The Responsive Window: A Simple System for Determining the Position and Characteristics of Knocks on Large Sheets of Glass

Joseph A. Paradiso, Che King Leo, Nisha Checka, Kaijen Hsiao

Responsive Environments Group MIT Media Laboratory 1 Cambridge Center, 5FL Cambridge, MA 02142 USA {joep,cheking,nchecka,kjhsiao}@media.mit.edu

ABSTRACT

We describe a system that locates the position of knocks and taps atop a large sheet of glass. Our current setup uses four contact piezoelectric pickups located near the sheet's corners to record the acoustic wavefront coming from the knocks. A digital signal processor extracts relevant characteristics from these signals, such as amplitudes, frequency components, and differential timings, which are used to estimate the location of the hit and provide other parameters, including the rough accuracy of this estimate, the nature of each hit (e.g., knuckle knock, metal tap, or fist bang), and the strike intensity. As this system requires only simple hardware, it needs no special adaptation of the glass pane, and allows all transducers to be mounted on the inner surface, hence it is quite easy to deploy as a retrofit to existing windows. This opens many applications, such as an interactive storefront, with dynamic content controlled by knocks on the display window.

Keywords

touch screen, interactive surface, acoustic tracking, intelligent window

INTRODUCTION

Although new display and projection technologies are enabling very large displays to become ever more inexpensive and commonplace, they are generally not interactive. In contrast, systems with smaller screens, such as video kiosks, are already quite common and often highly interactive through the use of a touch screen or other local interface. When large-screen systems in public spaces are made responsive, however, they naturally tend to open collaborative interaction. Since the display output is no longer localized or private, the use of such attains some characteristics of a spectacle or performance, naturally drawing in others. In order to explore these possibilities, various systems are being designed to make large displays and surfaces interactive [1]. This paper concentrates on one particular approach that instruments a sheet of glass through a simple retrofit of adhering acoustic pickup transducers at its four corners. By processing the signals in real time with fairly simple algorithms running in a DSP, we can estimate the location of a knock or tap on the glass, attaining the functionality of a touch screen. In addition to determining the hit's location, we also extract the amplitude and the spectral characteristics of the strike, indicating its intensity and whether it was a knuckle tap or hard-object (e.g., metal) impact. This enables the interface to become responsive to possible nuance or affect in the knock (e.g., responding more "forcefully" for hard impacts, etc.). Glass is now a very common construction material, often used as clear walls for room dividers or large windows bordering urban buildings. The techniques described in this paper aim at enabling these surfaces to become interactive. For example, information displayed on a projection or monitor on the inside of the glass can be determined by knocking appropriately on the outside. A straightforward application of this niche is an interactive storefront, where passers-by can navigate through information on the store's merchandise or explore special offers by knocking appropriately.

OTHER APPROACHES

In general, most techniques used in conventional touch screens [2] don't scale gracefully to large displays. Resistive sandwiches, the most common technique, made from sheets of conductive plastic that are shorted together when pressed by the user's finger, are now starting to be manufactured in sizes large enough to be used with plasma monitors. They are not yet produced in sizes to cover large windows (e.g., 2×2 meters), and because of their operational principle, would need to be mounted on the window's outside (active) surface, where they would be subject to potential damage over time. Active acoustic touch screens [3] detect the absorption of ultrasound launched into the surface of the glass when a finger is in

contact. Although they don't require any material to be adhered to the glass (aside from the piezoceramic transducers at a corner), they do need the glass' surface to be properly patterned with etched reflectors along the edges that direct the acoustic energy towards the middle of the glass plate. These systems also become error-prone as the surface scratches and deterioriates, plus attenuation of ultrasound in the glass can become an issue for large panes. Capacitive techniques are able to measure hands through the glass, but require a matrix of transparent electrodes to be patterened or adhered across the entire sensitive surface. Infrared LED curtains, featuring dense arrays of opposing IR transmitters and receivers, measure the hand position when the corresponding optical beams are occluded. Although these systems can be scaled up to a larger screen, their expense increases accordingly, plus, as the IR array must line the perimeter of the screen's outer surface, they can be subject to damage, dirt, and deterioration, especially if they're used outdoors.

Other techniques, such as video tracking [4] or reflection of IR illumination by proximate objects [5] have been used to make large screens interactive. Although they are steadily improving, such vision-based approaches can be slow and are often sensitive to image clutter, target reflectance, and changes in background lighting. Time-of-flight laser rangefinding has also been used to scan the surface of large displays [1,6] and find hands, but the potentially expensive laser scanner must be mounted outside the window, leading to reliability difficulties for outdoor operation.

