



The Importance of Parameter Mapping in Electronic Instrument Design

Andy Hunt , Marcelo M. Wanderley & Matthew Paradis

To cite this article: Andy Hunt , Marcelo M. Wanderley & Matthew Paradis (2003) The Importance of Parameter Mapping in Electronic Instrument Design, Journal of New Music Research, 32:4, 429-440

To link to this article: <https://doi.org/10.1076/jnmr.32.4.429.18853>



Published online: 09 Aug 2010.



Submit your article to this journal [↗](#)



Article views: 563



View related articles [↗](#)



Citing articles: 48 View citing articles [↗](#)

The Importance of Parameter Mapping in Electronic Instrument Design

Andy Hunt¹, Marcelo M. Wanderley² and Matthew Paradis³

¹Media Engineering Group, Electronics Department, University of York, Heslington, York, UK; ²Sound Processing and Control Laboratory, Faculty of Music, McGill University, Montreal, Canada; ³Music Technology Group, Music Department, University of York, Heslington, York, UK

Abstract

This paper presents a review of a series of experiments which have contributed towards the understanding of the *mapping* layer in electronic instruments. It challenges the assumption that an electronic instrument consists solely of an interface and a sound generator. It emphasises the importance of the *mapping* between input parameters and sound parameters, and suggests that this can define the very essence of an instrument. The terms involved with mapping are defined, and existing literature reviewed and summarised. A model for understanding the design of such mapping strategies for electronic instruments is put forward, along with a roadmap of ongoing research focussing on the testing and evaluation of such mapping strategies.

1. Introduction: Electronic instruments and the mapping layer

In an acoustic instrument, the playing interface is inherently bound up with the sound source. A violin's string is both part of the control mechanism and the sound generator. Since they are inseparable, the connections between the two are complex, subtle and determined by physical laws. With electronic and computer instruments, the situation is dramatically different. The interface is usually a completely separate piece of equipment from the sound source. This means that the relationship between them has to be defined (see Fig. 1).

The art of connecting these two, traditionally inseparable, components of a real-time musical system (an art known as *mapping*) is not trivial. Indeed this paper hopes to stress that by altering the mapping, even keeping the interface and sound source constant, the entire character of the instrument

is changed. Moreover, the emotional response elicited from the performer is shown to be determined to a great degree by the mapping. Whereas the input devices establish the physicality of the system and the synthesis methods govern the sound quality, the mapping somehow affects *how the player reacts psychologically and musically* to the instrument.

2. The importance of mapping

In this section we emphasise the dramatic effect that the style of mapping can have on “bringing an interface to life”. We focus on our own experience in designing digital musical instruments and comment on several previous designs. An extensive review of the available literature on mapping in computer music has been presented by the authors (Hunt, 2000; Wanderley, 2001) and (Hunt & Wanderley, 2002). A special issue of the journal *Organised Sound* has also been recently devoted to Mapping Strategies in Computer Music (Wanderley, 2002) and we refer the reader to it for a varied review of different approaches to mapping.

2.1 Informal observations

The first author has carried out a number of experiments into mapping. The more formal of these have been presented in detail (Hunt & Kirk, 2000; Hunt, 2000), and are summarised later in this paper. We begin with some rather simpler, previously unpublished, observations that originally sparked interest in this subject. We have retained the first person writing style to denote that these are informal, personal reflections.

Accepted: 26 February, 2003

Correspondence: Andy Hunt, Media Engineering Group, Electronics Department, University of York, Heslington, York YO10 5DD, United Kingdom. Tel.: +44 1904 432375; Fax: +44 1904 432335; E-mail: adh2@york.ac.uk

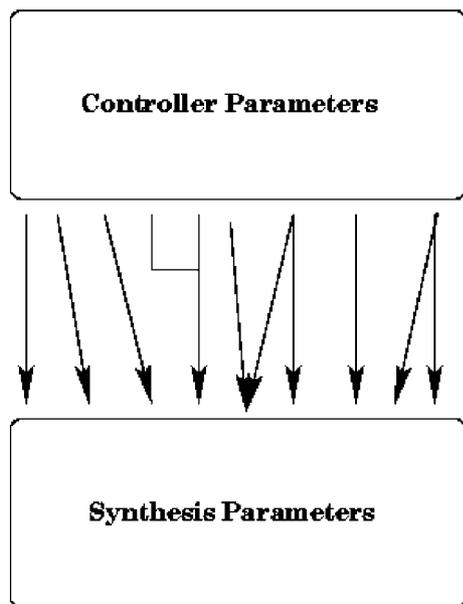


Fig. 1. A generic model of the mapping problem.

2.1.1 The accidental theremin

Several years ago I was invited to test out some final university projects in their prototype form in the lab. One of them was a recreation of a Theremin with modern electronic circuitry. What was particularly unusual about this was that a wiring mistake by a student meant that the “volume” antenna only worked when your hand was moving. In other words the sound was only heard when there was a rate-of-change of position, rather than the traditional position-only control. It was unexpectedly exciting to play. The volume hand needed to keep moving back and forth, rather like bowing an invisible violin. I noted the effect that this had on myself and the other impromptu players in the room. Because of the need to keep moving, it felt as if your own energy was directly responsible for the sound. When you stopped, it stopped. The subtleties of the bowing movement gave a complex texture to the amplitude. We were “hooked.” It took rather a long time to prise each person away from the instrument, as it was so engaging. I returned a week later and noted with disappointment that the “mistake” had been corrected, deleted from the student’s notes, and the traditional form of the instrument implemented.

2.1.2 Two sliders and two sound parameters

The above observation caused me to think about the psychological effect on the human player of “engagement” with an instrument. To investigate this further, I constructed a simple experiment. The interface for this experiment consisted of two sliders on a MIDI module, and the sound source was a single oscillator with amplitude and frequency controls. In the first run of the experiment the mapping was simply one-to-one, i.e., one slider directly controlled the volume, and the other directly controlled the pitch (Fig. 2).

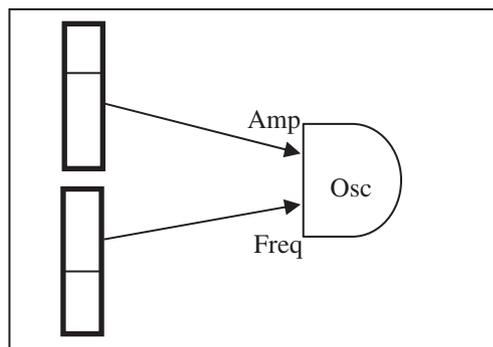


Fig. 2. Simple Mapping for Experiment 1.

