Elective Music Students Experiences with Jam2Jam

Andrew Johnston
Creativity and Cognition Studios
School of Software
University of Technology Sydney
Australia
andrew.johnston@uts.edu.au

James Humberstone
University of New South Wales
Sydney, NSW
Australia
aftertrace@me.com

ABSTRACT

This paper presents findings from a trial of the interactive music software Jam2Jam in a classroom music setting. Jam2Jam is software which allows musical novices to control generative music in real time. It has an interface which enables users to control multiple audio-visual parameters with a single gesture – an approach intended to facilitate complex, conversational interaction. Examination of students experiences with Jam2Jam indicates that students find Jam2Jam attractive and that it has considerable potential. However, a number of issues for improvement, particularly a need for increased transparency of operation are identified. Extensions to Jam2Jam which would enable students to incorporate more of their own material into the music and visual they create during jam sessions are also proposed.

Keywords

Music, performance, Jam2Jam

1. INTRODUCTION

This paper presents findings from a trial of Jam2Jam conducted at the MLC School, a private girls school located in suburban Sydney, Australia. We examine students' experiences with Jam2Jam in order to evaluate how effective it was in its current form and identify what changes could be made in order to improve it. We were particularly interested in identifying which aspects of Jam2Jam's design helped students engage creatively with the software, and in identifying barriers which prevented students from fully engaging with Jam2Jam and unlocking its potential as a teaching and performance tool.

2. CONTEXT

MLC School is a large private girls school in Sydney, Australia. Our study involved year 9 and 10 students (14–15 year olds) who participated in the elective music program. Thus, the students were musically experienced, able to read and write music and play acoustic instruments.

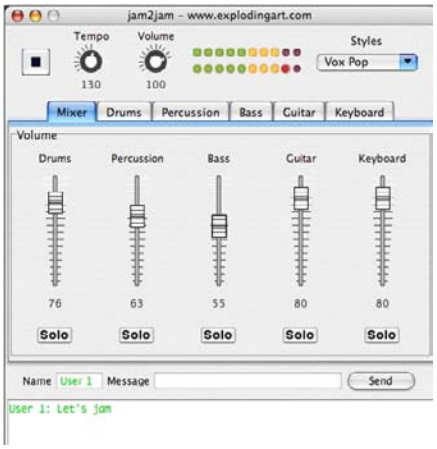

Figure 1: Jam2Jam Gray screenshot.

2.1 Jam2Jam

Jam2Jam is software which allows musical novices to control generative music in real time [3]. The music is generative in the sense that the computer uses algorithms to compose music based on parameters controlled by the user. Thus, as well as adjusting the volume or mix of instruments that are playing, the user is also able to change the parameters that affect the automated compositional processes applied by the computer.

Basic parameters that can be manipulated include the volume of individual instruments in the mix and the overall tempo of the music. In addition, compositional parameters for each instrument include density (number of notes or beats per time unit), timbre, pitch and articulation style. Each of these parameters is adjusted using on-screen controls.

There are two distinct versions of Jam2Jam: Jam2Jam Gray and Jam2JamAV. In this paper we focus on Jam2Jam AV. When we refer to 'Jam2Jam' this is the version we are referring to. The older version will be explicitly referred to as 'Jam2Jam Gray'.

Jam2Jam Gray uses a more basic user interface based on sliders (figure 1). Users switch between controlling the mix of individual instruments using a standard mixing desk metaphor and controlling the generative music parameters by selecting on-screen tabs.

Figure 2: Jam2Jam AV screenshot.

Jam2Jam AV, on the other hand, provides support for generative video as well as audio and uses a more complex interface (figure 2). In addition to controlling the audio and compositional parameters described above, users are able to mix videos based on pre-recorded video files or live video from a webcam, and control a range of video effects and settings such as kaleidoscope, brightness, ghosting, frame rate and saturation.