PASSIVE ACOUSTIC IMPACT TRACKING

Although the original motivation for this system was to track the location of knocks on a virtual fishtank [7], the first implementation of this technique was realized in collaboration with Hiroshi Ishii's Tangible Media Group for the PingPong+ (PP+) augmented ping pong table [8]. Here, four contact electret microphones recorded the impact of the ping pong ball at the four corners of each contiguous half of the table. A simple 8-bit "PIC" microcontroller was able to adequately time the leading-edge of each microphone signal as it arrived, and the difference in arrival times (processed through a linear least-squares fit) produced the Cartesian coordinates of the impact, essentially in real time. The knock of a knuckle on glass, however, is markedly different from the impact of a ping pong ball on wood. The ball provides a consistent, highamplitude waveform with a steep rising edge, requiring very little signal processing. The knuckle taps, however, are much more difficult - each tap is different (depending on anatomy, intensity, style of knock, etc.), and generally of lower frequency, hence exhibit a much softer edge, which is much more difficult to time. The wave propagation characteristics of the glass add even more Knuckle taps launch what are termed complexity. "bending" or "flexural" waves [9] into the glass: the glass bends into "ripples" that propagate across its surface, out from the point of impact. These are structural acoustic modes that travel much more slowly (e.g., by roughly an

order of magnitude) than the sound velocity in glass. Although this helps for resolving position, flexural waves are highly dispersive, causing the wavefront to increasingly break apart as it travels further through the glass and high-frequency compents arrive earlier. Different propagation velocities and analysis techniques must be used for knuckle taps and hard-object (e.g., "metal") taps, which travel nearly a factor of three times faster through the glass. These complications prevented the use of the simple constant-threshold discriminator as used with PP+; the signals in this case needed to be digitized and processed algorithmically.



Figure 1: Hardware configuration for knock tracker system

HARDWARE SETUP

Figure 1 shows the layout for our acoustic tap tracker Although the contact microphones that we system. originally used [10] were made from strips of simple PVDF piezoelectric foil [11] laminated onto the surface of the glass, we have shifted to piezoceramic transducers made by Panasonic (the EFV-RT series), which provide considerably more sensitivity across the full bandwidth of various knocks and taps (e.g., 100 Hz - 5 KHz). Figure 2 shows photos of both transducers attached to the glass (the electronics for the Panasonic mic are shielded from interference by a copper housing). Both systems employ a gain-of-10 voltage amplifier mounted at the transducer, enabling a long cable to be driven to the signal conditioning electronics, which filter out unwanted noise while applying additional gain. An inexpensive digital signal processor (the ADMC 401 from Analog Devices, designed for motor control applications) then extracts relevant parameters from the digitized waveforms and ships them across a serial link to a standard PC, which classifies the knock, estimates its coordinates, and runs appropriate interactive content.

Figure 3 shows a photograph of a prototype window under test in our laboratory with the microphones attached and Figure 4 shows the finished unit containing all signal conditioning and DSP electronics. Although we generally adhere the microphones to the glass surface with standard "crazy glue" adhesive, long-term installations may prefer a more permanent epoxy joint.



Figure 2: Contact mics - original PVDF w. preamp (left) and current piezoceramic unit w. assembly adhered to glass (right)



Figure 3: Test Window, with mics adhered at 4 corners

Although the piezoceramic transducers respond well to knocks and metallic hits, they don't give much signal when the window is banged with a fist. This primarily introduces low frequency vibration, down below 50 Hz, where the pickups aren't as sensitive. Accordingly, we have attached another pickup to the window to detect these "bash" events; it is an inexpensive electrodynamic cartridge, with the diaphram epoxied to the window glass. As such bashes are of very low frequency, they don't tend to provide good timing resolution and tracking, hence one transducer is sufficient, placed near the edge of the glass where it is out of the way. As this pickup responds very strongly to bash events (and very weakly to knuckle taps and metal hits), it detects them with essentially no ambiguity.



Figure 3: Unit housing analog electronics and DSP

Because the piezoceramic transducers are strongly adhered to the glass, they are essentially contact microphones and give very little response to signals coming from the air and not generated by taps on the glass. Certain sounds, however, such as a very loud clap or "snap" produced near the window, couple enough energy into the system to trigger a false event. To prevent this from occurring, we have provided for another transducer in our system; this is a simple crystal microphone, not adhered to the window, but listening to ambient sounds made in its vicinity. The peaky high-frequency response of the crystal mic makes it respond strongly to the clap sounds that instigate false triggers, hence any events that exhibit a significant amplitude from this channel are vetoed as external sounds, not assumed to be generated in the glass.

