The Birl: Adventures in the Development of an Electronic Wind Instrument

Jeff Snyder

1 Introduction

This is the story of how the instrument I call the Birl morphed from a large, strange electromechanical contraption into a miniature wind controller. The current version of the instrument is arguably completely unrelated to the original design. Only the name has carried over, and the explanation of the name no longer makes sense with what the instrument has become. The convoluted story of the instrument's development gives some insight behind the scenes at the various design problems, creative inspirations, and unplanned discoveries that guide the creation of new instruments.

I approach instrument design with a few things in mind I want to achieve, but many of those ideas do not end up in the final product. I don't consider this a failure of the design goals, but a gift of the process. One of my favorite parts of instrument design is when ideas emerge from accidents and surprises along the way. The Birl is an example of how sometimes the resulting object evolves from the process, as much as—or even more than—vice versa.

2 Origins of the Birl (2008)

In 2008, I formed a band with fellow composer and technological adventurer Victor Adan called the Draftmasters. We had both gotten excited about the musical and visual possibilities hidden in 1980s pen plotters, those large mechanical drafting machines that print images by moving a real pen around on a page. Victor and I were collecting plotters via Ebay bids, and, encouraged and guided by fellow plotter enthusiast Douglas Repetto, we were experimenting with controlling the plotters

Jeff Snyder
Princeton University, 310 Woolworth Center, Princeton, NJ 08542, USA e-mail: josnyder@princeton.edu

live, treating them as musical instruments as well as drawing tools. We found an X/Y plotter (meaning the paper stays stationary while the pen moves in both X and Y dimensions) that seemed perfect for the job, called the Roland DXY-1100. It was big but still portable, and quick to respond to serial commands sent from Python or Max/MSP over a USB-to-serial converter, so we could control it live without much trouble. We wanted our stage act to integrate the visual and audio elements of the plotter. As we drew an image, the sound produced would be an amplification of the motor noises generated by the instrument as it followed our drawing instructions. We experimented with placing electromagnetic pickups against the stepper motors inside the plotter to get a stronger audio signal by capturing sounds directly from the electromagnetic field the motors gave off as they turned. It worked beautifully, and produced a gritty, intense sound that combined the bass frequencies of the rotary motion with the digital hissing and white noise of the drive signal being sent from the motor controller ICs. A contact microphone on the pen-up/pen-down solenoid completed the instrument, and we drilled holes in the plotter bodies to install 1/4" jacks so we could simply show up with our plotters as though they were electric guitars. Video of our performance is available(24).

While I had one of our plotters open to repair a pickup, I accidentally pushed the plotter arm while the pickup on the motor was connected to an amplifier, and was surprised by the beautiful, clear glissando that erupted from the speakers. The tone color of the plotters in our live performance was naturally harsh, evoking a sort of robot apocalypse, but this sound was sweet and subdued. The difference was that there was no power applied to the plotter, so I was hearing only the electromagnetic waveform generated by the motor's motion (indirectly through the body of the motor), without the interference from the noisy PWM drive signal. I immediately

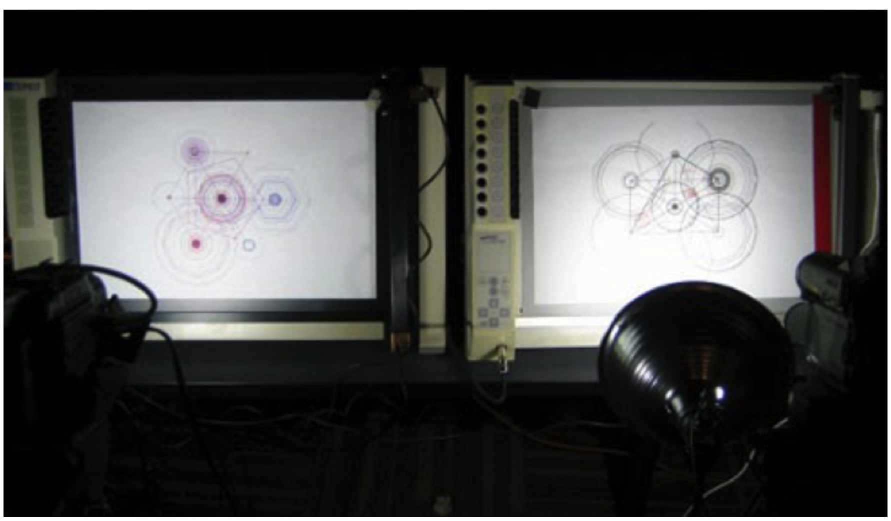

Fig. 1 The Draftmasters pen plotter band

began to ponder how I could harness that sound in a new instrument and be able to control it musically.

The primary challenge was how to turn the motor at a precise speed without driving it electronically. I built a small test rig with two stepper motors mounted on an aluminum plate, and experimented with ways to mechanically couple the rotors. If I drove one stepper motor electronically, I could use a friction belt and pulleys to drive a second passive stepper motor at the same speed. I soon realized that in this configuration I could dispense with the electromagnetic pickup, since the passive motor was acting as a generator and I could simply connect the unused leads from one of the electromagnets inside the motor itself. This resulted in an even more pure signal, approaching a sine wave in timbre. Soon, I had a working prototype that allowed me to accurately produce desired pitches within a range of a two octaves. Going above the usable pitch register resulted in a loss of torque, stalling the motors. Going below the usable pitch register produced a waveform that got more rugged as the steps in the motor became audible and the rotor inertia could no longer smooth out the tone. Swapping out the motors for different stepper models moved this pitch range around, but didn't manage to expand it much. I did find I could raise the register easily by varying the pulley sizes, for instance a pulley size ratio of 4:1 produced the expected pitch shift of two octaves up.

I later noticed a description of a similar idea in Handmade Electronic Music by Nic Collins(2), although he uses a DC motor instead of a stepper. In terms of historical precedent, the Hammond organ is also based on a related principle(1), with a spinning metal tonewheel being sensed by an electromagnetic pickup, however in the organ the pitch changes are produced by switching between several pickups pointed at tonewheels with different numbers of teeth, rather than by changing the speed of rotation.

In my early experiments, I was controlling the driver motor with MIDI signals sent to an AVR microcontroller brain. I found myself wondering how this new instrument should most naturally be controlled. Was there an instrumental interface better suited to this sound production method than any other? The sound properties of the stepper motor synthesizer were the following:

- It was monophonic. There was no possibility for polyphonic sound without creating multiple identical mechanisms, and this seemed unnecessary to me at the time.
- The tone color was very dark, with a strong focus on the fundamental and first harmonic.
- The amplitude could be controlled electrically, with a VCA, and had no "natural" envelope (such as a plucked string or percussion envelope).
- The amplitude was also somewhat coupled to the pitch, as higher frequencies produced higher amplitude signals, perhaps through inertia of the rotor.
- There was a natural vibrato to the sound, caused by slight inaccuracies in the pulleys.
- There was a natural portamento to the sound, due to the need for speed ramping in the motor control to avoid stalling. When moving between nearby pitches it

was inaudible, but when going from a very low note to a high note a ramp of more than 10ms was usually necessary.
- There was brief but noticeable overshoot to the pitch contour when changing speeds, due to the stretching of the rubber friction belt.

