

The importance of parameter mapping in electronic instrument design

Andy Hunt

Music Technology Group,
Electronics Department
University of York,
Heslington, York
YO10 5DD
email adh2@york.ac.uk

Marcelo M. Wanderley

Faculty of Music,
McGill University
555, Sherbrooke Street West
H3A 1E3 - Montreal - Canada
email mwanderley@acm.org

Matthew Paradis

Music Technology Group,
Music Department
University of York,
Heslington, York
YO10 5DD
email mdjp100@york.ac.uk

Abstract

In this paper we challenge the assumption that an electronic instrument consists solely of an interface and a sound generator. We emphasise the importance of the mapping between input parameters and system parameters, and claim that this can define the very essence of an instrument.

Keywords

Mapping Strategies, Electronic Musical Instruments, Human-Computer Interaction

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. 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 psychological and emotional response elicited from the performer is determined to a great degree by the mapping.

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 in [6], [16] and [17].

Informal Observations

The first author has carried out a number of experiments into mapping. The more formal of these have been presented in detail in [5] and [3], and are summarised later in this paper. Let us begin with some rather simple, yet

interesting, observations that originally sparked interest in this subject. We have retained the first person writing style to denote that these are informal, personal reflections.

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 the 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 in a week's time and noted the irony that the 'mistake' had been corrected, deleted from the student's notes, and the traditional form of the instrument implemented.

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.

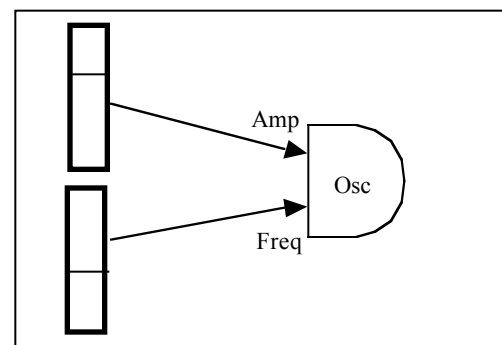


Figure 1. Simple Mapping for Experiment 1

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 (cf. Figure 1).

I let several test subjects freely play with the instrument, and talked to them afterwards. In the second experimental run, the interface was re-configured to emulate the abovementioned '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 (cf. Figure 2).

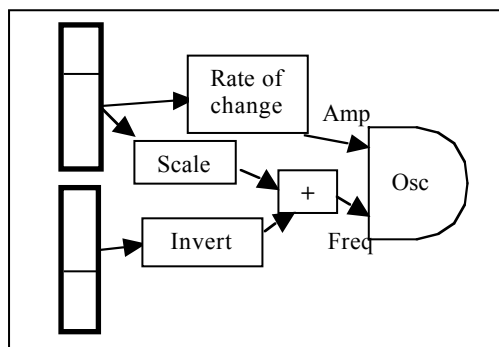


Figure 2. Complex Mapping for Experiment 2

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 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.

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.

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 [3]. The main part of this study took the form of 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 [15], and the details of the theory, experiments and results have been published [5]. 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 (cf. Figure 3) represented a typical computer music editing interface with on-screen sliders connected one-to-one to each sound parameter.




Figure 3. The 'mouse' interface

The second (cf. Figure 4) involved physical sliders (on a MIDI module) again connected in a one-to-one manner to the synthesis unit.




Figure 4. The 'sliders' interface

The third interface (cf. Figure 5) 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.




Figure 5. The 'multi-parametric' interface

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. These quantitative 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 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 parameter simultaneously. Surprisingly perhaps, they were nearly all extremely frustrated and angered by the 4 physical sliders. Comments abounded such as "I should be able to do this, technically, but I can't get my mind to split down the sound into these 4 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 - but not at the very beginning. At first it seemed counter-intuitive to 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.

Focusing on the Effect of Mapping Strategies

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 decreased reliance on visual

feedback, and provided 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 to focus entirely on the user's reaction to a change in mapping strategy.

These tests utilised three contrasting mapping strategies, with a fixed user interface and synthesis algorithm. The mappings were;

- a) simple one-to-one connections between input and output,
- b) one-to-one requiring the user's energy as an input. This was implemented by requiring the user to constantly move one of the sliders in a 'bowing'-like action
- c) many-to-many connections from input to output, but also requiring the user's energy as in b).

