

An interaction study of a digital music stand

November 8, 2007

Alexey Blinov

`abl35@student.canterbury.ac.nz`

**Department of Computer Science and Software Engineering
University of Canterbury, Christchurch, New Zealand**

Supervisor: Assoc. Prof. Dr. Timothy Bell

`tim.bell@cosc.canterbury.ac.nz`

Abstract

Digital music stands help performing musicians manage their music scores and assist them during performance by providing a way of turning pages without interruption. Many music stand systems exist at the moment, both on the market and in research labs. However, very few of those systems address the issue of page turning carefully and no systems consider animated page turning methods. In this report we propose and evaluate three animated page turning techniques. We also propose and evaluate three methods of interacting with the music stand using foot pedals. Our evaluations indicate that animation could be a very useful feature for beginning musicians; we also discover that using one control for issuing many commands to the digital music stand is quite difficult and requires careful consideration.

Acknowledgments

I would like to thank Tim Bell for great supervision, the Honours class of 2007 for a great year and the ability not to take all this too seriously, Phil Quinn for his help with .NET, and last, but certainly not the least, the musicians, who participated in the experiment.

Contents

1	Introduction	1
2	Background Information and Related Work	3
2.1	What is a digital music stand?	3
2.2	Technologies related to digital music stands	4
2.3	Issues associated with digital music stands	4
2.3.1	Issues associated with digital music	5
2.3.2	Screen issues	5
2.3.3	Page turning	6
2.4	Existing systems overview	8
2.5	Animated page turning	8
2.5.1	Benefits of animation	8
2.5.2	Potential drawbacks of animation	9
2.5.3	Proposed techniques	9
2.5.4	Parameters	10
3	Implementation and Pilot Evaluation of Page Turning Animation	13
3.1	Espresso music stand	13
3.1.1	Issues encountered while implementing the system	14
3.2	Pilot evaluation of animation	15
3.3	Music Stand Interface	16
3.3.1	Design considerations and available hardware	16
3.3.2	Possible techniques	17
4	Evaluation of the Interface and Animated Page Turning	19
4.1	Experiment design	19
4.2	Experimental findings	19
4.2.1	Feedback from participants around the world	19
4.2.2	Results of the evaluation	20
5	Discussion, Future Work and Conclusions	23
5.1	Discussion	23
5.2	Improvements and future work	24
5.3	Conclusion	24
	Bibliography	26
A	Appendix	27
A.1	Post-experiment questionnaire given to the participants	27

1

Introduction

The music industry is the one with traditions that span centuries. Music notation systems are well established and for every musician the notation system is like a second language, which they spend months, if not years, learning and mastering. Music scores are a standard way of communicating music information, and dealing with music scores is a ubiquitous part of any musician's life.

Beginning musicians rely on scores probably more than any other musician. Professional performing musicians usually play the music from memory and rarely use scores during performance; however rehearsing the piece is usually very rigorous and involves a lot of interaction with the score. Show musicians can be asked to perform any song at any time and they cannot usually hold all the songs in their head. Thus, they have a big pile of paper scores to manage and be able to retrieve the one required at the moment quite quickly. In orchestras the process of producing the scores for a concert is not a cheap one; sometimes the orchestra hires the scores from some library or, perhaps, another orchestra. Hired scores usually have some pencil markings, which might have to be erased, and new markings might have to be put in. All of this adds up to a very mundane and laborious process, which is normally done by a dedicated person (librarian), but can involve a lot more people. Computers have automated or made many of these tasks much easier, removing the routine from the life of a musician. Many technologies available nowadays help the musicians in all areas of musical activities: creating, publishing, annotating, managing and performing music.

There is no strict definition of what a digital music stand is. It can support any of those activities, however, its main function is to assist the musician to manage their music.

Formally, the first digital music stand appeared when computers were able to display images, which happened in 1980's, although, it took about a decade before the resolution and the quality of the computer displays became adequate enough for reading music comfortably. However, image or document displaying software by itself makes a poor music stand: the tasks that musicians do with music scores are very different from the tasks performed on a digital document or an image. An ideal digital music stand system would support managing digital music in some sort of a library; composing, editing and annotating music; and provide the musician with a way of performing without interruption to turn the page. One of the first such systems was described in 1996 [13]. It was a concept that was never made into a complete system, but served as an inspiration for many systems to follow, both commercial and research: the music stands available nowadays support many of the desired features and offer many more, such as score collaboration. However, few systems address the performance side of the digital music stand as attentively as it should be addressed. Namely, no systems use animation for page turning.

In our opinion this is a serious shortcoming, as animation can help the musician to orient themselves in the digital score better: animation makes clearer what events are occurring in the system at the moment and makes the experience of using the music stand more similar to the real world. This report focuses on evaluating animation as a means to aid musician's performance while using a digital music stand.

We propose three methods of animated page turning. We discuss general issues regarding the music stands in sections 2.1–2.3 and present an overview of currently available music stand systems in section 2.4. The benefits and potential drawbacks of animation are discussed in section 2.5; this section also introduces the proposed methods and various parameters associated with the animation.

The techniques were implemented into Espresso digital music stand. A small feasibility study of animation techniques was then conducted. During the experiment more was discovered than was intended. The relationship between a forward and reverse animations was not all obvious with the kinds of animation that we were using, which forced us to look into this and posed another question: what is the best way to

provide an interface to the system given the constraints that the performing musician has to work under.

We then propose three different methods of controlling the music stand during performance with the aid of pedals. This is discussed in detail in section 3.3. To evaluate the proposed animation techniques and the control methods the system was made public and musicians interested in the system were informed about it. Also, the system was tested on performing musicians from the University of Canterbury. The results of the evaluation are presented in chapter 4.

2

Background Information and Related Work

2.1 What is a digital music stand?

Most performing musicians use music scores for rehearsal and some use them for performance too. The most common storage medium of the music scores is paper: the scores are published in books or are simply printed out from a computer — there is no shortage of digitised scores available in online music shops. Paper music scores have their obvious disadvantages:

- paper is very prone to physical damage: tear, wear, general soiling. If not looked after well, paper needs to be replaced every so often;
- large amounts of paper quickly become unmanageable. A professional musician might have to carry tens if not hundreds of sheets of paper with them; searching through these could be cumbersome.

The digital age brought many technologies that solve these problems and assist musicians in many areas of musical activities. These technologies include optical music recognition (OMR), composition software tools (such as Sibelius¹, Finale² or Music Notepad [12]), automated score following [24], etc. (some of these technologies are discussed in section 2.2).

A digital music stand is a software or hardware tool that replaces a common music stand. It can support many of the technologies mentioned above, however its main purpose is displaying music. A music stand works with digital files, so the paper scores don't have to be carried around anymore. It does not need an external light as the computer screen emits the light, which could be very useful during some theatrical performances, where the light emitted by the musicians' lamps is unwanted (a screen is not completely dark, but can be far less noticeable than a lamp). The most important advantage of a digital music stand, however, is the fact that it can be completely silent and could provide a way for a musician to play without being interrupted to turn the page. In orchestras the music scores are printed so that the page turn is situated in a place where the musician is not playing; however, this is not always possible and it is not an option for a solo artist or a small group of artists. The result is that there is an audible dip in sound at the place of a page turn [13]; during recording in the studio the musicians have to be extra careful too — noise produced by the page turn could spoil the whole recording.

Hardware music stands usually come in the form of a tablet computer (a laptop computer, with a touch sensitive display, which either folds up into a pad or does not have the keyboard part of the laptop at all), sometimes stripped of most of tablet PC functionality apart from that required for the music stand. Standalone software tools that can be installed on a tablet computer, a laptop or even a desktop computer are available as well. An overview of current tools, both commercial and from research labs, is given in section 2.4.