Several of these properties suggested that a control paradigm based on a wind instrument model would make sense. The most obvious was its monophonic nature, which is commonly a property of wind instruments not shared by most string instruments, keyboard instruments, or percussion instruments. Also, the dark tone color immediately reminded me of a recorder or flute, and many people commented on its "birdlike" character, which brought to mind whistles, ocarinas, and other wind-powered instruments. I quickly started working on a wind-style interface for the new electromechanical oscillator.

3 The First Birl (2009-2010)

While I worked on the interface, the instrument took the name "the Birl". The word "birl" seemed appropriate in multiple ways: It is an old English or Scottish word for "rotate with a whirring sound," a type of bagpipe ornament, and slang for "to carouse." "Birling" is also the name of the sport where lumberjacks run on logs in a river, a connotation that delighted me.

The first Birl was a large instrument, held between the legs and connecting to the floor with a cello pin (see figure 2). The top of the instrument was a wind controller, with mechanical momentary push buttons arranged for the fingers and thumbs of the left and right hands. The base of the instrument was a large wooden resonator, in keeping with the concept of "acoustic electronic music" I had developed in my dissertation work(13). The body of the instrument was made from 1/4" birch plywood cut on a laser cutter, and the front was made from a thin spruce board with internal spruce X-bracing, like the top of a guitar or harp. The body housed the motor and pulley system, which was screwed into the inner right side of the instrument, and a hole in the center of the body was decorated with a laser-etched lute rose. Screwed into the spruce front from the inside was a Rolen Star vibration transducer, which resonated the top-plate to create the instrument's acoustic sound. By resonating the electronic sound through an acoustic body rather than from a speaker cone, I could achieve both a more natural radiation of the sound in space, and I could get an individualized color for the instrument. Each wooden resonator imparts a unique sonic filter onto the electronic sound passed through it, emphasizing certain frequencies and attenuating others. This idea extends back to instruments like the Ondes Martenot, an early electronic instrument with several acoustic resonators, and in my case was influenced by David Tudor's installation piece, ***Rainforest IV***. As for breath input, by 2010, a mouthpiece with a breath pressure sensor was fitted to the top of the instrument, but for the first performance in 2009, a Yamaha BC1 breath controller was used, since a more tailored custom solution hadn't yet been completed.

The instrument was self-contained, except for the power amplifier needed to drive the vibration transducer. Due to the inefficiency of the wood top when compared to the paper cone in a standard speaker, more watts were needed to get stronger sound levels than a small amplifier that could fit inside the body would have been able to provide. Therefore, the signal path was:

- The pattern of pressed pushbuttons for the keying system is sensed by a microcontroller, resulting in a frequency being sent to the motor controller.
- The motor controller controls the ”drive” motor, which then spins the passive motor via a friction belt.
- The electrical signal generated by the passive motor is sent to a voltage controlled amplifier (VCA). The amplitude of the VCA is directly controlled by the voltage from the breath pressure sensor.
- The audio signal from the VCA is sent to a power amplifier, which sends an amplified signal to the vibration transducer.
- The vibration transducer mechanically vibrates the spruce top-plate on the front of the instrument, producing the acoustic sound the performer and audience hear.

This was the system used for the first public Birl performance, a Wet Ink Ensemble(26) concert where my dissertation piece, ***Concerning the Nature of Things*** (13), was premiered. I had written two parts for the Birl in the composition, having determined the most useable pitch range and knowing the basic timbre

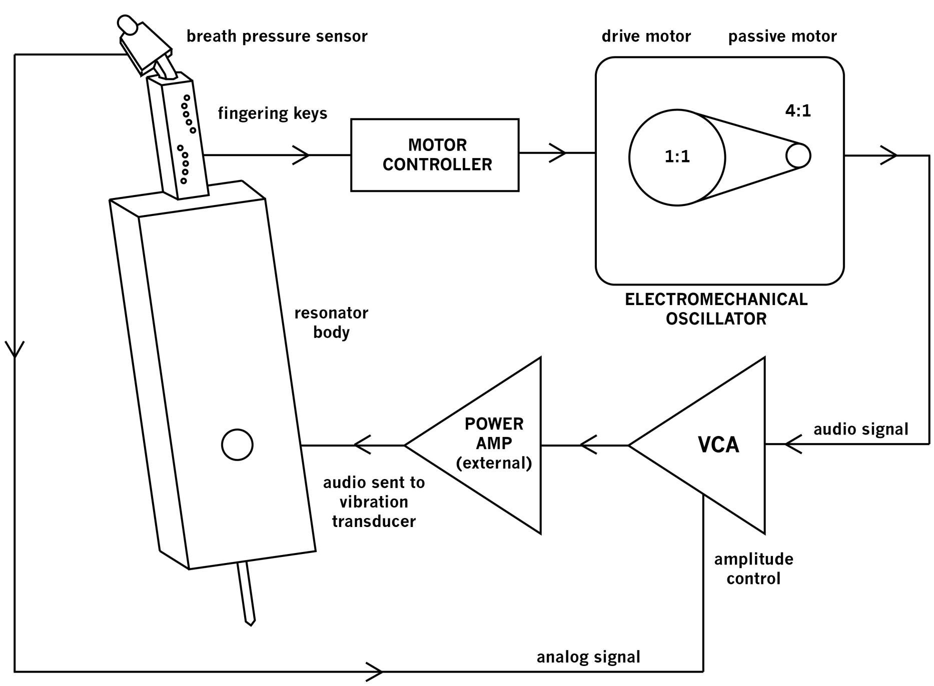

Fig. 2 The First Birl, diagram

the instrument would have, but not having actually finished the instruments' construction. About a month before the performance, worried about getting both Birls functional in time, I decided to focus on finishing one instrument, and cut the second Birl part from the score. Erin Lesser, the flautist for Wet Ink, learned to play the new instrument and provided feedback during the design and development phase. I had built the fingering system with only four buttons per hand, one for each finger (not counting the thumb buttons), so we had to work together to design non-standard fingerings for the pitches that weren't well served by the simple recorder-based keying system (such as a low C, C#, and F#). On the back of the keying system, I had added two buttons for the left thumb, for octave up and octave down, and three buttons for the right thumb, allowing for maneuvering within my Adaptable Just Intonation(13) tuning system. The fingering-to-pitch mapping was implemented as a lookup table, with specific patterns of open and closed buttons resolving to a particular note. Lesser tackled the unfamiliar instrument with enthusiasm and managed a very expressive performance even with the limited rehearsal time resulting from the instrument being completed barely a month before the premiere. However, after the performance, I was left with a considerable list of design problems Lesser had discovered with the instrument. It should be noted that I don't play any wind instruments, so I was heavily reliant on information from Lesser and other musicians who tested the prototypes.