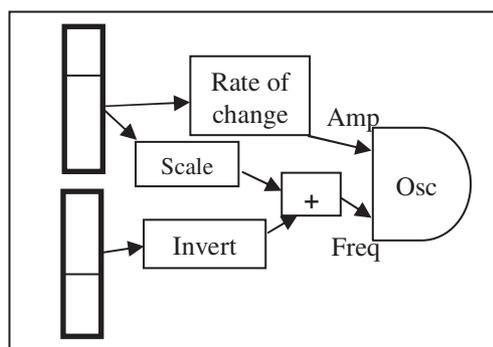


Fig. 3. Complex Mapping for Experiment 2.

I let several test subjects play freely with the instrument, and talked to them afterwards. In the second experimental run, the interface was re-configured to emulate the above-mentioned “accidental Theremin.” One slider needed to be moved in order to make sound; the rate of change of movement controlled the oscillator’s amplitude. But I decided to complicate matters (on purpose!) to study the effect that this had on the users. The pitch, which was mainly controlled by the first slider, operated “upside-down” to most people’s expectations (i.e., pushing the slider up lowered the pitch). In addition the second slider (being moved for amplitude control) was used to mildly offset the pitch – i.e., it was cross-coupled to the first slider (Fig. 3).

A remarkable consistency of reaction was noted over the six volunteers who tried both configurations. With Experiment 1, they all commented within seconds that they had discovered how the instrument worked (almost like giving it a mental “tick”; “yes, this is volume, and this is pitch”). They half-heartedly tried to play something for a maximum of two minutes, before declaring that they had “finished.” Problem solved.

With Experiment 2, again there was a noted consistency of response. At first there were grumbles. “What on earth is this doing?” “Hey – this is affecting the pitch” (implied cries of “unfair,” “foul play”). But they all struggled with it – interestingly for several more minutes than the total time they spent on Experiment 1. After a while, their bodies started to move, as they developed ways of offsetting one slider against

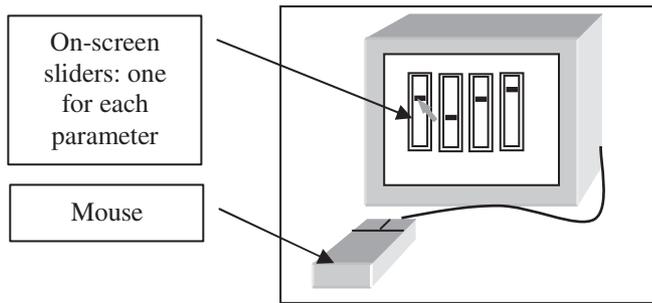


Fig. 4. The “mouse” interface.

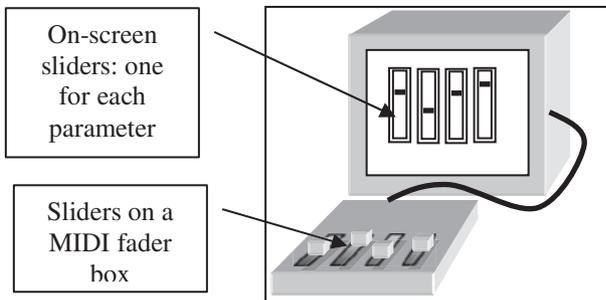


Fig. 5. The “sliders” interface.

the other, while wobbling the first to shape the volume. Nearly all the subjects noted that somehow this was rewarding; it was “like an instrument.” Yet in both cases the interface (two sliders) and the sound source (a single oscillator) were identical. Only the mapping was altered, and this had a psychological effect on the players.

3. Mapping experiments

Several formal investigations have been carried out by the authors in order to explore the essence and the effect of this mysterious mapping layer.

3.1 Complex mapping for arbitrary interfaces

The first author carried out an investigation into the psychology and practicality of various interfaces for real-time musical performance. The main part of this study took the form of a major series of experiments to determine the effect that interface configuration had on the quality and accuracy of a human player’s performance. The full thesis is available for download online (Hunt, 2000) and the details of the theory, experiments and results have been published (Hunt & Kirk, 2000). They are summarised here, in order to give an overview of their implications for mapping strategies.

Three interfaces were used, and these are now described. The first interface (Fig. 4) represented a typical computer music editing interface with on-screen sliders connected one-to-one to each sound parameter.

The second (Fig. 5) involved physical sliders (on a MIDI module) again connected in a one-to-one manner to the synthesis unit.

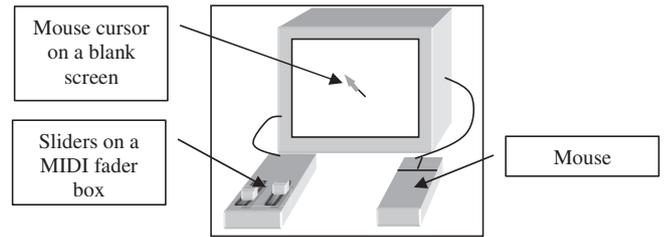


Fig. 6. The “multi-parametric” interface.

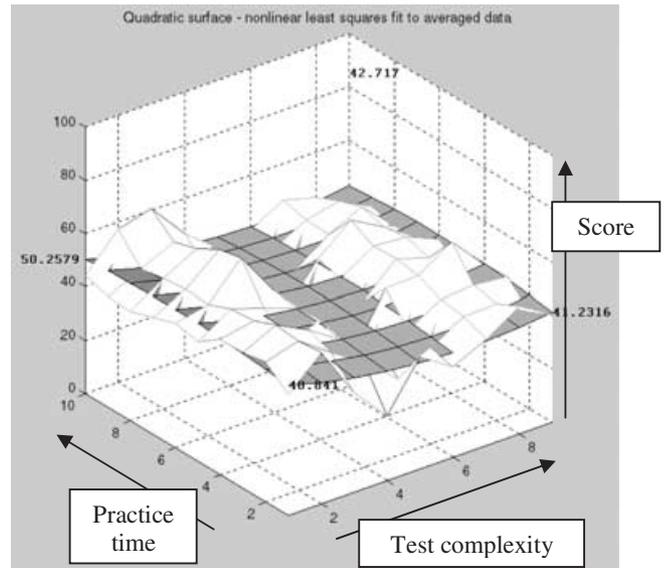


Fig. 7. Results for the “mouse” interface.

The third interface (Fig. 6) consisted of a series of multi-parametric cross-mappings, and, like the accidental Theremin mentioned above, required constant movement from the user to produce sound.