Naturally, a digital music stand possesses a number of disadvantages, namely:

- it requires external power. Using a laptop also adds the disadvantage of limited battery life and the fact that it can be dropped easily if not secured well;

¹<http://www.sibelius.com>

²<http://www.finalemusic.com/>

- like any computer system, it can crash at any moment for many reasons not necessarily to do with the music stand software. Obviously, this is very undesirable during a live performance and this fact alone could serve as a reason for a musician to not even try a digital music stand;
- musicians who are not comfortable with technology can find operating a digital music stand difficult. Also in a case when something goes wrong (e.g. operating system error) they could find it hard to recover;
- a computer screen is a completely different medium from the paper and some people might find it quite hard to read the music off a screen (issue of reading music off a computer screen is discussed in detail in section 2.3.2).

In addition, there is a sentimental side of the question: many believe that digitising art removes the “vestiges of author’s presence from the work” [16]. To some extent this is true: the expressiveness that the authors have at their powers when they are creating their works with ink and paper (or any other set of ‘analogue’ mediums) is much greater than that of a computer tool, which can only offer a set of standards. However, we believe that in the domain of music (and specifically music scores) the advantages that computers offer outweigh this disadvantage. Besides, the colour screen of the computer is able to represent the original work in full colour — we can simply take a digital photo of the work, while it is rare for a printed score to simply be a replica of the original work. Some musicians can also find the look of a computer inappropriate in a concert environment.

2.2 Technologies related to digital music stands

A digital music stand is an amalgamation of many technologies and standards, some specific to music and some more general. This section gives a brief overview of some of them.

PDF Portable Document Format (PDF) is the state of the art format for distributing digital documents. It is vector-based, which makes it very suitable for music scores that were created using computer. PDF is a standard way of exchanging digitised music and is a popular format in electronic music score stores;

MIDI Musical Instruments Digital Interface is a standard protocol for communicating musical data. The protocol consists of a number of events, which are purely semantical, i.e. only represent the note played, its loudness and length. How the MIDI information is processed is up to the client software. MIDI is used to connect several digital music instruments (such as synthesisers) together. In the context of a digital music stand MIDI can be used in conjunction with the score to assist score following: the stand can ‘listen’ to what’s being played and compare it to the information in the MIDI file;

JPEG Joint Photographic Experts Group image format is the most common format for exchanging photographic images. It is raster-based and it uses a lossy compression method. Jpeg is the output format of the great majority of modern digital cameras.

PNG Portable Network Graphics is a raster image format. Its advantage over jpeg is that it supports lossless compression and transparency.

OMR Optical music recognition refers to transforming a digital image of music into a semantically rich representation. This can mean recognising either individual notes or staff lines. Although quite high in accuracy, OMR systems are still not perfect, with average accuracy rate ranging 89–96% [1].

2.3 Issues associated with digital music stands

Switching to a digital music stand

2.3.1 Issues associated with digital music

One of the first steps of switching to a digital music stand is obtaining digital scores. Many classical scores are covered by copyright, which makes it illegal to produce digital copies of them. Thus, the musicians either have to create their own scores or obtain them from digital stores or free websites. Fortunately, there is no shortage of sources for obtaining digital scores on the Internet.

Once the number of scores starts becoming substantial, it will be more and more difficult to manage them. File system organisation (e.g. storing scores in directories by composers' names) has a major disadvantage: one score can belong to many directories. Thus, some sort of library is required; many document organisation solutions exist, some created purely for that purpose (e.g. Yep³), and some featuring 'accidental' functionality (e.g. iTunes⁴, whose purpose is to manage audio files, but it is capable of managing PDF documents as well). However, ideally the digital music stand would feature a library of its own.

2.3.2 Screen issues

Computer screen is another issues related to digital music stands.

A computer screen, be it a CRT screen or a modern LCD display, is a completely different medium from paper. First of all, it emits light rather than reflecting it. Secondly, the images on the computer screen are produced by the means of the screen refreshing itself regularly (i.e. flickering). At low refresh rates this flicker might be perceived by the user, which could be very tiring for the eyes. Hibbard has established that the screen refresh rate threshold at which no flicker is perceived on the eye is 60 Hz [15]. The problem of display flicker, however, mostly relates to CRT displays due to the fact that on every screen refresh only every second line of the screen is 'drawn' with the rest of the lines being drawn on the next refresh. LCD displays, on the other hand, refresh the whole screen at once. Since our main hardware target is laptop computers, we are not considering refresh rate a problem.

The question of reading from paper versus reading from a computer screen is not one lacking attention from researchers. Numerous papers discuss various issues related to the problem (e.g. [23]). Dillon provides a comprehensive overview of the research of different aspects of reading text from screen and from paper [11]. The main conclusion, however, is that neither of the media can be proven better or worse: there were a number of experiments speaking in favour of paper and an equal number of experiments disputing the results of the former. The only sound conclusion deduced was an obvious one: the quality of the text and imagery displayed on the computer screen drastically affects readability.

Regarding the music scores, two major tasks can be distinguished: performance (sight reading) and proof-reading. Picking [25] conducted an experiment in which he compared paper based music with music displayed on computer screen for proof-reading. As well as simply displaying static music on computer screen, Picking evaluated how different animated visual cues could aid the reader's performance. The visual cues included a jumping stave marker, a smoothly moving stave marker and a note marker, which were showing the current place in music. The experiment showed that paper outperformed the static computer display, but only marginally; nonetheless, the paper was outperformed by music displayed on a computer screen with animation cues. Picking concluded that "there is a strong indication that animation can offer very powerful reading assistance, especially to less able readers of music". McPherson's study [22] was evaluating animation with a different task: sight-reading music. A red marker was indicating current position in music (Figure 2.1) to assist musicians' orientation while they were performing the music. The study showed that "musicians did not find the [bar marker] useful" and some musicians found that "any movement at all became distracting". This is probably due to the fact that following the score is one of performing musician's main tasks and years of practice have made every musician an expert in the task and animation cues serve only as a distraction. The two experiments clearly show the difference of the nature of two score-related tasks and how useful animation cues are for musicians of different experience: while beginners can find guidance useful, professionals perceive it purely as a distraction.

Another limitation of screens is their size. When displaying music on a computer screen we aim to make music sufficiently large for the musicians to be able to read it comfortably. An ideal situation would be if we could display two conventional, non-scaled, sheets of music next to each other, but this would

³<http://www.ironicsoftware.com/yep/index.html>

⁴<http://www.itunes.com>




Figure 2.1: A screenshot from McPherson’s experiment showing a red bar marker that indicates current position in music

require a rather big stationary screen and a performing musician would often rather invest their money in a better instrument. Besides, one of the desired features of the digital music stand is its mobility.

A truly mobile music stand would be a sheet of electronic paper: it is flexible and requires far less power than a laptop. However the technology does not yet offer an affordable e-paper option. The best tool currently available on the market in terms of price/mobility trade-off is a tablet computer, followed by a laptop. A pianist can easily place a tablet PC on the music stand that each piano has, however, fitting a laptop to it might be somewhat awkward; a guitar (or any other instrument that allows for space around the performer) player would be able to place a laptop computer on a stool or a small desk in front of them.

A standard laptop would feature a display of size between 12 inches and 17 inches (diagonal), 14.1 inches being the most popular. Tablet computers usually come with either a 12 inch or 14 inch display. A typical size of a paper music score is 12.5 inches by 9.5 inches, measuring 15 inches in diagonal. Even though the computer screen might have the same diagonal measure as a sheet of music, the aspect ratio of the sides might be different from that of the paper score. Smaller screens usually feature 16:9 screen aspect ratio, while older 14 inch displays usually have 4:3 aspect ratio. Any conventional laptop screen is too small for displaying an equivalent of two pages of music to be read easily. McPherson’s experiment [22] showed that some of the musicians find even the 17 inch screen too small to accommodate for two pages of music. In fact, we will need a 17 inch screen to display a page of music in its original size: a 12 or 14 inch display is just not enough to display the whole page at once. Therefore, either the music needs to be reduced in size (80–90% scaling depending on the score) or less music needs to be displayed. An experiment by Bell et al. [4] compared different sizes of music for a proof-reading task. The participants were given scores of two sizes—one larger than conventional music and one smaller—with intentional errors and were asked to find errors in the scores, while listening to the piece. The experiment revealed almost no difference in performance, however, the participants in the post-experiment questionnaire expressed a preference for larger music.

