TouchKeys: Capacitive Multi-Touch Sensing on a Physical Keyboard

Andrew McPherson
Centre for Digital Music, School of EECS
Queen Mary University of London
Mile End Road, London E1 4NS, United Kingdom
andrewm@eecs.qmul.ac.uk

ABSTRACT

Capacitive touch sensing is increasingly used in musical controllers, particularly those based on multi-touch screen interfaces. However, in contrast to the venerable piano-style keyboard, touch screen controllers lack the tactile feedback many performers find crucial. This paper presents an augmentation system for acoustic and electronic keyboards in which multi-touch capacitive sensors are added to the surface of each key. Each key records the position of fingers on the surface, and by combining this data with MIDI note onsets and aftertouch from the host keyboard, the system functions as a multidimensional polyphonic controller for a wide variety of synthesis software. The paper will discuss general capacitive touch sensor design, keyboard-specific implementation strategies, and the development of a flexible mapping engine using OSC and MIDI.

Keywords

augmented instruments, keyboard, capacitive sensing, multi-touch

1. INTRODUCTION

There are many excellent reasons to use an iPad or other touch-screen device as a musical controller, among them flexibility, continuous gesture recognition, and direct relationship between image and touch input. However, no current touch-screen device can replace the tactile feedback of a traditional instrument. Tactile feedback is crucial in keyboard performance, since pianists generally play by feel rather than by sight.

This paper presents a system of capacitive multi-touch sensing which attaches to the surface of a physical keyboard. Each key contains sensor pads and a controller which measures the location and contact area of fingers on the key surface. The complete system, consisting of up to 8 octaves, communicates with a computer by USB. The touch measurements transform the keyboard into a continuous multidimensional control surface.

1.1 Related Work

The use of electronics to enhance the capabilities of traditional instruments dates back over a century [9]. In the past 20 years, several authors have explored continuous extensions of the keyboard. In 1990, Moog, Rhea and Eaton created a touch-sensing piano keyboard [1, 6] that is the most direct antecedent for the present work. The project was never commercially produced and remained a work-in-progress at Moog's death in 2005 [7], though Eaton has used it in several performances.

The Seaboard [3] uses a keyboard-shaped silicone surface and force-sensing resistors to provide multidimensional measurements of each touch. The Haken Continuum [2] extends the concept of the keyboard to a generalized mechanical control surface measuring the three-dimensional position of each finger. Most recently, the Evo keyboard[1] measures front-to-back touch position along a segment of each key.

Like the Moog-Eaton design, this work maintains the traditional feel of the keyboard, while providing more detailed data in the form of multiple touches and finger contact area across the entire surface of each key.

2. CAPACITIVE SENSING

Capacitive touch sensing allows high-precision tracking of a user's finger motion with no electrical contact between user and device. A conductive plate forms a capacitor with the surrounding free space and ground layers. Objects which are conductive or have a substantially different dielectric constant than air, when brought into proximity with the plate, will change its capacitance [8]. Capacitance values are typically measured either by charging the plate to a known voltage and measuring discharge time in an RC circuit, or by measuring its frequency response in a resonant circuit.

The capacitance of a single sensor can be read as a continuous value which roughly corresponds to the proximity and size of nearby objects. To measure *position*, an array of discrete sensors are required (Figure 1 bottom). Sensors are measured one at a time, with the remaining sensors tied to ground. A finger touch will activate several adjacent sensors, from which a centroid value (weighted average) can be calculated (Figure 1 top). Because the sensor values are continuous, position resolution can far exceed the number of sensor elements in the array.

Though more complex to implement than resistive position sensors, capacitive sensing has the advantage of requiring no finger pressure (indeed no contact at all) to operate. With certain sensor configurations, multi-touch capability is also supported, where resistive sensors are limited to at most one or two points of contact. Capacitive sensing can be combined with existing pressure (aftertouch) keyboard systems, and unlike aftertouch, both pressed and unpressed keys can be read.

2.1 Keyboard Sensor Design

The TouchKeys use the PSoC 1 'CapSense' series of ICs from Cypress Semiconductor.[2] The chips contain dedicated

NIME'12, May 21 – 23, 2012, University of Michigan, Ann Arbor.