First, there were tuning problems in the upper octave. I had switched shortly before the concert from metal pulleys to plastic pulleys, since the reduced weight

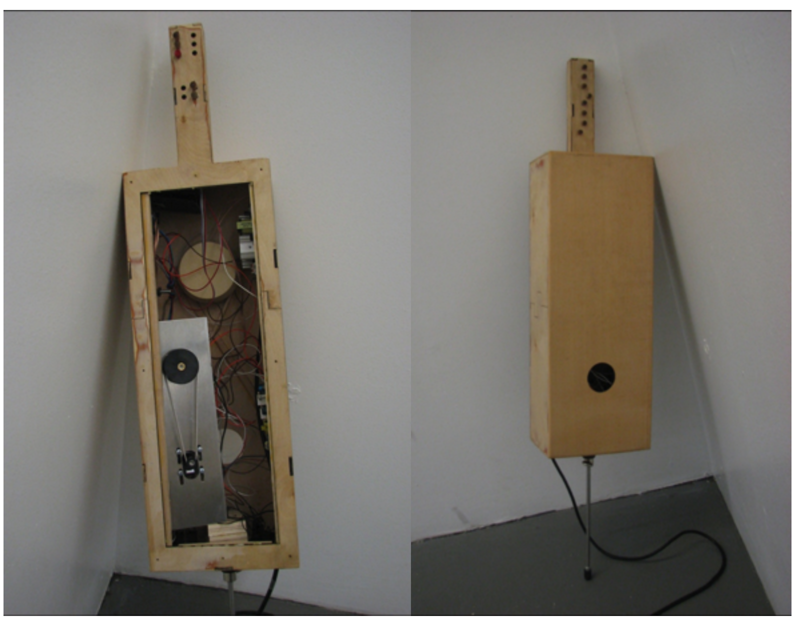

Fig. 3 The First Birl, prototype

allowed me to lower the ramp times for the motors. The plastic pulleys were not as precisely sized, though, and the difference caused seriously flat pitches in the high register that I didn't recognize until the day of the performance. In my music, this is especially problematic, since a great deal of attention has gone into precise Just Intonation(3). This was easily fixed by switching to precision steel pulleys from SDP/SI. But the added weight meant I needed to find motors that could handle more torque.

Another serious problem was the acoustic sound caused by the keys. I had used tactile pushbuttons because they had a satisfying click response when actuated, but when the pushbuttons were mounted in the resonator body, they were acoustically amplified to an unacceptable level. I liked the key click sounds in principle, but they made truly quiet playing untenable. Lesser also noted that the actuation force required for the pushbuttons was far above what was normal for a flute or other wind instrument, and was tiring for her fingers. I decided to redesign the button system.

Fig. 4 Stina Hasse plays the First Birl

Most problematically, the electromechanical tone generator system produced unintended acoustic vibration noise in the resonator body in addition to the intended electrically amplified signal. This sound was not unpleasant in itself, as it was in tune with the electrical signal and changed pitch with the notes being played, but since it was mechanical in nature it could not be electrically attenuated by the VCA. Therefore, whenever the motor was spinning, the instrument was humming, even if the VCA had attenuated the volume to ”off.” Since the motor in the Birl changed speed with every new pitch, the humming sound seemed unnatural, as though the instrument didn’t really stop sounding the note when the performer ceased to blow into the instrument. The motor couldn’t be stopped between notes because ramping up from a standstill would create a dramatic glissando on every attack. This definitely had to be solved, and a solution was not immediately obvious.

There were also more minor issues I hoped to address in the next iteration. The higher pitches from the motor were naturally louder for reasons I didn’t entirely understand. This is also the way many woodwind and brass instruments operate - it is difficult to play quietly in the high registers as it takes more breath to overblow the notes -and Lesser was able to compensate by reducing her breath pressure for higher pitches, but it was very difficult in faster passages with leaps. It seemed worthwhile to build a more automatic compensation system into the instrument. It was difficult to minimize pitch glitches when changing many keys at once, such as going over the ”break” in the instrument, where the fingering changes most drastically from one note to the next. This is also a problem for acoustic wind instruments, but it seemed much less forgiving in this instantaneously calculated digital version. Lesser told me that the majority of her practice time was spent working to minimize these glitches.

4 The Second Birl (2011)

The first issue I dealt with in the second iteration was the unintended acoustic vibration noise. I removed the stepper motor system from inside the resonator and made a prototype board that combined a custom power amplifier with the stepper motors and motor controllers, the analog VCA, and the power supplies for all the circuitry. This worked well and sounded much better - I recorded a studio version of ***Concerning the Nature of Things***(13) using this prototype, with the stepper motor board in an isolation booth to keep the mechanical noise away from the resonator and microphones. However, it was very messy, fragile, and not really useable in live performance. Having a board with the motors on it backstage seemed impractical, so I decided to try to build an enclosed box to acoustically isolate the motors.

I designed a box about 12 inches square, made from birch plywood. The front was an aluminum panel for controls and jacks, and the back had an acrylic window that made the motor system visible. When the first Birl had hidden the pulleys inside the resonator I had been disappointed that they weren’t part of the visual signature of the instrument, and, thinking back to the Draftmasters, I wanted to give the audience more of a view into the unusual workings of tone generator. Inside the box, the motors were mounted on a thick aluminum plate, and the plate was suspended off the base of the box with rubber vibration isolation mounts. The box itself was isolated from the table or floor with large rubber feet. I lined the inside of the box with vibration damping rubber-lined foam, intended to muffle sounds from boat engine rooms.

I designed a printed circuit board (PCB) with the stepper motor driver circuitry, and another PCB for the audio processing of the electrical signal from the passive generator motor, designed to stack with a PCB for the control panel components. The goal was to get everything that had been on the messy prototype board into a nice, neat box that could be on stage next to the controller/resonator body and connected with a short MIDI cable. I left out the power amplifier once again, due to space, heat, and weight considerations, but I decided to expand upon the shaping of the audio signal.

In the original instrument, the audio path was simple. The waveform generated by the passive motor went directly through a VCA for amplitude control and was converted into acoustic sound through the driver transducer. In the time since I designed the first Birl, though, I had started to see the tone generator as an oscillator for a system that could be a more complete synthesis voice. Therefore, I chose to build into the new Birl some extended functionality that allowed further shaping of the sound.

First, I added to the motor mount the ability to drive two simultaneous passive generator motors. The drive shaft of a single driver motor was affixed with two pulleys, and these pulleys drove the two passive motors with different pulley size ratios, one at a 2:1 ratio and the other at a 4:1 ratio. This meant I could mix two resultant oscillator signals, one an octave higher than the other. On the audio PCB, I made an "oscillator" section that allowed for crossfading between these - like an 8' and 4' stop on an organ. After the oscillator crossfader, the mixed signal went through a waveshaper that could add high harmonics to the signal, essentially a distortion

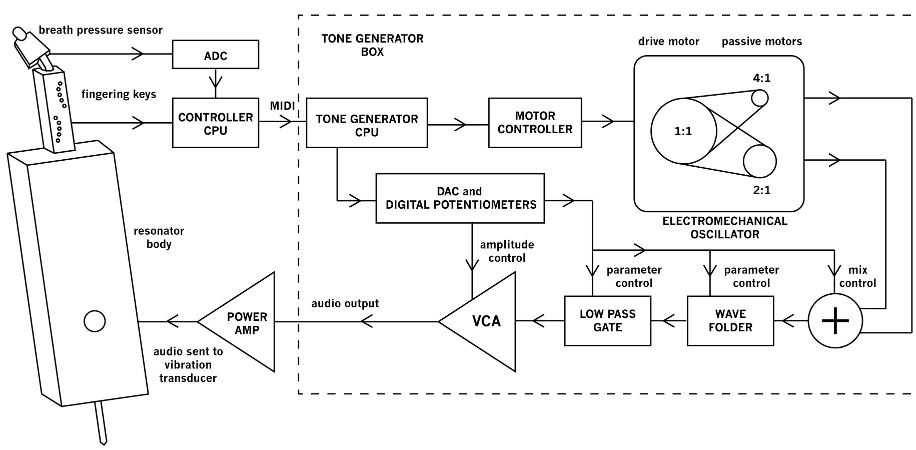

Fig. 5 The Second Birl diagram

circuit. This allowed for more timbral possibilities than the original ”natural” waveform. After the waveshaper, the signal passed through a Low Pass Gate, based on Don Buchla’s design from the 1970s(5) - a vactrol-controlled lowpass filter acting as a VCA. The signal then went through a final VCA and finally to an output jack on the box. Aiming for maximum flexibility, I designed the whole ”voice” as a semi-modular system, with patch points for each input and output, and voltage control inputs for all parameters. I also added digital-to-analog converters (DACs) and digital potentiometers to allow computer or MIDI control of the analog functionality.

