

Cutting the cord - In-circuit programmable microprocessors and RF data links free the performer from cables.

D.G. Malham
Department of Music, University of York

Abstract

A versatile combination of an in-circuit programmable microprocessor and standard licence-exempt R.F. modules is investigated as a replacement for the cumbersome and restricting cable harness usually required when using sophisticated, performer mounted sensors.

Introduction.

This paper describes the development of sensing and data link technologies for performer-computer interaction. This work was pursued within a practical context - the realisation of *construction 3* for soprano saxophone and multiple media within the context of the RIMM project (RIMM, 2001). The guiding principals were:

- to establish performer control of the work;
- to eliminate physical connections between the performer and computer, establishing performer independence;
- to implement a functional (playable) hardware interface that was sensitive to the performers instrumental technique;
- to define appropriate mappings of performer data for sound/graphic control and generation;
- to allow a flexible interpretation of the composition during performance;
- to allow performer control of the surround-sound presentation of the work.

The sensing system and mapping strategies in *construction 3* establish a relationship between performer and computer by combining several data sensing and communication strategies simultaneously. This is an attempt to develop a simple model that incorporates some of the inter-performer communication channels which exist during performer-performer interaction, combining communication gestures, aural information and visual cues.

Analysis

Video tapes of the performer were analysed to determine the methods used to sense and derive control information. This analysis generated a specification for the sensing devices selected and their application. Following the analysis, the instrument's movement, the movement of feet(toes) and arms were defined as performer data sources in addition to the acoustic audio signal. All of these data sources were considered independently and in combination to provide information about the performers movements and gestures.

System Specification

In order to accommodate the requirements indicated above, a wireless data acquisition system was developed which integrated a variety of hardware sensing devices: accelerometers, digital compasses, force sensing resistors, switches. The data acquisition system was designed to allow the use of a flexible word length to represent data from analogue sensors. This provides an improved dynamic range for continuous controllers over MIDI-based controller devices.

In order to achieve the aims of the project, it was regarded as necessary to;

- give the performer the fullest possible control over the realisation of the piece
- minimise changes to natural performing style.
- maintain imperceptible delays between performer action and system response

A number of sensor technologies were investigated for their suitability for the task. The candidate technologies were;

- Video tracking
- Force sensing resistors
- Switches
- Accelerometers
- Electronic compass

Video tracking

Two members of the team, Dr. Tony Myatt and Dr. Rob Fletcher, already had extensive experience in the use of video tracking technologies as a result of their involvement in the *Chaotic Constructions* project (Fletcher, 1996). The tracking for *Chaotic Constructions* successfully used standard colour video cameras and specially written colour 'fingerprinting' software to track coloured patches on a kinetic sculpture. However, the existing system had a number of problems which had to be overcome.

In the original form, resolution was limited to 90(h) x 70(v) pixels due to the relatively low performance of the machine used, an Silicon Graphics Indy. The resolution was increased for the current project, at the cost of increased response times, by means of a running average of the positional values. This route was seen to be acceptable since, with a minimum delay of 40mS resulting from the frame rate of the camera (25fps), the response was already too slow for some purposes, such as triggering scene changes. Accordingly, video tracking was only used for less time critical control parameters.

Previous experience suggested that the colour fingerprinting software was less robust than was desirable when using passive, reflected light only targets in conditions of variable lighting. The initial thought was to use coloured LED's as targets to get better colour stability and separation. However, when it was noted that even low-cost cameras were sensitive to infra-red light it was decided to use infra-red LED's coupled with an infra-red filter over the camera lens to reject a large portion of the ambient light in the scene. This was found to greatly increase the robustness of the system. A number of high brightness surface mount LED infra-red emitters were mounted in a 3-d array in order to get a reasonably bright omnidirectional light source. This was attached to the saxophone so that its movements could be tracked and used to control the positions of graphic elements, audio mixing and, via the gesture recognition subsystem, sound source trajectories. Powered by a small battery pack and with a visible LED connected in series to act as a fault indicator, the omnidirectional light source thus formed will run for many hours before needing battery replacement.