There is a very delicate trade-off between the amount of music displayed on the screen and its legibility. Evidence points to the fact that although some musicians don’t mind the music to be smaller than usual, most still prefer it to be standard size or even larger when displayed on a computer screen. We suspect that this choice is affected by many factors, including musician’s age. Younger musicians might prefer having more music on the screen, while older ones might have trouble reading small music. Also, if the musician is not experienced with the piece, they would like the score to be bigger. However, there is a limit to how big the music should be to remain readable: too big a score would mean very little music will fit in musician’s field of view. In any case, a bigger score means that we cannot always display as much music on the screen as we would like to. This brings us to the problem of page turning.

2.3.3 Page turning

Having less music on the screen means that the piece will be split into more parts, which leads to more page turns. The paper scores are normally printed in such a way that the page turn happens during the musician’s rest time. However this is not always possible. Therefore, a digital music stand needs to provide a way of continuous display of a musical score that is otherwise too large to fit onto one screen. This is very hard to do with paper: a possible option is a book that is printed so that turning a page changes the contents

of only one side of the book; also if the score is printed on individual sheets of paper, they can be slid to the side instead of being turned. When displaying music on the screen, however, we can employ all kinds of different techniques. We will refer to the issue as a whole as page turning, even though some of the methods discussed aren't strictly turning pages.

The following methods of turning from old page to a new page are possible:

Static page turns Here the screen of music is simply replaced with the new music. This method has an obvious disadvantage of 'surprising' the musician. Even if the musician issues the command to turn the page to the music stand themselves, a small delay in the system (that could be caused by memory repagination, for example) combined with a short loss of attention (even blinking) could confuse the musician. Besides, since the next page is going to look quite similar to the previous page, it may not be obvious that it has changed; the musician can also initiate two page turns instead of one and not notice it. Furthermore, if the scores have been created hastily (e.g. by taking pictures of the scores with a digital camera), it is quite likely that the pages will not be aligned with each other, so the musician would not know where to expect the music on the next page.

Semi-static page turns With this method the old music is replaced with the new music in steps, part by part. When displaying one page of music at a time, the top half could be replaced first and the bottom half later, or the old staff lines could be gradually replaced by the new ones. When displaying two pages of music, one page could be replaced at a time [13]. This method is an improvement over completely static methods, however it still features the same disadvantage: the musician might not notice that the top of the page has in fact changed. If we will be replacing the old page with a higher granularity (e.g. one staff line at a time), the likelihood of musician noticing the change is even slimmer, so we need to provide some sort of indication of the change. This could be either a static marker (like the number of the staff line or a small shape in the margin) or an animated cue. As mentioned before, musicians find animated guidance distracting. Static guidance would increase the musician's mental load, as they will have to remember either what the different colours represent or what staff number they were up to.

Animated page turns With this method the screen of music is gradually replaced with the new music with the aid of animation. Animation is a strong visual cue, it is much less likely to be unnoticed and it makes it clearer what has changed.

McPherson researched some of the page turning techniques [22]. The researched methods included displaying a single page of music and gradually replacing the old staff lines with the new ones, displaying two pages of music and replacing one page at a time, and vertical and horizontal scrolling and jumping of the score. The methods were evaluated on a small group of jazz school students. It was discovered that jumping was "too difficult [for the musicians] to follow" and the method preferred by most of participants was the vertical scrolling, where the pages of music were joined into a continuous sheet, which was scrolling from the bottom to the top of the screen. This was followed by the single page method, where the old staff lines were gradually replaced with the new ones, so that when the musician gets to the bottom of the page, they can continue reading from the top. The results give an indication that animation could be a very useful feature in a digital music stand.

The vertical scrolling method that participants in McPherson's experiment favoured the most has a major disadvantage: it requires constant attention of the musician. Continuously moving music will substantially hinder a musician's ability to find current place in music: if the musician looks away from the score, they can not longer rely on their spatial memory to find the place they were up to. Usually during performance there is no lack of distraction factors, especially if the musician is playing in an orchestra: the musicians have to look at the conductor, their instrument, and fellow musicians. In fact, some (rather expressive) solo musicians seldom look at the score. Thus, we believe that a page turning system should keep the music itself static, but provide an animated way of moving from one page to another.

In the following sections we will give an overview of the digital music stand systems currently available, what page turning methods they offer and how they can be improved.

2.4 Existing systems overview

Muse is a concept system developed by a group of researchers from Carnegie Mellon University in 1996 [13]. Muse is an all-around system that features a specifically designed computer with pre-installed software on it. The authors discussed the problem of page turning in the context of automation and considered a semi-static method, where two pages are displayed and they are refreshed individually.

MICON, a Music Stand for Interactive Conducting [7] is a music stand that was developed in addition to Personal Orchestra [8], a museum exhibit that let the visitors to control a recorded piece of music through conducting gestures. The system used a complete emulation of a book of scores, that was rendered in 3D with OpenGL. The music is displayed on the pages of the book and is static. A page turn would “start when the last beat of the right page was reached, and end with the first beat on the new page”; this was found “unnaturally slow”. The system was tested on music students. No specific evaluation details are available from the authors, and the only conclusion that was drawn from the evaluation was that “the students liked the appearance of MICON”.

While MICON was developed with non-musically trained users in mind, Multimodal Musical Stand (MMMS) [2] was “intended for expert interaction of trained instrumentalists who wish to perform with a computer interactively”. This was a research system, the main goal of which was to introduce gestural methods of interaction with the system. Page turning was not addressed in the design.

Music Object-Oriented Distributed System (MOODS) is a system that is designed to be used by the whole orchestra [5]. The system “frees composers and conductors from mundane, time-consuming clerical tasks” by providing a way for the musicians and the conductor to collaborate on the score. Each musician’s stand is a computer, which is networked with a central controller, which manages the music. The controller also initiates the page turns for all the musicians. The authors used a semi-static page turning technique, where the old staff lines are replaced by the new ones once they have been played. An evaluation of the system showed that the users preferred the method used to both static page turns and scrolling score.

Of all research systems, the system developed by Kosakaya et al. [19] has received the most attention in terms of page turning. The authors chose a semi-static page turning method: the screen was updated in two passes. The first pass was initiated by the musician and updated some part of the screen; the second pass updated the rest of the screen automatically after a certain delay. A study was conducted to determine what ratio of the screen update and what delay values were preferred by musicians of different experience. It was found that professional musicians were able to perform with the refresh ratio in the range of 70–100% and delay values of 0–1 seconds, i.e. with most of the old page being refreshed on the first pass and the rest of the page refreshed after one second delay. Amateur performers, however, were more comfortable with lower refresh ratio values and higher delay values.

Commercial music stand systems available in the market include eStand⁵ and MusicPad Pro⁶. eStand is sold as a dedicated device with a touch screen and pre-installed software (Figure 2.2(a)) and as a stand-alone software package. The software uses a static page turning technique, refreshing the whole page at once. MusicPad Pro (Figure 2.2(b)) is available as a hardware device, very similar to eStand. However, MusicPad Pro has an optional half-page preview method for page turning.

2.5 Animated page turning

None of the systems discussed above are considering using animation for page turning. Research systems are focusing on features such as collaboration and innovative input methods; commercial systems put the emphasis on score annotation and music libraries. We believe that animation has potential of improving musician’s experience with the music stand and propose three animated page turning methods. This section describes the general advantages and disadvantages of the animation and the methods proposed.