These mappings were used to control the parameters of a stereo FM synthesis algorithm, including amplitude, frequency, panning, modulation ratio and modulation index. The input device used was a MIDI fader box. Users were asked to play with each 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 accuracy of musical control that could be achieved.

Whilst experimenting with the first mapping test (one-to-one) many users noted that the simple division of parameters was not very stimulating. Users tended to learn the parameter associations very quickly but then struggle to achieve any improvement in their performance or expressive output.

The second test generated a range of comments, which suggested that the process of injecting energy into a system presented a much more natural and engaging instrument. However, due to the proximity of sliders on the interface they found it difficult to control other sliders whilst providing the required 'bowing' action. However, this problem lessened over time as the user practised.

The third and final user test (many-to-many mappings) provided some interesting results. Most of the test subjects noted that the appeal of this instrument was that it was not instantly mastered but required effort to achieve satisfactory results. The instrument presented a challenge to the user, as one would expect from a traditional expressive instrument.

These tests highlighted the differences between a general-purpose interface, (such as the mouse) which has simple mappings but allows the user to begin working instantly, and an interface with more complex mappings which must be practised and explored in order to achieve truly expressive output.

Learning from Acoustic Instruments

In [12] the second author and collaborators discussed the fact that by altering the mapping layer in a digital musi-

cal 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 [2], 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 exactly the same mechanically (one had two separate controllers and the other one two-dimensional controller), so the situation is not the same as when using the same input device with different mapping strategies, even if the results are similar.

In [12], a Yamaha WX7 wind controller was used as the input device, and sound was generated 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 compared 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. However, it was decided to challenge this assumption and to solely work on the mapping layer between the controller variables and the synthesis inputs (for a complete description see [12]).

Another point became clear in this process: even if the WX7 was a faithful model of a saxophone providing the same types of control variables (breath, lip pressure and fingering), these variables worked totally *independently* 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 since the equivalent "reed" in the controller was a plastic piece that did not vibrate, and moreover was not coupled to an air column, variables were simply independent.

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), such 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.

Using these three different models, the system was performed by different musicians and non-musicians. Results indicated that wind instrument performers tended to stick with complex cross-coupled mappings similar to the single reed behaviour (the third mapping strategy used), whereas beginners initially preferred simpler mappings (easier to play and produce stable sounds).

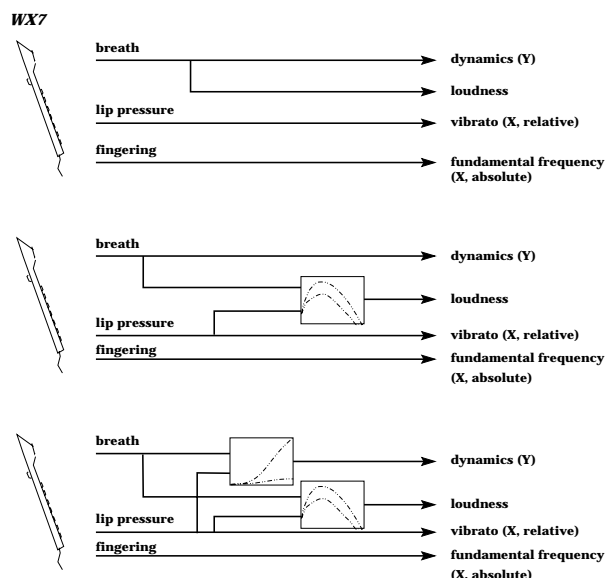


Figure 6. Several mappings used in the clarinet simulation presented in [12].

The two most important consequences of this work were:

- 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.

- 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. Possible options could include:

- complex (cross-coupled) control of loudness with one-to-one control of timbre,
- one-to-one loudness and complex timbre controls, or,
- complex loudness and timbre controls, such as in the real instrument.

Even if these results supported the essential role of mapping (and the importance of devising mapping strategies other than one-to-one during the design of digital musical instruments), they could not be easily extrapolated to more general situations. In fact, in the above specific case, there *did exist* a model of complex mapping to be followed, since the controller was a model of the acoustic instrument. So what about mappings in general digital musical instruments using alternate controllers, those not based on traditional acoustic instruments?