Force sensing resistors

After considering the possibility of positioning Force Sensing Resistors (FSRs) (Steadlands, 2001) as 'switch pads' on the body of the saxophone, a decision was taken in conjunction with the performer to use them as foot, or rather, toe, operated switch devices instead. A set of thick insoles was modified by cutting out an area under the big toes and replacing it with a force sensing resistor. This approach initially produced good results but with increasing use problems of false triggering appeared, resulting from normal movements of the performers feet. Attempts to overcome this problem by varying the threshold or by mounting the FSR's in the upper part of the shoe above the toe failed to meet the required repeatability standards. Discussions with the performer showed that this was largely because of the lack of any form of mechanical feedback. As the main use for these devices was to initiate scene changes, this level of repeatability was unacceptable.

Switches

The initial solution to this problem was to replace the FSRs with dome switches, such as are in membrane keyboards. Although not much thicker than the FSRs, these have a significantly greater mechanical travel and a positive 'click' or snap action which the performer was easily able to feel when placed under the big toe. The repeatability of these switches was better, but an extra layer of protection was added in the control software by only looking for switch changes within a window around the time of any expected scene change. However, long-term reliability of these switches is not good, so the switching mechanism may be replaced with fsr's whilst keeping the snap action dome for haptic feedback.

Accelerometers

Recently, low cost dual-axis surface micromachined integrated accelerometers have become available (Analog Devices 1999) which exhibit sufficient sensitivity and dynamic range to enable the tilt of an object in a standard 1G gravitational field to be measured with high accuracy. Moreover, they output a simple, pulse width modulated digital data stream, rather than an analog voltage, thus facilitating connection to the computer. The smallest available unit, in a 5mm

square package, was mounted, with its ancillary components on a small (12mm x 18mm) PCB. This only needed a simple a five wire connection and could be mounted almost anywhere. Various positions were tested including in finger rings, on the arm (typically above the elbow) and on the saxophone itself. The finger ring proved too inhibiting to the performer during normal playing and the saxophone mount proved to need movements which were awkward for the performer if they were to be distinguished from normal performance motions. The above elbow mount proved least obtrusive whilst still providing a good degree of control and a high level of repeatability. It was found in practise that there was a significant degree of 'jitter' in the raw output of the device, caused by inadvertent small movements and vibration pickup so running averages were used to smooth the data and further enhance resolution. Since the raw data rate was so high (and unlike the situation in the case of typical video tracking systems) the additional delays were not a significant problem.

Electronic Compass

The low response speed of most video tracking systems set limits to their usability. Preliminary investigations have established the feasibility of using magnetoresistive (MR) sensors to sense position within the earth's magnetic field directly. Although current devices are physically larger than the accelerometers and have a higher power consumption, it is possible to use one on the saxophone together with an accelerometer, thus enabling the position of the saxophone bell to be sensed in three dimensions with a significantly higher update rate than via video tracking. This system replaces the earlier video tracking system, thus removing a layer of complexity whilst at the same time increasing performance.

Sensor Interfacing and Communication

Signal Conditioner

The accelerometers, switches and compass are connected to a microcontroller within the **Signal Conditioner (Figure 1)** where their outputs are conditioned into a form which can be sent to the main system via a radio data link operating on the license free 433MHz frequency band to the European EN 300-220 standard (Radiometrix 2001). The microcontroller used for this task was the *In-Circuit Programmable* version of the

MicroChip PIC unit (Microchip 2001), since it offered a simple instruction set coupled with low power consumption (essential for what was to be a battery powered system) and a good range of development tools. Perhaps more importantly, though, the device can easily be reprogrammed via a serial link, enabling the unit to be reprogrammed 'on the spot' to meet the needs of a particular project. All the software developed for this project is modularised to assist in this. What follows describes the system as developed for *construction*

3

The switch inputs require little in the way of processing except inverting their outputs since depressing a switch produces a low level, releasing it a high one. The data from the two switches is transmitted in both normal and inverted form in order to provide an extra error checking mechanism, if the true and complemented forms do not match on reception, the data is rejected.