Of particular interest is that the interface for Jam2Jam AV is designed to enable users to control multiple audio-visual parameters with a single gesture. As figure 2 shows, icons surround the screen. Each of these is a 'parameter selection' icon which corresponds to an audio or visual parameter that can be adjusted by the user. Audio parameter selection icons (instrument volume, density, etc) are arranged on the left and bottom of the screen and visual parameter selection icons (saturation, frame rate, etc) are on the right and top of the screen. Users are able to select parameters for adjustment by clicking on the icons.

In the centre of the screen are five 'parameter adjustment' icons. Four of these represent instrumental roles in the sound mix: drums, bass, lead (guitar icon) and chords (keyboard icon). Moving these icons around the space in the middle of the screen adjusts whichever parameters have been selected using the parameter selection icons around the screen edge. The fifth icon (webcam) is used exclusively to adjust the selected video parameters.

For example, if the user selects the volume icon on the left of screen. then moving the drum kit icon up and down will adjust the volume of the rhythm track, moving the bass guitar icon will adjust the volume of bass in the mix, etc. If the volume parameter alone is selected for adjustment in this way, then Jam2Jam AV is effectively working like a basic mixing desk - moving instrument icons up and down changes their volume in the music being generated.

Where things get more complex is when more than one parameter adjustment icon is selected. It is possible, for example, for the user to select the volume parameter selection icon on the left of screen *and* the timbre parameter selection icon on the bottom of screen. This means that moving an instrument icon up and down will adjust its volume and moving it left and right will adjust its timbre.

This feature enables the two dimensional gestures made with the mouse to be 'cross coupled': one gesture affects two parameters. Cross coupling is a key feature of acoustic instruments - increasing volume on a brass instrument also affects timbre and pitch, for example - but, because of the separation of control interface and sound production inherent in electronic music, it is often not present in computer based instruments. Experiments, most notably those conducted by Hunt [6, 5], have shown that cross coupling is an effective strategy for enhancing the expressiveness of computer-based musical instruments.

Cross coupling is a strategy to enable musical interactions which are 'conversational' [7, 8]. In conversational interactions the user and computer share control over the music in much the same way that musicians in a more traditional ensemble do. While the user may at times seek to explicitly direct the performance, at other times they will surrender control and let the computer shift the performance in new, possibly unexpected, musical directions. Past studies involving professional musicians have indicated that if a computer's musical response to a given input is completely predictable it is unlikely to trigger conversational interactions [7]. Obviously though, a computer response which has no apparent relationship to the actions of the user is equally problematic. The key to effective conversational interactions is finding a balance between controllability and complexity.

Jam2Jam provides additional scope for conversational interaction by allowing multiple users to share control of the interface. When two users are running Jam2Jam on a network, they are able to link together, which enables joint control of the music and visuals. When one user moves an icon on their screen the corresponding icon moves on the other user's screen and vice versa. Thus conversational interaction can occur between each user and Jam2Jam, and also between users *through* Jam2Jam.

3. METHOD

This paper is primarily focused on examining students' experiences with Jam2Jam. This is part of a larger program of creative work and research in the area of designing for musical expression that the first author has been conducting for several years. As part of this work, a research framework aimed at linking design and evaluation has been developed. In this section we will describe this framework in order to place the Jam2Jam evaluations in context.

At a high level our approach draws on action research and design science, in that we attempt to improve understanding of the nature of musical expression by actively developing strategies and techniques to support it. By carefully examining the impact of our work in real-world contexts we are able refine our designs but also, and perhaps more importantly in the long run, improve understanding of musical expression more broadly.

A simplified view of an action research approach is shown in Figure 1. This very simple process involves drawing on literature and past experience to develop a theory and plan of action. This plan of action is implemented and the effects

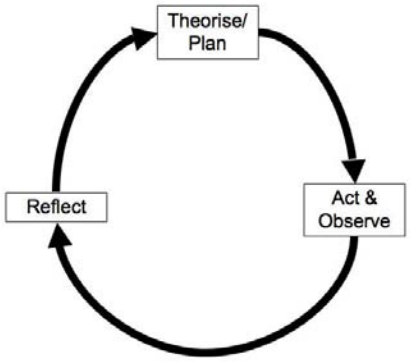

Figure 3: The Action Research cycle

carefully observed. Reflection on the impact of the actions taken lead to refinement of the theory and development of further plans informed by past actions.