2.5.1 Benefits of animation

The main problems with static and semi-static techniques were outlined in section 2.3.3:

- Static techniques may be confusing, since there is no flow between the pages;

⁵<http://www.estandmusic.com>

⁶<http://www.freehandsystems.com>



Figure 2.2: Commercial music stand systems

- Semi-static techniques could be somewhat confusing too: musician can simply forget whether they turned half of the page or not;
- Any delay caused by the system may prompt the musician to trigger the page turn once more and it can be hard to tell how many page turns have occurred.

Animation offers better awareness of what is happening in the system. It is much harder not to notice a moving object than an object that changes its position at discrete intervals. The factors that influence human attention are motion, size, shape, colour and contrast of the objects [28]. Two different pages of music from the same score are highly unlikely to be different in any of these factors; animated page turns, however, will provide the factor of motion, which means it is much less likely to be unnoticed.

Studies also show that users prefer animation when navigating within a digital document and perform better in search tasks when the navigation is aided with animation [18]. It has also been shown that resemblance with real world improves user experience, both in the domain of music stands [7] and applied to other computer simulations, such as virtual books [9].

2.5.2 Potential drawbacks of animation

As mentioned previously, the musicians find animation distracting during performance. Although this was proven only for guiding animation (i.e. a marker indicating the current place in the score), the same could hold true for page turning animation.

Since the effect of animation is not instant, it could be frustrating for the musician if they have made a timing error and executed a page turn too late.

There are also purely technical disadvantages to the animation: smooth animated imagery requires higher processor load, which means higher power consumption and potentially higher processor fan load, which would also produce undesired noise.

2.5.3 Proposed techniques

For the reasons discussed in section 2.3.3 the techniques that were moving the music itself were not considered, despite the fact that some of them are physically more realistic (e.g. sliding the page sideways). Since a page of music is always read from top to bottom all the animation methods replaced the page starting from the top.

The following techniques are proposed:

Page fold technique This technique emulates a page peeling off a stack of pages (Figure 2.3).

Page roll technique Emulates a page rolling down towards the bottom of the screen, revealing the next page (Figure 2.4).

Simple line technique This technique is an approximation of the previous method. It simply draws a straight line between the parts of the old page and the new page (Figure 2.5).

All of these techniques reveal the upcoming music at the top of the page, while keeping the current music visible and not moving.




Figure 2.3: Page fold technique at different stages




Figure 2.4: Page roll technique at different stages

2.5.4 Parameters

There are a number of parameters associated with the animations. The main parameter is the speed of animation, i.e. how long it takes for a page turn to complete. Too fast an animation may cause problems when triggered too soon: the old music will be obscured before the musician will finish playing the music



Figure 2.5: Simple line technique at different stages

on the page. If the animation is too slow, then it might be that it will not reveal the staves on the new page by the time musician gets to them. This, of course, heavily depends on when the musician will trigger the page turn. An experienced user will know when they need to turn the page if they are aware of their current speed setting; however, a novice user might forget to trigger the page turn on time.

Stopping the animation half-way through the page turn is another option. This means that the page turn will have to be done in two steps, but will give the musician the choice of triggering the page turn earlier. This method will simply be enhancing the half-page preview technique.

To reduce the number of commands needed to be issued to the system back to one, but retain the freedom of making a slight timing error, we could set different speeds for different page halves. Thus the musicians would be able to set a really fast first half animation and a slow second half animation, which would produce a good compromise between stopping half way through the transition and wiping the whole page in one transition. The opposite (slow first half and fast second half) could also be of potential use for scores done in such a way that going fast over the first half of the page would mean obscuring the music currently being played.

The page fold technique is the most realistic, but has a major disadvantage of obscuring the music at the bottom of the page too early. While this could be fixed by setting a rather low speed of the first half of the animation, we could also limit the maximum height of the page fold.

There is no standard way of paginating a score: one transcriber could distribute a score across 20 pages, while another would be able to fit it in 10 by putting more music on one page. While testing the system we have seen many scores that weren't formatted in a standard way, including a landscape oriented score. This means that there can be no standard settings for the animated page turns — every score would require individual set of settings. In fact, individual pages in the score could require different settings. The settings would depend not only on the amount of music that each page contains, but on the difficulty of the score. For example, the musician might find that due to an especially tricky bit of music at the end of one page they forget to initiate a page turn and only remember to do it just before the last notes on the page, therefore they might choose to use the transition with very fast first-half-page wipe.

All the described techniques could be adapted for horizontal animation, which would be suitable for scores with landscape orientation, such as a conductor's score.

3

Implementation and Pilot Evaluation of Page Turning Animation

3.1 Espresso music stand

The animation methods proposed in the previous section were implemented into Espresso, a research digital music stand system [3]. In its current version Espresso featured an OMR system, which was capable of recognising staff lines. The page turning techniques implemented in the system were static page turns and some of the animated page turns, such as page wipe and page scroll. Unfortunately the animated page turning techniques were implemented as proof of concept and were not very easily modifiable, so the page drawing logic had to be rewritten.

The system was implemented on .NET 2.0 platform, using the C# language. Microsoft's GDI+ library was used for animation drawing. The system used PDF files as input; it then used ImageMagick library¹ to resample the pages of the PDF document into a set of PNG images. The images then were drawn on the screen.

All the animation methods described in section 2.5.3 were implemented in the system. Figure 3.1 shows different stages of page fold animation.

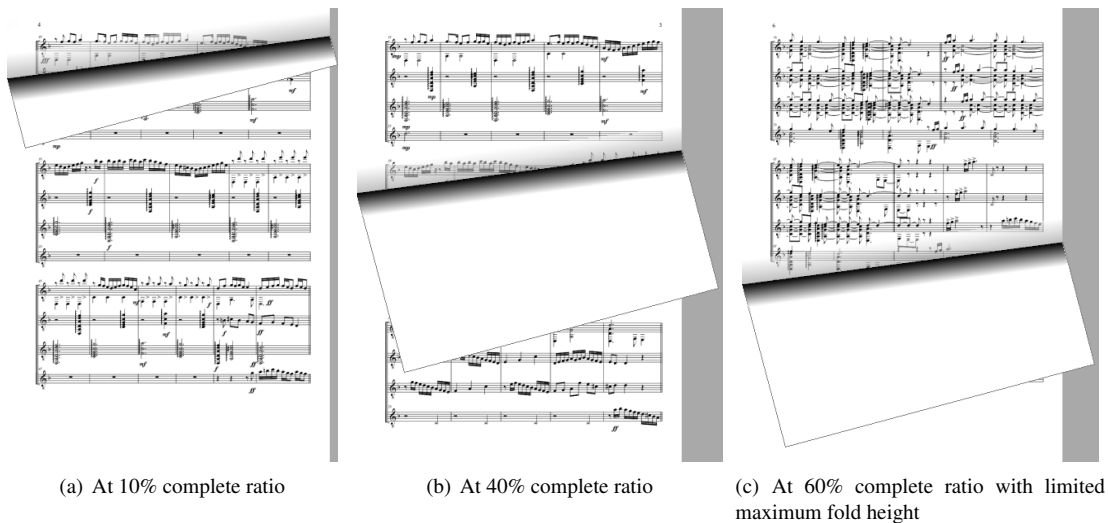


Figure 3.1: Page fold animation at different stages

When implementing the system, the following design considerations were made:

- Resampling PDF documents takes a considerable amount of time, so to aid faster document open-

¹<http://www.imagemagick.org>

ing the image files were cached on the hard drive once they were rendered. PNG was chosen as the file format for storage due to its lossless compression capabilities. This way once the file has been opened, subsequent opens will take much less time. This is quite important in performance environments, where a piece has to be opened in a matter of seconds;

- All the settings had to be saved between the program runs. This was done by writing them out to an XML file before closing the program and reading back the file when the application started;
- A compromise between the quality of the resampled image and the amount of resources required for its processing had to be made. Low quality image would be very fast to process, but could hinder readability; high fidelity images could cause delays in display, which is very undesired. Resampling resolution of 100 DPI was chosen as the compromise. Later this was found inadequate; the situation is explained in the next section.

