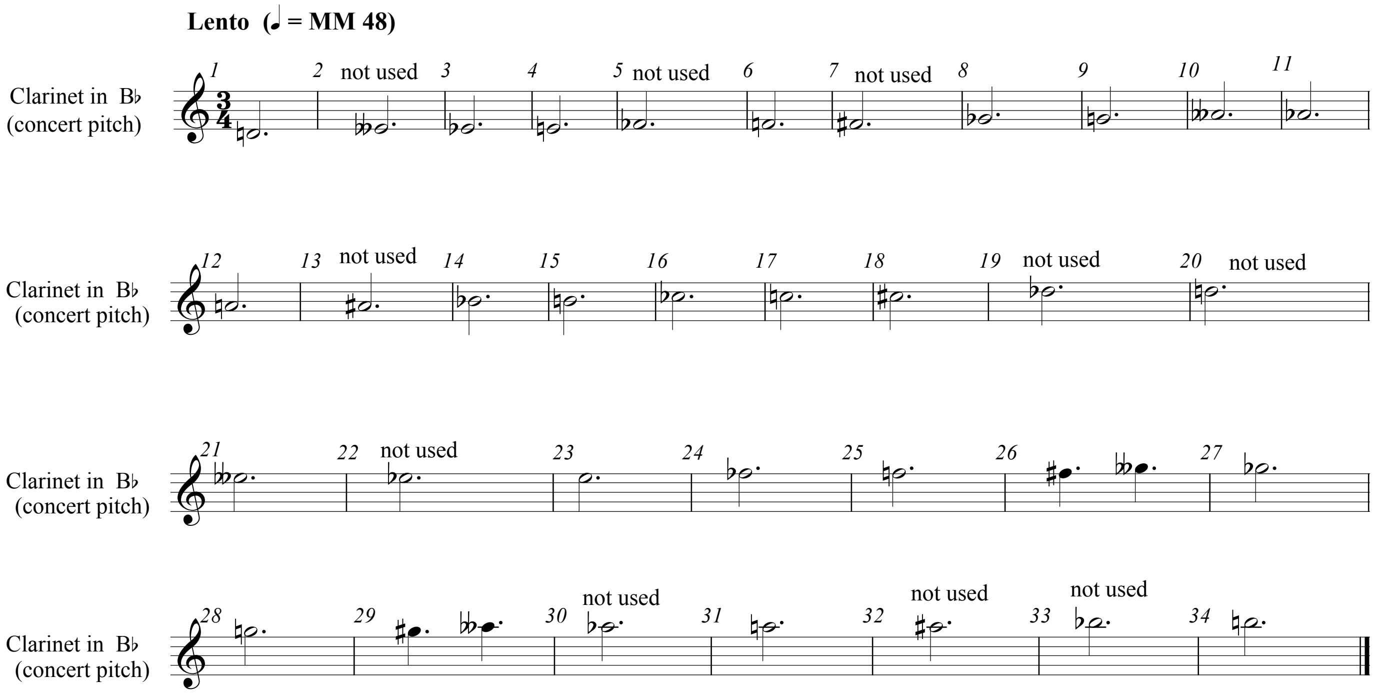

Example 9: The clarinet gamut

The *Rosegarden Codicil* software makes it possible to hear the actual composition in 19–ET, as well as to perform along with it. In the process of learning *Ash,* Pearson found it useful to play in unison with the computer and to check intonation against the keyboard, as well as to set up her ‘fingering gamut’ in advance. The visual representation presented on the computer screen provided a further valuable means of correcting intonation during rehearsal, in preparation for the performance with the intended microtonal harmony.