Part of the reason for adding comprehensive digital control was the need to compensate for the higher volume in the upper octave of the instrument. With the final VCA controlled digitally, I could easily program curves to apply to the amplitude based on the frequency of the oscillator, allowing for a more even response.

Once I had assembled the audio and motor control PCBs, I installed them in the box that could now be controlled from the Birl wind controller or using MIDI from a computer.

While the new instrument had a unique sound, there were design problems preventing it from being as useable as I had intended. The mechanical noise had been solved - the box was very quiet and no longer caused any acoustic issues, but electrical noise problems arose. The noise from the motor controllers was audible in

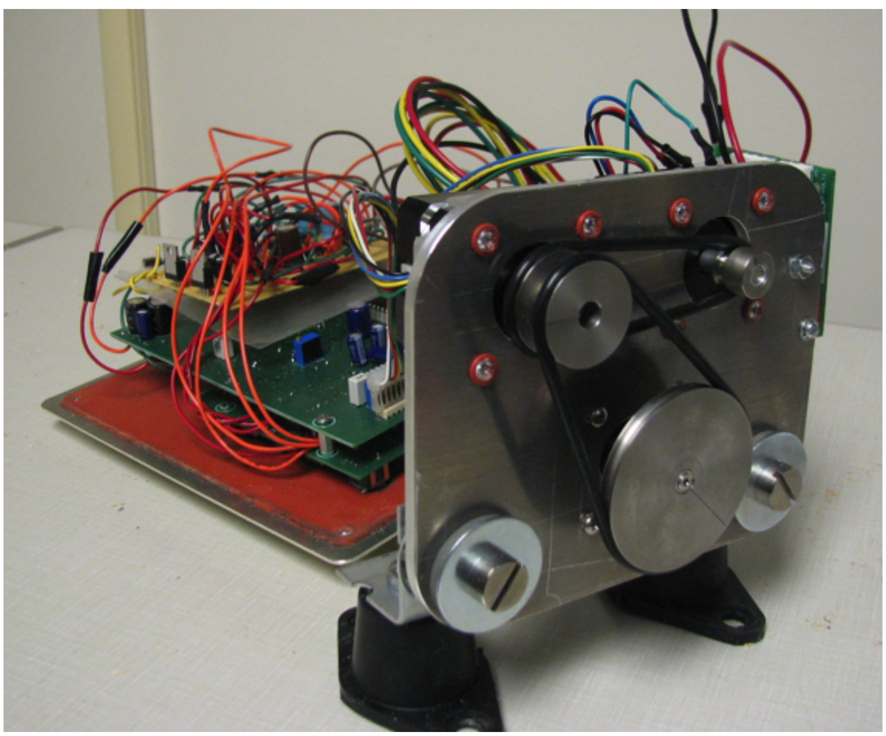

Fig. 6 The Second Birl PCBs and motors, taken out of the box enclosure

the audio circuits, despite carefully isolated ground planes and independent power supplies. The physical proximity of the circuits was just too close to expect the high currents of the motors to not interfere with the audio. This problem was similar to the acoustic sound problem, in that the injected motor noise continued even when the VCA was off, and the sound was related to the frequency of the spinning motors, so it couldn't be easily ignored.

Even more seriously, the acoustic insulation had come at the price of heat insulation. The motors generated heat, and despite heatsinks and the 1/4" thick aluminum mounting plate, once I had sealed off the air transfer from the inside of the box to the environment, the heat had nowhere to go. After about a half hour, the box was hot to the touch and needed to be turned off to cool down. This also limited the instrument's ability to be used in a full-length concert. I've since discovered the existence of heat pipes, a technology designed to handle this very problem, but I have yet to try that solution.

Testing this system with composer and flautist Natacha Diels, we decided that despite the problems, the tone generator sounded quite good. The controller, however, was still awkward, and the discrete nature of its pitch control was not ideal. The pushbuttons needed to be replaced with something more comfortable and mechanically quiet, ideally something that sensed a continuous change on each key rather than simply an on/off event. Also, travelling to events in far away locations convinced me the controller needed to be separated from the resonator for porta-

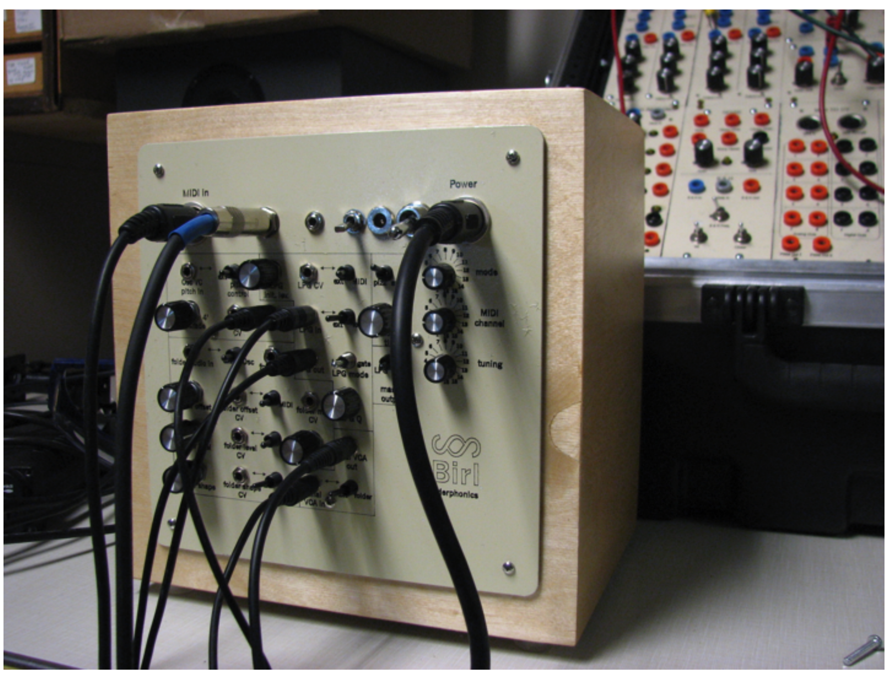

Fig. 7 The Second Birl tone generator box control panel

bility. This led me to focus the next phase of research on improving the controller portion of the instrument.

5 The Third Birl (2012-2015)

Now that I had decided to break the instrument into three separate pieces—controller, resonator, and tone generator—the first order of business was improving the finger sensing on the controller. I wanted the instrument to be more like open-holed wind instruments such as the bansuri and the tin whistle. I built prototypes to test both infrared (IR) reflectance sensing and capacitive sensing, and found the response for the capacitive sensing fit my needs better, so I moved forward with that.

