The Cilia: Aspiring to Virtuosic Controller Design.

Abstract
This paper profiles some conclusions from several years of performing with and programming for the Cilia, a new flute controller. One condition for virtuosic music controller performance is able and responsive hardware. However, it has been proposed before that control data routing and network design plays a more crucial role in meeting the objective of creating a controller which can facilitate high degrees of dexterity and nuance in controller performance. This paper endeavors to support and amplify this proposition in two ways. First, reference to the control domain of acoustic instruments is instructive for articulating design strategies and criteria for control network design. Second, novel concepts of control are thereby idenitfied, contextualized within a theory of wholes and parts, and shown in their implementation in the Cilia application software. This paper contributes to broader discussion by providing a conceptual tool box, tools useful in both descriptition and prescriptition.

Keywords
Novel controllers and interfaces for musical expression, Novel musical instruments, Musical mapping strategies, Musical human-computer interaction, Performance analysis

Introduction
The Cilia is a vertical flute controller. It consists of a physical flute controller (currently with a PIC microcontroller) which directly produces control data and a software application (MaxMSP) which implements the contrl data in generating musical sound. It is intended primarily as an event controller (by definition requiring simultaneity between control gesture and sonic generation). It aspires to facilitate virtuosic performance, following a high standard set by acoustic instrument makers. If it can do this, then to function as a process controller (requiring control operators for sequences, algorhythms, sonifications, or pre-recordings) is simple. By definition, flutes split an air column produced with an open *embouchure* (lip technique). The Cilia mouthpiece splits the air column and derives control data from an analysis of the split air column's dynamics. All flute wind techniques can be used. Five high-resolution (0-6000) three-dimensional (x, y, z) track pads enable the use of numerous finger techniques found on keyed flutes as well as non-keyed flutes. The physical performance gestures of such a non-keyed flute, the Japanese shakuhachi, served as a first design paradigm.

Figure 1. The Shakuhachi and the Cilia.

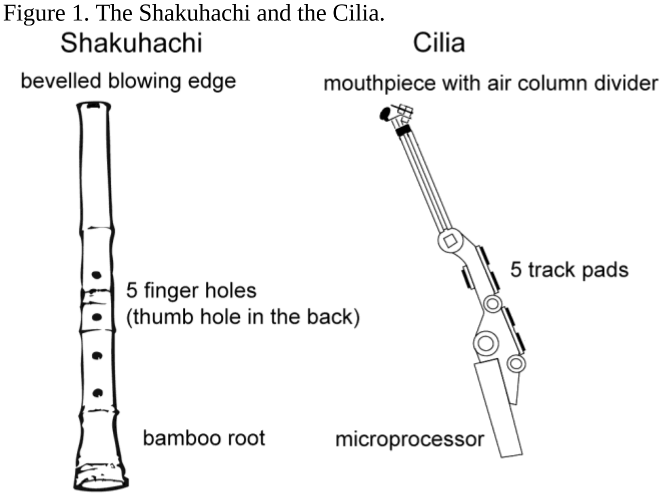

The Cilia was first conceived as a composer's instrument, facilitating 'orchestral' like improvisation and performance. Not intending to imitate the sound of the shakuhachi, the Cilia rather enables dexterity in manipulation of several timbre types (overtone structure, harmonic color, density, and texture). It performs monophonically, heterophonically, homophonically, polyphonically, and cacophonically; terms that colletively describe a 'compositional virtuosity.' Its sound generating routines have tended away from pitch-centrificity, an aesthetic choice. Although virtuosity in composition-performance is the end, the practical design problem begins specifically with the task of facilitating athletic virtuosity. That is the focus here.

This paper traces the first years of an ongoing design process. It begins with concepts of control that are relatively familiar to electronic musicians, with which it then enacts a reverse analysis of the 'control-dimension' of the shakuhachi. Additional categories of control were developed which prestructure control data network design, and influence synthesis design such that it be capable of receiving and using such control. The disucssion here will be limited to the former.

Control and Abandon.

Control data networks are complex. Complexity is ever an implicit topic. In performance practice, complexity begins with a contradiction that resides within the term controller, between control and abandon. All experienced musicians know this, whether improvising or interpreting. There is a feedback continuum between the immediate production of sound and the responsive adjusting of technique. There is a moment of abandon as the player puts the sound *out there*; then in the response a moment of control as the resistances of the instrument are negotiated. The resistances of good instruments are complex, always providing some nuance beyond the player's control. Experienced players negotiate quickly and efficiently. There is a poise that is part of the athlete's psychology that enables intense focus on movement and reaction. The back and forth feedback happens so fast and constantly, a listener can't parse it. The sound shimmers, and even if one can't parse it, this shimmering is still audible. Shimmering is another way of saying temporally complex meaning that past, present and future are equi-primordial.