Users attempted to copy (using each interface) a series of sounds produced by the computer. The accuracy of reproduction was recorded for each user, over several attempts, spread out over a number of weeks. Results were gathered numerically, and plotted on a series of graphs to compare the effect, over time, of each interface.

The graphs are in the form of three-dimensional plots which summarise the results for the three interfaces, by averaging together the responses from all users. What results is a representation of the overall performance trend on each interface—plotted over time, and against test complexity.

The test scores are generated by a human marker (moderated by another marker) and are plotted on the vertical axis, with a “perfect” score of 100% being at the top. The axis labelled Test Complexity represents the test number in increasing order of parameter complexity. The axis labelled Practice Time is the session number and thus represents increasing user contact time with the interface.

The test results from all the human test subjects were averaged into the following three graphs (Figs. 7, 8, and 9).

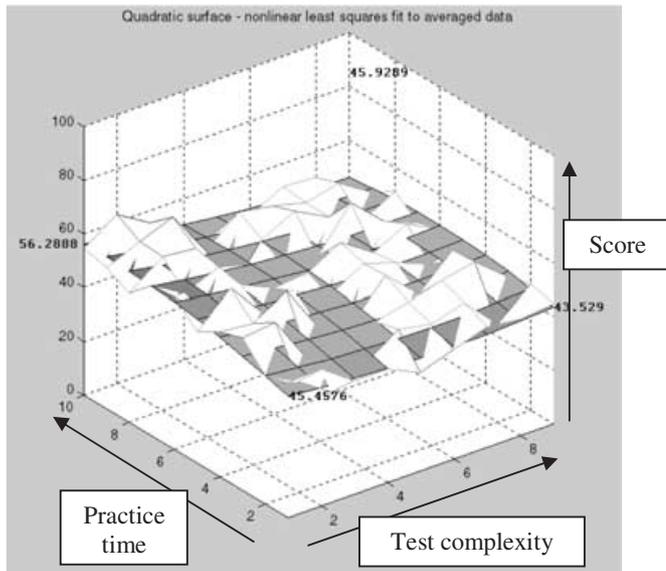


Fig. 8. Results for the “Sliders” interface.

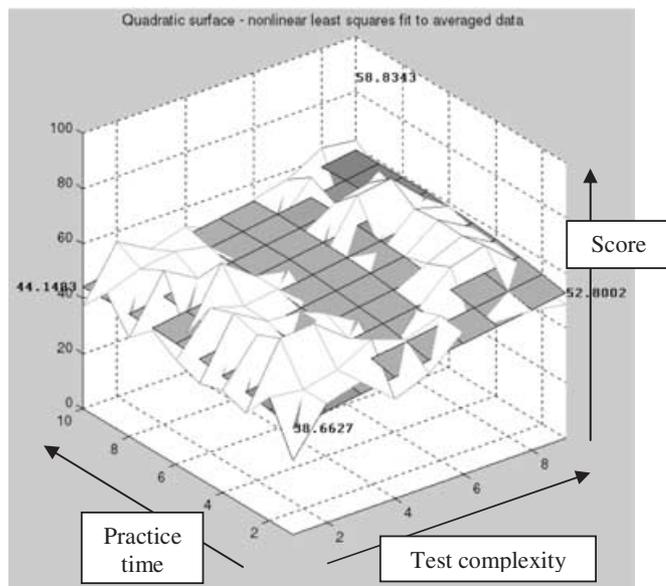


Fig. 9. Results for the “Multiparametric” interface.

Figure 7 shows the first of the 3D plots. This graph encapsulates all the data collected for the *mouse* interface during the longitudinal tests. This graph shows that the mouse does indeed give a reasonably “flat” response over all the tests. There are signs of a very small improvement over time. Note the small upward bend in the plane which indicates that the best scores are for the simplest few tests.

Figure 8 shows a similar plot for the *sliders* interface. The angle of the surface allows some immediate comparisons to be made with the mouse interface plot above.

- For the simplest tests the initial scores are lower for the sliders than for the mouse.

- There is a rapid learning curve for the simpler tests which means that at the final session the score is much higher than for the mouse.
- The sliders perform slightly better than the mouse for the more complex tests.
- The learning curve is only slight for the more complex tests.

This indicates that the “sliders” interface is more of a challenge than the “mouse” to start with, but a significant amount of learning can take place for the simpler tests only.

Figure 9 shows the same type of plot for the multiparametric interface.

This shows a very different response. The angle of the surface indicates that a very different experience occurred for all users with the multiparametric interface. The following points of comparison with the previous two graphs are noted:

- For the simplest test the scores are always lower than those for the mouse or sliders, but they improve over time.
- The scores get better for more complex tests and are much higher than the other two interfaces.
- There is a good improvement over time across all test complexities.

The upward tilt of the plane towards the far-right corner is the most striking feature of the graph. It demonstrates that, on average, the multiparametric interface allows the best performance on the complex tests whilst also allowing a general all-round improvement on tests of all complexities. Neither of the other interfaces had this characteristic.

Studies taken qualitatively (by interviewing each candidate after each test) show clearly that users actively enjoyed this interface, and commented that it was “more like playing an instrument” than operating a conventional computer interface.

The above results can be summarised for the multiparametric interface as follows:

- The test scores in general were much higher than those for the other two interfaces, for all but the simplest tests.
- There was a good improvement over time across all test complexities.
- The scores got better for more complex tests!

This last result may seem rather counter-intuitive at first sight; that people performed better on the harder tasks. However, this brings into question the definition of a “hard task.” If an interface allows the simultaneous control of many parameters, maybe it really is easier to perform the more complex tasks, and harder to accurately isolate individual parameters. A similar finding was presented by Jacob et al. (1994), as explained in section 4 below.

A range of qualitative results was also gathered by interviewing the test subjects to establish their subjective experience of using each interface. They all concluded that the “mouse” interface was the most limited – as they could see how impossible it would be to operate more than one para-

meter simultaneously. Surprisingly perhaps, they were nearly all extremely frustrated by the four physical sliders. Comments abounded such as “I should be able to do this, technically, but I can’t get my mind to split up the sound into these four finger controls.” Some users actually got quite angry with the interface and with themselves. The multi-parametric interface, on the other hand, was warmly received – although not at the very beginning. At first it seemed a difficult interface for most users, but they rapidly warmed to the fact that they could use complex gestural motions to control several simultaneous parameters without having to “de-code” them into individual streams. Many users remarked how “like an instrument” it was, or how “expressive” they felt they could be with it, but that it needed practice.

3.2 Focusing purely on the effect of mapping

In the above experiment several factors may have affected the results. For instance, the multiparametric interface used cross-coupled parameters in addition to the user’s energy. It also reduced reliance on visual feedback, and provided a two-handed input, all of which may have contributed in varying degrees to the interface’s effectiveness.