Although the signals from the four piezoceramic transducers of Fig. 1 appear to provide sufficient information to detect bashes and veto claps on their own, this would require considerable additional signal processing and tweaking - the additional two transducer channels make this determination much more simple and robust.

The DSP provides 26 MIPs worth of processing and samples all 6 input channels into 12 bits at 50 kHz. It continuously samples the input signals, triggering when any rise above a preset threshold. We currently retain 8 ms worth of data from the 4 piezoceramic transducers and the clap veto transducer (5 ms before the trigger and 3 ms afterwards, nicely capturing the incoming wavefront before the glass begins steady-state modal oscillation). As its low-frequency signal comes somewhat later, we retain 20 ms from the bash transducer (5 ms before trigger and 15 ms afterwards).

Figure 5 shows data from one transducer triggered by the signal from another for a series of knuckle taps running

between them atop a 1 cm-thick piece of glass; the pickups are separated by 0.9 meters. The progressive time delay, upon which this technique is based, is quite evident, as is the dispersive nature of the glass.



Figure 5: Increasing delay as knock moves away from pickup

DATA PROCESSING

Our earlier attempts [1] ran on a Hitaschi SH-1 embedded microcontroller, which digitized and processed the transducer data. Our initial strategy was centered around triggering the system when the amplitude of a transducer rose well above noise, then walking up the attack of the knock transient to determine the point at which the signal emerged from the noise floor. Although this attained a basic degree of performance, it could be quite erratic, as background noise and the low-amplitude high-frequency signals arriving earlier through dispersion could often have considerable (and variable) influence on the timing. Our next steps [10] digitized the signals with an data acquisition card, allowing us to explore more complex algorithms We attained much more robust under MATLAB. performance by cross-correlating the signals across pairs of sensors and extracting the differential time from the correlation peak. Because of the distortion encountered when propagating through the glass, however, the correlation could sometimes become ambiguous, producing two or more significant peaks. In these cases, the redundancy in the system (provided by the additional sensor - only 3 are needed to specify position in a plane) and data from a calibration procedure were used to select the appropriate peak. Details are provided in Ref. [10].

We have recently improved this approach by limiting the number of cycles of each waveform used in the crosscorrelation, as the signal after the transient passes is dominated by multipath and modal oscillation, hence has little dependence on the strike position. We likewise use a heuristically-guided chi-square matching scheme to pick the segments of the waveforms from each sensor that are most similar to one-another, then cross-correlate these to determine the time delay. These measures generally result in a much clearer correlation maximum with less ambiguity.



Figure 6: Histogram of knuckle and metal tap frequency

This correlation technique is used for knuckle taps, which are generally at much lower frequency (e.g., several hundred Hz) than hard taps with a metallic object, which can create significant components in the vicinity of 2 KHz. As the metal taps exhibit a much sharper transient, we default to the original technique, where we walk back along the waveform after triggering to detect where the first arrival of the signal disappears below the noise floor. Likewise, because of dispersion and the different acoustical modes launched by the hard, metallic impact, the propagation velocity for metal taps is generally 2-3 times faster than for knuckle taps. Although one could discern knuckle taps from metal taps by examining the frequency distribution obtained from a FFT, we accumulate a coarse but adequate estimate of frequency by counting the number of times the sensor waveforms cross zero across a fixed interval. As seen in Figure 6, which shows this quantity for several knuckle taps and metal taps, the two distributions are nicely separated and this distinction is quite reliable.

The operations described above are currently all executed in the DSP, and all relevant parameters are sent over a serial link to the attached PC, which then determines the impact position. In our earlier work [1,10], we derived the (x,y) impact coordinates by running the timing results through a third-order polynomial that was determined via a linear least-squares fit to data collected on a regular grid. To avoid this lengthy calibration process (that was prone to overfitting), we have recently shifted to a deterministic algorithm. As the position ambiguity curves for any pair of sensors form a hyperbola [12], we calculate the closest intersection between a pair of hyperbolas, each derived from the data from two different sensors. The amount of miss between the hyperbolas reflects the quality of the data, hence the resolution of the position estimate.