---

[1] `http://www.endeavour.de`

[2] `http://www.cypress.com/?docID=25608`

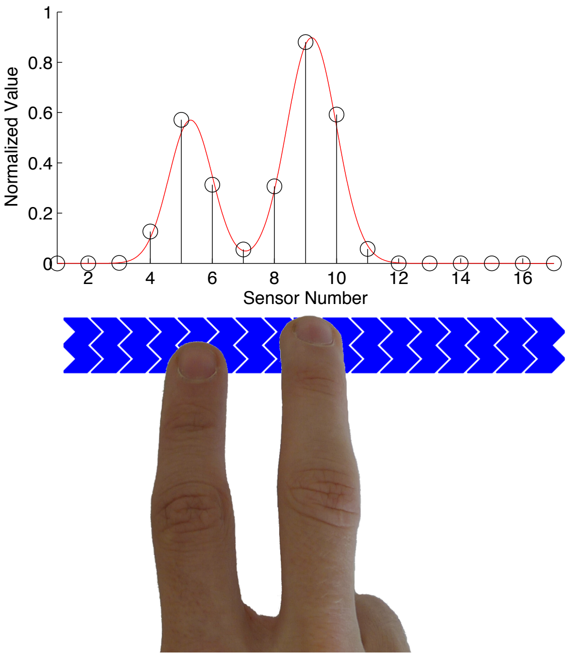

Figure 1: Simluation of multi-centroid calculation from individual sensor readings.

analog hardware, including a sigma-delta ADC, to measure capacitance values on each pin. Because of this dedicated hardware, measurement time and sensor resolution significantly exceed the performance of software implementations such as the one used in the Atmel AVR (and Arduino) controllers. The Cypress chips include a microcontroller core in addition to the touch-sensing hardware; however, its performance is limited, so the TouchKeys implementation handles most data processing at a later stage.

The sensors are made from 0.8mm printed circuit boards routed to the shape of each key. Figure 2 shows a two-octave set of keys, which is fabricated as a single board with scoring that allows each key to be separated once assembled. Figure 3 shows the keys installed on a five-octave MIDI keyboard. Plastic spacers laser-cut around the electronic components are placed underneath the circuit board to create a flat mounting surface. The sensors are attached to the keyboard with adhesive tape which creates a secure but removable bond.

The black keys have 17 discrete sensor pads arranged in a single row. They are capable of sensing finger position along the lengthwise axis of the key. The white keys have 25 sensor pads, divided between a single row in the back and a grid of rows and columns in the front (Figure 4). The front of the white keys senses finger postion in two dimensions. Positions are calculated as the centroid of adjacent values [5], and up to three touches can be measured on each key.

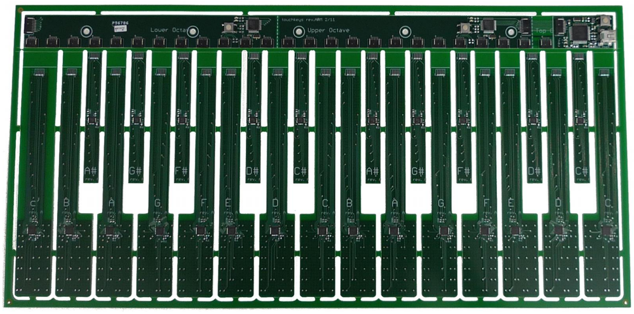

Figure 2: TouchKeys in two-octave scored panel.

A complete scan of 25 sensors, centroid calculations, and communication of the results via an I2C bus takes approximately 4ms. Scan rates up to 250Hz are thus possible, though because of timing constraints elsewhere in the system, 125Hz operation was found to be the most reliable. Further information on aggregation of data from multiple keys and transmission to the computer can be found in [5].

2.2 PCB Design Guidelines

Figure 4 shows the printed circuit board layout for one key (A). The board has four layers. The top layer contains the sensor pads; the layer beneath contains traces connecting rows of pads in the two-dimensional grid; the third layer is a hatched ground plane, and the bottom layer contains the components and most signal traces.

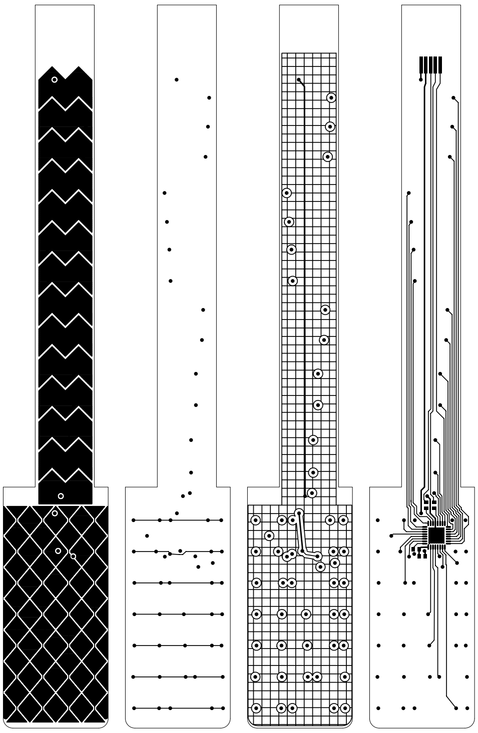

Figure 4: PCB layers of one key, from top (at left) to bottom.

