

Parthenope: A Robotic Musical Siren

Michael Sidler
Worcester Polytechnic
Institute
100 Institute Road
Worcester, MA 01609
msidler@wpi.edu

Matthew Bisson
Worcester Polytechnic
Institute
100 Institute Road
Worcester, MA 01609
mcbisson@wpi.edu

Jordan Grotz
Worcester Polytechnic
Institute
100 Institute Road
Worcester, MA 01609
jtgrotz@wpi.edu

Scott Barton
Worcester Polytechnic
Institute
100 Institute Road
Worcester, MA 01609
sdbarton@wpi.edu

ABSTRACT

Parthenope is a robotic musical siren developed to produce unique timbres and sonic gestures. Parthenope uses perforated spinning disks through which air is directed to produce sound. Computer-control of disk speed and air flow in conjunction with a variety of nozzles allow pitches to be precisely produced at different volumes. The instrument is controlled via Open Sound Control (OSC) messages sent over an ethernet connection and can interface with common DAWs and physical controllers. Parthenope is capable of microtonal tuning, portamenti, rapid and precise articulation (and thus complex rhythms), and distinct timbres that result from its aerophonic character. It occupies a unique place among robotic musical instruments.

Author Keywords

Musical Robots, Aerophones, Sirens

CCS Concepts

- Applied computing → Sound and music computing;
- Computer systems organization → Robotics;

1. INTRODUCTION

Sirens are often thought of as "noisemakers" that have limited utility as musical instruments [4]. Part of the reason for this attitude is due to the fact that many sirens are unidimensional in terms of pitch and rhythm and are hard to control. When considering instruments to add to the robotic ensemble of WPI's Music, Perception and Robotics Lab (mprlab.org), we were intrigued by the possibility of harnessing the unwieldy utterances of the siren through computer control of electromechanical actuators. There are few robotic instruments that have ventured down this path, thus an opportunity to explore new territory in musical expression was presented.



Figure 1: Parthenope

2. BACKGROUND AND PRIOR ART

A siren has a simple design: a pitch is created by passing air through concentric holes arranged radially on a spinning disk. The frequency produced is equal to the number of holes multiplied by the rotational speed of the disk. A disk with 44 holes spinning at 10 rotations per second will result in a frequency of 440 Hz, producing the pitch A4 [4].

$$f[\text{Hz}] = n[\text{holes}] * v[\text{rps}] \quad (1)$$

Sirens are most often recognized for their use as alarms (e.g. air raid sirens). They are recognized by their sweeping portamenti and piercing timbre. As a result, their use in musical pieces is often for textural or representational effects. They helped illustrate the clamour and intensity of an industrializing world in the early 20th century in works such as Edgard Varèse's *Ionisation* or the *Ballet Mécanique* by George Antheil. The use of the siren is also heard in the popular domain: the song *Ridin' the Storm Out* by REO Speedwagon opens with the sweeping sound of a siren that commands the attention of listeners.

Variations of the siren's design have been made to try to achieve the functionality of more traditional musical instruments, which can produce a larger range of pitches and musical gestures. A number of siren machines have been built that feature a motor that spins a perforated disk. None of them are fully autonomous though in that each requires some degree of human involvement.

The Helmholtz double siren, as built by Rene Bakker in 2008, has two spinning disks that are capable of sweeping portamenti [1]. The tone is changed by altering a secondary disk, which limits the flow of air. The timbre produced



Licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0). Copyright remains with the author(s).

NIME'20, July 21-25, 2020, Royal Birmingham Conservatoire Birmingham, UK.

is reminiscent of an accordion or harmonica, and can be changed by moving a case over the disk. The instrument is limited by its complex design, which makes it harder to reproduce musical gestures.

The Loman Siren Organ was patented by Abraham Loman in 1915 and interpreted by Rene Bakker in 2008 [2]. This instrument has a series of six disks that can produce two tones in semitone increments over the span of an octave. The disks are encased, displayed in the familiar visage of an organ, which affects the timbre produced by the instrument. It is played manually with an octave wide keyboard. The range can be changed by pulling a stop that increases the speed of the disks.

In 2014, a group from Stanford created the "Siren Organ", which operates by actuating valves to allow air to flow through the three disks on the instrument [3]. The speed of the disk and the flow of air is human controlled. However, this control relies on human input, limiting the range of musical expression. Precision is difficult when physically moving a valve or potentiometer, echoing the issues of musical gesture reproduction with the Helmholtz double siren.