The peak amplitude is also captured for each sensor across the data acquisition interval. For the four contact pickups in the corners, these quantities reflect the intensity of the hit. Although some techniques use differential amplitude to determine position of impact (e.g., Max Mathew's Radio Drum [13]), since the wavefront is attenuated as it travels through the glass, we have found that the timing data is much better behaved. The peak amplitude of the other two sensors are used to determine a "bash" event (if the bash sensor amplitude dominates), or a "clap" veto (if the clap sensor amplitude dominates).



Figure 7: Estimated positions for knuckle taps at 5 locations



Figure 8: Estimated positions for metal taps at 5 locations

PERFORMANCE

We have used this system thus far with primarily two different types of glass, a 1/4" thick pane of tempered window glass and a 1-cm thick pane of shatterproof roomdivider glass. In the 1-cm pane, we observe a propagation velocity of 670 m/s for the knuckle tap and 1700 m/s for the metal tap; the speeds are roughly 20-30% slower in the thinner glass. Our present systems have been used with an active area between the sensors of roughly a meter and we are currently instrumenting a 2-meter by 2-meter window, with results expected shortly. The results presented here in Figures 6,7,8 are all taken with the 1-cm glass, with sensors spaced at the corners of an 80-cm square.

Figure 7 shows the (x,y) reconstructed points for knuckle knocks at 5 locations (25 knocks at each site), and Figure 8 shows the analogous data for metal taps (the "+" symbols in the corners denote the sensor positions). Although the net scatter is similar (we see a resolution of = 3.5 cm for the knuckle taps and = 3.3 cm for the metal taps at each position), they have a totally different structure. While the metal taps are uniformly distributed (with a few stragglers in the tails), the knuckle taps appear to cluster into 3-4 tight clumps at each strike location. This is due to the peakpicking and waveform aligning logic associated with the cross-correlation process in the data analysis - the clumps are essentially offset by integral delays of a waveform cycle. The tight nature of these clumps indicates that a significantly better resolution may be able to be obtained by picking a better correlation match, perhaps through an iterative dispersion-canceling algorithm.

Because of systematic inaccuracies in the hyperbola definitions, this data shows some systematic skews (averaging 8 cm) between the actual hit locations (\mathbf{X}) and the centroids of the reconstructed points (\mathbf{O}). This is an artifact of the way in which the hyperbolas were defined, however, and can be easily calibrated out.

The analysis was able to provide good coordinates for essentially all knuckle taps (only one tap was declared inconsistent out of 125 candidates), and the vast majority of metal taps (where 13 out of 125 candidates were failed).

The window upon which these sensors was mounted extended for over a meter above and below the region plotted here (it was grabbed along its sides by a rubber bumper to quickly damp resonances and avoid rattling). As a result, a user was able to tap both above and below the region bounded by the sensors. The differential timings are able to determine this condition, however, and still produce coordinate estimates. As the time differences in the vertical direction show little change with vertical position in these regions, the resolution here was poor; the horizontal coordinate was still usable, however, reaching the order of = 5.5 cm when tapping a foot above or below the sensor-bounded perimeter.

The system responded quite quickly. The DSP generally produced parameters within 50 ms of an impact, and a PC (a 1 GHz Pentium 3) produced x,y coordinates circa 15 ms later, yielding a net system latency within 65 ms. Tighter coding could speed this up considerably.

The trigger thresholds were placed sufficiently low to respond to a soft knock and not activate sporadically with room noise. When sharp room sounds did occur, the "clap" veto signal successfully prevented the system from falsely responding. The "bash" detector reliably discriminated between knocks and even modest fist "bangs".



Figure 9: A Knock-Activated Browser



Figure 10: Responsive Window at the Ars Electronica Center

Figure 11: The Responsive Window in action

APPLICATIONS

As this system has evolved, we have explored its use with increasingly sophisticated applications. All involved graphics projected onto a screen behind the glass; the user interacts with the information by knocking, tapping, and banging on the front surface. In all cases, the pickups were mounted on the rear of the glass.

Our first application [1] was a simple diagnostic that plotted circles centered at the location of the knock, with radius proportional to the estimated coordinate accuracy. After perfecting our hardware and analysis with this tool, we built a more user-relevant application, illustrated in operation in Figure 9. This is a "knock-driven" browser, where a user navigates group projects by knocking on the corresponding picture, launching relevant web pages and/or playing associated video clips. The targets (bounded by the interpicture gaps) were sufficiently far apart to make this system quite usable, even with a coarse resoultion of several centimeters. This system was shown at several large sponsor meetings at our laboratory, where its novelty made it quite popular.