An additional experiment was subsequently carried out by the third author, designed to focus purely on the mapping layer of electronic musical instruments.

These tests utilised three contrasting mapping strategies, with a fixed user interface (a MIDI slider box) and synthesis algorithm (a simple FM instrument). Users were asked to explore 3 different instruments for as much or as little time as they wished. In each phase of the experiment the interface *looked* exactly the same to the user, however the sonic response to the user’s interaction was different in each case.

The mapping employed for each phase was as follows:

1. A simple one-to-one mapping where input parameters were directly linked to audio controls
2. A one-to-one mapping as in instrument 1 but also requiring the user to input energy to the system in order to make the instrument sound. This was implemented by requiring the user to constantly move one of the sliders in a bowing-like action. Users experienced with traditional instruments would be familiar with this concept of having to constantly inject energy to a system.
3. A many-to-many mapping in which parameters are cross coupled and modified to build complex relationships between certain input gestures and audio output. This mapping also required the user’s energy as in 2.

These mappings allowed the user to control the parameters of a stereo FM synthesis algorithm, including amplitude, frequency, panning, modulation ratio and modulation index.

In each phase of the experiment, users were asked to play with the interface until they felt they had a good sense of how to “drive it” to perform musical gestures. No time limit was given to this process; the users were encouraged to explore the possibilities of each set-up. Data was collected on the

users’ (subjective) views on the comparative expressivity and learnability of each mapping, and the quality of musical control that could be achieved.

Whilst users were experimenting with these instruments an interesting trend was observed. The first (one-to-one) test did not seem to hold the users’ attention. They treated it as a “task” to complete rather than an opportunity to experiment with a musical instrument. The second test received a more favorable response. The third test seemed to inspire the users to experiment and explore the interface to the point of composing and performing their own short melodies and musical gestures.

Users were asked to fill in a questionnaire relating to their feelings towards each experiment. The questionnaire asked each user to score the instruments in different areas such as:

“How easy did you find it to control individual parameters?”

“How expressive could you be with the interface?”

“How did you feel your performance with the interface improved over time?”

Most of the users found the first interface to be very “dull,” or not very stimulating, and were not willing to play it for more than a few minutes before moving on to the next experiment. Users commented that the simple division of parameters was not musically interesting.

Comments from the second user test suggested that the incorporation of user energy presents a more engaging natural instrument, allowing the user to feel that they have connected with the instrument. The simple division of parameters no longer seemed to be such an issue.

The final test presented much more of a challenge to the users, with some of them spending several minutes just trying to make it produce sound. As time progressed the users became more and more immersed in what they were doing musically and forgot about the technical aspects of the experiment. The key to this interface’s appeal seemed to be that it felt more like a traditional instrument and promoted holistic thinking rather than forcing the users to analyze exactly what each of the controls did. This allowed the users to step back from the technical issues relating to the instrument and concentrate on their musical output from the system.

3.3 Learning from acoustic instruments

In terms of designing mapping strategies, much can be gained from the analysis of the way existing acoustic instruments work. Several acoustic instruments present complex relationships between physical control variables and the resulting sound. These relationships can serve as inspiration for the design of expressive new digital musical instruments.

Rovan et al. (1997), used a Yamaha WX7 wind controller as the input device, and generated sound using additive synthesis models of clarinet sounds in IRCAM’s FTS environment (later in jMax). The idea behind the project was simple; many wind instrument performers complained that MIDI wind controllers tend to lack expressive potential when com-

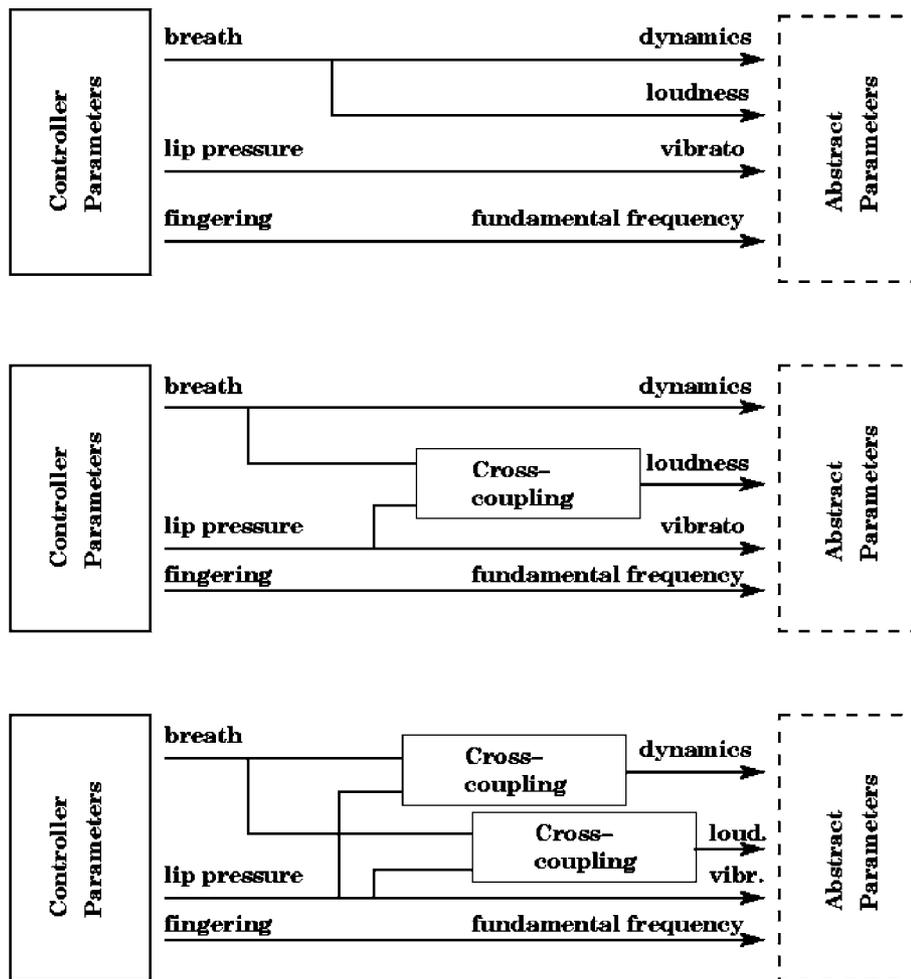


Fig. 10. The three mappings used in the clarinet simulation presented in Rován (1997).

pared to acoustic instruments such as the clarinet or saxophone. A common path to solving this problem involves improving the design of the controller by adding extra sensors. Rován et al. (1997) discussed the fact that by altering the mapping layer in a digital musical instrument and keeping the interface (an off-the-shelf MIDI controller) and sound source unchanged, the essential quality of the instrument is changed regarding its control and expressive capabilities. Previous studies, notably by Buxton (1986a), presented evidence that input devices with similar characteristics (e.g., number of degrees of freedom) could lead to very different application situations depending on the way these characteristics were arranged in the device. In that study, however, the devices were not the same mechanically (one had two separate one-dimensional controllers and the other one two-dimensional controller), so the situation is different to using exactly the same input devices with different mapping strategies.

