New Renaissance Institute®
Technology White Paper

THE NRI® RICH TOUCHPAD
Operation and Applications

New Renaissance Institute®
P.O. Box 128
Belmont, California 94002
www.newrenaissanceinstitute.com
Email: inquiries@newrenaissanceinstitute.com

Updated June 18, 2007.

ABSTRACT

This document describes the user-level operation of, and applications for, the NRI® rich touchpad, a novel touchpad controller, based on U.S. Patent 6,570,078, which can be used for a wide variety of real-time applications. The touchpad is very powerful, provides an unprecedented capability to enter large amounts of information at high speed and is extremely flexible in the kinds of input it can process, the kinds of output it can produce and in how it is configured. The touchpad creates images of the pressure exerted on its surface and can recognize images created by contact with different parts of the hand (e.g. a fingertip, flat finger, palm, wrist). It can process multiple regions of contact simultaneously and can extract the values of a large number (typically three to six) of continuous parameters from each region of contact. The touchpad is simple to use, and its capabilities can easily be extended. It can be favorably compared to conventional computer pointing devices, such as the mouse, trackball and conventional touchpad, which typically provide control of only two continuous parameters at any one time and can process only a single region of contact.

The general-purpose nature of the touchpad permits the same basic system to be used in a wide range of applications. They include CAD/CAE workstation control, real-time machine control, human-machine interfaces for the physically disabled and electronic musical instruments. The touchpad can also be used in intelligent machine sensing and robotics applications.

The touchpad can be implemented in a variety of ways. In one implementation, it incorporates a two-dimensional pressure-sensor array, a data acquisition and compression stage, an image processing and recognition stage, and an application interface. Special hardware and algorithms permit the data processing and image processing to be carried out in real time. The system can be modularized to support partitions of the sensor array into functionally discrete regions and aggregations of sensor arrays to form larger arrays.

This document is based on U.S. Patent 6,570,078 and related issued and pending patents, all licensable from NRI®. Detailed hardware and software reference designs can be discussed under negotiable terms. All financial or in-kind proceeds from such arrangements will be used to fund pure academic research at NRI®. Contact inquiries@newrenaissance-institute.com *for more information.*

Table of Contents

1 Introduction

This document, based on U.S. Patent 6,570,078 [1], describes the user-level operation of, and applications for, the NRI® rich touchpad, a novel controller that can be used for a wide range of data entry and real-time applications. The touchpad was originally envisioned as a human-machine interface, though it can also be used in machine sensing and robotics. When used by the hand, the touchpad provides an unprecedented capability to enter large amounts of information at high speed.

The touchpad incorporates a pressure-sensor array for gathering information in the form of real-time images of the pressure exerted on it [2]. When used as a human-machine interface, these images are typically produced by contact with the user's hand, though they can be produced using other parts of the body, such as the foot. The pressure images are presented to a data acquisition and compression stage, which sends its output to an image processing and recognition stage. The image processing and recognition stage is used to identify the shape of particular types of pressure images (e.g. images of a fingertip, flat finger, thumb, palm, wrist). The touchpad can process multiple regions of contact simultaneously and can extract the values of a large number (typically three to six) of continuous parameters from each region of contact. An application interface assigns these values to control signals, which can control arbitrary external systems. Special hardware and algorithms enable the data acquisition, image processing and recognition, and the derivation of parameter values to be carried out in real time [1,2].

The rich touchpad can be favorably compared to conventional computer pointing devices, such as the mouse, trackball, stylus tablet and conventional touchpad. Conventional computer pointing devices typically provide simultaneous control of only two continuous parameters and can process only a single point of contact. In addition, the way in which the rich touchpad is operated is much more natural and intuitive than the way in which conventional pointing devices are operated.

The touchpad can be used as a human-machine interface in a wide variety of applications, including CAD/CAE workstation control, real-time machine control, human-machine interfaces for the handicapped and electronic musical instruments. In addition, as mentioned earlier, the touchpad can also be used in machine sensing and robotics applications. The touchpad naturally lends itself to metaphors useful in a wide range of user interface applications and is simple to use, and its capabilities can easily be extended. Because of the touchpad's general-purpose nature and flexible reconfiguration capabilities, one basic system can be adapted for a wide range of applications.

This whitepaper is organized as follows. Section 2 discusses the user-level operation of the touchpad. Section 3 provides examples of metaphors that can be employed to make the touchpad easier to use. Section 4 describes some ways in which the basic features of the touchpad can be extended. Though a few applications for the touchpad are discussed in these sections, section 5 provides a number of additional examples of applications.