All of the parameters described in section 2.5.4 were implemented and were made user settings, which were accessible through File menu (Figure 3.2).

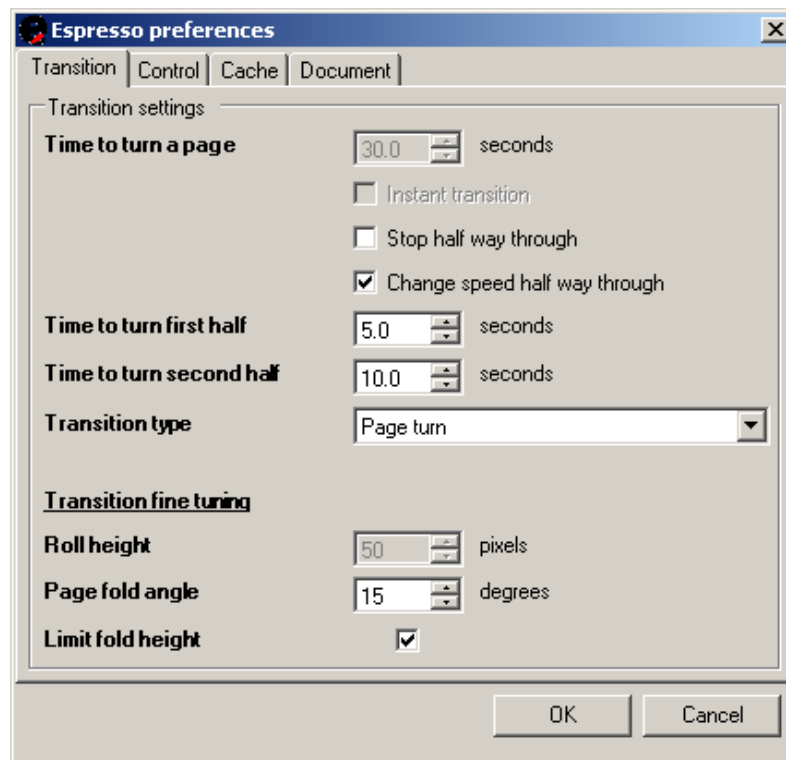


Figure 3.2: Espresso preferences window

3.1.1 Issues encountered while implementing the system

The resolution of 100 DPI turned out to be not high enough for some of the scores. Due to the resampling algorithm the scores in which the staff lines are not horizontal, but are slightly skewed, were rendered with 'white rivers' across the pages. Figure 3.3(a) shows such score in its original form and figure 3.3(b) shows the same score resampled at 100 DPI. In the end it was decided to have the user select the resolution at which the documents were resampled. The upper limit was chosen to be 300 DPI as this proved adequate for every score tested.



Figure 3.3: Low resampling resolution aliasing defect

3.2 Pilot evaluation of animation

A small feasibility study was conducted to evaluate different animated page-turning techniques. The goals of the study were to find and eliminate possible bugs that were not encountered during initial testing. It was also intended to get as much feedback as possible from the participants before sending the system to the testers, so that possibly some of the methods could be ruled out, some improved, and possibly new ones implemented.

Espresso was installed on a Compaq tc4400 tablet computer with 12.1 inch display running Windows XP with 1024x768 pixels resolution. Two musicians, both pianists, have used the system for a period of 30 minutes by playing an unfamiliar piece of music. The page turns were triggered by a synthesiser sustain pedal, which was connected to the laptop via a USB cable.

The first musician was already familiar with digital music stands: they used Adobe Acrobat software for displaying the scores before. They found that they liked the page roll technique better than others. They also really appreciated the animation stopping half way through and commented that “it makes absolute sense what is going on”.

The second musician had no previous experience with computer-based music stands. Overall, they found their experience with Espresso very enjoyable and thought that “it was a great idea”. They found that they have no preference over which animation technique they preferred and said that “a simple line is good enough”. The piece of music that the musician was playing had a repeat in it, i.e. a place in the score, where they had to return to a previously played staff line and play it again. The start of the repeat happened to be on the previous page. While the system provided a way to return to the previous page via the keyboard and a page-up/page-down button on the side of the tablet computer, there was no way to invoke the *previous page* command with the pedal. Furthermore, only the static backward transition was implemented (i.e. the previous page simply replaced the current page).

The repeats (and D.S. and D.C. notations, which are other kinds of repeats) are the main reason the musician would need to travel back a page during performance. During a rehearsal navigation in the music document is quite vital, but it could be done via other controls that are more suitable for the task, such as the keyboard or the drop-down menu. A possible solution to having to navigate within the score in order to play repeats is to use an optical music recognition system, which finds the repeats and automatically ‘unfolds’ the music, i.e. identifies the chunk of music that needs to be repeated and inserts a copy of it after itself. However, most OMR systems require a high quality image to work successfully [20]; this means

that the music has to be scanned at a reasonable resolution and without artifacts such as page curl due to binding in case the score was in a book. However, often the musician would like to create a copy of the score as quickly as possible. For example, imagine a group of musicians who have decided to have an impromptu playing session and only one of them has the score. The others then could simply take photos of the scores with a digital camera (it doesn't have to be very high quality camera either) and import them into Espresso.

Nevertheless, the evaluation posed two interesting questions: what is the best way to indicate a page turning back and how should it be initiated? Since going back one page usually means the user is returning to the place of a repeat, the music on the previous page will need to be accessible as quickly as possible: the location of the repeat is unknown. We were not considering instant transition for the reason discussed in section 2.3.3; using animated methods meant that the animation would have to be quite fast and independent of the animation speed selected for the forward transition. A value of one second was chosen as the compromise between user's ability to notice and interpret the motion of the page and the speed adequate enough to reveal the music on previous page. Furthermore, to aid user's metaphor of traveling within the document we have made the animation reverse, so that the action of returning to a previous page is not perceived as another *turning the page* action, but rather as an inverse of the *going forward* action.

The question is how to initiate the *go back a page* command? Given that the musician controls the page turning with a pedal, a few options are available. They are discussed in the next section.

3.3 Music Stand Interface

Dillon in his overview comparison of paper and computer screen as media for working with documents writes that "... the most obvious difference between reading from paper and from screens is the ease with which paper can be manipulated and the corresponding difficulty of doing so with an electronic text" [11]. One of O'Hara's implications for design of a digital document manipulating system is the "need to support quicker, more effortless navigation techniques" [23]. The digital music stand as a page-turning tool is the system where efficient navigation within the document is vital. In this section we will briefly explain why navigation is important and will introduce a few options of interface implementation.

3.3.1 Design considerations and available hardware

Navigating within the music score has two basic use cases: during performance (to turn pages) and during a break (to find a required page or to return to the beginning of the document). While not performing, the musician is not so constrained, thus any conventional interface can be used: we can employ a mouse or a keyboard, which can allow for highly efficient navigation. A performing musician, however, does not have their hands available to them and is very constrained in movement. Also, the cost of user error is much higher — if the musician accidentally turns two pages instead of one, they will have to stop performing to realise what the error is and then fix it.