Another point became clear in this process; even if the WX7 was a reasonably faithful model of a saxophone providing the same types of control variables (breath, lip pressure and fingering), these variables worked totally *inde-*

pendently in the MIDI controller, whereas they are *cross-coupled* in acoustic single-reed instruments. This natural cross-coupling is the result of the physical behaviour of the reed, and variables were simply independent since the equivalent “reed” in the controller was a plastic piece that did not vibrate, and moreover, was not coupled to an air column.

Based on these decisions and facts, the authors proposed different mappings between the WX7 variables and the synthesis parameters. The first was basically a one-to-one relationship, where variables were independent. The second was a model where the “virtual airflow” through the reed (loudness) was a function of both the breath and lip pressure (embouchure), as in an acoustic instrument. The third was a model that took into account both the “virtual airflow” and the relationship between spectrum content to breath and embouchure; a model that would match even more closely the real behaviour of the acoustic instrument.

The two most important consequences of this work were:

1. By just changing the mapping layer between the controller and the synthesis algorithm, it was indeed possible to completely change the instrumental behaviour and thus

the instrument's feel to the performer. Depending on the performer's previous experience and expectations, different mappings were preferred.

2. By deconstructing the way that the reed actually works, it was noted that the choice of mapping could be important as a pedagogical variable. Indeed, in stark contrast with acoustic instruments where the dependencies between parameters are unchangeable, cross-coupling between variables can easily be created or destroyed in digital musical instruments. This means that performers could focus on specific aspects of the instrument by explicitly defining its behaviour.

4. Analysis of experiments

In this section we set the results of the above music-interface experiments in the context of work by other authors and from mainstream Human Computer Interaction literature.

The experiments in section 3.1 reveal results which are at once challenging, and yet understandable. Experience with acoustic musical instruments, and their in-built complex cross-couplings of control parameters, would tend to support our results that the complex, multi-parametric interface performed significantly better. Our natural expectations of having to practice hard to play an instrument are reflected in the multiparametric interface's slow start; it takes time to get used to a complex mapping. However, it is notable that nobody questions the need to practice an instrument, yet in the same breath everybody demands "easy-to-use" computer interfaces. Much of the HCI literature concentrates on making interfaces more direct, and less hard to learn. At first glance this would seem a laudable goal, even an obvious one. But are we perhaps throwing out the proverbial baby with the bathwater (in this case the bathwater represents interface complexity and learning time, and the baby represents the rewards of a flexible interface)?

So, why should a complex and cross-coupled interface completely outperform a more standard sliders-only one? Most users of the multiparametric interface describe a moment when they "stopped thinking" about the interface and began to "just play" with it – almost as if their body was doing it for them, and their conscious mind was somehow "not in control." This is remarkably close to what acoustic musicians describe as the "flow" experience. In a fascinating web-site for musicians, violinist and psychologist A. Burzik (2002) describes the four principles of practicing in flow as:

1. A special physical contact with the instrument;
2. The development of a subtle feeling for sound;
3. A feeling of effortlessness in the body;
4. A playful and free-spirited handling of the material studied.

"Flow" is the term coined by M. Csikszentmihalyi (1975) to describe a holistic state of consciousness where actions happen continuously, accompanied by sharpened senses

acting in a unified manner. It is also associated with feelings of success and being uplifted – which certainly concurs with the experiences of the users of the multiparametric interface.

This contrasts greatly with the anger and frustration experienced by users of the parallel sliders interface. This may come as a surprise to some, especially considering that such banks of physical sliders are known the world over as the basis of the mixing desk. Fitzmaurice et al. (1997) conclude that slider banks (and other space-multiplexed graspable interfaces) outperform time-multiplexed interfaces, such as the mouse. Does our work contradict their findings? No, but it does qualify them. One possible interpretation of Fitzmaurice et al. is that we should abandon the mouse and provide separate graspable interfaces, but our results show this is not necessarily true. Where a mouse is used to time-multiplex (e.g., move sequentially between different on-screen areas such as icons and menus) then it will be less efficient than an interface where the same parameters are physically graspable. Our work supports this, as for the simplest tests – the physical sliders outperform the time-multiplexed "mouse" interface. However, in our final interface the mouse is used in a much more creative way – to provide *continuous multiparametric* control. Users of the sliders interface had to mentally split down the holistic sound into its four sonic components before being able to derive the motor controls for each individual slider. Users were happy to do this when only one parameter changed. For two simultaneous parameters it was much more difficult, but still just possible. With three and four parameters most people simply gave up, or tried until they were angry with frustration. This would seem to imply that space-multiplexed systems begin to fail when more than two parameters need to be changed simultaneously.

Somehow the multiparametric interface is aiding the users to cope with handling simultaneous sonic variables. The work by Jacob et al. (1994) on Integrality and Separability of input devices sheds light on how this may be happening. Jacob's work shows that an interface system work best where the structure of the input devices matches the perceptual structure of the task being performed by the user. In other words it is not just about which device is used (e.g., a mouse) but *how* that device is mapped onto the system parameters, and how well that helps the users to think about the task. In our case the users were hearing a holistic sound, yet the sliders interface forced them to decode the sound into four separate control streams – a very hard and frustrating task. In contrast the multiparametric interface allowed users to think spatially about the simultaneous control of pitch, timbre and volume. Users reported listening to a sound and "just knowing what shape that was" in terms of their required movements on the interface.

A further contribution to the success of the multiparametric interface could be the way in which the user's two hands are used. Kabbash et al. (1994) describe how a well-designed bimanual input can greatly outperform a single-handed control, but one hand is usually the dominant control

and the other more of a modifier (Guiard, 1987). This is exactly the situation in the multiparametric interface, where the main timbral parameters are all available to the right hand, and the left hand controls sliders which offset them and control the panning. Kabbash et al. also describe the converse situation, where an inappropriately designed two-handed interface can be significantly *worse* than one hand. This is what we believe was happening with the physical sliders interface. It also seems that the structure of the multiparametric interface was allowing the user to break down the sound into more useable sized “chunks” (Simon, 1996; Buxton, 1986b) of information, whereas the task of breaking down the sounds into a form usable by the physical sliders interface was just too much for the users.