The sustained action of shimmering requires *athletic virtuosity*. The task of how to encourage virtuosic performance can be unsderstood as that of encouraging shimmering. How can one facilitate shimmering in design? We propose next to translate it into exacting control domain terms.

Unidirectional Control Types.

Discrete control means direct implementation of triggers (e.g., simple switches or sequence tracking routines) and continuous control of data streams (e.g., from capacitive sensors such as track-pads). Both instantiate *direct* control. They directly affect a signal path. This has been indicated by others as one-to-one mapping, but would also include 'divergent mapping' or isomorphic parsing (Rovan et al. 1997). Direct routing includes one-to-many and many-to-one mappings.

There are four types of fingerings on the Cilia. They are all second order direct discrete controllers as they are derived from primary controllers (raw pad data) read for discrete *state* control (is the finger off or on the pad?). The combinations of states from the different pads yield the second order controllers. Non-vented and vented fingerings are woodwind players' terms. The former describe continuous closures from the top of the tube down to the bottom. The latter describe non-contiguous closures. Numeric fingerings describe the number of pads whose state is on. They are meaningful only on an electronic controller. When non-vented fingerings are combined with ratio threshold values from the mouthpiece (there are two thresholds, therefore three possible discrete ratio values), what is produced is analogous to a 3-octave range. Four non-vented fingerings become a series of 12 basic notes.

Two of the track pads' continuous controllers are all that is needed to divide the pads into quadrants, providing four state detectors per track pad. The Cilia is thus capable of playing homophonically (i.e., non-parallel chords sequences). Each combination of quadrant state values maps onto intervallic structures (how chords are defined). It should be noted that multi-tasking of the pads on several levels is the norm.

Dynamic control happens when controllers directly control other controllers. This is a simple example of what has been called 'convergent mapping' (Rovan et al. 1997). A formal theory of types would admit of discrete and continuous dynamic controllers, and of first and second orders.

Indirect dynamic control occurs when controllers indirectly control controllers by way of another mediating controller. In the Cilia work, the mediating controller is invariably a dynamic controller or another indirect dynamic controller. As with dynamic controllers, other sub-distinctions can be anticipated formally. They have not yet been implemented in the Cilia programming.

Figure 2. Demonstration of Unidirectional Control Types.

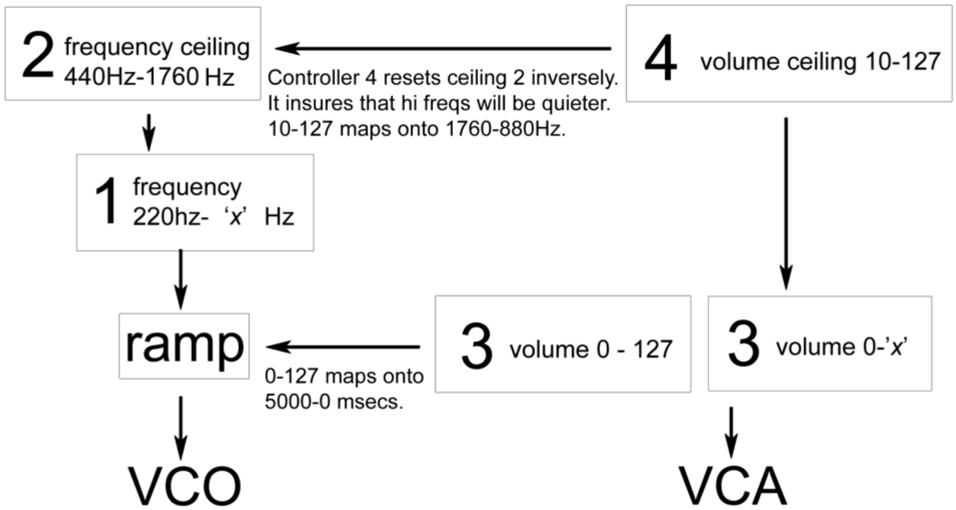

Consider the control category implementations in figure 2. If controller 3 went directly to the VCA it would demonstrate simple direct continuous control. As controller 4 modulates its upper limit, it functions as a dynamic continuous controller. Similarly controller 1 is a dynamic continuous controller as controller 2 is setting its upper limit. Controller 4 parses and because its values are scaled, provides an example not only of complex control, but also of indirect dynamic control. It affects controller 1 through manipulating its dynamic controller, controller 2. Complex control is demonstrated in the parsing of controller 3. When not modulated by dynamic controller 4, it routes directly into a time function, “summing” with the already modulated frequency to make frequency ramps (glissandi).