User bandwidth characterises the rate at which users process information [21]. Bandwidth is measured in bits/second, which indicates the amount of information that can pass through a channel in given amount of time. If using interface A it takes 5 seconds to complete a task, and it takes 8 seconds to complete the same task using interface B, then interface A is said to have higher bandwidth. With a design goal of making the musician's performance uninterrupted the interface to the digital music stand is one with very limited bandwidth — the great majority of music instruments occupy both musician's hands, leaving only the head and the feet free. Foot pedals seem like a natural choice for control device in this case: head-mounted control systems are rare, could be awkward to control, would require the musician to look at the score all the time, and would look inappropriate in the concert environment. There are two types of pedals: push pedals and continuous controller pedals. The former allow only one action to be performed on them: a simple push; continuous controller pedals are usually made in the form of a seesaw and allow to control a variable setting (such as volume), which value is determined by the recline angle of the pedal. The controlled variable could be the completeness ratio of the page turn, i.e. to turn the page the musician would need to bring the pedal from fully reclined in one direction to fully reclined in the opposite direction at the rate they feel comfortable with.

The disadvantage of the push pedal is that it could be difficult to find. Musicians who stand up could

also find it awkward to operate, since they need to balance on one foot and do the push with the other. Dannenberg proposes a solution to these issues — a pressure sensitive pedal [10]. The pedal appears as a flat surface and triggers an action only when enough pressure is applied to the surface. Although a nicer solution ergonomically, it is still a switch that allows user to do only one action.

In the following section we propose different techniques of using push pedals to control page turning.

3.3.2 Possible techniques

Existing digital music stand systems use either buttons [13] or pedals [19] as the control devices. Some of the systems use automatic page turning, controlled by a person in charge (usually the conductor) [5, 6]. Fully automatic page turning, using score following techniques [26] is a possibility, although it is prone to errors. Some of the experiments shows that musicians accept the automated page turning with enthusiasm. However, these systems would work only in orchestras, and would not be suitable for small bands and solo musicians. Furthermore, a system, in which the behaviour is not triggered by the user brings a degree of surprise and may be found confusing by some users [22].

Our first method uses two pedals instead of one: one pedal for the forward turn and the other for the backward turn. The advantage of this technique is that the metaphor of right representing ‘next’ and left representing ‘previous’ is a well established one among most computer users, so it shouldn’t be too hard to get used to; the disadvantage, however, is the addition of another control. This means that it will take slightly longer to find the pedal; besides, some musicians (like pianists and organists) already have enough pedals to manage, so fewer additional controls there are, the better. Introducing more controls also increases the total cost of the system. Thus, we also tried employing just one pedal for the task.

In computers the left mouse button is usually used for two actions: click and double click. Our second method employs the same technique for the pedal. The disadvantage of the method is the inevitable delay in executing a single push. To execute a double push, we need to make a single push first. This means that either the action that a single push is assigned to must be executed and then cancelled on the second push, or the system has to wait for the user to do a second push, and if it is not made within a certain time limit start turning the page. Turning the page immediately after the user pushed the pedal once may cause confusion. In fact, if the musician uses static page turning techniques, we definitely don’t want to turn the page if we know that a second push might follow: in this case the user will first be taken to the next page and then, after the second push, immediately back two pages, which will almost certainly confuse them. Thus, we need to wait for the user to make the second push, and if none were made, we can start turning the page forward. Regarding the double-click speed value, there is a trade-off between the responsiveness of the interface and ease of executing a double push.

But what if more than two commands are needed? For example, the musician might want to go the last annotation they made or the repeat could be more than one page back. The problem is very similar to interfaces for physically impaired people, but to a smaller extent: the number of commands that need to be issued to the music stand during performance is very limited, while universal interfaces usually control the whole operating system. The problem to solve is allowing users to execute many commands using just one action. Most of the research in the field of interfaces for the disabled is concerned with eye tracking (e.g. [17]), however software solutions, which do not require any additional hardware exist as well. An example of a software universal interface is a *scanning menu* [27, 21]. A scanning menu steps through a list of items to select from at a given rate: the next item is chosen automatically with a preset delay. When the desired item is chosen the user pushes a button and the selection is made. We have adapted this method with an important alteration. A musician has to pay attention to the music they are playing and cannot afford to spend time to look at the menu and wait for the right item to be selected; besides, the rhythm of the jumping menu selector may interfere with the rhythm of the piece, confusing the musician. Thus, our menu works in the following way:

- A push of the pedal brings up the menu, selects the first item and starts the timer; the menu is semi-transparent, so that if it obscures something important, the effect is not as harmful;
- When the timer expires the current item is selected;
- Any subsequent push while the timer has not expired selects the next item in the list and restarts the timer;

- After the selection is made, the menu disappears from the screen.

The menu allows us to have more than two commands that can be sent to the system having only one pedal at the cost of time to select the item needed (which is actually not going to be that high since most of the time the users will be selecting the first item anyway). An example of the third command would be *go to the previous annotation*, but since at the time of the experiments the software did not support annotations, the third command we have put into the menu was *jump to previous page*, which turned to the previous page immediately. Figure 3.4 shows the menu with the first item selected.




Figure 3.4: Page turning menu with the first item selected

4

Evaluation of the Interface and Animated Page Turning

To evaluate the new version of Espresso we have made the system available publicly and asked musicians, who previously expressed interest in the music stand to try it. Furthermore, we have tested the system with musicians from the University of Canterbury.

4.1 Experiment design

The software was made public and e-mails were sent out to people who have expressed their interest in previous versions of Espresso. They were asked to use the software for a few weeks, trying different page turning methods. After a few weeks the participants were sent out a questionnaire, where they had to rank different page turning methods and comment on their usability.

The software was also evaluated with several musicians from the University of Canterbury. All the musicians were male pianists. They were asked to play various scores using Espresso (Figure 4.1). The chosen scores were unfamiliar to the musicians. In the first part of the experiment the musicians were shown different page turning methods; they tried using each method several times with different scores. They were told about all the parameters associated with each method and were free to choose whichever values they wanted for any parameter. In the second part of the experiment they were given a score with a repeat in it that went one page back. Here they were told about different control methods and were asked to try each one. During the experiment their feedback was recorded, and after the experiment they were asked to fill out the same questionnaire that remote participants were given.

4.2 Experimental findings

4.2.1 Feedback from participants around the world

The first feedback that we have received from the participants was the request of the system for Macintosh platform. This was requested by at least four people out of 22. This confirms the popular view that Macintosh seems to be the platform choice for professional musicians (much professional music software is developed exclusively for this platform). However, nowadays the owners of Mac computers have a choice of running Windows operating system as well using either virtual machine or dual-boot software, which is what we suggested to the users.

Several features were requested by the distant users. One of the participants used Adobe Acrobat as their music stand and kept their music scores in large (up to 90 pages) pdf documents, so that they only had to open one file during the performance and just navigate between the pages to get to different scores. Since Espresso was resampling each page of the document when it was being opened for the first time, it took far too long to open a 90 page document, and the participant thought that the system had crashed. The problem was fixed by introducing a user setting that was the length of the document that was considered too long and checking that the document being opened was not longer than this limit. If it was longer, then the user was asked which pages of the document they would like to open.

A few bugs were discovered by some of the participants. The bugs were fixed within a day and fixed software was released and participants were notified.

Unfortunately, when e-mailed with the post-experiment questionnaire, none of the participants replied,



Figure 4.1: A participant performing music using Espresso

although it was known that at least seven of them have used the software. The participants said that they were either too busy at the moment or have not used the software enough to be able to answer the questionnaire. This suggests that when creating a software system, targeted at professionals, any obstacle that the user encounters while using the software (especially if it is new to them) simply draws them away from the software and increases the likelihood of them abandoning the software altogether.

4.2.2 Results of the evaluation

None of the three participants in the local experiment (further referred to as participants A, B and C) found reading the music from computer screen difficult, rating the difficulty as 1 out of 5 (not difficult at all). Static page turning was immediately disfavoured by one of the participants (who has never used the music stand before), but was not found too difficult to use by the others. Half-page preview technique was an improvement over a fully static page turn for two of the participants, however, participant B found it quite difficult to use due to a degree of surprise: “Something’s changed and you subconsciously think: ‘what’s changed?’ ”.