Our next deployment, shown in Figures 10 and 11, was somewhat more sophisticated. It is a semipermanent installation at the Ars Electronica Center that was debuted in the 2001 Ars Electronica Festival as the "Responsive Window" [14]. It is the system on which the data of Figs. 6-8 were taken. Figure 10 shows the setup from a distance - it is a piece of 1 cm glass 1 meter in width and 3 meters high upon which our sensors are affixed as mentioned previously. A "holoscreen" [15] is mounted behind the glass and located within the perimeter enclosed by the transducers. The holoscreen is a holographic diffuser that gives a rather "ghostly" appearance to the graphics. Nothing is projected (and the screen is mildly translucent) until the glass is knocked, at which point the images appear.

The graphics that we are running resulted from a collaboration with Ben Fry from the MIT Media Lab's Aesthetics and Computation Group. It's essentially a simple knock-driven drawing program that illustrates the full responsiveness of the system. One knock generates a circle centered at the impact coordinate that oscillates as a function of knock intensity. Two and three knocks extrude a rotating 3D object, with dimensions extending out to the knock coordinates. The rotation rate and fill intensity are similarly a function of the knock intensity. A forth knock sends the object spiraling into the screen (with rate depending on knock intensity), vanishing at the impact location. The color of the object is a function of the spectral content of the sensor waveforms, with knuckle knocks producing bluish objects and metal taps producing reddish objects. Fist bangs flash the entire screen with a bright red transient, with large rotating rings centered at the estimated impact location. Finally, hits above and below the sensor perimeters launch circular waves into the screen, centered on the assumed horizontal location of the hit.

We are now in the process of adapting this technique to a much larger piece of glass, measuring 2×2 meters, for an interactive art exhibit that will be exhibited at the Kitchen in New York City in association with the San Francisco artist JD Beltran. It is another "knock-screen" application, where a desk full of objects is projected upon the glass; when the images of relevant objects are knocked, a corresponding video stream is launched. If the user knocks during the video, a relevant still image will briefly flash up, centered at the knock location.

Although dispersion effects can be more significant as the impact wavefront travels through a larger amount of glass, we will be using the 1/4" tempered sheet here, which appears to have better propagation characteristics and provide a better resolution than the 1 cm glass used in the results presented here.

We are considering deploying this system next in a large outdoor storefront on a busy street to explore applications in interactive retail experience. Acoustic sensors are no stranger to this venue, being commonly used here to detect glass breakage and deter theft. Now we will encourage (hopefully less violent) knocks as people explore interactive content, perhaps bringing something new to the art of window-shopping. Although our sensors are made to be adhered to the inner surface of a glass sheet, they will probably not function with air-insulated double-paned windows, as the rear surface is acoustically isolated from the front, where the impact occurs.

USER EXPERIENCES

As mentioned above, this system has recently emerged from the lab and has been now used by thousands of people at the MIT Media Lab and the Ars Electronica Center. The Media Lab event was monitored by students, who directed visitors to knock on the screen. Once this began, a crowd would generally gather and spontaneously knock-navigate through the layers of content, feeding off on-another's actions. Aside from the novelty of the knocking action, this metaphor is directly related to already ubiquitous touch screen interactions, hence once people start knocking, they had little difficulty using the system. Because of the MATLAB software, the latency in this demo was considerable, however (on the order of a half second!), which was somewhat detrimental to the interaction.

In contrast, the Ars Electronica installation responded essentially promptly, hence latency wasn't a problem. Although there were attendants in residence there, people were for the most part left alone to explore the installation. There was considerable variation in styles, but a common thread was often established in the way in which they approached the interaction. Being used to touch screens, visitors initially started touching the glass. When no interaction resulted, they began lightly tapping. At this point, they quickly exceeded the DSP's trigger threshold, and graphics began appearing. Seeing this response, they began knocking harder, producing more response. Many attendees then started knocking with rings, pens, or other hard objects, launching metal-tap events. Several then graduated all the way, as they began banging on the window with their fists. Fortunately this system is robust enough to take (and aptly respond) to such vigorous usage.

In general, people seem captivated by the full response to the knocking interaction - a very natural gesture between people across the physical boundaries of windows and doors but still rare across the boundary between the real and the virtual. Knocking does have its drawbacks, however; after several hours of using the system, one's knuckles can feel the accumulated strain, as glass is a notoriously hard material. This is certainly an interface for particular niches, not common interaction.