Though they are touched on only briefly in this whitepaper, the touchpad can be used for many music and performance applications. These are described in detail in a companion music and performance applications whitepaper [3], as well as in U.S. Patents 6,570,078 [1] and 6,689,947 [4], and in pending pre-grant patent publications U.S. 2004/0099131 [5], U.S. 2004/0074379 [6] and U.S. 2004/0069125 [7]. For information about the touchpad technology, see the companion technology and applications whitepaper [2], as well as U.S. Patent 6,570,078 [1]. Although not discussed in this document, of related interest may be patents of Westerman and Elias [8,9], as well as the associated proximity detection products made by Finger Works™, Inc.

2 Rich Control

As mentioned earlier, a principal difference between the rich touchpad and conventional pointing devices is that the rich touchpad provides simultaneous control of many more continuous (as well as discrete) parameters. This section briefly describes the pressure-sensor array employed by the user to operate the touchpad and then goes on to illustrate and discuss the touchpad's user-level operation.

Pressure-Sensor Array

An array of pressure-sensors, illustrated in Figure 1, is a central component of the touchpad. The user operates the touchpad by touching the surface of the array. The dimensions of the array range from several inches on each side to a little more than an inch or even less, depending on the application. The spatial resolution of the individual pressure sensors is 1 to 2 square millimeters, and the number of gradations in pressure the sensors can measure ranges from 16 to 256, also depending on the application. The sensors produce pressure measurements, which are passed to an image processing stage. The image processing stage extracts values of various parameters, which are then assigned to signals used to control an external device.

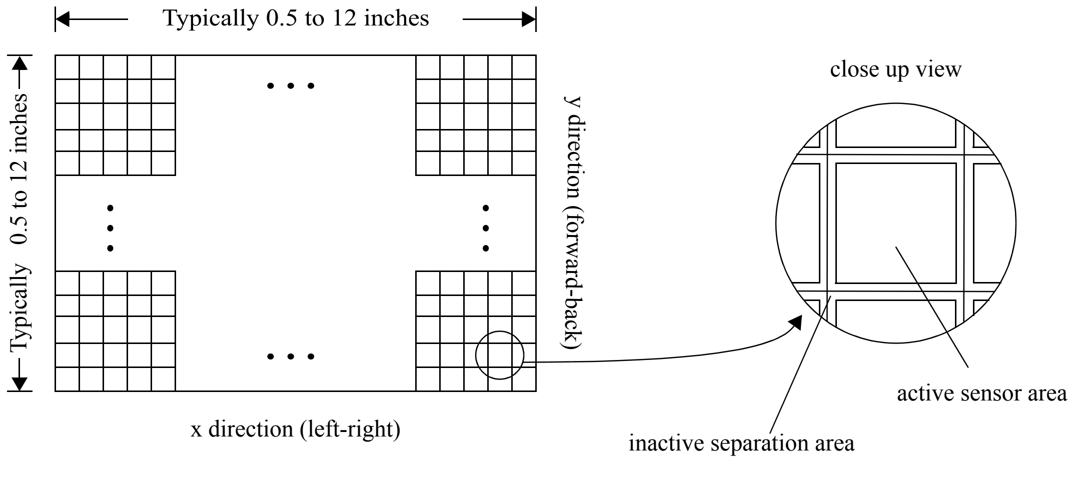

Figure 1. Pressure-Sensor Array

Pressure Images and Regions of Contact

Though the touchpad can process multiple, non-overlapping regions of contact, it will be useful to start by considering a simple case in which it is operated using a single region of contact. Assume the sensor array is contacted by the end joint of a single finger, as suggested in Figure 2a. The pressure image produced by such contact will be similar to that shown in Figure 2b. Note that the darker the image, the higher the pressure.