Animation was preferred to static techniques by participants A and B, but not participant C. Participant B found that they were now able to use half-page preview enhanced with animation, since animation was “reassuring that something was happening”. Participant C, however, commented that they would still prefer to use static half-page preview over animated methods because of its immediateness — the music on the new page can be read right after the pedal push and there is no need to wait for the animation to uncover the upcoming music.

However, the most favoured setting for the page turn among participants A and B was changing the speed of the transition half-way through the page. They found it easier than animated half-page preview, since it required only one push of the pedal; the speed settings chosen were really fast (1.5–3 sec.) first half-page and substantially slower (around 9 sec.) second half-page wipe. Participant A tried reversing the settings (slow first half, fast second half), but found that it was much harder to notice the animation.

The type of transition preferred by participants A and B was the page roll, since it provided a better clue than a simple line, but did not have the cost of obscuring the music that page fold had. It was much

easier for participants to find the suitable speed settings for the page roll than for the page fold.

Controlling the music stand with two pedals was found quite easy by all the participants and received a rating ranging 1–3 (1 being very easy to control). All the participants required two or three attempts to get used to going back one page initially, after which they had no trouble returning to a repeat. Using one pedal and double push was found to be frustrating by all the participants because of the rhythmic differences between the music and the double push. When the musician is performing they hold the rhythm of the music in their head; the rhythm of the piece can be much slower than double push of a pedal at the lowest speed allowed, which means that the musician has to switch their attention from the music to the control pedals and then back, which, as participant B stated, adds “a whole another layer of rhythmic ambidexterity”. Using a menu with a longer delay was welcomed by the participants as a better option, but was still rated worse than the two pedal option, due to the delay that was introduced for the *turn a page forward* command. Participant A said that they would prefer the menu to appear at the lower right corner, because that is what their eyes are focused on when they get to the bottom of the page.

Regarding returning back a page, participant A commented that they have never had to return more than one back. However, participant B, who plays at music and theatre shows quite often, said that it is quite often the case that the score, instead of being 20 pages long, is made 10 pages long, but with a big repeat that spans 5 pages or so.

The pedal used in the experiment triggered a transition on the way up. Some of the participants commented that they are more used to pedals reacting to pedal-down action. However, both methods have their advantages. Pedal-down has no response delay, it triggers a command as soon as the musician touches the pedal; with pedal up there will be a slight delay between when the musician reaches the pedal and when the command will be executed. Furthermore, the musician, being immersed in playing, might push the pedal down and forget to let go of it. The advantage of the system reacting to pedal-up event is that there is far less risk that the musician will miss the pedal: they could put their foot on it in advance. This could be important for those using the music stand for the first time; with experience, however, the musicians will get used to having extra pedals and would be able to find the pedals easily. In fact, participants A and B commented that they wouldn't mind having to reach for a button on the music stand to initiate a page turn.

5 Discussion, Future Work and Conclusions

5.1 Discussion

Unfortunately the number of participants does not allow us to draw statistically meaningful conclusions from the experiments. However, we have received some invaluable feedback from our participants, and some implications for design can be drawn.

None of the participants found reading from the screen tiring, which means that a modern 12 inch display is actually good enough for displaying a scaled down A4 page of music to be read comfortably.

Importance of responsiveness of the system was something all the participants agreed on. Everyone noted that even a slightest delay in system response caused them to wonder if the command has in fact been issued. Participant C put an especially strong emphasis on the fact that any command in the system should be immediate. They found that they preferred to trigger the page turn right at the end of the page and immediately jump to the top of the next page. Thus, the technique they preferred was the static one. However, all other participants agreed that an animated page roll method with faster top half animation and slower bottom half animation was the best one. This is probably due to the fact that out of all the participants participant C was the one for whom playing piano was a major part of their occupation, while for the others it was an addition to their day-to-day activities. This indicates that the animated page-turning techniques are quite useful for amateur and beginning players, while professionals might find that they don't need the advantages offered by the animation, opting for a simpler and quicker alternative. This reinforces the results obtained by Kosakaya et al. [19], who found that professional musicians were more comfortable with refreshing most of or the whole page at once, while less able musicians preferred having some room for timing error, which is what the animation and half-page preview techniques provided. It was also important that the animation was configurable, since different musicians will have different animation speed preferences.

The participants who preferred animated page turning chose page roll as the animation technique. The simple line was not always obvious enough as the cue, and the page fold was good, but the animation was not perceived as a folding page, but simply as "something moving across the screen". Although the animation for the page fold was just an approximation, it can be concluded that high fidelity is not necessary in providing visual cues for the musicians.

It turned out to be extremely difficult to employ one foot control for several music stand commands. Double pushes are quite hard to do while playing music: the musician always plays to the tempo, consciously or subconsciously (depending on experience) counting the beats in their head. When performing a double-push, they will naturally time the pushes with the music beats, which most of the time is going to be too slow for a double-push. A solution would be to lower the speed of the double push, but this would make the system very unresponsive, since we will have to wait longer after a single push to confirm the *turn a page forward* command (for discussion on this see section 3.3.2).

Of all the participants only participant B found the pedal operated menu useful and said that it is a feature they would require often, which suggests that the technique is worth looking into, however, it should not be the only technique available, as it will only be found useful by show performers, for whom it is not a rare case to have a repeat that goes more than one page back.

5.2 Improvements and future work

Our experiment results point suggest that the animation is found useful by some musicians. However, a much more rigorous study with more participants is required to make reliable conclusions. The evidence points to the fact that amateur musicians would appreciate the animation more than professionals, so the study should distinguish between the musicians of different experience.

A potentially useful feature would be a static marker that would remind the musician when to turn the page. This feature would require an optical music recognition system, which would split the score either into staff lines or bars, depending on how good the system is. The musician then would need to enter the timing and rhythmic information of the piece (beats per minute, bars per line, etc.). Given the animation settings, the application will then be able to calculate at which bar the musician will need to trigger the page turn and place a small marker above that bar. The system could also use audio input to compare what's being played to the score and place a marker at the place where the page turn will need to be initiated approximated based on the current animation settings.

Annotations can be seen on almost any paper score and for most of musicians the process of annotating music is ubiquitous and is strongly associated with the paper scores themselves. The annotations can be regarding the music specifics, such as volume, tempo or mood; or they can be purely technical, such as bowing instructions for a violin. A comprehensive classification of most common annotations is available [29]. Computers changed the way people exchange documents and collaborate on them. There has been some research on how paper documents are used by professionals in a commercial organisation [14], but no similar studies have been conducted in the field of music documents. Tablet computers make creating annotations on digital documents extremely similar to writing on a piece of paper, but with the advantage of being able to erase what's written. Typing the notes on the keyboard is not as natural as writing them with a stylus, but is certainly a more affordable and readily available option. More research needs to be done in the area of score annotations to discover how the musicians interact with paper scores, what part annotations play in the process, and how these processes can be supported by the digital music stands.

Portable devices (PDAs, iPhones, etc.) are becoming more and more available. Although they are not very suitable as a primary display for a music stand, they could be sufficient enough for pieces, where the score need not be of very high detail and only serves as a reminder of what to play, e.g. jazz improvisations or scores that the musician is very familiar with. We have tried emulating a portable screen and found it to be legible for displaying music; obviously, more research needs to be done in this area.

5.3 Conclusion

We have presented a digital music stand system with novel animated page turning techniques. Our experiments show that animation can be a very useful feature in a page turning digital music stand. Our hypothesis that the animation has a potential of offering much better awareness of the events occurring in the system has been confirmed by most of participants. The only musician who preferred static page turning technique was the professional musician, who rarely relies on scores during performance, only using them during rehearsal.

Skills that have been mastered for years are difficult to forget. Professional musicians can be very reluctant to switch to a digital music stand for the performance assistance, as they have been dealing with paper music scores all their lives and are very used to them. It takes a perfectly functioning system to make the musician interested in it. Our experiments show that any minor flaw in the system can make the musician lose interest in the system completely. Due to minor bugs in our software we were not able to receive as much feedback as we were hoping for, so a more thorough experiment is required to obtain sound conclusions about the system.