Processing for the accelerometers is considerably more complex but fortunately a code template was available from the manufacturers (Analog Devices 2001/2). This, however, had to be modified extensively to:

1. Allow more than one accelerometer to be used on a microcontroller
2. To change the arithmetic from sign/magnitude to two's complement
3. Install an 'accelerometer not present' routine to stop the microcontroller becoming locked up
4. Add a routine to perform a running average, based on the formula

$$\text{Accel}_{\text{ave}} = (1-n)\text{Accel}_{\text{ave}} - (n)\text{Accel}_{\text{new}}$$

where n is a value between zero and one. In this case, n is set to 1/16.

Wireless Links

The wireless data links used separate transmitter and receiver modules with a maximum data rates of 120 kBits/second. However, the link was run at 80 kBits/second as this proved to be more reliable in operation. For this class of radio link, continuous operation is not normally an approved procedure and assuming that a 10% duty cycle may become compulsory, a considerably reduced rate of around once every 30mS is obtained. However, as the duty cycle rating is based on an average over 5 minutes, the rate can be tailored to be faster when necessary,




Figure 1. System block diagram

for instance when a switch change is detected. Using separate transmitter and receiver modules, rather than the more usual combined (*transceiver*) unit has a number of advantages. Firstly, with transceivers, antenna choices are limited by the requirements of license free operation, to one of three low gain types in order to keep effective radiated power within legal limits (Radiometrix 2001). However, if the receiver does not share an antenna connection with the transmitter, only the transmitter need use the standard aerial and the receiver can use a higher gain, more directional type such as a Yagi. Not only does this improve the received signal/noise ratio, but the directionality helps to reject the reflections which can cause pattern sensitivity problems. A second advantage is that since the performer unit need only contain the transmitter module, current demand from the battery pack is reduced by at least 33%.

Signal Conditioner

Hardware

The unit is contained in a small belt mounted box containing two PCBs and a battery pack. The main PCB holds the microprocessor, power regulator, clock and connectors for the sensors. Extensive precautions were taken to avoid EMC problems. The transmitter module was mounted on a second PCB connected to the first with a ten way ribbon cable. As well as the transmitter, this board contains a power switching circuit to enable the transmitter to be turned off when not needed. The antenna, a 155mm length of wire, is attached

directly to the transmitter output pin in order to meet licensing conditions.

Transmission Protocol.

The transmission protocol chosen was one in which the data is split into 4 bit nibbles which are in turn used to look up a 6 bit *data transmission unit (DTU)* in a pre-computed table. The 6 bit DTUs form a DC free code, each symbol of which has three ones and three zeros (see **table 1**). This is necessary in order for the adaptive data slicer in the receiver to function properly but it also provides a certain amount of error correction capability since only 16 out of 64 possible codes are valid. A further layer of error detection is provided via a *checksum* which is calculated in real-time from the contents of all the DTUs in a *data burst (Figure 2)* during transmission. This is transmitted at the end of the data burst. The transmitted checksum is compared with one calculated in the *interface/communications controller (ICC)* on reception, errors being detected by the ICC when the calculated and received checksums are different. No error correction capability has been implemented, the ICC simply holding the last correct values when an erroneous transmission is received. Since data bursts occur every 30 mS, testing has shown that a correct burst will, in general, be received within one or two iterations (30-60mS), providing the link has not become hopelessly degraded. With line-of-sight operation over a few metres, detected error rates were a few per hour even without using a higher gain, more directional antenna on the receiver with undetected



Figure 2. Data Burst Structure

rates falling well below this. With non-line-of-sight operation, for instance, with the transmitter at the end of a 10 m corridor and the receiver in a room at right angles to it, detected error rates have proved to be about one in a hundred, with undetected errors in the region of 1 in 10,000. This represents an extreme condition which should not occur in normal use. Should it be necessary to improve on this for performances venues with particular difficulties, directional antennas could be used as a first step, with the possibility of using two (or more) ICC's spaced apart to form a *dual diversity* reception system which is a well known technique for increasing reliability in radio transmission systems.

Data Burst Structure

Each data burst starts with a quarter data rate square wave section then a further four cycles of half data rate which acts as a sync character. The first section acts as a 'primer' for the receiver's internal self adjusting data slicer with the quarter data rate section being employed to set the internal clock rate of the software data recovery system in the receiver. In the present implementation, the first six bits after the sync character are used to encode the switch data nibble, followed by 12 bits encoding the X accelerometer data, then more blocks of 12 bits containing Y,U and V channels of accelerometer data plus a spare block for use later. This is followed by eighteen bits containing the checksum (three nibbles). The transmitter is turned off after this to conserve battery power. The data bursts are normally transmitted every 30mS but if any change of state occurs in the switches in between normal transmissions, the sequence is re-timed so that a new burst is transmitted immediately, thus ensuring that no more than a maximum delay of 6 mS occurs for the most time critical data.