Based on results and experiments from this design and guidance from the Cypress datasheets[3], I suggest the following guidelines for capacitive touch design:

- Locate the controller IC centrally. The *parasitic capacitance* (capacitance when no finger is present) depends on the complete conductive area, including pad and trace. Long traces should be avoided. This was a particular challenge in designing the white keys.

- A ground plane is required to avoid stray interference. In a two-layer design this can be shared with components and signal traces on the bottom, but a dedicated inner layer in a four-layer design is helpful. Cypress recommends the ground plane be hatched to avoid excessive parasitic capacitance. The TouchKeys notably do not follow the recommendation that a ground plane surround the sensors on the top layer. Its absence does not seem to affect performance.

---

[3] `http://www.cypress.com/?docID=28734`

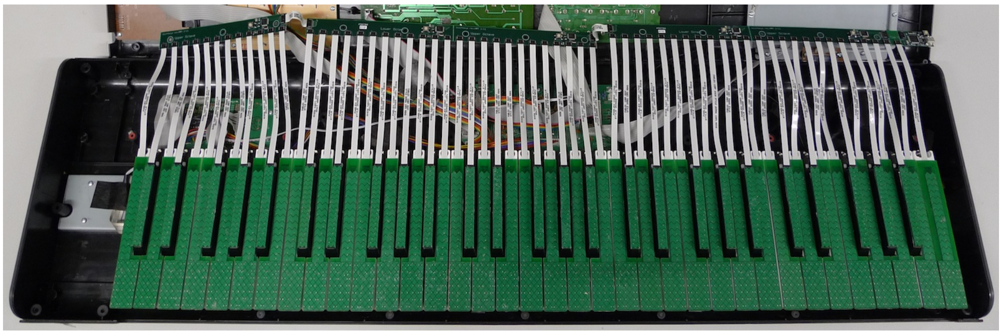

Figure 3: TouchKeys on a 61-key MIDI keyboard (open to show connections). Experiments with surface coating found the bare soldermask to produce the best feel and highest-resolution data.

- Signal traces should not cross underneath a sensor pad unless a ground plane separates the two.
- Pads in an array should be V or zigzag-shaped. This ensures that a touch partially activates several adjacent pads. Similarly, pads should be close enough together that a finger activates several at once. Minimum spacing will likely be constrained by number of pins available on the controller.
- Keep communication lines clear of sensor traces.
- In a situation such as a physical keyboard, any action that moves the key will ideally also register as a touch location. This requires extending sensor pads as close as possible to the edges of the part.

2.3 Surface Coating

The initial design of the TouchKeys [5] used a thin plastic laminate on top of the circuit board. The intention was to more accurately simulate the look and feel of the traditional keyboard. Many types of plastic were tested, including polypropylene, PETG, Delrin, acrylic, teflon, nylon and polycarbonate. Enamel and epoxy paints were also tested. Experimentally, it was found that the laminate could be no thicker than 0.5mm on the black keys, and that on the front of the white keys, even a laminate of 0.25mm reduced performance in the two-dimensional sensor area.

Unexpectedly, many pianists indicated that the raw soldermask coating of the circuit board (an insulating layer applied during fabrication) produced a better feel than the various plastics, many of which were felt to be too sticky on the fingers. The copper sensor pads are slightly raised with respect to the etched parts of the circuit board, so the keys have a texture that was initially thought to be a drawback. However, some pianists observed that ivory keys also have a textured surface, which is not a problem in performance. The next design revision will use white and black soldermask to maintain the standard look of the keyboard.

In general, the designer can optimize any three of the following quantities at the expense of the fourth: sensor pad size, coating thickness, measurement speed and measurement resolution. The size of the TouchKey pads are constrained by the geometry of the keys, and speed and resolution were prioritized over coating thickness.

2.4 Data Aggregation

The controller on each key is responsible for scanning all the sensors in sequence, calculating up to three centroid locations and sizes, and transmitting this data via I2C to a second "octave" controller. The octave controller aggregates the data from an octave of keys and routes it to a "host" controller which communicates to a computer via USB [5].