3. DESIGN

The above examples suggested that there was musical possibility in robotic sirens that had yet to be discovered. As roboticists, many of the machines that we work with create noise as a by-product. The fundamental idea of Parthenope was to create what is typically seen as a noisemaker into a musical voice through computer control of electro-mechanical actuation. Such an approach could surpass the limitations of the previous works mentioned.

These motivations resulted in a number of specific design objectives for Parthenope. We sought to create an instrument that could play continuous gestures and discrete notes in ways that humans could not. The disk's rotations were to be controlled to ensure speed of intonation and accuracy of pitch. The timbre produced should be variable, compelling, and in combination with a wide dynamic range, could be either the primary or supporting voice in an ensemble. We sought to minimize system latency so that the machine could be used in real-time interactions. We wanted to make our instrument accessible to less-technical musicians, thus it was to be played through Digital Audio Workstations (DAWs) via the MIDI protocol.

3.1 Disk Design




Figure 2: Diagram of one of Parthenope's disks

The design of Parthenope's disks dictates the structure of the rest of the instrument. The cast acrylic disks have been laser cut to create four concentric sets of holes. A note is produced by passing air through any one set of holes.

The number of holes in each of the four concentric sets affect the way that Parthenope generates sound. The control algorithm prioritizes the smallest change in angular velocity when choosing what holes to send air through. This helps to minimize the portamento and time the instrument takes

to settle on the correct pitch.

To determine the range and speed optimizations of a particular disk, the following calculations are used: First, the number of holes in each set are chosen (e.g. ranging from inner most to outer most set: 11, 12, 14, 15). Next, a starting note is picked (e.g. C4). The disk speed in rotations per second is calculated for the innermost hole set using equation 1. The note C4 has a frequency of 261.626 Hz, which when divided by 11 holes gives a speed of 23.784 rps. This disk speed is used to find the frequency that each of the other hole sets would play. For example, the second hole set would have a frequency of $12 * 23.784 = 285.408 Hz$. This note is closest to the pitch D4 (293.665 Hz). In practice, the motor speed would be changed to account for this difference. The nearest notes of the remaining two hole sets are calculated according to this process (in this example, E4 on the third hole set and an F4 on the fourth hole set). A spreadsheet aids with this calculation process.

A disk with 11, 12, 14, and 15 holes in each set will have a range of five semitones in the pattern W W 1/2 between the four notes. In Parthenope's design, the inner hole set of the second disk will have a fixed offset from the inner hole set of the first disk.

This offset is given as an integer number of semitones and is referred to as the *disk semitone offset*. The offset in speeds of the two disks can be used to optimize various configurations. The disk semitone offset can be set to seven to produce a configuration for Parthenope that is optimized to play a diatonic major scale. (Figure 3). This optimization only exists over a one octave range as there are only eight sets of holes. Any notes played outside of this range will require the motor to make a greater speed change, resulting in a pitch glide up or down.

This procedure can be followed to create any pattern that the composer desires. A range of eight octaves is obtained with two disks, 48 semitones apart, each with a hole pattern of 1,2,4, and eight holes (Figure 4). A consequence of this configuration is larger pitch glides to non-optimized frequencies. This will cause a decrease in accuracy as tempo increases. Additionally, the intonation time will increase, as the instrument has a greater speed change to overcome. A disk semitone offset of six could be chosen to alter this configuration to have a three and a half octave range and a tighter grouping of notes (Figure 5). The desirability of such behavior will depend on the music being played.

3.2 Motor and Motor Controller

In order to precisely control and change the rotational velocity of the disks, a high power motor must be used. For this task we chose the RCRunning 3650 brushless motor and an ODrive v3.4 24V motor controller. A brushless motor was chosen for its quiet operation and high power. The ODrive controller was chosen for its implementation of field oriented control (FOC). FOC is a method for driving the three phases of a brushless motor using sinusoidal waves instead of square waves. This prevents the undesirable high pitched whining noise that is customary of standard brushless motor controllers.