With the resonator no longer being a necessary part of the controller, I had to reconsider the visual design of the instrument. The resonator on a cello end pin had given the instrument a striking and unusual visual presence, and with the breath controller and key system removed from the resonator base, I worried about the instrument taking on a ”typical” soprano saxophone visual style, like the Akai EWI(15) and the Yamaha WX(21) instruments. I decided to try to keep the vertical orientation of the instrument, and enforce that by creating a mouthpiece to angle the instrument more like a bass clarinet or tenor saxophone. I imagined the instrument would be played seated, and tried to design it so that it would be comfortable to play with the base of the instrument resting against the player's leg. I found the laser-cut plywood design of the first Birl somewhat ugly, so I aimed to create a design that could be milled out of solid wood, to have a more beautiful presentation.

I had only limited time access to a 3-axis CNC mill, so I designed the wooden body to be millable without flipping the part, to avoid wasting time with realignment

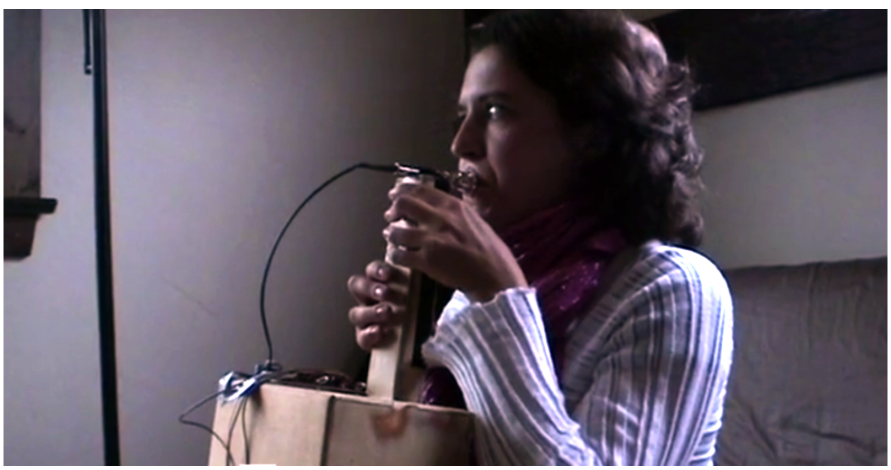

Fig. 8 Natacha Diels tests the Second Birl

during the milling process. I made it with a clamshell type of design, so that I could mill both parts from only one side, and then screw the two parts together around the circuit board to house the electronics. Before I milled the wooden version, I tested the design using a 3D printed white plastic model. The shape of the new enclosure was highly influenced by Scandinavian design; I had recently visited the Danish Museum of Art & Design(19) and I found the sharp corners and clean lines of Jacob Jensen's(20) designs for Bang and Olufsen(17) to be particularly inspiring. This led me to a relatively boxy design, somewhat reminiscent of a 1970s Volvo automobile, but unusual for wind instruments, which are typically cylindrical or conical in shape. In a way, the melodica is a closer visual reference than a recorder or flute.

The controller's circuit board had undergone several revisions since the first Birl. Now that continuous data from the fingers was going to be possible, I wanted to improve the data throughput from the device, so I switched from standard serial MIDI to OSC over Ethernet. I designed the brain of the controller around an AVR32(16) microcontroller because it could easily send Ethernet information. The capacitive sensing for the keys was handled by a Cypress PSoC(18) microcontroller using the CapSense CSD library. The keys themselves were simply aluminum standoffs, since any metal object can be a sensor using that technology. The initial revision of the board added multiple capacitive sensors for embouchure sensing; I was hoping I could retrieve some reasonable data on lip position by doing machine learning on the data from sensors touching the top, bottom, and sides of the mouth.

While I still at this point considered the electromechanical oscillator to be an important part of the instrument, it was slowly beginning to seem more optional, rather than essential. As I made plans to lend Birl prototypes to musicians for "in-the-field" user feedback, the impracticality of the motor system for anyone's use but my own became more and more apparent. For instance, in transporting the Birl for to show at an event at the Mass MOCA museum, the motor connections were damaged, and I had to do some emergency surgery on the instrument that a musician couldn't be expected to do. I decided that the instrument needed to have the option of a simpler digital voice for musicians who weren't as interested in the strange electromechanical oscillator, but found the wind controller useful. I designed an internal digital synthesis circuit to live inside the birl controller itself, which could

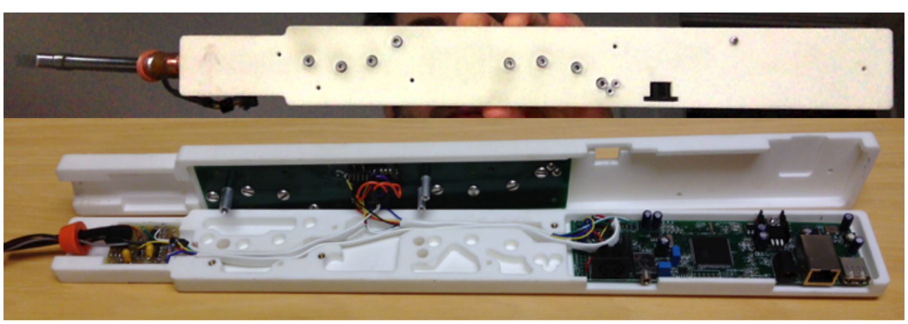

Fig. 9 Third Birl early prototype, 3D printed with vacuum cleaner mouthpiece

be used instead of an external computer, although in practice those who used the instrument in this next iteration always used it with an external computer, to allow themselves more flexibility in synthesis options.

One major hurdle was the question of how I would map data from multiple continuous key sensors into a single pitch output. I wanted expressive pitch bends to be intuitive for players to execute with their fingers, but it wasn't immediately clear how that mapping should work. The first idea I had was to use machine learning. My colleague Rebecca Fiebrink had written a program called Wekinator(4), which allowed for easy experimentation in applying machine learning techniques to digital music making. It could take OSC data in and send OSC data out, which was perfect for my new Ethernet-ready instrument. Fiebrink, who happens to be a flautist as well, joined me to try some tests where we trained a neural net on a simple flute scale. The results were astounding. The trained model made choices that were surprisingly intuitive, and the resulting system allowed Fiebrink to bend each pitch up to the next. With this encouraging solution to the pitch-mapping problem showing the way, I started testing the Birl with other professional performers.

The primary testers of the third Birl in 2013 and early 2014 were jazz saxophonist David Schnug, and avant-rock saxophonist Sam Hillmer. Both were most interested in using the instrument as a controller for digital synthesis. At the time they first tried it, it had a vacuum cleaner attachment for a mouthpiece, and the embouchure sensors were pieces of copper wire covered in heat-shrink tubing and wire-tied to