The host microcontroller implements a USB serial device. MIDI, even in its native USB implementation, is ill-suited for TouchKey data since controls are limited to 7 bit resolution (compared to 10 bits or more for touch position data) and Control Change messages are channel-wide rather than specific to each note. The serial data is unpacked by the computer into OSC messages which can in turn be dynamically mapped to MIDI. The mapping is discussed in the next section. The software provides a real-time visualization of key touches (Figure 5) which can be used in live performance or for debugging.

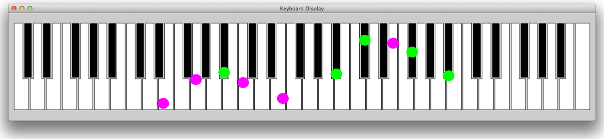

Figure 5: Real-time display of touch position.

3. MAPPING TOUCH DATA

Open Sound Control is the native output of the TouchKey system, with messages for the following actions:

- Touch onsets and releases. Distinct from MIDI note onsets and releases, this indicates when a finger touched or left the key surface.
- Changes of position and contact area for each touch.
- Pinch and slide gestures involving two or three fingers.
- Raw data frames of all touch locations and sizes.

The complete control system consists of touch data correlated with MIDI data from the underlying keyboard. This gives a picture of both activity on the key surfaces as well as physical key motion. Where the keyboard supports aftertouch, a form of three-dimensional sensing is available on pressed keys. All OSC messages are tagged with a MIDI note number so a synthesis program can easily correlate touch sensor and keyboard data.

3.1 Dynamic MIDI Mapping

Though OSC messages can be sent to any program, most commercial software synthesizers are implemented as VST or AudioUnit plugins whose parameters are set by MIDI Control Change messages. To control these plugins, a dynamic MIDI mapping system was developed (Figure 6). Each touch parameter can be mapped to a different MIDI controller, to the pitch wheel, or to aftertouch. Input and output ranges are adjustable, and controls can be sent as absolute values or as relative values with respect to the original touch or note onset (see playability discussion below).

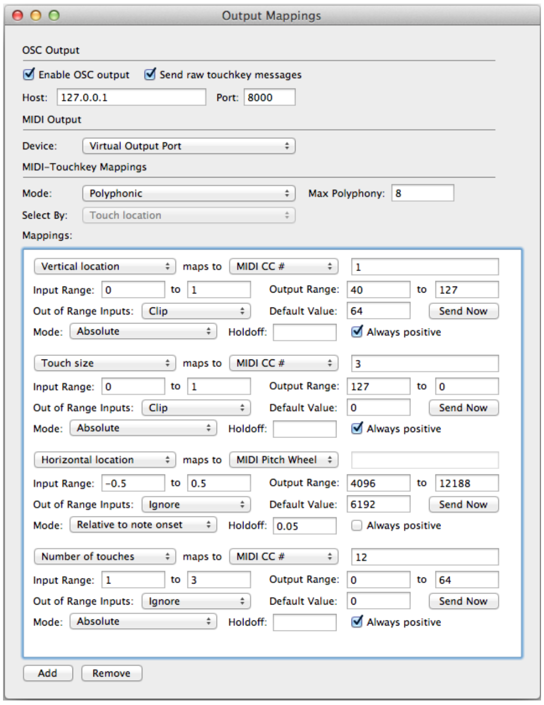

OSC Output

☑ Enable OSC output ☑ Send raw touchkey messages

Host: 127.0.0.1 Port: 8000

MIDI Output

Device: Virtual Output Port

MIDI-Touchkey Mappings

Mode: Polyphonic Max Polyphony: 8

Select By: Touch location

Mappings:

Vertical location maps to MIDI CC # 1
Input Range: 0 to 1 Output Range: 40 to 127
Out of Range Inputs: Clip Default Value: 64
Mode: Absolute Holdoff: ☑ Always positive

Touch size maps to MIDI CC # 3
Input Range: 0 to 1 Output Range: 127 to 0
Out of Range Inputs: Clip Default Value: 0
Mode: Absolute Holdoff: ☑ Always positive

Horizontal location maps to MIDI Pitch Wheel
Input Range: -0.5 to 0.5 Output Range: 4096 to 12188
Out of Range Inputs: Ignore Default Value: 6192
Mode: Relative to note onset Holdoff: 0.05 ☐ Always positive

Number of touches maps to MIDI CC # 12
Input Range: 1 to 3 Output Range: 0 to 64
Out of Range Inputs: Ignore Default Value: 0
Mode: Absolute Holdoff: ☑ Always positive

Figure 6: Mapping editor assigns parameters to multi-channel MIDI messages.