Finally, there is the result common to all of the above experiments that the users report a better response when continuous input of *energy* is required from them. Again, we should perhaps expect this from our experience with acoustic musical instruments, but this is often ignored because computers are self-powered devices which do not require our energy to operate. This point has been studied by C. Cadoz, who proposed the notion of an *energy continuum* between performer actions and the sound being produced in most acoustic instruments (Cadoz, 1999). In fact, part of the muscular energy spent by the performer is transformed into acoustic energy in the resulting sound through *excitation gestures*. This is known as the *ergotic function* of the gestural channel (Cadoz & Wanderley, 2000).

We seem to be discovering that we need to use our energy to operate effectively. The response from nearly every user in all of the above experiments was that the interface perceptually transformed into an instrument when the user’s continuous energy was required. Maybe this is another requirement of “flow”? – that we need to continuously put energy into a system, and to steer its flow through the system. Returning to the “accidental theremin” of 2.1.1, once the energy requirement was removed and the system played continuous sound without movement, the users who were so engaged, lost interest rapidly. Could it be that there is an in-built requirement of human beings to use their own energetic control to transmit that energy into the world around them?

Taking the above thoughts, it is important to think about *how* to design mappings that will engage the user, and also be perceptible by an audience. In the following section we review the work done on mapping and contribute to the debate about how to design creative instruments for live performance.

5. Models and guidelines for mapping

Since there will not always be ready models for inspiration when designing new digital musical instruments, the design of mapping strategies becomes more complex. We therefore need to propose guidelines for mapping strategies and devise

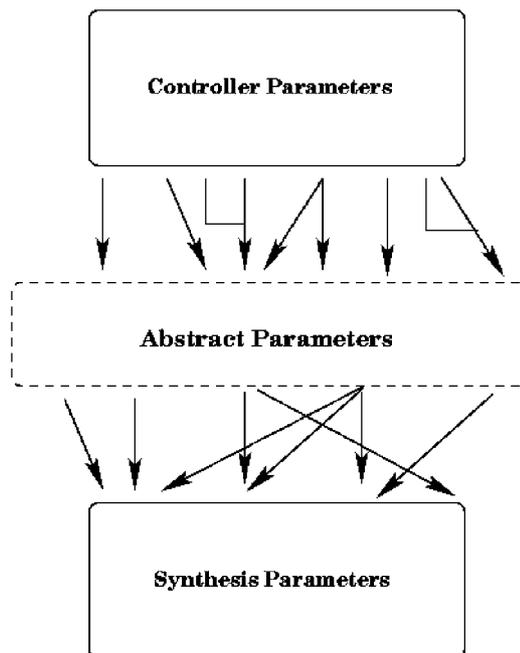


Fig. 11. Two-layer mapping model showing intermediate layer of abstract parameters.

models that can facilitate mappings other than simple one-to-one relationships.

In trying to answer this question of how to extend a specific mapping solution to a more general case, a model of mapping for digital musical instruments was proposed in (Wanderley et al., 1998). It was based on the separation of the mapping layer into two independent layers, coupled by an intermediate set of user-defined (or “abstract”) parameters (see Fig. 11).

This model was presented in the framework of a set of extensions to jMax later known as ESCHER (actually, a set of objects developed at Ircam to perform interpolation using additive models).

This idea is based on previous works, such as those of Mulder et al. (1997a), Métois (1996), and Wessel (1979). A similar direction was presented by Mulder et al. (1997b), Mulder and Fels (1998), and later by Garnett and Goudeseune (1999), and Goudeseune (2002). Basically, all these works have used higher levels of abstraction as control structures instead of raw synthesis variables such as amplitudes, frequencies and phases of sinusoidal sound partials. The main point made in (Wanderley et al., 1998) was to explicitly think about two separate mapping layers and the strategies to implement these, and not on the choice of intermediate parameters themselves, whether perceptive, geometrical or “abstract” (Wanderley & Depalle, 1999).

The intrinsic advantage of this model is its flexibility. Indeed, for the same set of intermediate parameters and synthesis variables, *the second mapping layer is independent of the choice of controller being used*. The same would be true

in the other sense: for the same controller and the same set of parameters, *multiple synthesis techniques could be used by just adapting the second mapping layer, the first being held constant*. Specifically in this case, the choice of synthesis algorithm is transparent for the user.

The two-layered model has been expanded to include three mapping layers (Hunt & Wanderley, 2002; Arfib et al., 2002). Figure 12 depicts a model containing these three mapping layers.

These works support the idea that, by using multi-layered mappings, one can obtain a level of flexibility in the design of instruments and that moreover, these models can indeed accommodate the control of different media, such as sound and video, in a coherent way. Section 5.1 describes in more detail the justification for a multi-layer mapping model.

5.1 One-to-one mappings using multiple layers

We have noted that there is a tendency for designers to make one-to-one mappings when constructing an interface. We can *use* this tendency to improve the mapping process if we utilise the many layered models outlined above. The following scenario may illustrate this:

Imagine a system whose interface inputs included “button 1,” “button 2,” “slider 1,” “slider 2,” “mouse x” and “mouse y.” Let us suppose that the synthesis system was a Frequency Modulation module with inputs such as “carrier frequency,” “carrier amplitude,” “modulation frequency” etc. Now consider the two possibilities below:

Case 1: let us consider a designer working to connect the above inputs to the above outputs. We are quite likely to see arbitrary connections such as “mouse x controls carrier frequency,” and “slider 1 controls modulation frequency” (similar to those shown in Fig. 1). These give us the often-encountered one-to-one mappings.

Case 2: let us imagine that a mapping layer has already been devised to abstract the inputs to parameters such as “energy,” “distance between sliders,” “wobble” etc. Also let us imagine that there is a mapping layer before the FM synthesis unit, providing higher-level control inputs such as “brightness,” “pitch,” “sharpness” etc. This gives us a situation similar to that found in Figure 12. Now we can picture the designer making a relationship such as “energy controls brightness.” On the surface this may appear to be yet another one-to-one mapping. Indeed it is – at the *conceptual* level. However, when you consider how “energy” is calculated from the given inputs, and how “brightness” has to be converted into the FM synthesis primitives, you will notice how many of the lower-level parameters have been cross-coupled.

Thus the many-level mapping models are a way of simplifying the design process, and of helping the designer to focus on the final effect of the mapping, as well as providing a convenient method of substituting input device or synthesis method.