The clarinet used for the current project is a 3/4 Boehm-system instrument, a design which is played almost universally, except in Germany and Austria. Experience with eighteenth- and nineteenth-century clarinets, both originals and reproductions, suggests that for each new instrument a player should devise his/her own fingering chart. Therefore, the fingerings shown in this paper represent only one player’s solution to realising the musical example. None the less, the Boehm clarinet’s seventeen keys and seven finger holes are useful in creating microtonally-compatible fingerings. For example, 12–ET fingerings can be modified by shading open finger holes, with the right hand, as shown by the fingerings for e′, f♯′ and a′. Similarly, depending on the desired pitch, keywork can be opened or closed to facilitate small adjustments, of which the fingerings for a and c’ are an example. Where pitch requirements could not be completely satisfied by designated fingerings, it was necessary to use embouchure pressure further to modify the tuning. The fingering chart shows instances of this feature for the notes a′, b′, d′′, e′′ and f♯′′. The performer’s ability to adjust intonation by such embouchure pressure is relatively constant across most of the dynamic range, and is more limited only at extreme dynamics especially *fortissimo*.

As we experimented with the *Codicil* at a workshop at the Royal College of Music in February 2007, both Pearson and Morrison found that the file of *hot* notes was more useful *after* the melodies of *Ash* had been learned, rather than as a learning tool: in other words, as a means of checking their intonation *post facto*. In fact, the file of *hot* notes was somewhat distracting in the early learning stages. More useful in the process of actually learning the melodies, they felt, would have been a file with the pitches of the melody elongated into a series of long notes in *cantus firmus* style: notes which are long enough to be listened to carefully in order to check the intonation of each. We surmised that this might have been because it is important during the learning process to practice moving between pitches of the melodies in their sequential order. Both performers said that getting the feel of each pitch and the feel of their succession – in the vocal tract, in the head, in the embouchure – was an important part of getting the intonation right. So the *cantus firmus* soundfile has now been incorporated into the list of rehearsal resources.

In the discussion following the workshop, Morrison raised two other issues, which may or may not involve challenges particular to her. One is that she has absolute pitch. In working with such bodies as the vocal ensemble *Exaudi* on works by composers including Xenakis, Ferneyhough or Sciarrino (as she had done immediately before our workshop sessions at the RCM), this may be especially useful. But it seems that it may be absolute in relation to 12–ET! The whole issue of categorical perception in relation to 19–ET is thereby raised. Can singers become intonationally ‘bilingual’, thinking in 19 categories or in 12 at will? Secondly, Morrison is normally accustomed to learning works quickly from score, unlike, for example, many opera singers, who may learn their parts essentially by ear, ie working with the help of a *repetiteur* at a much slower learning pace. Thus, for example, seeing an F♭ in the score, may produce an immediate sonic image of the note in the ‘mind’s ear’ which then has to be corrected, because, in 19–ET, an F♭ is only 63 cents lower than an F♮, and may feel more like some variety of F♮ than like a *real* F♭. In fact, in 19–ET, only the A♮ in each octave is exactly the same frequency in 19–ET and 12–ET; moreover, in *scordatura* notation only a′ in 19–ET is identical to a′ in 12–ET. Morrison therefore wondered if learning the score by ear may not be more effective than working from score in the way that performers of contemporary music normally do, and as our group was doing in our RCM workshop.

To sum up, we assert that the combination of these rehearsal resources, based on the capabilities of what we have called the *Rosegarden Codicil*, viz the *singalong* soundfile, the *scordatura* keyboard and the pitchtracker for the tuning file, the *elongated* file stretching the melody out in *cantus firmus* style, and the file of *hot* notes, forms a combined resource with which the singer or the clarinettist can build up confidence in establishing correct intonation for the performance of 19–ET scores. A final additional resource must be mentioned: the keyboard accompaniment. In the scores which were developed for this particular project, the accompaniment was designed to support the voice and clarinet in establishing correct intonation, by embedding some of the principal tones of the melody into the harmony, even though that harmony is not, by and large, triadic let alone functional.