CONCLUSIONS

We have demonstrated a technique of easily retrofitting common windows for contact interactivity by measuring the position of a knock or tap, determining the type of impact, and estimating the impact intensity. Because of the complicated nature of the various impacts and nonideal propagation characteristics of the glass, this is not a precise pointing device, yielding resolutions on the order of 3.5 cm across roughly a meter. Its accuracy is adequate, however, for several applications, such as navigating through content at interactive storefronts. Better signal analysis could extract more reliability and accuracy from this system; some possibilities could involve locating the arrival transients with a wavelet analysis or compensating for the dispersion in the glass by using coarse estimated range information to remove the expected dispersion from the signals. For larger windows, more transducers can be redundancy added. providing increased in the measurements and keeping the minimum propagation distance small. Also, more sensitive pickup sensors could be investigated, such as wideband accelerometers, providing reliable response across larger areas. We have brought this system out to the public, where it is received with interest because of its novelty - once users recover from their initial confusion with touch screen operation (one must hit rather than touch), they generally adapt well to it.

ACKNOWLEDGMENTS

We acknowledge our artistic collaborators (Ben Fry of the MIT Media Lab and J.D. Beltran of San Francisco), Steve Francis of Analog Devices for assistance with the DSP system, and our colleagues Josh Lifton and Warit Wichakol for their help with early versions of this system. We thank our laboratory sponsors for supporting this project.

REFERENCES

- Paradiso, J.A., Hsiao, K., Strickon, J., Lifton, J. and Adler, A., "Sensor Systems for Interactive Surfaces," *IBM Systems Journal*, Vol. 39, No. 3&4, October 2000, pp. 892-914.
- 2. Quinnell R.A., "Touchscreen technology improves and extends its options," *EDN*, vol.40, no.23, 9 Nov. 1995, pp.52, 54-6, 58, 60, 62, 64.

- 3. Kent, J. (1997). Acoustic touch position sensor using higher order horizontally polarized shear wave propagation. US Patent No. 5591945.
- 4. Maes, P. et al. 1997, "The ALIVE System: Full-body interaction with animated autonomous agents," *Multimedia Systems*, vol.5, no.2, March 1997, pp.105-112.
- Rekimoto, J., Matsushita, N., "Perceptual Surfaces: Towards a Human and Object Sensitive Interactive Display. *In Workshop on Perceptual User Interfaces* (*PUI-97*), Banff, Published by Microsoft, October 1997, pp. 30-32.
- Strickon, J., Paradiso, J., "Tracking Hands Above Large Interactive Surfaces with a Low-Cost Scanning Laser Rangefinder," *CHI '98 Conf. on Human Factors in Computing Systems: Summary*, ACM Press, New York, 1998, pp. 231-232.
- Galyean, T., et al., "Virtual FishTank," SIGGRAPH 98, Conference Abstracts and Applications, ACM Press, New York, 1998, p. 116.
- Ishii, H., Wisneski, C., Orbanes, J., Chun, B., and Paradiso, J., "PingPongPlus: Design of an Athletic-Tangible Interface for Computer-Supported Cooperative Play," in *Proceedings of Conference on Human Factors in Computing Systems (CHI '99)*, Pittsburgh,

Pennsylvania, May 15-20, 1999, ACM Press, NY, pp. 394-401.

- Cremer, L., Heckl, M., and Ungar, R.R., Structure-Borne Sound, Second Edition, Springer-Verlag, New York, 1990.
- Checka, N., "A System for Tracking and Characterizing Acoustic Impacts on Large Interactive Surfaces," MS Thesis, MIT Department of Electrical Engineering and Computer Science and MIT Media Lab, May, 2001.
- Paradiso, J., "The Interactive Balloon: Sensing, Actuation, and Behavior in a Common Object," *IBM Systems Journal*, Vol. 35, Nos. 3&4, 1996, pp. 473-487.
- Reynolds, M., "An Immersive, Multi-User, Musical Stage Environment," ACM Computer Graphics, SIGGRAPH 2001 Proceedings, ACM Press, NY, 2001, pp. 553-560.
- Mathews, M.V., "The Conductor Program and Mechanical Baton," in Mathews, M.V., and Pierce, J.R., eds., Current Directions in Computer Music Research, MIT Press, Boston, MA, 1989, pp. 263-282.
- Paradiso, J., "The Responsive Window," in Take Over, Proc. of the 2001 Ars Electronica Festival, Springer Verlag, Linz, Austria, September 2001, pp. 261-263.
- 15. See: http://www.holopro.de/