Soundstone: A 3-D Wireless Music Controller

Adam Bowen
Stanford Product Design and CCRMA
Building 610, 447 Santa Teresa
Stanford, CA 94305
650-387-0298
abowen@stanford.edu

ABSTRACT

Soundstone is a small wireless music controller that tracks movement and gestures, and maps these signals to characteristics of various synthesized and sampled sounds. It is intended to become a general-purpose platform for exploring the sonification of movement, with an emphasis on tactile (haptic) feedback.

Keywords

Gesture recognition, haptics, human factors, force, acceleration, tactile feedback, general purpose controller, wireless.

1. INTRODUCTION

Controlling sound with gesture has long held interest for musicians and electronics experimenters. The Theremin is probably the most prominent example of a device that allows gestural control of music. Other instruments have followed, and these have employed a variety of sensing methods (IR, capacitive, etc). In most cases, they require the performer to be stationed at the instrument.

To capture broader and more dynamic gestures, it would be useful to have a device that is mobile and that senses movement relative only to the performer's body. For a course at Stanford's CCRMA center, my goal was to create a truly "untethered" experience – allowing performers to move freely and control music, without being tied to a particular position or reference frame. The idea for a small, handheld accelerometer-based controller was born.

This project was also an opportunity to explore the use of haptics in music controllers. As the device would be free-moving, with no physical "stops", it seemed valuable to include a feedback mechanism that would allow for the performer to sense the result of his/her gestures. Perceptually, the human sense with lowest latency is touch, so a simulated tactile means was employed: a small vibrating motor, like those found in cellular phones.

2. IMPLEMENTATION

The general overview of the system is that a small, battery-powered controller (the stone) senses 3-axis acceleration and sends this data over a wireless link to the host computer running Max/MSP.

Once the data enters Max/MSP, it is filtered and interpreted to provide several signals: tilt, velocity, and jerk in all dimensions. Specific Max patches determine which of these signals are utilized.

As an added dimension of control, a force-sensitive resistor (FSR) is incorporated into the device. This analog "button" can be mapped to amplitude (intensity) or other aspects of the various sound patches. For feedback, a tri-color LED is housed in the handheld device, as well as the vibrator motor.

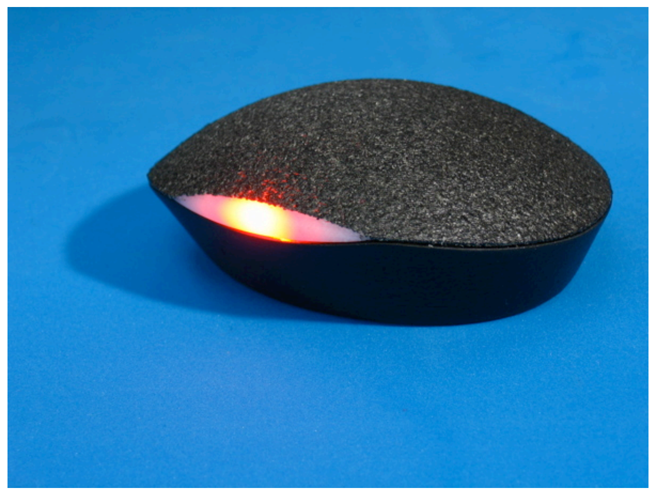

Figure 1. Handheld device

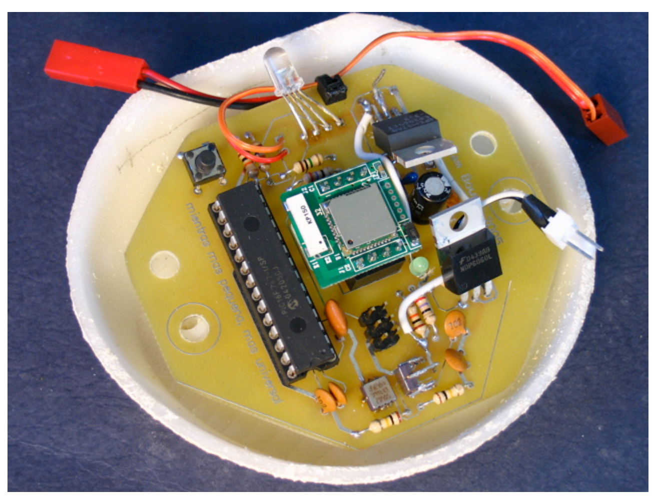

Figure 2. Electronics and enclosure

A PIC microcontroller is programmed to run onboard and sample the analog signals, format the data, and send it over a Bluetooth wireless connection. Bluetooth was chosen due to its high bandwidth (~800kbits/s), and to eliminate the need for special receiver hardware (a Bluetooth transceiver comes pre-installed on some Apple computers, such as the PowerBook G4 used in this case).

The high transmission speed of this scheme allows not only for high sample rates (~1kHz), but for low-latency responses from the host computer. For instance, if a virtual "wall" or limit is reached by the controller, a particular Max patch might return a command to quickly pulse the vibrator motor, providing tactile feedback. If a percussion patch is invoked, it can provide a tactile tap response when a drum is "hit."

The overall system is as shown in Figure 3.

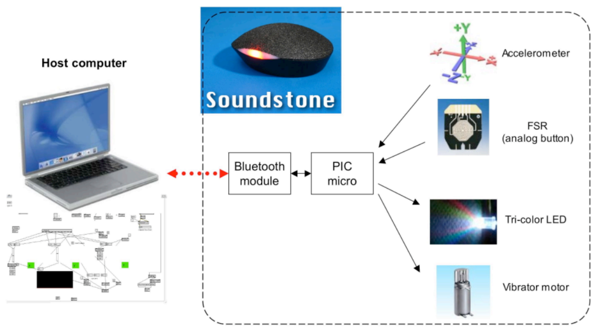

Figure 3. Functional diagram

The handheld unit is a controller in the truest sense – sampled data is sent to the host computer, with no interpretation or synthesis on the device itself. This may change as I look at ways to improve feedback latency.

The form of the device was carefully studied to be somewhat generic, such that it might be grasped in a variety of ways that seem comfortable and natural.

Figure 4. Soundstone in use

3. FIRST SHAKES – EVALUATION

Since the hardware and communications were finished, it has become a platform for many experiments in gestural control. Modes include percussive, continuous control, complex gesture recognition, and combinations thereof. For percussive applications, transient impulses in acceleration are detected. For continuous control, tilt data is used. Gesture recognition combines both.

Current patches incorporate wavetable and additive synthesis as well as physical modeling. In one patch, the device becomes a virtual canteen: tipping it over starts the sound of a dripping faucet. Increasing tilt increases the frequency, until gradually the sound transforms into an undulating waterfall.

In another patch, a basic drum kit is implemented, where jabbing the control along different axes triggers distinct samples.

From initial experiences of using the device, it has shown to be intuitive, consistent, and expressive in a variety of modes.

4. NEXT STEPS

As future work on this project, I will continue developing new modes and patches to further explore the device's potential. It might also be of interest to augment its sensing means with gyros, additional buttons, etc.

5. ACKNOWLEDGEMENTS

Thanks to Bill Verplank and Matt Wright of CCRMA, Ed Carryer, Pascal Stang, and the PD Loft community.