MODELS AND GUIDELINES FOR MAPPING

Since there will not always be ready models for inspiration when designing mapping strategies for new digital musical instruments, the task then becomes one of proposing guidelines for mapping and also, if possible, devising models that can facilitate the implementation of mapping strategies 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 [13]. 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. This model was presented in the framework of a set of extensions to jMax later known as ESCHER (actually, a set of objects developed by Norbert Schnell to perform interpolation using additive models).

This idea is based on previous works, such as those of Mulder et al. [9], Métois [8], and Wessel [18]. A similar direction was presented by Mulder and Fels in [10] and later by Garnett and Goudeseune [7]. 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 [13] 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” [14].

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 original two-layered model has recently been expanded to include three mapping layers in two independent performance works by Hunt and Myatt [11], and by Arfib and collaborators [1]. 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.

One-to-one Mappings – 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”. These give us the oft-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. 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.

FUTURE DISCUSSION OF MAPPING

From the evidence presented above in both informal and controlled experiments, there is definitely a need to come up with 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 [4]. A further effort is currently being carried out in the form of a special issue on “Mapping Strategies for Real-time Computer Music” guest-edited by the second author [17] to appear as volume 7, number 2 of the journal *Organised Sound* later this year.

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

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 non-trivial task.

We are 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.

REFERENCES

- [1] Arfib, D. Personal Communication. (2002)
- [2] Buxton, W. There's more Interaction than Meets the Eye: Some Issues in Manual Input. In Norman, D. a. and Draper, S. W. (eds), *User Centered System Design: New Perspectives on Human-Computer Interaction*, Hillsdale, N.J.: Lawrence Erlbaum Associates, 319-337. (1986)
- [3] Hunt, A. *Radical User Interfaces for Real-time Musical Control*. DPhil thesis, University of York, UK. (1999)
- [4] Hunt, A. and Wanderley, M. M. (eds.) Mapping of Control Variables to Musical Variables. Interactive Systems and Instrument Design in Music Working Group. (2000) Website: <http://www.notam02.no/icmd/interactivesystems/mapping.html>
- [5] Hunt, A., and Kirk, R. Mapping Strategies for Musical Performance. In M. Wanderley and M. Battier (eds.) *Trends in Gestural Control of Music*. IRCAM – Centre Pompidou. (2000)
- [6] Hunt, A., Wanderley, M. M., and Kirk, R. 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. (2000)
- [7] Garnett, G., and C. Goudeseune. Performance Factors in Control of High-Dimensional Spaces. In *Proc. of the 1999 International Computer Music Conference*. San Francisco, CA: International Computer Music Association, pp. 268 - 271. (1999)
- [8] Métois, E. *Musical Sound Information: Musical Gestures and Embedding Systems*. PhD Thesis. MIT Media Lab. (1996)
- [9] Mulder, A., S. Fels, and K. Mase. 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. (1997)
- [10] Mulder, A., and Fels, S. Sound Sculpting: Manipulating Sound through Virtual Sculpting. In *Proc. of the 1998 Western Computer Graphics Symposium*, pp. 15-23. (1998)
- [11] RIMM The Real-time Interactive MultiMedia project. (2001) Website: <http://www.york.ac.uk/res/rimm/>
- [12] Rován, J. B., Wanderley, M. M., Dubnov, S., and Depalle, P. 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. (1997)
- [13] Wanderley, M. M., Schnell, N. and Rován, J. Escher - Modeling and Performing "Composed Instruments" in Real-Time. In *Proc. IEEE International Conference on Systems, Man and Cybernetics (SMC'98)*, San Diego, CA, pp. 1080-1084. (1998)
- [14] Wanderley, M. M., and Depalle, P. Contrôle gestuel de la synthèse sonore. In H. Vinet and F. Delalande (eds.) *Interfaces Homme-Machine et Creation Musicale* - Hermes Science Publishing, pp. 145-163. (1999)
- [15] Wanderley, M. M. (ed.) Interactive Systems and Instrument Design in Music Workgroup. (2000) Website: <http://www.notam02.no/icmd/interactivesystemswg.html>
- [16] Wanderley, M. M. *Performer-Instrument Interaction. Application to Gestural Control of Sound Synthesis*. PhD Thesis. University Paris VI, France. (2001)
- [17] Wanderley, M. M., ed. Mapping Strategies for Real-time Computer Music. Special Issue. *Organised Sound* 7(2). To appear in August 2002.
- [18] Wessel, D. Timbre Space as a Musical Control Structure. *Computer Music Journal*, 3(2):45-52. (1979)