3.3 Air Supply

An Eagle Silent Series 2-HP 20-Gallon air compressor is used to supply 30 PSI air through the instrument to the Air-TAC 4V210-08 solenoid valves. These are five port, two way pilot operated valves, with one of the channels closed off. This prevents air from flowing unless one of the solenoids is purposefully actuated. The solenoids' valves direct the release of air through nozzles. These nozzles are of different sizes to produce different amounts of airflow and therefore




Figure 3: Optimized frequencies of a disk with 11, 12, 14, and 15 holes and a semitone offset of 7




Figure 4: Optimized frequencies of a disk with 1, 2, 4, and 8 holes and a semitone offset of 48




Figure 5: Optimized frequencies of a disk with 11, 12, 14, and 15 holes and a semitone offset of 6

different volumes. Diameters of 0.2mm, 0.6mm and 1.0mm were chosen as they produced a wide dynamic range. Three nozzles (one of each diameter) correspond to a hole set on a disk. With four sets of holes on each disk, one rotational speed can produce four different pitches at three different velocity levels each (see figure 2).

3.4 Instrument Control

Parthenope is controlled by an Adafruit Grand Central M4 Express. This is an ARM Cortex M4F development board chosen for its large pinout, floating point processing support, and configurable serial ports. The microcontroller can interface with a computer or any other device over ethernet through the Adafruit Ethernet FeatherWing. This is an ethernet interface built around the WizNet WIZ5500 chip. Parthenope accepts OSC messages over the ethernet interface, which can be used to play notes and adjust instrument parameters. The primary way to play Parthenope is through a DAW such as Ableton Live. A Max for Live patch was created to translate the MIDI information from the Ableton Live Set into OSC messages that are sent to the microcontroller. The patch also includes buttons that can start/stop the disks.

3.5 Aesthetics and Visualizations

The enclosure for Parthenope was inspired by the analog synthesizers of the 1960's and 70's. The two have a timbral connection that we sought to reference in Parthenope's physical design. We envisioned Parthenope as a product of modern technology that is infused with the spirit of early experimental electronic music.

The casing of the instrument is made from a translucent black acrylic that allows for a glimpse at the internal mechanisms when lit internally with signal LEDs. The case also displays analog voltage and current indicators to visualize the power of the instrument. The clear spinning disks are prominently displayed above red modules that couple each system together. We find the design simple and engaging.

4. EVALUATION AND RESULTS

Our technical evaluation of Parthenope focused on intonation time, pitch accuracy, timbral envelope, dynamic range, and system latency. One of the ways that Parthenope produces different pitches is by changing the rotational speed of the motor, which takes time. We measured the time it took for the instrument to accurately produce all intervals

within a chromatic scale relative to a tonic (in this case the pitch C5), both ascending (e.g. $1 > 2, 1 > 3, \dots 1 > 12$) and descending ($12 > 11, 12 > 10, \dots 12 > 1$). After a note message was sent, a timer was started and the estimated pitch produced by the instrument was measured using Max's fzero object (which uses wavelet transforms). When the produced frequency reached a threshold $\pm 5\%$ of the target frequency, the pitch was considered accurate and the timer was stopped. The experiment was repeated twice for a total of three blocks. The resultant values were averaged, which can be seen in Figure 6.




Figure 6: Intonation time of ascending and descending pitch intervals in the chromatic scale. The y-axis represents time (msec) and the x-axis indicates pitch interval size (semitones).

Intonation times tended to rise with interval size (more with ascending than descending intervals), though this relationship was variable. Intonation times were generally in the 100-250 msec range, which, in practice, would affect the amount of pitch "glide" produced at the start of a note.

To assess pitch accuracy, we collected reported frequencies for each pitch produced once they reached the specified threshold. The absolute values of these readings were averaged to give a measure of pitch accuracy. The pitches produced were quite accurate with most values deviating from ideals by 1% or less.

The timbral envelope of a tone produced by Parthenope is characterized by a percussive attack (produced by the solenoid valves) with energy at 200-400 Hz and 700-800 Hz, which lasts roughly 300 msec. The lower partials of the sound develop quickly, starting about 3 msec after the percussive beginning, and are followed by the higher partials

emerging by 15 msec. The spectrum is strongly harmonic with surprisingly little energy at the fundamental, with emphasis on harmonics 4-7 and considerable range that extends into the 7 kHz region. These details can be seen in Figure 7.




Figure 7: A Spectrogram of Parthenope playing the pitch C5 (log frequency scale).

To test the dynamic range of Parthenope, we placed a LotFancy Sound Level Meter 1ft from the front face of the instrument. With the instrument powered off in a quiet room the sound level was measured at 54.4dBA. With the motors powered on, the sound level was 59.3dBA. Next, the instrument played a two octave C major scale starting at C3 at each of its three volume levels. The dBA reading was taken for each note at each volume level. The readings for each volume level were averaged. The quietest was 80.5dBA, the middle was 95.7dBA, and the loudest was 98.2dBA.