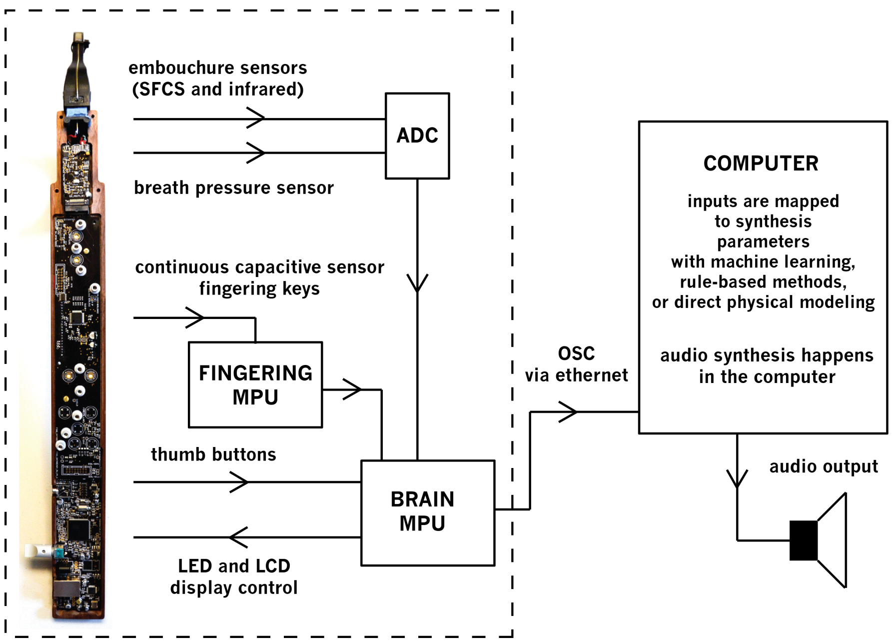

Fig. 10 Third Birl diagram

the mouthpiece. I tried using neural nets to map data from the embouchure sensors, but the values drifted significantly during testing, possibly due to the loose wire-ties allowing the sensors to move. I still got a surprisingly useful result, as can be seen in an online video(22). I mapped the embouchure sensors to timbre parameters on both a simple FM synthesis patch and a wind instrument physical model (the ”blotar” by Dan Trueman and Perry Cook).

Both Schnug and Hillmer found the pitch bends reasonably easy to control. Interestingly, Hillmer was entirely focused on the pitch bending possibilities - in the software I had left an option to turn it off to allow easier discrete fingering, and he commented that he would never use this feature. He enjoyed the strange, yet controllable bends that were possible, and even invented an extended technique in his first performance with it - wearing rubber gloves while playing to reduce the sensitivity of the sensors and make the bends more intense, eliminating the ability to achieve exact pitches but exaggerating the slippery weirdness of the pitch mapping. He used this technique in a live installation performance in NYC called Apparition(23).

Unlike Hillmer, Schnug was interested in both the continuous and the discrete fingering options. What Schnug found most exciting was when I mapped the embouchure sensor data to physical modeling synthesis parameters. He was extremely interested in the possibilities opened up by this mapping, which allowed for very unusual sound transformations controlled entirely with the mouth. As a musician who specializes in experimental and free jazz, he wanted a way to have a wild range of sounds under predictable embouchure control, to provide a world of timbres that could rival the saxophone’s variability.

The feedback from these performers started me thinking in another direction. When the original idea for the Birl was forming, the choice of a wind control paradigm was simply because I wanted a ”winds” section in the ensemble of invented instruments I was building, and the properties of the electromechanical oscillator I was experimenting with suggested that wind control would be a good match. It had a wooden resonator because it was intended to be used for my own chamber music, specifically a suite of pieces I am writing in just intonation for electronic instruments with acoustic resonators. While the electromechanical oscillator and wooden resonator of the earlier Birl designs made sense in my original context, they weren’t ideal for the needs of these professional musicians. Though these musicians were very experimental and open-minded, touring internationally would not be easy with those contraptions in tow. However, both of these performers found the possibility of a more expressive wind controller aimed at the needs of an experimental musician to be very interesting. While existing commercial wind controllers were certainly useable in those contexts, there seemed to be room for something that sought to fill those needs more directly - a wind controller designed with a player like Evan Parker in mind.

Rather than considering the Birl controller as just part of a whole that I would later reconnect to the oscillator and resonator, I began to imagine what the controller could be without those more mechanical components. For the past few years, the controller has become separated from electromechanical ideas that initiated the design process, and both the oscillator and the resonator have been placed on the back burner while I solve the problems posed by these new design goals. They may rejoin the Birl someday, but it will be as accessories rather than as the heart of the instrument.

After a few adjustments to the dimensions of the body and the shape of the mouthpiece, I milled several bodies out of walnut and maple during a residency at the Haystack Mountain School of Crafts. I also made a new CAD design for a mouthpiece that I had 3D printed.

I wrote a paper for the New Interfaces for Musical Expression (NIME) conference(14) about the application of neural net machine learning to the pitch-mapping problem. While I was showing the Birl at NIME, a colleague questioned why machine learning was really an improvement over a rule-based approach, and my answer didn't really satisfy me. I was interested in the fact that machine learning would easily allow the user to modify or completely alter the mappings by entering new training examples, but it was true that a rule-based approach to the mapping would be significantly more efficient, and, although it would be more difficult for the user to change the behavior to their preferences, the inner workings of the algorithm could be understandable to an expert user, unlike the more black-box internals of a neural net. The following year I had a student develop a rule-based pitch mapping algorithm that worked well for situations where user-defined fingerings were not necessary.

In addition to the neural net and rule-based pitch mapping options, I was curious about using a physical modeling approach. Gary Scavone had written several papers

Fig. 11 Leila Adu and Dave Schnug test the first prototype of the Third Birl

describing physical models of woodwind toneholes(8)(12)(9)(11)(10) that could be continuously varied from open to closed. Since I was already exploring digital synthesis options as an alternative to the electromechanical oscillator, that suggested another idea for the mapping: rather than trying to get a ”pitch” parameter from the array of floating-point values from the key sensor readings, one could instead generate the synthesis directly using a digital physical model of a tube with holes of the right size in the right places. A student and I managed to create a model of a full tube with continuous toneholes for every key on the Birl. Once that was completed, I needed to figure out how to put the virtual toneholes in the right places, and with the right radii. The student wrote a solver allowing the user to enter a scale in cents. From this user-defined scale, the solver designs a virtual tube with the correct tonehole placement. It worked great, although the tuning was not completely

Fig. 12 Dave Schnug and Pedro Eustache try later prototypes of the Third Birl

accurate - more work needs to be done on that front. One interesting advantage of this method is that extended techniques like multiphonics arise naturally out of the system, and the results of half-holing will usually be intuitive. One downside is that, unlike the machine learning or rule-based systems, the available pitches are limited by the number of keys, making this approach better for creating instruments similar to those with only a few open holes, such as the shenai, rather than instruments with complex keywork, like the oboe.

Throughout the 2014/2015 school year, I was lucky to have four wind players in PLOrk, the Princeton Laptop Orchestra, which I direct. We were working up a program of 15th and 16th century music arranged for electronic instruments, so the Birl was a perfect addition to our electronic orchestra. This was a fantastic opportunity to get real-world test data on the problems of the current prototype.

The PLOrk musicians made several requests. First, they found the finger holes (at the time just holes drilled in the enclosure with standoffs inside the holes) to be too small and hard to locate without looking at their hands. I experimented with inserting thumbscrews into the standoffs to make physical touch plates for the keys, similar to the capacitive keys on an EWI. With that change, some of them noted that it was now hard to avoid accidentally activating keys when you just wanted to rest your finger near them. I added 3D printed plastic guards around the keys, intended to give fingers a place to rest when not pressing the keys.