Thus, in the case of Jam2Jam, theories of musical expression, interaction design and software development lead to the development of software which is in a sense an embodiment of these theories. By carefully observing Jam2Jam in use we are able to reflect and draw conclusions that lead us to refine our theories of creative interaction and improve the design.

Other authors have provided a rationale for the design of Jam2Jam [3, 4]. Our focus in this paper is on examining its impact in a particular context and reflecting on the implications the experiences of particular users have for its future development.

3.1 Human Computer Interaction

Several authors have recognised the potential of human-computer interaction (HCI) techniques to investigate the experiences of performers who use musical interfaces. In general, the approach has been to use quantitative techniques from HCI which tend to equate interface effectiveness with efficiency. Wanderley and Orio [13], for example, propose a series of "musical tasks" which might be used in order to evaluate how effectively an input device can support expressive performance. These tasks are intended to create a kind of benchmark which will make it easier to compare one interface device with another. The intention is that these benchmark figures, derived as they are from formal studies of users doing prescribed musical tasks, might complement traditional technical measures of device capabilities such as output rate and precision.

This would certainly be worthwhile. However, this approach is very much focussed on the devices and their ability to efficiently translate the intentions of the user into parameters for the computer. The experiences of the users who use the devices, being hard to quantify, are comparatively neglected.

Kiefer, et al [9] draw on Wanderley and Orio's task-based approach but extend it by gathering qualitative data during interviews with study participants. They found that analysis of the interview data provided significant insights, surfacing a number of unexpected issues and suggestions from users for alternate uses of the controllers they were evaluating. Of particular interest is their observation that the findings based on quantitative data "seemed to be a limited measure of the device compared to the subtlety of the participants' observations".

Stowell, et al [12] present a qualitative evaluation method based on Discourse Analysis which they applied to a voice-based interface they developed. Unlike the more restricted task-based approach, the participants in their study were not given specific musical tasks to complete but were instead encouraged to explore the interface in their own way, at least initially. Following the free exploration, participants were played some example recordings created using the interface and asked to use these as inspiration for creating their own work. Finally, a semi-structured interview was conducted in which video recordings of the participant's session with the interface were sometimes used to stimulate discussion. Of particular interest was the use of discourse analysis techniques to analyse the interview transcripts. The use of structured qualitative analysis methods such as these is one way to maintain rigour without compromising relevance by forcing study participants to only perform actions which are easily measurable.

The work described by Kiefer et al and Stowell is significant because it broadens the scope of what constitutes 'evaluation' in this context, recognising that while ergonomics and efficiency are important, they are not the primary determinants of the quality of a musical interface. This thinking is reflected in the broader field of HCI, where there has been recognition that the task-based approach alone is inadequate, particularly when considering software intended to support creative work. A number of HCI researchers therefore have turned their attention to the 'user experience' (eg. [1, 10]).

Just how to approach this is a difficult issue. As we have discussed, 'traditional' approaches have focused on measuring user performance when carrying out various well-defined tasks such as navigating a web-site or entering figures into a spreadsheet. Software designed to facilitate musical expression presents a problem in this context as it is difficult to formulate tasks to assign to users which are measurable but also meaningful [13]. If the aim was to produce a general-purpose musical instrument for performing traditional music, then evaluation would be simpler. Tasks such as playing a scale, trilling, etc. could be assigned and measurements made to ascertain how successfully users were able to execute them. The benefit of this approach is that it would be possible to somewhat objectively compare two different musical instruments in terms of this restricted definition of playability. However, where a new 'instrument' is intended to create new and unusual sounds - to explore new languages of composition and performance - this approach is problematic. Part of the rationale for creating new musical interfaces is that they disrupt performers' ways of thinking about music so that they are stimulated to try new ways of playing and composing. Attempting to determine how effectively they enable performance of current styles of music might be

interesting, but it would not facilitate learning about how to design instruments which encourage divergent thinking.