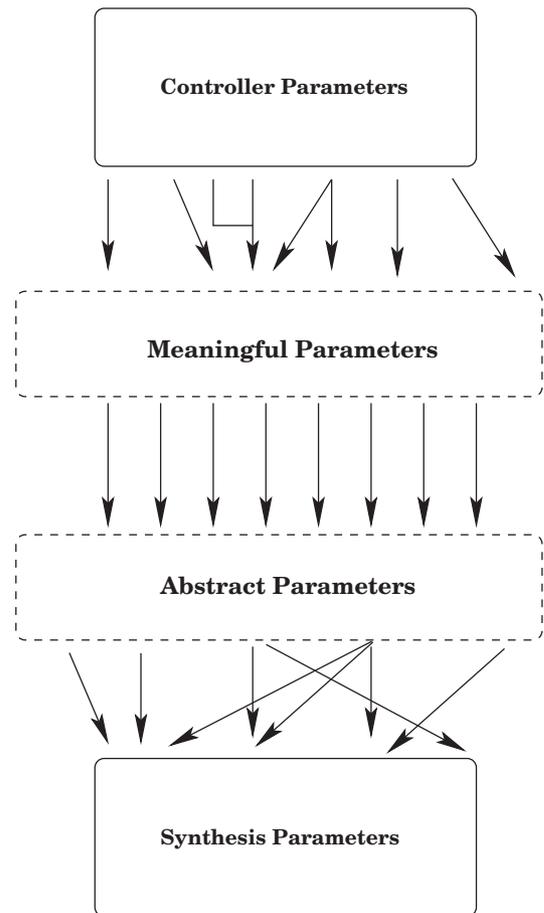


Fig. 12. Three-layer mapping model.

6. Towards designing, testing and evaluating mapping strategies

The experiments outlined above have proven the importance of the mapping layer in electronic musical instruments. However the question remains, how do we efficiently and effectively design, test and evaluate a mapping strategy?

The development of mapping strategies for “known interfaces” (those which inherit their properties from existing interfaces) has a head start as there are plenty of successful examples already widely used in the form of traditional musical instruments. However, fewer examples exist for novel interfaces – newer designs that have no traditional counterparts.

Novel sensor based interfaces, MIDI-slider boxes and graphics tablets all have the potential to be employed as very expressive musical devices, as long as the mapping is implemented correctly.

The third author is carrying out ongoing research in an attempt to further our knowledge of the mapping layer. The goal of this research is to develop techniques and models which allow researchers to evaluate the performance of a given mapping strategy, and to provide focussed environ-

ments for the construction of such mappings by answering the following questions.

How do we define the mapping layer?

In order to fully understand the workings of the mapping layer it must firstly be fully defined. This will give researchers a clear structure and vocabulary with which to refer to elements of a system. Current research concentrates on the mapping layer as a high level entity. Strategies are defined (one-to-one etc.) along with layers which can carry out specific functions. It is proposed that a mapping layer should be studied at a much lower level in order to fully understand the implications of a given parameter configuration.

What effect does a specific treatment of parameters have on the user experience?

It is important for interface designers to understand how specific operations affect the overall user experience that a system offers. For example (in a musical context) is there a different psychological effect from adding parameters together rather than averaging parameters? One would expect that this would create a very different interface, but which method is best for the task in hand?

At a lower level, a mapping layer can be examined in great detail. As with any complex system certain points in the layer will contribute to the overall output of the system in different ways. It is proposed that these “areas of activity” can be thought of as mapping *nodes* which are examined and classified according to their function within the network of parameter relationships. At this level of detail it becomes possible to use specific mapping operations such as addition, division, averaging, rate of change or acceleration (to name a few) to build up nodes of parameter modifications which can influence the interaction model in different ways.

How can a system be designed to give the appropriate balance of accuracy in real time control and expressive intuitive interaction?

Different interaction tasks require different interaction models. A task such as entering text into a word processor requires a very accurate interface, but it is not concerned with creative expression. Inaccuracies cannot be tolerated as they instantly render the output from the system unusable. A musical instrument however may not require the same degree of accuracy in order to create usable output. Examples of this can be found most traditional instruments. Players have to rely on skills built up over many years of practice before they can control an instrument accurately and this in turn promotes a greater understanding of the interface by the user. Recent informal experiments indicate that accurate interfaces do not tend to promote expressive output from a user and vice-versa. In order to understand this relationship, a set of tests has been proposed as ongoing research to enable the

collection and analysis of data over time relating to the accuracy, speed, expressive control and suitability of interfaces for specific tasks.

How can we measure the performance of a mapping strategy?

How can the subjective data taken from a user’s experience of an interface be represented with actual performance data (data that can be recorded whilst a user is carrying out a specific task)? Can trends be observed in objective data to represent Expressivity, or the enjoyment that a user experiences in their interaction?

In all interface design the end result is ultimately a user experience. Mapping strategies can provide a tailored experience dedicated to the task in hand, rather than using a global parameter association as can often be found with devices such as the mouse. Simply trying to extend the position of a user’s hand into a cursor on the computer screen is neglecting a great deal of the sophisticated interaction that users are capable of in an attempt to provide a global solution for all computing tasks.

7. Future discussion of mapping

From the evidence presented above in both informal and controlled experiments, there is definitely a need to find better-designed mappings than simple (engineering style) one-to-one relationships. General models of mappings have been proposed and expanded to incorporate multimedia control, but also to fit several levels of performance, from beginners to highly skilled players.

One attempt to foster the discussion in this direction has been initiated in the context of the ICMA/EMF Working Group on Interactive Systems and Instrument Design in Music (Wanderley, 2000). Readers may also wish to refer to the special issue on “Mapping Strategies for Real-time Computer Music” guest-edited by the second author (Wanderley, 2002).

We therefore welcome comments and criticism on issues related to mapping so as to push the discussion on this essential, although often ignored, topic.

8. Conclusions

The mapping “layer” has never needed to be addressed directly before, as it has been inherently present in acoustic instruments courtesy of natural physical phenomena. Now that we have the ability to design instruments with separable controllers and sound sources, we need to explicitly design the connection between the two. This is turning out to be a complex task.

This paper has indicated that mapping can define the very essence of an instrument and that complex mappings can provide quantifiable performance benefits and improved

interface effectiveness. We experimented with a multi-layer model to help designers to implement complex but usable mappings, and are carrying out a programme of research to look into the testing and evaluation of mapping strategies.