Power-up Mode Selection

Apart from its normal operating mode, the SC can be set into either test mode or calibrate mode by setting different switch combinations during power-up. modes if the initial state of the foot switches is selected appropriately. In a test mode, a preset data packet with the same structure as a normal data packet but with fixed contents is transmitted continuously. This is used by the ICC to check for undetected errors in the link by comparing the received data with what it knows should have been

sent. In calibrate mode both accelerometers are calibrated for rest position output.

Packaging

The whole unit fits into a small belt pack (**Figure 3**) with a 9 volt battery as power source. Battery life is tens of hours.




Figure 3 Belt pack and accelerometer

Interface/communications controller

The ICC accepts commands from the main computing system over the rs232 interface and responds by either sending the appropriate data back or by entering or leaving test mode, during which the collection of link statistics based on the test sequence is performed

Hardware

The same microprocessor board is used as in the SC, but a board is added with an rs232 physical layer interface on it. The receiver module is also on a separate board, like the transmitter module in the SC. The units are built into a fully screened metal box, with a nine pin 'D' type connector for the serial link to the main computer and a BNC connector for connecting an external antenna.

Reception

On reception of a signal, the microprocessor in the ICC switches to monitoring the *received data* line from the receiver for pulses with a nominal

length of two bit periods with a 25% margin of error. If it does not detect a pulse of this length within sixteen bit periods (200µS), the process is aborted. If a pulse of this length is received, it measures the length of the first full block of eight pulses. The timing is re-measured for each data block and is averaged with all previous bit period measurements. The bit period timing derived from this is used to set the actual received clock. All further time based operations on the input data block are derived from this measurement which, as a result of its averaging of a number of input pulses, enables the clock to be set sufficiently accurately to avoid problems caused by the normal manufacturing tolerances (+- 0.01%) of the crystal oscillators used for timing. This procedure has a significantly lower software overhead than the more frequently used software phase locked loop system.

Once the preamble has been received the pulse timing subroutine starts looking for the one bit length pulses which form the sync character. When an appropriate sequence of pulses is received the data bit reception routine is entered. For the expected duration of the actual data portion of the received data burst this routine samples the received data line eight times in the middle of each *bit cell*, ignoring the edge regions where noise would have most effect. The logic level in each cell is derived by applying majority logic - if there are more 'one' results than 'zero' results from a bit cell, the value in that cell is deemed to be 'one', otherwise it is 'zero'. This imparts a considerable degree of noise immunity to the process.

The received data bits are assembled into blocks of six, thus reconstructing the original DTU. A checksum is computed from them before they are stored in memory for later decoding. When the whole data burst has been received, the received DTUs are decoded to recover the original data. If during decoding any of the DTUs fall outside the accepted range of the dictionary used in data encoding an error is flagged. Likewise, if the computed checksum does not match the transmitted one, an error is flagged.

The two error types increment different counters and these are used to supply link error statistics. If no errors are detected, the received data is used to update the stored values of data. If any error is reported, the old values are retained and the new ones discarded. The switch status has a further

level of error detection since they are used to effect major scene changes which makes them particularly critical. The stored values are returned, on request, to the main computer over the rs232 link. The various command codes that can be sent to the ICC are detailed in Table 2.

In test mode, a further error counter is updated whenever the received data does not match the preset pattern **and** neither of the other error detection mechanisms have reported an error.

Operational experience.

As noted above, very low error rates and overall good performance are possible with this system. It was found, however, that in a highly reflective RF environment such as was present for the first performance at SIMK (the concert hall has steel mesh reinforced concrete as a significant part of its structure), the single receiver approach became highly sensitive to antenna positioning. In view of this, implementing multiple diversity reception became a high priority for further development in order to make the unit easier for unskilled personnel to use. This said, the system performed faultlessly in performance.