We found that the sensitivity curve on the keys wasn't ideal, as there was a jump in the value when the sensors went from touched to not touched. I experimented with applying clear vinyl stickers to the keys to eliminate this threshold, then eventually found a durable solution in epoxy spray. Some of the players didn't like the custom mouthpiece, and instead attached a clarinet mouthpiece, which worked fine since the

Fig. 13 PLOrk performs with Birls

embouchure sensors were not yet functional in these prototypes anyway, but doing so did bring the instrument's profile back into soprano saxophone territory. I had made some versions of the circuit board with capacitive sensors for the thumbs as well as the fingers, but the PLOrk members greatly preferred mechanical switches for the thumb buttons. The PLOrk members also complained of the difficulty in avoiding glitches in the pitch output of the instrument when changing octaves. This problem remained from the first Birl, and it remains unsolved as yet. We performed several pieces using the Birls at the final concert of the year; video is available(25).

I worked on further developing the embouchure sensors, trying a swept frequency capacitive sensing technique(7) as a way to get more reliable data from the lip sensors. It worked well, and I designed a new embouchure sensor circuitboard to integrate the new technology. However, I noticed that although the new lip sensors did a decent job of detecting lip position, they did not register the inside of the mouth, which began to reveal itself as an important element of embouchure. I tested several methods of sensing the space inside the mouth, including acoustic sensing (via a microphone) and infrared, and found infrared to be remarkably reliable.

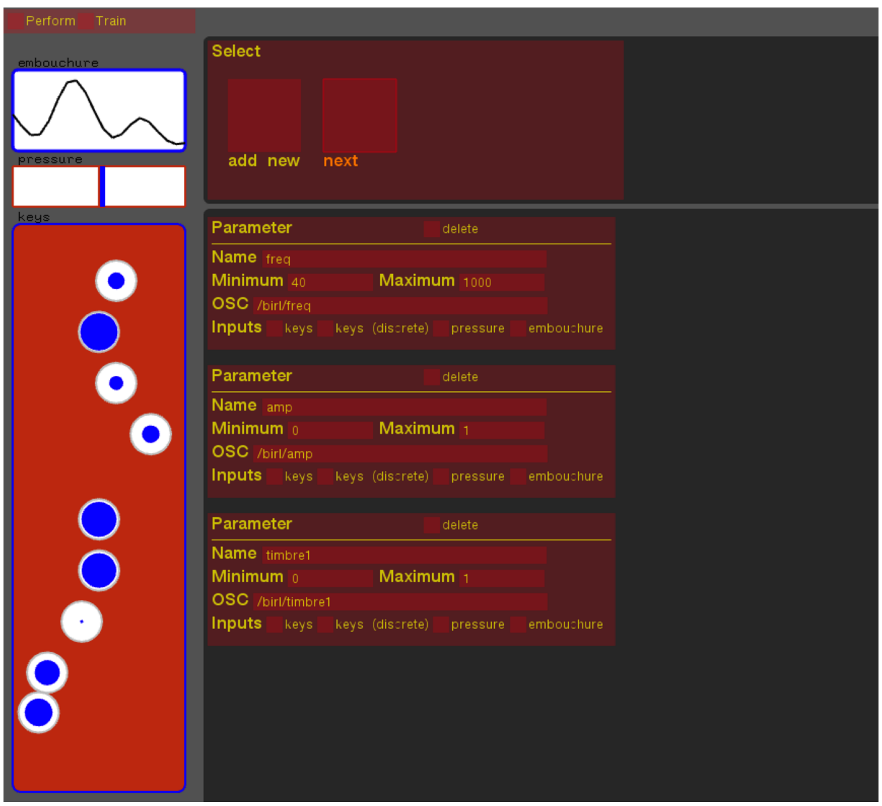

Fig. 14 Neural net training software for the Third Birl by Gene Kogan

I was now considering the controller as its own independent instrument, not needing either the resonator or the electromechanical tone generator. The new prototype was also designed in such a way that it wasn't out of the question to affordably make it in multiples. I began envisioning the release of a future version of the instrument as a product.

I had released the Manta, a hexagonal grid touch controller, as a product in 2008, but I wasn't as confident of an existing market for the Birl, since electronic wind instruments have a bit of a cheesy reputation, and there are several existing wind controllers on the market. However, the continuous key sensing and general feel of the instrument seemed different enough for a successful product. A friend suggested that I develop a simplified version of the instrument, focused on making it affordable. Thus I started the design of the fourth Birl, the MiniBirl.

6 The Fourth Birl (2015-present)

In designing the MiniBirl, I decided to eliminate the most expensive parts: embouchure sensing and embedded sound synthesis. I needed to reduce the cost of manufacture, simplify the user's experience of communicating from the instrument to a computer, and improve the portability.

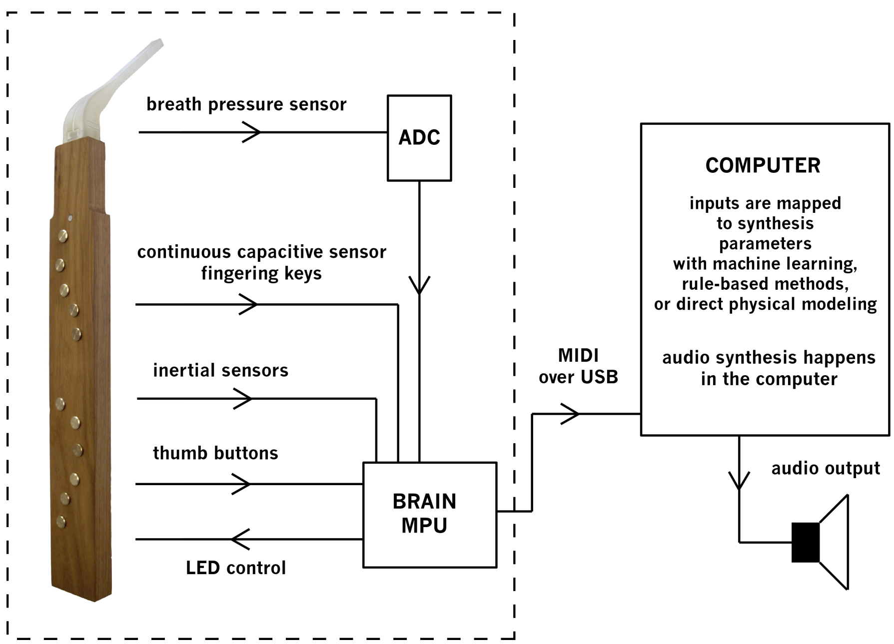

Fig. 15 Diagram of the Fourth Birl

I reduced the cost by stripping the design down to a single 2-layer circuitboard and removing all the complex parts. Instead of a separate brain microcontroller, the key sensor microcontroller would also be the brain. In addition to removing the embouchure sensing and the internal synthesis, I removed the LCD display. I replaced the Ethernet jack with a USB jack. Using USB as the only communication meant that I no longer needed a separate power cable, which made the instrument more elegant. What remained were the continuous keys sensing, the breath pressure sensor, and an indicator LED for the breath pressure. I added two new features that didn't add significant cost: an X/Y touchpad for the right thumb and an accelerometer/gyroscope IC to detect orientation of the instrument. I replaced the combination of standoffs and thumbscrews that I was using as keys with Chicago screws, which integrated several parts into one and lowered the cost of each key.

Once I had removed the Ethernet jack and the LCD screen, I realized that I could reduce the thickness of the instrument. I simplified the design from the wooden clamshell to a single piece of wood, milled from the back, with a fiberglass panel covering the cavity. With the reduction in circuit complexity, I was also able to shorten the length of the instrument considerably.