Researchers in the broader field of Human-Computer Interaction have recognised the limitations of task-focussed approaches and are proposing new ways of thinking about 'evaluation' in the context of systems which have uses that are open to a range of interpretations. Sengers and Gaver [11], for example, argue that interaction designers are becoming less concerned with designing software which unambiguously conveys and supports a clearly defined 'purpose'. They propose that HCI needs to support interactions in which users may have multiple interpretations of what a system is for and how it works. 'Evaluation' in this context goes beyond identifying whether users' interpretations of a system's purpose and behaviour matches the designer's anticipated interpretation. Rather, "evaluation shifts from determining whether an authoritative interpretation was successfully communicated to identifying, coordinating, stimulating, and analyzing processes of (evaluative) interpretation in practice" [11], p.105.

3.2 Approach

Our approach has been to focus on two key areas:

1. Examination of students' experiences with Jam2Jam in the context of the meaningful engagement matrix [4] and modes of interaction [7, 8].
2. Identifying opportunities for improvement in the design of Jam2Jam.

In order to study students' experiences with Jam2Jam, we conducted a series of interviews with students as they used it. The author, equipped with a small digital video camera, moved about the classroom observing (and recording) the students as they explored Jam2Jam. Students thoughts and opinions were actively solicited during this time. In general, open questions were used to encourage students to verbalise their experiences and to help reduce the effects of interviewer bias. However, the aim of the study was to get as rich a picture as possible of the students' experiences with Jam2Jam and this was prioritised over consistency of procedure. In essence, the process was more akin to a user dialog than usability testing [2].

In addition, an interview with the classroom music teacher, James Humberstone, was conducted after the students' sessions with Jam2Jam. The aim was to get additional perspective on the use of Jam2Jam in the music program at MLC and to identify areas for improvement.

The software Transana [15] was used to facilitate analysis of the 4.5 hours of video that was gathered. Transana is open source software for conducting qualitative analysis of video and audio data. In particular, it facilitates the annotation of video with keywords and comments as well as tools for grouping related video clips together based on these keywords. This enables the researcher to build hypothesis and gather evidence through detailed examination of both the verbal comments and responses of the musicians and their behaviour while using the software.

4. FINDINGS

4.1 Usability Issues

Students were initially asked to simply play with Jam2Jam and see what they could do with it. During this session the researcher moved about the room with a small handheld video camera, observing what the students were doing and asking students about their experiences.

Students were attracted to Jam2Jam and understood the underlying idea that it allowed them to control various musical properties by moving icons around the screen. The fact that the selection icons around the outside of the screen determined which audio/visual parameters the moving icons adjusted was intuitively understood by most students.

However, beyond this, several characteristics of the interface prevented them from really grasping the details of specifically how to establish control over what Jam2Jam was doing. The main aspects which inhibited students' ability to establish control over Jam2Jam were:

- Lack of clarity of whether icons around the outside of the screen were selected or de-selected.
- Difficulty recognising the internal state of Jam2Jam. Changing some parameters had immediately perceivable impact on either sounds or visuals but others were more subtle and hard to pick up.
- The link between the icons around the outside of the screen and the movement of the movable icons was unclear. Students at times selected an icon on the bottom of the screen for example and then moved a movable icon in the vertical direction. Intuitively they felt it should have an effect but because there was no horizontal movement the music produced by Jam2Jam was unaffected.
- Because they didn't realise which musical parameters the parameter selection icons referred to - and were not always able to identify this through experimentation - students needed guidance (from the teacher or from the Jam2Jam documentation) to get started. Comments from students included: "With the manual it makes it completely understandable..." and, "I realised that if you moved them around it would change but I didn't know why."

Mostly, these issues are easily addressed and do not relate to the core functionality of Jam2Jam. Subsequent releases of Jam2Jam for example have improved the visual design to some degree so that it is now much clearer whether icons are selected or deselected.