Multiple Diversity Implementation

Following the experience gained in the *construction 3* concert in Berlin, research into the use of multiple diversity reception was commenced. The low cost of the receiver/ICC units meant that rather than applying a traditional approach employing specially designed receivers, a simpler and more easily extensible approach to implementing a multiple diversity system could be taken. The communications protocol between the ICC and the main computer was modified so that a flag could be sent from the ICC to the computer, utilising an unused bit in the 'switch' response. This flag is set if the data in the data burst contains errors. The main computer can then use this information to make decisions about how to treat the data it has been given. Using two or more separate receiver/ICC combinations, each connected to a different rs232 port, it is possible to select which data stream to use, based on the contents of the flag. In extreme circumstances, many reception channels could be added and the main computer could be used to provide a yet further layer of error protection. By comparing all

the data streams that show 'no error' status further undetected errors can be picked up from the presence of differences between the data streams from different reception points. Selection of the most reliable data is based on voting and/or statistics about error rate trends in each of the channels.

Conclusions

The combination of technologies used and, in particular, the freedom and intimate acoustic relationship which was established by the use of wireless technology appeared to aid the integration of the performer with the performance of the work. The unit's small size, low weight and ease of concealment within the performer's suit rendered its operation essentially transparent to the audience helping to minimise impact on the performer's natural performing style. This will be further enhanced when the RF link is moved to the higher 868-870 MHz band where aerials are even smaller. Moreover, the great bandwidth available means that multichannel operation will be significantly eased, leading the way to ensemble playing.

References

Analog Devices 1999, "Dual Axis, Low g, Fully Integrated Accelerometers"
<http://www.analog.com/library/analogDialogue/archives/33-01/accel/index.html> accessed 04/March 2003

Analog Devices 2001, "Using the ADXL202 With a Microcontroller"
http://www.analog.com/publications/magazines/accel_news/issue7/3.html, accessed 04/October,2001

Fletcher, Dr. R. P., 1996 "Monitoring Colour Change using on-line Video",
<http://www.uauug.org.uk/journal/v2n4/fletcher/rfpaper.html>, accessed 04/March 2003

Microchip, 2001 <http://www.microchip.com/>, accessed 04/March 2003

Radiometrix, 2001, "TX2-RX2" datasheet, downloaded from <http://www.radiometrix.co.uk> 26/May 2001

RIMM, 2001 <http://www.york.ac.uk/res/rimm/>, accessed 03/March 2003.

(The RIMM Project (IST - 1999 - 21022) was supported by the European Commission 5th Framework programme Information, Societies, Technology programme as an Accompanying Measure Trial under key action 3, Action lines: 1.1.2 - 3.2.1.)

Steadlands, 2003 "Pressure Sensing Technology",
<http://www.steadlands.com/pst.htm> accessed 04/March 2003

	<u>000</u>	<u>001</u>	<u>010</u>	<u>011</u>	<u>100</u>	<u>101</u>	<u>110</u>	<u>111</u>
<u>000</u>								
<u>001</u>				<u>x</u>		<u>x</u>	<u>x</u>	
<u>010</u>						<u>x</u>	<u>x</u>	
<u>011</u>			<u>x</u>		<u>x</u>			
<u>100</u>				<u>x</u>		<u>x</u>	<u>x</u>	
<u>101</u>		<u>x</u>	<u>x</u>		<u>x</u>			
<u>110</u>		<u>x</u>	<u>x</u>		<u>x</u>			
<u>111</u>								

Table 1 DC free code book

<u>x</u>	<u>send X axis value from first accelerometer</u>
<u>y</u>	<u>send Y axis value from first accelerometer</u>
<u>u</u>	<u>send U axis value from first accelerometer</u>
<u>v</u>	<u>send V axis value from first accelerometer</u>
<u>b</u>	<u>send both X and Y values from first accelerometer</u>
<u>s</u>	<u>send switch status</u>
<u>f</u>	<u>send a string containing all parameters</u>
<u>T</u>	<u>set test mode, starts collection of link statistics based on the test sequence is started</u>
<u>t</u>	<u>clears test mode, revert to normal operation</u>
<u>r</u>	<u>return link statistics (will include 'undetected error count' from last test mode session)</u>

Table 2 Command options for Interface/communications controller