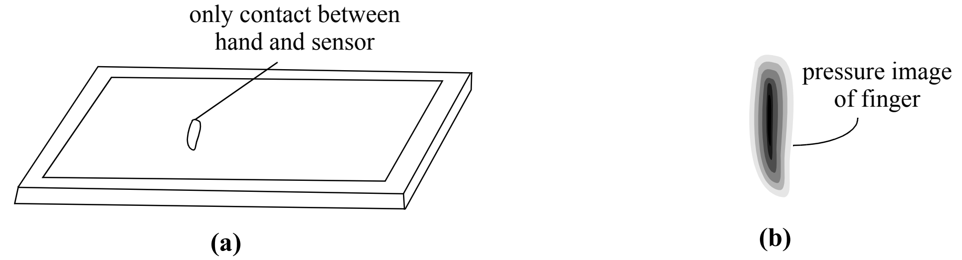

Figure 2. Contact of Finger with Pressure-Sensor Array and Resulting Image

The finger can contact the sensor array anywhere on its surface, so the geometric center of the image determines two parameters -- the center's x and y coordinates -- that are easy to control. The roughly elliptical form of the image has a measurable angular orientation, which the user can vary by moving her wrist so the finger pivots around its point of contact with the sensor array. A moment's experimentation by the reader will show that these three parameters (the ellipse's orientation, and the x and y coordinates of the geometric center of the image) are easy to control independently of one another. It is also possible to control two dimensions of the finger's tilt, the degree to which pressure is concentrated with respect to the left-right axis and the front-back axis. So far there are five continuous parameters that are easily controlled independently of one another. Finally, the user can readily learn to keep all these parameters relatively constant while varying the average or total pressure exerted, giving a sixth independent parameter. Thus the touchpad enables a single finger to simultaneously control six independent parameters in a way that is very easy to learn. By varying the values of the parameters, the user can vary corresponding parameters of an arbitrary external device. Figure 3 illustrates the six degrees of freedom just described.

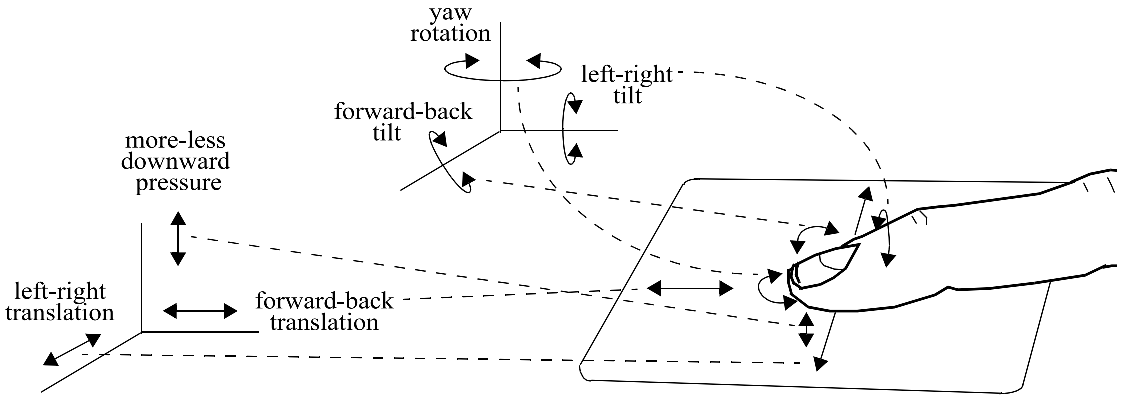

Figure 3. Six Parameters Can Be Controlled Simultaneously with One Finger

Many different kinds of contact with the sensor array can be made using the hand, some of which are illustrated in Figure 4. Figure 4a illustrates the hand position used in the case just described, in which the end joint of the finger touches the sensor array. But, to give just a few examples, the touchpad can also be operated using a fingertip (4b), a flat finger (4c), a thumb end joint (4d) and a wrist (4e).