Making the state of Jam2Jam more readily perceivable is a design challenge relating to both the visual design and the algorithms used to manipulate the generated audio. One aspect of Jam2Jam's interface which caused recurring confusion with students, particularly in the early stages, was that the link between the icons around the outside of the screen and the movement of the movable icons was unclear. Students at times selected an icon on the bottom of the screen

for example and then moved a movable icon in the vertical direction. Intuitively they felt it should have an effect but because there was no horizontal movement the music produced by Jam2Jam was unaffected.

This comparatively minor problem is compounded though when movement in the *correct* direction doesn't result in audible change. This is usually because the instrument that has been selected is not sounding at the time. For example, if the keyboard icon moved in the horizontal direction while the 'timbre' parameter is selected on the bottom row, the timbre of the keyboard sound will normally be altered. However, if the keyboard is not sounding at the time then moving this icon will appear to have no effect – at least until the keyboard begins playing again. It should be noted that while students were often confused by Jam2Jam in the early stages, they remained interested and did not appear to become frustrated:

> Researcher: So, before you figured that out, were you getting frustrated because you can't figure [it] out?
> Student 1: No.
> Student 2: We were just enjoying the sounds, dragging them around. We knew that they were changing.

While these students were unable to articulate exactly what was happening and how their actions were linked to what Jam2Jam was producing, they could hear that there was a link and were happy to enjoy them without necessarily needing to control them.

4.2 Engagement

Dillon et al [4] categorise users' interactions with Jam2Jam into a set of categories they call the 'meaningful engagement matrix'. This matrix is made up of five different 'modes of creative engagement' which are linked to three aspects of musical meaning. The modes of creative engagement are:

- "Appreciating – listening carefully to music and analysing music representations.
- Selecting – making decisions about musical value and relationships
- Directing – managing music making activities
- Exploring – searching through musical possibilities and assessing their value
- Intuiting – participating in intuitive music making." [4], p. 6

The aspects of musical meaning which complete the matrix relate the actions of the user to the socio-cultural context within which they take place. As the students in our study primarily interacted with Jam2Jam as a stand-alone (rather than networked) instrument, we did not examine these social aspects of their jamming in detail.

We did, however, note that the students primarily interacted with Jam2Jam AV in either the 'appreciating' or 'exploring' mode. We did not see students move into the other modes, primarily because they failed to establish a sufficient degree of control over the performance. Dillon et al [4] describe engagement with Jam2Jam Grey as being primarily in 'exploring' and 'appreciating' (see p.6). They described more advanced users as moving into 'selecting' and 'directing' modes at times. The students in our study (using the newer Jam2Jam interface) didn't appear to make this jump.

The comments by this year 9 student illustrate how a lack of instrumental control over Jam2Jam may have contributed to this:

> "'Just then we actually changed the music. Which was the first time that happened. Not just speeding it up or slowing it down - the music actually changed. And that was the breakthrough. But I don't know how we did it!"

Likewise, Dillon et al [4] describe students using Jam2Jam Gray gaining insights into styles by noticing for example that hip hop slowed down sounds like reggae. The newer Jam2Jam AV interface did not seem to provide this same insight as the changes made to the interface seemed to have more complex, unpredictable effects. Because they could not establish clear links between their actions and the resulting musical changes it was hard for them to gain musical insights of this type.

Interestingly, Jam2Jam Gray was trialed by a few students who found it installed on the lab machines. They found that while less immediately visually appealing, it was simpler to understand and manipulate:

> "...It is a lot simpler, because you got all the things in one. Because you can see which way the things are going."

One student in particular articulated a key difference between Jam2Jam AV and Jam2Jam Gray:

> "It depends like what you want to do with the software. Like if you just want to play with things I guess this [Jam2Jam AV] is alright. But like if you actually want to like perfect music with it then you should do the sliding bars [Jam2Jam Grey]. Because that would be a lot easier."

This student identifies how the playfulness of Jam2Jam AV comes at the expense of controllability. The transparency of Jam2Jam Gray tends to encourage an instrumental (or 'directing') approach in which the music is 'perfected'. The lack of control she felt with Jam2Jam AV tended to lead her to a more 'playful' approach.