This pedagogical patch enables performance with a sine wave. The aesthetic objective is to avoid having very loud activity at higher frequencies, to do so in a way that is not entirely predictable, and in a way that a performer can control. Complex multiple referent controllers perform deterministically, even if statistically within ranges of values. Such *behavior*, such inexact predictability aids in making a performance lively.

Complexity begins in control data configuration with the parsing of a single control stream, and *in its non-isomorphic implementations*.

Reverse Analysis of the Control Dimension of a Shakuhachi.

Reverse analysis here means use of electronic music concepts to understand the control dimension of the shakuhachi anew. It begins with common practical knowledge that any player can attest to.

Breath and finger actions are two types of physical gesture used on the shakuhachi. There are four breath actions parameters. The *angle* of the vertical flute to the lips changes the ratio of air in and out of the tube. *Force* is the amount of pressure applied by the stomach muscles to push the diaphragm and force air from the lungs. *Air volume* is the shape of the *embouchure* hole. The more horizontally oblong the hole is, the more breathy the sound: the less so, the denser the timbre. *Velocity* is the size of the *embouchure* hole. The smaller the faster. Fingerings change not only the length of a standing wave, but also the impedance configuration, the resistance that create instabilities in the standing wave. Any combination of closed and open holes counts as a fingering. Covering only part of a hole (shading) enables many other techniques. The percentage of area covered is one parameter. So is the shape of the shading, whether it is covered from the top of the hole down, or from the side.

Figure 3. Reverse Analysis of the Control Structure of a Shakuhachi.

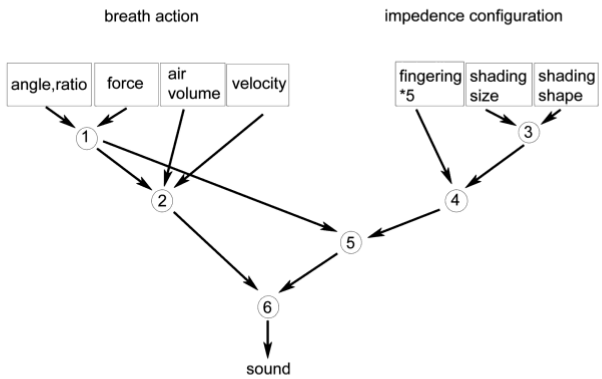

At node one (fig. 3), the breath parameter angle acts on force, or conversely breath with consistently less force enables greater angling (greater pitch and timbre spectrum). If the angle or ratio value is closer to 0 (the breath is almost all in the flute or out), then less force will be possible. This is an example of continuous dynamic control.

This is repeated at node 2, except with greater complexity as two more parameters are acting as dynamic controllers.

Direct continuous control is seen in the movement of the fingers over the keys (affecting pitch) (node 3), and in the increasing of breath affecting both amplitude and timbre in a continuous spectrum (force).

At node 5, velocity as a discrete threshold controller acts on the fingerings by changing the frequency register, octave shifts. The fingerings function progressively differently from the norm of non-vented fingerings the higher the register. This demonstrates discrete dynamic control.

If the value of any breath parameter is 0, then all are zero-ed. Whatever else is going on at node 6, this shows multiple dynamic control and interdependence.

All fingerings produce unique frequencies in combination with timbre variations. In this way the shakuhachi demonstrates direct discrete control.

The one instance of indirect dynamic control can be seen by tracing force from node 1 to node 5. Insofar as it effectively changes the interior of the instrument (by over-blowing), it affects the impedance configuration.

This descriptive also reveals limits in the paradigm. Indirect dynamic control is hardly present; the control routing cannot become as complex as one can easily imagine with electronic controllers. Furthermore, the control routing configuration is fixed on an acoustic instrument. The Cilia is an interface for a universal machine and need not be so limited.

What design lessons are here?