Bibliography

- [1] BAINBRIDGE, D., AND BELL, T. The Challenge of Optical Music Recognition. *Computers and the Humanities* 35, 2 (2001), 95–121.
- [2] BELL, B., KLEBAN, J., KUCHERA-MORIN, J., OVERHOLT, D., PUTNAM, L., AND STURM, B. THOMPSON, J. Interactive musical performance using the multimodal music stand. *NSF Integrative Graduate Education and Research Traineeship annual meeting* (January 2007).
- [3] BELL, T., BLIZZARD, D., GREEN, R., AND BAINBRIDGE, D. Design of a digital music stand. In *ISMIR* (2005), pp. 430–433.
- [4] BELL, T., CHURCH, A., MCPHERSON, J., AND BAINBRIDGE, D. Page turning and image size in digital music stand. In *International Computer Music Conference* (2005).
- [5] BELLINI, P., FIORAVANTI, F., AND NESI, P. Managing music in orchestras. *Computer* 32, 9 (1999), 26–34.
- [6] BELLINI, P., NESI, P., AND SPINU, M. Cooperative visual manipulation of music notation. *ACM Transactions on Computer-Human Interaction (TOCHI)* 9, 3 (2002), 194–237.
- [7] BORCHERS, J., HADJAKOS, A., AND MÜHLHÄUSER, M. MICON a music stand for interactive conducting. *Proceedings of the 2006 conference on New interfaces for musical expression* (2006), 254–259.
- [8] BORCHERS, J., SAMMINGER, W., AND MUHLHAUSER, M. Engineering a realistic real-time conducting system for the audio/video rendering of a real orchestra. *Multimedia Software Engineering, 2002. Proceedings. Fourth International Symposium on* (2002), 352–362.
- [9] CHU, Y., WITTEN, I., LOBB, R., AND BAINBRIDGE, D. How to turn the page. *Proceedings of the third ACM/IEEE-CS joint conference on Digital libraries* (2003), 186–188.
- [10] DANNENBERG, R. B. New interfaces for popular music performance. In *NIME '07: Proceedings of the 7th international conference on New interfaces for musical expression* (New York, NY, USA, 2007), ACM, pp. 130–135.
- [11] DILLON, A. Reading from paper versus screens: a critical review of the empirical literature. *Ergonomics* 35, 10 (1992), 1297–1326.
- [12] FORSBERG, A., DIETERICH, M., AND ZELEZNIK, R. The music notepad. *Proceedings of the 11th annual ACM symposium on User interface software and technology* (1998), 203–210.
- [13] GRAEFE, C., WAHILA, D., MAGUIRE, J., AND DASNA, O. Designing the muse: A digital music stand for the symphony musician. *CHI* 96 (1996), 436–441.
- [14] HERTZUM, M. Six roles of documents in professionals’ work. *ECSCW 99* (1999), 41–60.
- [15] HIBBARD, W., LEVKOWITZ, H., HASWELL, J., RHEINGANS, P., AND SCHROEDER, F. Interaction in Perceptually-Based Visualization. *Perceptual Issues in Visualization, IFIP Series on Computer Graphics* (1995), 23–32.
- [16] ISHII, H. Reflections: “The last farewell”: traces of physical presence. *interactions* 5, 4 (1998).

- [17] ISTANCE, H., SPINNER, C., AND HOWARTH, P. Providing motor impaired users with access to standard Graphical User Interface (GUI) software via eye-based interaction. *Proceedings of the 1st European Conference on Disability, Virtual Reality and Associated Technologies (ECDVRAT'96)* (1996), 109–116.
- [18] KLEIN, C., AND BEDERSON, B. B. Benefits of animated scrolling. In *CHI '05: CHI '05 extended abstracts on Human factors in computing systems* (New York, NY, USA, 2005), ACM Press, pp. 1965–1968.
- [19] KOSAKAYA, J., TAKII, Y., KIZAKI, M., ESASHI, A., AND KIRYU, T. Research and evaluation of a performer-friendly electronic music stand. In *Proceedings of the 2005 International Conference on Active Media Technology* (May 2005), pp. 11–15.
- [20] LOBB, R., BELL, T., AND BAINBRIDGE, D. Fast capture of sheet music for an agile digital music library. *International Computer Music Conference* (2005).
- [21] MANKOFF, J., DEY, A., BATRA, U., AND MOORE, M. Web accessibility for low bandwidth input. *Proceedings of the fifth international ACM conference on Assistive technologies* (2002), 17–24.
- [22] MCPHERSON, J. Page turning — score automation for musicians. *Honours project report, Department of Computer Science, University of Canterbury, Christchurch, NZ* (1999).
- [23] O'HARA, K., AND SELLEN, A. A comparison of reading paper and on-line documents. *Proceedings of the SIGCHI conference on Human factors in computing systems* (1997), 335–342.
- [24] ORIO, N., LEMOUTON, S., AND SCHWARZ, D. Score following: state of the art and new developments. In *NIME '03: Proceedings of the 2003 conference on New interfaces for musical expression* (Singapore, Singapore, 2003), National University of Singapore, pp. 36–41.
- [25] PICKING, R. Reading music from screens vs paper. *Behaviour & Information Technology* 16, 2 (1997), 72–78.
- [26] PUCKETTE, M., AND LIPPE, C. Score Following in Practice. *Proceedings of the ICMC* (1992), 182–185.
- [27] SIMPSON, R., AND KOESTER, H. Adaptive one-switch row-column scanning. *IEEE Transactions on Rehabilitation Engineering* 7, 4 (1999), 464–473.
- [28] SUNDSTEDT, V., DEBATTISTA, K., LONGHURST, P., CHALMERS, A., AND TROSCIANKO, T. Visual attention for efficient high-fidelity graphics. *Spring Conference on Computer Graphics (SCCG 2005)* (2005), 162–168.
- [29] WINGET, M. Heroic Frogs Save the Bow: Performing Musician's Annotation and Interaction Behavior with Written Music. *Proceedings of the 7th International Conference on Music Information Retrieval - ISMIR 2006* (October 2006).

A Appendix

A.1 Post-experiment questionnaire given to the participants

1. Have you used a computer-based music stand before?

Yes No

2. If yes, was it using static or animated page-turning?

Static Animated Both

3. Did you find it difficult to read music from the computer screen?

1 2 3 4 5
No, not at all Yes, somewhat Yes, very

4. How usable did you find static page turning?

1 2 3 4 5
Very unusable Neutral Very usable

If you want to, explain why:

5. How usable did you find the half-page preview technique?

1 2 3 4 5
Very unusable Neutral Very usable

Why?

6. Please rate the three page-turning methods on usability

Page wipe (simple line)

1 2 3 4 5
Very unusable Neutral Very usable

Page turn

1 2 3 4 5
Very unusable Neutral Very usable

Page roll

1 2 3 4 5
Very unusable Neutral Very usable

7. Did you find controlling the music stand in general distracting?

1 2 3 4 5
 Not at all Somewhat Quite distracting

easy to use?

1 2 3 4 5
 Not at all Somewhat Yes, very

prone to errors?

1 2 3 4 5
 Not at all Somewhat Yes, very

8. Please rate the three control methods on ease of use

One pedal

1 2 3 4 5
 Easy Neutral Difficult

Two pedals

1 2 3 4 5
 Easy Neutral Difficult

One pedal with menu

1 2 3 4 5
 Easy Neutral Difficult

9. Which of the following page-turning settings did you find works best for you?

Constant animation speed

Constant animation speed stopping half-way through the turn

Faster first half animation, slower second half animation

Slower first half animation, faster second half animation

Why?

10. How would you improve Espresso music stand?