Following the workshop at the RCM, we made up a file with the pitches of the melody elongated into a series of long notes in *cantus firmus* style, exactly as described in the previous paragraph. In April 2007, in the Laboratory of the Glasgow University Centre for Music Technology, we recorded Morrison singing the pitches of the melody of *Ash* in this *cantus firmus* form. First she sang it assisted by the keyboard accompaniment and the pitch-tracker, then assisted by the pitch-tracker alone, and then finally without assistance either from accompaniment or pitch-tracker. At the time of going to press, we are in the process of measuring these performances. After experimenting with a Vector Quantisation method, we are using a Shifting-Window Piecewise Linear Segmentation method. An excellent review of the various methods and applications of Vector Quantisation is given in Gray 1984, and for the particular method we tried see Linde, Buzo and Gray, 1980. For Shifting-Window Piecewise Linear Segmentation methods, see Keogh, Chu, Hart and Pazzani, 2001.

This measuring process foreshadows one of the directions in which we intend to develop this research beyond the 'pilot project' stage described in this article, namely, to measure objectively performances of microtonal music by a number of different singers, of varying levels of experience.

Acknowledgements

The authors gratefully thank the *Rosegarden* development team for their invaluable assistance in adding microtonal functionality to their program, and M. S. Cuthbert for permission to use his *ClarFinger* font in the clarinet fingering table.

Bibliography

Bailey, 2001: N.J. Bailey, 'Why Free Software Happens: A Survey of Motivation in the Open Software Community', in *Circus 2001: New Synergies in Digital Creativity Conference for Content Integrated Research in Creative User Systems*, ed C. Boehm, D. Garcia and H. Schuter (Glasgow: University of Glasgow), pp. 35–44.

Barbour, 1951: J Murray Barbour, *Tuning and Temperament* (East Lansing, MI: Michigan State College Press).

Berend, 1997: A. Berend, Personal Communication, National Film and Television School (April).

Blackwood, 1985: E. Blackwood, *The Structure of Recognisable Diatonic Tunings* (Princeton, New Jersey: Princeton University Press).

Blackwood, 1991: E. Blackwood, 'Modes and Chord Progressions in Equal Tunings', Perspectives of New Music, 29/2 (Summer), pp. 166 – 200.

Bolton, 2005: Jered Bolton, *Gestural Extraction for Musical Audio Signals*. Unpublished PhD dissertation, Glasgow University (CMT.gla.ac.uk/Publications/Theses.html).

Bosanquet, 1875: R. H. M. Bosanquet, *An Elementary Treatise on Musical Intervals and Temperament* (London: Macmillan).

Bowen, 1968: M. Bowen, ‘Harry Partch’, *Music and Musicians* 16, pp. 20 – 25.

Dolson, 1986: M. Dolson, ‘The Phase Vocoder: A Tutorial’, *Computer Music Journal* 10/4, pp. 14–27.

Duke, 1985: R. A. Duke, ‘Wind Instrumentalists’ Intonational Performance of Selected Musical Intervals’, *Journal of Research in Music Education*, 33/2, pp. 101-112.

Emre, 1992: Y. Emre, *The City of the Heart: Yunus Emre’s Sufi Verses of Wisdom and Love*, ed S. Faiz (Nicosia: Near East University Press).

Gray, 1984: R. M. Gray, ‘Vector Quantization’, *IEEE ASSP Magazine* (April), pp. 4–29.

Karrick, 1998: B. Karrick, ‘An Examination of the Intonation Tendencies of Wind Instrumentalists Based on their Performances of Selected Harmonic Intervals’, *Journal of Research in Music Education*, 46/1, pp. 112-127.

Keberle, 1999: D. S. Keberle, *Fast and Easy Microtonal Fingerings for the Clarinet and ‘Four to Go’ for flute, clarinet, violin, cello, piano and percussion*, Unpublished PhD Dissertation, University of Pittsburgh.

Keogh, Chu, Hart & Pazzani, 2001: E. Keogh, S. Chu, D. Hart & M. Pazzani, ‘An Online Algorithm for Segmenting Time Series’, *Proceedings of IEEE International Conference on Data Mining*, pp. 289–296.

Kopiez, 2003: ‘Intonation of Harmonic Antervals: Adaptability of Expert Musicians to Equal Temperament and Just Intonation’, *Music Perception*, 20/4, 383-410.