1/ The control that we see demonstrated on the shakuhachi involves the summing of several control sources to create other levels of control (dynamic and indirect dynamic). *Richness of control is achieved by such summing.*

2/ At node one, force is parsed and routes to two different purposes; it contributes to the breath signature, and to the impedance configuration. *The degree of complexity of control yields a higher measure of nuance and depth in the sonic outcome.*

3/ Combinations of the five fingers on the shakuhachi yield fingerings, each of which yield a new sonic possibility. Higher order controllers can be developed through tracking of combinatorial outcomes of simpler controller mechanisms. *Higher orders of control enable manipulation of higher orders of musical content, e.g., scales.*

4/ The shakuhachi demonstrates how certain control parameters are bound together, mutually dependent. It is important for instrumentalists to have thresholds and limits. *Designing such dependencies with consistency aids the instrumentalist in developing a performance practice.*

5/ Electronic controllers offer far greater complexity of control design than is possible on acoustic instruments. *The ease with which the control limits of acoustic instruments can be surpassed in controller design suggests a significantly undeveloped field.*

Reflexive Control.

Reflexive control types are not unidirectional. A preliminary typology of this type has not been attempted. However a more generalized implementation enables the following account.

The next diagram (fig. 4) charts a complex sub-routine within a much larger application for the Cilia. The intention is first to bear witness to the potential complexity of control routing, and then through a reduction of a sub-routine to introduce a novel control category.

There are 7 original control sources; two microphones functioning at sub-gain levels as amplitude sensors in the mouthpiece, and five track pads outputting three continuous control streams each. They are indicated at the top of the diagram. Solid circles indicate points of dynamic control; empty circles, indirect dynamic control. Signal processing does not begin until just above the numbers 3 and 4. The final signal is output after 6. Everything above 3 and 4 is control processing and routing. Control streams split and at some later point in the configuration flow rejoin with themselves. The numbers 1-6 indicate points where this rejoining or 'self-referencing' in the control path happens. The original control shape undergoes modification and delay after parsing. Still, something of its original shape is retained at the time of rejoining.

Figure 4. Complex Control.

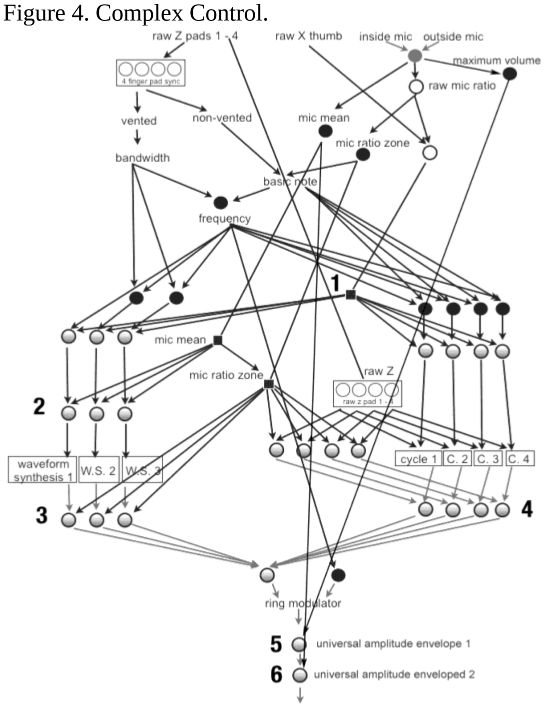

The next diagram (fig. 5) reduces and exposes these reflexive moments in the control routing. Starting from the simple split air column, four of the derived control streams are traced. They eventually all meet again but after modification and variable delay.

Figure 5. Reflexive Control.

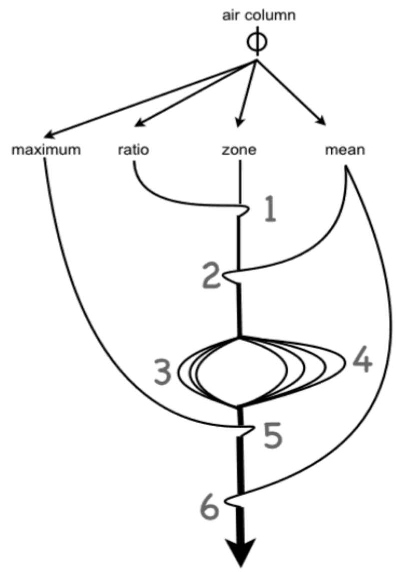

Something of the original air column gestalt is preserved in all these derivations indicated in the diagram. They split off from the original, but in various ways continuously modify the original by interacting with other control streams or by being modified through implementation in a signal path. A sedimentation of variation accrues, the variants also being marked by different delays. It is very much a *reflexive control structure.* It finds a sonic analogy in heterophony.