We are still in the early stages of understanding the complexities of how the mapping layer affects the perception (and the playability) of an electronic instrument by its performer. What we know is that it is a very important layer, and one that must not be overlooked by the designers of new instruments.

Acknowledgements

Many thanks are due to the following people for their collaboration in the work described in this paper: Ross Kirk, Butch Rován, Shlomo Dubnov, Philippe Depalle, and Norbert Schnell.

References

- Arfib, D., Couturier J., Kessous L., & Verfaillie V. (2002). Strategies of mapping between gesture data and synthesis model parameters using perceptual spaces, *Organised Sound*, 7, 127–144.
- Burzik, A. (2002). A Holistic Method of Practice for all Instrumentalists. Available: <http://www.practising-in-flow.de/> Last visited on 10/04/03.
- Buxton, W. (1986a). There's more interaction than meets the eye: Some issues in manual input. In: D. Norman, & S.W. Draper (Eds.), *User Centered System Design: New Perspectives on Human-Computer Interaction*. Hillsdale, N.J.: Lawrence Erlbaum Associates, pp. 319–337.
- Buxton, W. (1986b). Chunking and phrasing and the design of human-computer dialogues. In: H.J. Kugler (Ed.), *Information Processing '86, Proceedings of the IFIP 10th World Computer Congress*, North Holland, pp. 475–480.
- Cadoz, C. (1999) Continuum énergétique du geste au son. Simulation multisensorielle interactive d'objets physiques. In: H. Vinet, & F. Delalande (Eds.), *Interfaces Homme-Machine et Création Musicale*. Paris: Hermès Science Publishing, pp. 165–181.
- Cadoz, C., & Wanderley, M. (2000) Gesture-Music. In: M. Wanderley, & M. Battier (Eds.), *Trends in Gestural Control of Music*. Paris: Ircam – Centre Pompidou, pp. 71–93.
- Csikszentmihalyi, M. (1975). Beyond boredom and anxiety: Experiencing flow in work and play. (reprint, Jossey Bass Wiley; 2000 ISBN: 0787951404).
- Fitzmaurice, G., & Buxton, B. (1997). An empirical evaluation of graspable user interfaces: Towards specialized, space-multiplexed input. *Proc. ACHCHI 1997* pp. 43–50. <http://www.acm.org/sigchi/chi97/proceedings/paper/gf.htm> Last visited on 10/04/03.
- Guiard, Y. (1987) Assymmetric division of labor in human skilled bimanual action: The kinematic chain as a model. *Journal of Motor Behaviour*, 19, 486–517.
- Hunt, A. (2000). *Radical user interfaces for real-time musical control*. DPhil thesis, University of York, UK. Available: http://www-users.york.ac.uk/~elec18/download/adh_thesis/
- Hunt, A., & Wanderley, M.M. (Eds.) (2000). *Mapping of control variables to musical variables. Interactive systems and instrument design in music working group*. Website: www.igmusic.org
- Hunt, A., & Kirk, R. (2000). *Mapping strategies for musical performance*. In: M. Wanderley, & M. Battier (Eds.), *Trends in Gestural Control of Music*. IRCAM – Centre Pompidou.
- Hunt, A., Wanderley, M.M., & Kirk, R. (2000). Towards a model for instrumental mapping in expert musical interaction. In *Proc. of the 2000 International Computer Music Conference*. San Francisco, CA: International Computer Music Association, pp. 209–211.
- Hunt, A., & Wanderley, M.M. (2002). Mapping performance parameters to synthesis engines. *Organised Sound*, 7, 103–114.
- Garnett, G., & Goudeseune, C. (1999). Performance factors in control of high-dimensional spaces. *Proc. 1999 International Computer Music Conference*. San Francisco, CA: International Computer Music Association, pp. 268–271.
- Goudeseune, C. (2002). Interpolated mappings for musical instruments, *Organised Sound*, 7, 85–96.
- Jacob, R., Sibert, L., McFarlane, D., & Mullen, M. (1994). Integrality and separability of input devices. *ACM Transactions on Computer-Human Interaction*, 1, 3–26.
- Kabbash, P., Buxton, B., & Sellen, A. (1994). Two handed input in a compound task. *Proc SIGCHI Conf. Human Factors in computing systems*, pp. 417–423.
- Métois, E. (1996). *Musical Sound information: Musical gestures and embedding systems*. PhD Thesis. MIT Media Lab.
- Mulder, A., Fels, S., & Mase, K. (1997). Empty-handed gesture analysis in Max/FTS. In *Kansei, The Technology of Emotion. Proceedings of the AIMI International Workshop*, A. Camurri (Ed.), Genoa: Associazione di Informatica Musicale Italiana, October 3–4, pp. 87–91.
- Mulder, A., Fels, S., & Mase, K. (1997). *Mapping virtual object manipulation to sound variation*. In: T. Rai, & R. Basset (Eds.), *IPSJSIG notes*, 97, 63–68.
- Mulder, A., & Fels, S. (1998). Sound sculpting: Manipulating sound through virtual sculpting. *Proc. 1998 Western Computer Graphics Symposium*, pp. 15–23.
- Rován, J.B., Wanderley, M.M., Dubnov, S., & Depalle, P. (1997). Instrumental gestural mapping strategies as expressivity determinants in computer music performance. In *Kansei, The Technology of Emotion. Proceedings of the AIMI International Workshop*, A. Camurri (Ed.), Genoa: Associazione di Informatica Musicale Italiana, October 3–4, pp. 68–73.
- Simon, H. (1996). *The Sciences of the artificial*. MIT Press; 3rd edition, ISBN 0262691914.
- Wanderley, M.M., Schnell, N., & Rován, J. (1998). Escher – Modeling and performing “composed instruments” in real-time. *Proc. IEEE International Conference on Systems, Man and Cybernetics (SMC'98)*, San Diego, CA, pp. 1080–1084.

- Wanderley, M.M., & Depalle, P. (1999). Contrôle gestuel de la synthèse sonore. In: H. Vinet, & F. Delalande (Eds.), *Interfaces Homme-Machine et Creation Musicale* – Hermes Science Publishing, pp. 145–163.
- Wanderley, M.M. (Ed.) (2000). Interactive systems and instrument design in music workgroup. Website: www.igmusic.org Last visited on 10/04/03.
- Wanderley, M.M. (2001). *Performer-instrument interaction. Application to gestural control of sound synthesis*. PhD Thesis. University Paris VI, France.
- Wanderley, M.M. (Ed.) (2002). Mapping strategies for real-time computer music. *Organised Sound*, 7.
- Wessel, D. (1979). Timbre space as a musical control structure. *Computer Music Journal*, 3, 45–52.