Linde, Buzo & Gray, 1980: Y. Linde, B. Buzo & R. M. Gray, ‘An Algorithm for Vector Quantizer Design’, *IEEE Transactions on Communication* (January), 28/1, pp. 84–95.

Loosen, 1993: F. Loosen, ‘Intonation of Solo Violin Performance with Reference to Equally Tempered, Pythagorean, and Just Intonation’, *Journal of the Acoustical Society of America*, 93/1, pp. 525-539.

McLaren, 2006: B. Mclaren, ‘Microtonality in the U S’ (http://daschour.club.fr/micromegas/mclaren.html, accessed 08/12/2006).

Mandelbaum, 1961: J. Mandelbaum, *Multiple Division of the Octave and the Tonal Resources of 19-Tone Temperament*. Unpublished PhD Dissertation, University of Indiana. (Ann Arbor, MI: University Microfilms).

Pardo & Birmingham, 2002: B. Pardo & W. Birmingham, ‘Improved Score Following for Acoustic Performances’ in *Proceedings of the International Computer Music Conference, Gothenburg, Sweden, September 16–20, 2002* (Los Angeles: ICMA).

Rakowski, 1990: A. Rakowski, ‘Intonation variants of musical intervals in isolation and in musical contexts’, *Psychology of Music*, 18/1, pp. 60-72.

Rehfeldt, 1994: P. Rehfeldt, *New Directions for Clarinet*, revised edition (Berkeley, Ca: UCP).

Portnoff, 1976: M. Portnoff, 'Implementation of the Digital Phase Vocoder Using the Fourier Transform', *IEEE Transactions on Acoustics, Speech and Signal Processing*, 24/3 (June), pp. 243–248.

Schroeder, 1968: M.R. Schroeder, 'Period Histogram and Product Spectrum: New Methods for Fundamental Frequency Measurement', *The Journal of the Acoustical Society of America*, 43/4, pp. 835–838.

Stembridge, 1992: C. Stembridge, 'Music for Cimbalo Cromatico and Other Split-Keyed Instruments in Seventeenth-Century Italy', *Performance Practice Review* 5, pp. 5–43.

Stembridge, 1993: C. Stembridge, 'The Cimbalo Cromatico and Other Italian Keyboard Instruments with Nineteen or More Divisions to the Octave', *Performance Practice Review* 6, pp. 33-59.

Stembridge and Wraight, 1994: C. Stembridge with D. Wraight, 'Italian Split-Keyed Instruments with Fewer than Nineteen Divisions to the Octave', *Performance Practice Review* 7, pp. 150–181.

Straus, 1990: J.N. Straus, *Remaking the Past: Musical Modernism and the Influence of the Tonal Tradition* (Cambridge: Cambridge University Press).

Sundberg, 1982: J. Sundberg, 'In tune or not? A Study of Fundamental Frequency in Music Practice', in *Tiefenstruktur der Musik*, ed C. Dahlhaus & M. Krause (Berlin: Technical University of Berlin), pp. 69–97.

Yasser, 1932: Joseph Yasser, *A Theory of Evolving Tonality* (New York: American Library of Musicology).

Websites

Network for Interdisciplinary Studies in Science, Technology and Music ('*n*-ISM'). http://www.n-ism.org

Centre for Music Techology, Glasgow University home page. http://markov.music.gla.ac.uk/Papers/PDTalk/conclusions.html

Rosegarden sequencer. http:// www.rosegardenmusic.com

Wikipedia home page. http://en.wikipedia.org/wiki/Main_Page

Jack audio connection kit. http:// jackit.sourceforge.net

Scala home page. http:// www.xstall.nl/huygensf/scala

Pure Data home page. http://crca.ucsd.edu/~msp/software.html

Fluidsynth home page. http://www.fluidsynth.com

Audiofile

Hair, Graham (2007) *Ash* for soprano, clarinet and harmonium. Performed by Amanda Morrison (soprano), Ingrid Pearson (clarinet) and Graham Hair (harmonium). (Recorded at the Royal College of Music, London on February 7, 2007).