If one were to conduct physics laboratory analyses on the standing wave phenomena inside the shakuhachi and were able to show how ones abstracted data charts of rudimentary control gesture data are affected by the sound generation, perhaps it would look something like fig 5. Would the micro-delays micro–echoes and micro–amplifications caused by asymmetries of resistance along an interior bamboo wall find successful simulation with reflexive control data configurations? Successful or not, it would presumably be as complex as what was modeled here speculatively and creatively. Additionally, if the basic conditions for resonance are delay and repetition, then we could claim that virtual resonance is achieved in the above instance.

Reflexive control is not another magnitude of mediation between control nodes, but describes a different type of control structure. A threshold is passed when a control data network becomes self-referential. Reflexive control is a condition for any simulation of resonance, a quality of all good acoustic instruments. It is also a common if intuitive measure by which players attribute independence to their instruments, as though the instrument itself speaks. The anthropomorphizing of ones instrument is not unimportant to shimmering: musicians frequently need to know their instruments as their other.

Control Destination Types.

A control data stream always resolves into a control destination. Control efficacy is determined by how well data is routed to the whole or a part of a sound generating routine. What can we say of such wholes and parts? Let us call parts that are dependent upon one another *elements*. Pitch, amplitude and duration are such parts. No sound generation is possible when one is absent. The first level of an independent part let us call an *individual*. A simple note wherein we can distinguish no demarcation is an example. When several such individuals are sounded as in a chord or density, let us call this a *plenum*. Control can be directed to any level of whole and part.

View the categories together (fig. 6) and it is clear that the design potential is enormous.

Figure 6. Classification of Control Categories.

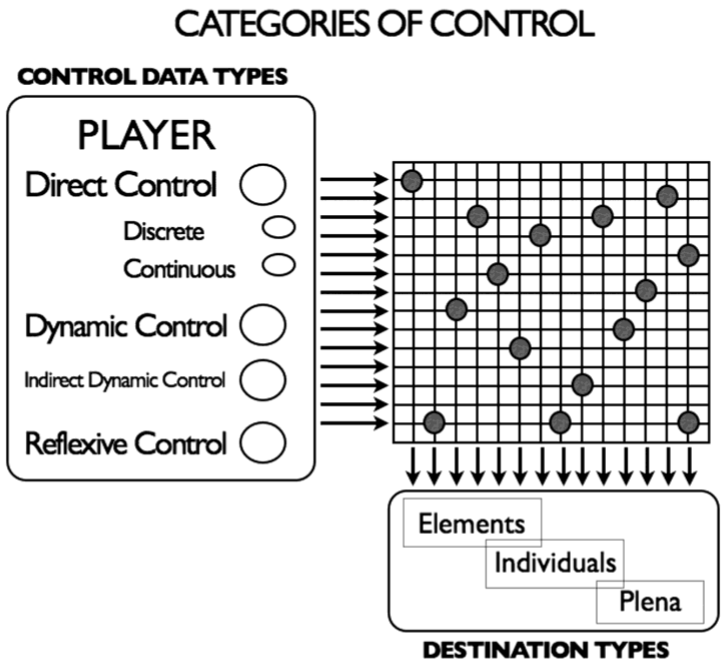

The above diagram could be described as a toolbox. It proposes control data categories to be used in digital instrument control network design. The use of a routing matrix indicates that the mapping potential is enormous. Such a conceptual toolbox is required in any systematic design approach.

It is perhaps a truism that good instruments are hard to play, taking years to master. Yet overcoming difficulty is a unambiguous measure of virtuosity. Performance complexity, nuance and hopefully, musical astonishment are the rewards at the end of such endeavor. Capable instruments facilitate virtuosity and in turn, complex control design enables the instrument. It is left to the Cilia and the like to further make the case in performance.

Reference

Rovan, J. B., Wanderley, M. M., Dubnov S. and Depalle P. 1997). *Instrumental Gestural Mapping Strategies as Expressivity Determinants in Computer Music Performance.* Hypermedia document. Paris: Mediatheque de l'IRCAM.