After switching to USB I had to decide on a protocol, and I chose USB-MIDI so as to be compatible with most music software. Sending the data as 7-bit packets would be limiting, but pairing bytes to create 14-bit messages made for reasonable resolution. I'm still working on the problem of how to encode the continuous pitch information in a way that standard music software can understand. Continuously varying pitch over a seven-octave range does not easily translate to note-ons, note-offs, and pitch bend data.

The new MiniBirl prototype had its premiere in April 2015, played by Sean Mac Erlaine of the group This is How We Fly. The following year, it was tested by professional wind players Pedro Eustache, Noah Kaplan, and Steve Lehman, who also provided valuable feedback.

Lehman was interesting in that he was almost entirely uninterested in the continuous finger sensing possibilities. His style emphasizes clean, precise, accurate playing - as he told me ”I've worked for years to sound like a computer on the saxophone!” To this end, he also found the lack of tactile feedback on the touch sensors to be a significant downside, not balanced by the continuous sensing possibilities. On the other hand, he really liked the X-Y touchpad for the right thumb and the unusual look of the instrument.

Kaplan is a saxophonist and composer who is deeply fascinated with alternate tunings, so he was excited by the possibilities in Gene Kogan's machine learning software. We spent a half hour training the system to recognize several unconventional fingerings as being either quarter-tone deflections or comma alterations. This was interesting to him, since the existing commercial wind controllers output only ”cooked” pitch information, not allowing access to the actual key-press data to reprogram unusual fingering options.

Eustache came to visit my New Jersey lab from California, where he works as a film soundtrack session musician, with credits on several films, including Pirates of the Caribbean. His feedback was incredibly detailed, since unlike any of my other test users, he has been regularly playing commercially available wind controllers for many years. He had very precise and helpful ideas about how the thumb buttons should be situated, and other physical layout details. He was most excited about the continuous fingering option, since he is a multi-instrumentalist (most of my test users had been either saxophonists or flautists) and he is regularly called upon to play a huge variety of wind instruments from various cultures, many of which use open-holed fingering techniques. He found the open-holed fingering options to function very well, although he wanted the keys closer to the sides of the instrument so that he could roll his fingers off of them as one does on a bansuri, and he thought the relatively sharp corners used in the design were counterproductive for this technique. He was also overjoyed about the small size of the instrument, especially the fact that it could fit in a backpack. He was theoretically very excited about the embouchure sensing, although it wasn't working when he visited, and he loved that the breath sensor was fast enough for fluttertongue effects. I was surprised by his enthusiasm for the visualization tool I had built with Gene Kogan, which allowed the user to view the data from the fingering sensors in real time. He wanted to be able to record and slow down the visualization so that he could analyze the synchroniza-

Fig. 16 The author testing out the Fourth Birl

tion mistakes made when leaping across octaves, imagining a tool something like Duncan Menzies's P-bROCK training program(6).

While refining the design, I started to see the MiniBirl not just as a simpler, cheaper cousin to the Birl, but potentially as the core of the Birl itself. Instead of making two Birls, one larger and one smaller, I could make the MiniBirl modular, so that if one wanted to add internal audio synthesis or embouchure sensing, one could do so by snapping on another piece, perhaps even a wooden resonator body or an electromechanical oscillator, bringing the instrument back to its conceptual roots. The upcoming circuit board iteration takes this idea into account, designing in connectors that can communicate with additional modules.

7 The Birl of the Future! (present-beyond)

I've been working on the Birl continuously for eight years. In this wandering adventure, I explored various things the instrument could be, and followed performer needs and my own varied interests where they took me. The future of the project is still very open, and I hope to continue to use it as a platform to follow whatever intriguing paths present themselves.

8 Reflections on the Process

It might seem unusual to call all of these instruments ”the Birl” when each has distinct design characteristics. However, in my mind they are one: I consider the Birl to be a gradually changing, continuous project. It has been exciting to undertake a project where the focus of the research is about the process, where continual reinvention flows from experiments I undertake and feedback from performers.

References

[1] Aldridge H (1996) Music's most glorious voice: The Hammond organ. The Journal of American Culture 19(3):1

[2] Collins N (2006). Handmade electronic music: the art of hardware hacking. Taylor & Francis

[3] Doty D (2002) The Just Intonation Primer. Other Music, San Francisco

[4] Fiebrink R (2011) Real-time human interaction with supervised learning algorithms for music composition and performance. Doctoral dissertation from Princeton University.

[5] Parker J, D'Angelo S (2013) A Digital Model of the Buchla Lowpass-Gate. In: Proc. 16th Intl. Conf. Digital Audio Effects (DAFx-13), pp.278 285

[6] Menzies D, McPherson A (2013) A Digital Bagpipe Chanter System to Assist in One-to-One Piping Tuition. In: Proceedings of the 2013 Stockholm Music Acoustics Conference/Sound and Music Computing Conference (SMAC/SMC)
[7] Sato M, Poupyrev I, and Harrison C (2012) Touché enhancing touch interaction on humans, screens, liquids, and everyday objects. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. pp. 483 492
[8] Scavone G(1997) An acoustic analysis of single-reed woodwind instruments with an emphasis on design and performance issues and digital waveguide modeling techniques. Doctoral dissertation from Stanford University.
[9] Scavone G(1999) Modeling wind instrument sound radiation using digital waveguides. In: Proc. of the ICMC (1999). pp. 355 358
[10] Scavone G, Cook P(1998) Real-time computer modeling of woodwind instruments. In: Proc. Int. Symp. on Musical Acoustics (ISMA-98) pp. 197 202
[11] Scavone G, Smith J O(1997) Digital waveguide modeling of woodwind toneholes. In: Proceedings of the 1997 International Computer Music Conference
[12] Scavone G, Van Walstijn M(2000) The wave digital tonehole model. In: Proceedings of International Computer Music Conference (ICMC-2000). pp. 465 468
[13] Snyder J (2011) Exploration of an Adaptable Just Intonation System. Doctoral dissertation from Columbia University. http://www.scattershot.org/full-dissertation.pdf. Cited 30 May 2016
[14] Snyder J, Ryan D(2014) The Birl: An Electronic Wind Instrument Based on an Artificial Neural Network Parameter Mapping Structure. In: Proc. New Interfaces for Musical Expression (NIME2014), pp. 585 588 http://www.nime.org/proceedings/2014/nime2014_540.pdf. Cited 30 May 2016
[15] http://www.akaipro.com/product/ewiusb. Cited 30 May 2016
[16] http://www.atmel.com/products/microcontrollers/avr/32-bitavruc3.aspx. Cited 30 May 2016
[17] http://www.bang-olufsen.com/en. Cited 30 May 2016
[18] http://www.cypress.com/products/32-bit-arm-cortex-m-psoc. Cited 30 May 2016
[19] http://designmuseum.dk/en. Cited 30 May 2016
[20] https://en.wikipedia.org/wiki/Jacob_Jensen. Cited 30 May 2016
[21] http://usa.yamaha.com/products/music-production/midi-controllers/wx5/. Cited 30 May 2016
[22] https://vimeo.com/86167177. Cited 30 May 2016
[23] https://vimeo.com/86461028. Cited 30 May 2016
[24] https://vimeo.com/4611451. Cited 30 May 2016
[25] https://vimeo.com/album/3382506. Cited 30 May 2016
[26] http://www.wetinik.org/. Cited 30 May 2016