At the end of one of the classes some of the students discussed which 'audience' Jam2Jam AV would be best suited

to. One student suggested that it would be best targeted at "Young kids...before (grade 4 to 5)". She felt that the colourful graphics and video would be appealing to this age group. However, she felt that "older kids" (their age), would find a program such as GarageBand more appealing because "...it has a lot more features in it. Like you can mess around with more sounds". This student also suggested that being able to load in her own songs and then use the Jam2Jam controls to remix and edit in real time would make the software more appealing.

Another student disagreed that Jam2Jam AV was best targeted at young kids "...because to understand how it works...you need to like have a basic understanding of music in general." She felt that understanding concepts such as density and articulation would require a higher level of musical skill. She felt that the visuals were appealing to both older and younger students but that younger students would lack the necessary musical understanding to meaningfully control the sound and musical parameters that Jam2Jam users control.

4.3 Design Issues and Suggestions

A number of issues and suggestions for further refinement of Jam2Jam AV emerged from the sessions. The most often repeated comment was that students would like to be able to upload their own music and videos into Jam2Jam and use its features to manipulate and vary that music. While they didn't mind the tunes generated by Jam2Jam, they felt that being able to provide their own music would enhance their engagement with the software. Several students mentioned that software such as GarageBand was more enjoyable for them because it provided them with features for manipulating their own material and thus provided more variety.

> Researcher: So in which ways is Garageband better do you think? What do you enjoy more about it?
> Student: The variety of sounds. Like the variety of like clips and you can upload your own like... little sound samples into it.

Later, this student went on to say that if Jam2Jam was able to be made more complex it would be more appealing to older students. She thought Jam2Jam in its current form was appealing particularly to younger children, but that it should be able to be reconfigured to provide greater complexity for more advanced students.

> "If you like added complexity to it, so in like separate files, so you can do a really complex version, and like give a simple version to kids."

We see this as a request for a musical application which has a "low 'entry fee' with no ceiling on virtuosity" [14]. Jam2Jam has many features which attract students. They like the idea of mixing music and visuals in real-time, but they didn't seem to make the step to using Jam2Jam as an expressive instrument.

Because of bugs (addressed in subsequent releases), the students experience of networked jamming was limited. Even when students did manage to connect Jam2Jam with other students in the room they found it frustrating that they were unable to determine who they were actually jamming with. If Jam2Jam were able to make the identity of the networked jamming partner more visible it could help students focus on collaborative jamming. The following quote illustrates a situation which often arose:

> Student: We changed the background as well but the actual melody seemed to go to a variation or something. Something we did changed it.
> Researcher: But you're not exactly sure what it was that changed it?
> Student: No I'm still not sure what's changing mine.
> Researcher: Is that frustrating for you?
> Student: Um, a little bit. Yeah because sometimes I don't know if I'm clicking it or if someone else is clicking it, cause I don't know who I'm with. If I knew who I was with it might be a bit more helpful.

This student went on to suggest that the webcam could be used to put other jammers' faces on or next to the icons they were moving that this would help make the behaviour of Jam2Jam more understandable.

4.4 Teachers' Perspective

After the sessions with the students, the author and the MLC instrumental music teacher, James Humberstone, met to reflect on what had occurred and discuss possible future directions for Jam2Jam. This discussion was video recorded and analysed with the help of Transana. In this section I will detail these ideas, several of which have already begun to be incorporated into more recent versions of Jam2Jam.

James' key concern was that Jam2Jam did not provide musical experiences which were complex enough to justify its inclusion in the musical curriculum at MLC beyond one-off special classes. He wondered whether a way could be found to allow students to create their own customised musical content for Jam2Jam as he sees this as a way to more fully engage students both compositionally and instrumentally. He believes this could be a way to integrate repertoire being studied by the class into Jam2Jam sessions. Using Jam2Jam to explore this repertoire could enhance student engagement with this material by using the Jam2Jam controls and interface as a vehicle to explore it in new ways. In order to achieve this though, Jam2Jam needs to provide better support for instrumental improvisational performance.