The main source of latency between receiving a message and sounding the appropriate note is the solenoid valves. This latency was defined as the time difference between the illumination of an LED, representing the power to the solenoid, and a thin piece of paper being moved by the air released from the nozzle. We recorded the solenoid actuating with a camera at 240fps. From the footage, an average of ten trials gave us a consistent six frames between the two events. This results in a solenoid actuation latency of 25msec +/- 2.4msec. The other potential source of latency is the microcontroller processing the message, however, this will be an order of magnitude smaller than the solenoid latency. This is due to the 120MHz clock speed and floating point processor core of the microcontroller.

5. MUSICAL APPLICATIONS

The instrument occupies a unique musical niche as a consequence of its physical design, means of sound production, and computer control. The quality of the sound produced is a cross between a sawtooth oscillator on an analog synth and acoustic brass. There is a rawness and a roughness to it that when sustained, can be grating (particularly when played in isolation). Parthenope finds a more comfortable home when articulating rapidly, which mollifies its unruly qualities. At high speeds, the jaggedness of the timbre provides a sharp attack, helping it cut through noisy or busy sonic environments. Thanks to the minimal activation time of the solenoids, Parthenope is able to play quite rapidly (it is comfortable at 50 msec intervals).

Parthenope excels in producing portamenti. By modifying the speed of the disks while air is flowing, continuous pitch changes are created. This allows the instrument to produce a range of gestures from long sweeps to abbreviated synth-like pitch glides that are typically found in electronic

dance music. Such variability in combination with the capacity to produce discrete notes is unique in the world of musical robots. Another strength is in articulatory variation. Contrasts can be created between legato phrases with portamenti only to transition to aggressive staccatissimo at the change of a note.

In addition to the pitches it produces, Parthenope's physical design results in sonic idiosyncracies, originating from the mechanical sounds of solenoids and hum of spinning disks. The sound of the disks' rapid rotation can intensify musical gestures or provide a background texture that can help congeal disparate sound sources.

Parthenope's configurable design allows it to support a potentially unlimited number of disks through a simple re-configuration of constants in its code. Each disk can be altered to contain different numbers of holes, which will in turn affect the amount of portamento relative to a particular tuning system. The instrument is thus an excellent candidate for microtonal explorations.

When composing with Parthenope, the instrument most often occupies the lead voice because of its piercing sound. However, its timbre allows for contrast and combination with other instruments to voice unique phrases. The third author composed a piece for Parthenope and (human-played) vibraphone that explored such sonic interactions. The sharp attack of the instrument allows for clear rhythmic articulations, which enables it to assume a percussive role if desired.

6. CONCLUSION AND FUTURE DIRECTIONS

Parthenope is successful as it allowed us to harness the energy of a noisy siren and transform it into a voice for musical expression. The controls afforded by the instrument and the idiosyncratic sounds that it produces inspire our musical creativity in new ways.

Despite this success, there are still areas for improvement. We would like to expand the number of playing modes. Currently the instrument can only play monophonic parts, but new control algorithms could be written to allow each disk to play independently, allowing Parthenope to produce polyphonic passages. Adding more disks to Parthenope could increase the range of the instrument, add to the number of notes that could be played in the polyphonic mode, or even be used to create auditory effects by placing disk modules in different locations relative to the composer or listeners. Parthenope could also be augmented with alternative controllers, allowing it to be played as a continuous instrument by a human, similar to a theremin. In addition to the above technical improvements, we hope that time will allow more musicians to explore Parthenope and discover creative ways to incorporate its unique sound into new musical works.

7. REFERENCES

- [1] R. Bakker. Helmholtz double-siren. <http://www.youtube.com/watch?v=xaBoC7tbAE0>, July 2008.
- [2] R. Bakker. Loman siren organ. <http://www.youtube.com/watch?v=gSn4EUNXz4M>, July 2008.
- [3] R. Collecchia, D. Somen, and K. McElroy. The siren organ. In *Proceedings of the International Conference on New Interfaces for Musical Expression*, pages 391–394, London, United Kingdom, June 2014.
- [4] B. Hopkin. *Musical Instrument Design*. See Sharp Press, Tucson, Arizona, 1996.