MIDI Control Change messages are limited by the fact that they apply to an entire channel, prohibiting polyphonic control. The mapping engine allocates a new MIDI channel for each Note On message, rebroadcasting it to one of several identical copies of a synth plugin, each listening on a different channel. Before the Note On is sent, Control Change messages are sent on the same channel based on the current touch information, ensuring that the controls take their proper position before the note begins. If a Note On arrives before any frames of touch data for the same key, retransmission is delayed for up to 10ms to allow touch data to come in. In practice, this is only necessary for the fastest strikes originating above the keyboard.

3.2 Control and Playability

On the piano, the location of each touch is partially constrained by fingering. Mappings based on the absolute position of each finger in each dimension will thus be challenging. Composer John Eaton remarked of the Moog touch-keyboard: "It's very difficult to play. But an instrument should be difficult to play. That's the only way to master musical materials, by overcoming these difficulties" [7].

Certain simple strategies can produce a more easily-learned instrument, including tracking deviation from initial touch position (motions which rarely occur in traditional technique) or mapping to parameters which are expressive but which do not require precise control to produce acceptable results (e.g. pluck position on a virtual string model). Detailed user evaluation of mappings is currently underway; preliminary results indicate that the two principles above produce usable results.

3.3 Example: Analog-Modeling Synth

The TouchKeys were configured to control the FXpansion Strobe[4] analog-modeling monosynth. Eight copies of the synth were hosted in Apple's AULab environment. Each copy was configured identically but on a different MIDI channel. The TouchKey mapping engine dynamically routed notes from the host keyboard to one copy of the synth, and the touch mappings in Figure 6 were used to send Control Change messages to specific channels:

- Vertical (front-back) position controlled the cutoff of a low-pass resonant filter.
- Horizontal position controlled the pitch wheel to bend notes up or down. This feature is presently only available on the white keys since they sense touch on two axes; it allows an intuitive "vibrato" motion (rocking finger on the key).
- Contact area was mapped to the level of white noise mixed into the main oscillator output. Normal technique (large contact area) creates a pure sawtooth tone, but special effects can be generated by playing on the very tip of the finger or the fingernail. When this is coupled with sharp attack and long release time, interesting percussive sounds can result.
- Number of fingers (1-3) was mapped the number of stacked oscillators, each of which was slightly detuned so that multiple touches produced a wider sound.

This is only one of many possible mappings, but it demonstrates how the TouchKeys can be configured to operate with standard VST/AudioUnit plug-ins.

4. CONCLUSION

This paper has presented the TouchKeys, a multi-touch augmentation of the traditional keyboard. The sensor system can be installed atop acoustic or electronic keyboards, where it provides continuous multidimensional control over each note while retaining the tactile feedback that is important to keyboard performance.

An important area of future exploration is the connection of the TouchKeys to the magnetic resonator piano (MRP), an electromagnetically-augmented acoustic piano developed by the author [4]. Electromagnets inside the acoustic piano can shape the amplitude, frequency and timbre of each note in real time, expanding the piano's musical vocabulary. Multidimensional keyboard control is a natural extension of this project, and indeed, performers frequently suggest finger motion along the keys as a means of note-shaping.

5. REFERENCES

[1] J. Eaton and R. Moog. Multiple-touch-sensitive keyboard. In *Proc. NIME*, 2005.

[2] L. Haken, E. Tellman, and P. Wolfe. An indiscrete music keyboard. *CMJ*, 22(1):30–48, 1998.

[3] R. Lamb and A. Robertson. Seaboard: a new piano keyboard-related interface combining discrete and continuous control. In *Proc. ICMC*, 2011.

[4] A. McPherson and Y. Kim. Augmenting the acoustic piano with electromagnetic string actuation and continuous key position sensing. In *Proc. NIME*, 2010.

[5] A. McPherson and Y. Kim. Design and applications of a multi-touch musical keyboard. In *Proc. SMC*, 2011.

[6] R. A. Moog and T. L. Rhea. Evolution of the keyboard interface: The Bösendorfer 290 SE recording piano and the Moog multiply-touch-sensitive keyboards. *CMJ*, 14(2):52–60, Summer 1990.

[7] NJ Star-Ledger. Waiting to be heard, Wanted: Someone to bring the successor to the Moog synthesizer to life. March 13, 2006.

[8] J. Paradiso and N. Gershenfeld. Musical applications of electric field sensing. *CMJ*, 21(2):69–89, 1997.

[9] C. Roads. Early electronic music instruments: Time line 1899-1950. *CMJ*, 20(3), 1996.

---

[4] `http://www.fxpansion.com/index.php?page=62`