> "To me I really want to take the element of improvisation and performance much further with the software. I feel that we'd certainly exhausted everything that the year 9's and 10's could do with it – and fair enough, they're quite old, musically literate kids. But...even for younger children, there's only a certain amount that you can do with the sliders that are given and the variation of audio and video. So I would really like to

> see future tools that would allow a higher level of improvisational control."

James felt that earlier versions of Jam2Jam (Jam2Jam Gray) were better in this respect. In the classes, several students found this application and felt that it was easier to feel a direct link between their actions and the sounds produced by the software.

> "And in the earlier version of Jam2Jam, which wasn't half as, um, attractive as this one. But you could do things like change the chord progression which was going on. So that gives you quite big control."

James went on to say that it was critical for students to feel ownership of the music they were creating. The move (in more recent versions of Jam2Jam AV) to allow students to bring their own music into Jam2Jam is a very positive step in this direction. Improved support for instrumental control would also help students feel ownership of their real-time performances as well as the compositions being used as source material.

> "I think the students have to be able to feel some ownership of it – at any age at some point. And as I say it's very engaging as it is and you know, you make the changes with the sliders. But to go to the next level, to make them want to use it again and again and to integrate it more completely into an ongoing music curriculum I think it has to have that next level [of support for bringing content into Jam2Jam]."

Another way to add depth to the students' experience with Jam2Jam would be to allow students to see the music being generated in traditional music notation. James argues that this would allow musically literate students to more readily perceive what Jam2Jam is doing:

> "For students that are musically literate it would enable them to analyse more closely the changes that they are making. As well as hearing them - and we've already said sometimes those changes aren't always easy to hear - they'd be able to see them."

A variation on providing a traditional notation view of the music Jam2Jam is producing could be to provide students with access to the Jam2Jam code that is generating the music. Jam2Jam is essentially an interface which allows users to adjust parameters for generative processes that produce music. If students were able to see, and perhaps modify, the underlying algorithms which generate this music, it would provide them with yet another perspective on music and composition, and add depth to their experience with Jam2Jam.

5. SUMMARY AND CONCLUSION

Based on our observations we have identified some key areas to address in the design of Jam2Jam. First, basic usability needs to be improved to remove unnecessary impediments to full engagement. Previous studies indicate that conversational interaction requires a degree of instrumental control. Performers want and expect a certain degree of autonomy from music software used in live performance - in the sense that its responses should be surprising at times [8, 7]. However, for the interaction to be conversational, these responses need to be clearly related to the input from the user. As we have discussed, the ability of Jam2Jam to cross-couple two audio and visual parameters should encourage conversational interactions by increasing the complexity of the audio-visual response to user input. However, this falls down when users are sometimes unable to perceive whether their input is having any effect at all.

This leads to the second area for improvement: making the internal state of Jam2Jam more evident. As we have pointed out, if students adjust parameters of an instrument that is not sounding at the time there will be no audible indication that their actions are having any effect. In addition, the only visible indication that they are having an effect is movement of the icon. In order to address this issue either the effect of the user's actions on the parameter must be made more visible (eg. by increasing or decreasing the size of a bar next to the icon) or the state of the icon should indicate whether an instrument is sounding or silent (eg. by altering transparency). That is, when the effect of the parameter change will not be audible, it should either be made visible, or the fact that the parameter change will have no effect should be conveyed.

A related area for improvement is in enhancing the visibility of the connection between Jam2Jam partners. Currently Jam2Jam does not identify the jamming partner in any way. As we have observed, this can lead to frustration because users see their icons being moved about but don't necessarily associate this with musical actions of another person. Simple measures such as overlaying the name (or better, a photo) of the networked jammer on the icon they are currently moving could help enhance the networked jamming experience.

The teacher and students involved in this study are keen to see extensions to Jam2Jam which would enable users to incorporate more of their own material into the music and visuals they create during jam sessions. As we have outlined, Jam2Jam does not play pre-recorded music but instead allows students to control parameters for algorithms which generate the music in real-time. Within this paradigm, students could for example begin to explore the creation of their own algorithms to create unique rhythms or melodies. This would provide more advanced students (such as those in the elective music program we studied) with scope for more in-depth engagement with Jam2Jam and potentially give students new perspectives on compositional processes.

In this paper we have examined the experiences of teacher and students at MLC with Jam2Jam. Our focus has been on how Jam2Jam could be improved and enhanced to provide users with a more expressive and flexible tool for collaborative audio-visual performance. While we have been critical

of some aspects of Jam2Jam, it is important to stress that we see the Jam2Jam in its current form is a great foundation upon which to build. Dillon et al [4] describe Jam2Jam as the "xylophone of computer music" and point out that using computers in this context seems intrinsically attractive to most students. Our observations back this up – the students were interested in Jam2Jam and immediately understood and were attracted to the notion of collaborative audio-visual jamming. We hope that the work we have described here encourages others to explore the use of Jam2Jam in the classroom and will enable those who design new technologies for musical expression to draw on the lessons we have learned.

6. ACKNOWLEDGMENTS

Our thanks to the students of MLC who participated in this study for their enthusiasm and insightful comments.

This research was conducted within the Australasian CRC for Interaction Design (ACID), which is established and supported under the Australian Governments Cooperative Research Centres Program.

7. REFERENCES

[1] M. A. Blythe, K. Overbeeke, A. F. Monk, and P. C. Wright. *Funology: from usability to enjoyment*. Kluwer Academic Publishers, Norwell, MA, USA, 2004.

[2] J. Buur and K. Bagger. Replacing usability testing with user dialogue. *Communications of the ACM*, 42(5):63–66, 1999.

[3] S. C. Dillon. Jam2jam: Networked improvisational musical environments. In M. Moore and B. W. Leung, editors, *School music and teacher education : a global perspective in the new century*, pages 19–30. Dept. of CAPE, the Hong Kong Institute of Education, : ISME International Office, Hong Kong : Nedlands, W.A. :, 2006.

[4] S. C. Dillon, B. A. Adkins, A. R. Brown, and K. L. Hirche. Communities of sound : examining meaningful engagement with generative music making and virtual ensembles. *International Journal of Community Music*, 2(1), 2009.

[5] A. Hunt, M. M. Wanderley, and R. Kirk. Towards a model for instrumental mapping in expert musical interaction. In *Proc. International Computer Music Conference*, 2000.

[6] A. Hunt, M. M. Wanderley, and M. Paradis. The importance of parameter mapping in electronic instrument design. *Journal of New Music Research*, 32(4):p429 – 440, December 2003.

[7] A. Johnston, L. Candy, and E. Edmonds. Designing and evaluating virtual musical instruments: facilitating conversational user interaction. *Design Studies*, 29(6):556–571, 2008.

[8] A. Johnston, L. Candy, and E. Edmonds. Designing for conversational interaction. In *Proceedings of New Interfaces for Musical Expression (NIME)*, 2009.

[9] C. Kiefer, N. Collins, and G. Fitzpatrick. Hci methodology for evaluating musical controllers: A case study. In *New Interfaces for Musical Expression (NIME)*, 2008.

[10] J. McCarthy and P. Wright. *Technology as Experience*. The MIT Press, 2007.

[11] P. Sengers and B. Gaver. Staying open to interpretation: engaging multiple meanings in design and evaluation. In *DIS '06: Proceedings of the 6th conference on Designing Interactive systems*, pages 99–108, New York, NY, USA, 2006. ACM.

[12] D. Stowell, M. D. Plumbley, and N. Bryan-Kinns. Discourse analysis evalutation method for expressive musical interfaces. In *New Interfaces for Musical Expression (NIME)*, 2008.

[13] M. M. Wanderley and N. Orio. Evaluation of input devices for musical expression: Borrowing tools from HCI. *Computer Music Journal*, 26(3):62–76, 2002.

[14] D. Wessel and M. Wright. Problems and prospects for intimate musical control of computers. *Computer Music Journal*, 26(3):11–22, 2002.

[15] D. Woods and C. Fassnacht. Transana v2.22. Madison, WI: The Board of Regents of the University of Wisconsin, 2007.