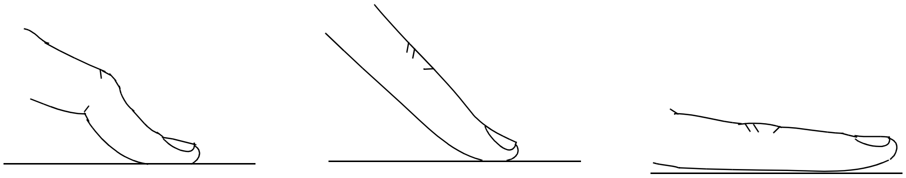

(a) first finger-joint **(b)** fingertip **(c)** flat finger

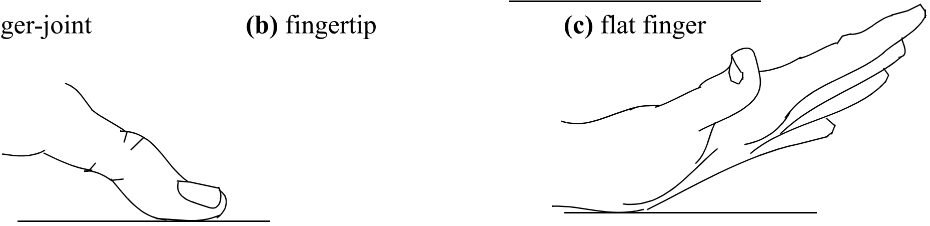

(d) first thumb joint **(e)** wrist

Figure 4. Examples of Hand Positions Used to Operate the Touchpad

The outlines of images created using these forms of contact, as well as some others, are illustrated in Figure 5. As can be seen in the figure, images created with the fingertip (5a, top), the flattened end of the finger pressed with moderate force (5a, middle), the end of the finger pressed with significant force (5a, bottom), the thumb end joint (5c), the wrist (5d) and the side of the wrist (5e) consist of a single contiguous region, while images created with an entire finger laid flat on the sensor array (5b) are compound images, images consisting of multiple, non-overlapping contiguous regions. In general, images created by the sensor array will consist of one or more contiguous regions of non-zero pressure surrounded by a “sea” of pixels measuring zero pressure. As is illustrated in the case of the end joint of the finger, the shapes of these images change somewhat but not significantly as pressure is varied, unless very light pressure is used. Under that condition, many images thin and can break into disconnected regions, and the three-segment image of the flat finger can become a two-segment image.

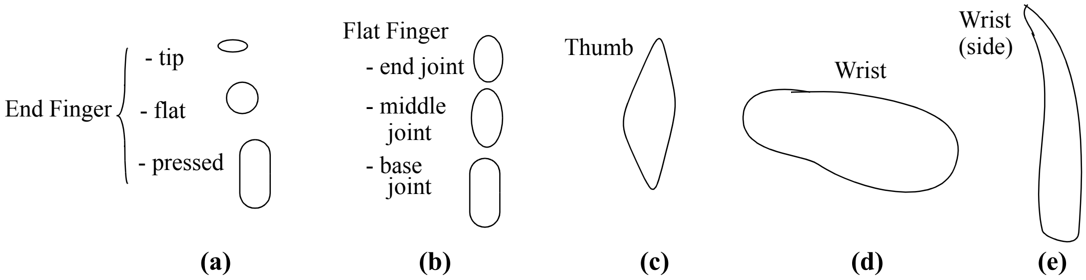

Figure 5. Outlines of Some Relatively Simple Pressure Images

Images created using the palm, the fist and the whole hand have more complex shapes, as illustrated in Figures 6a-c. Figure 6a shows the outline of an image created using a palm, Figure 6b the outline of one created using the fist, and Figure 6c the outline of one using a flat hand. The image of the fist is roughly curved, and the degree of curvature can be varied by the user. As can be seen from the image of the flat hand, contact with two or more digits creates a gap between them (S_1, S_2, S_3, S_4) that can be varied by the user. Such variations can be exploited to vary corresponding parameters of an external device controlled by the touchpad.

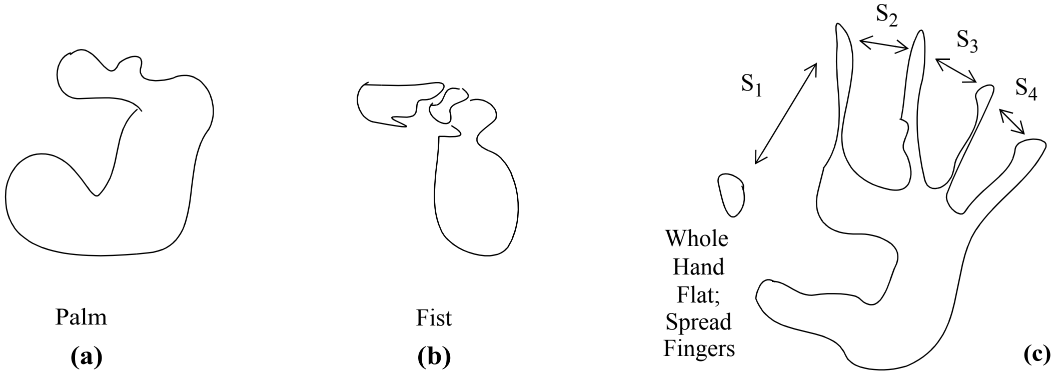

Figure 6. Outlines of Some Relatively Complex Pressure Images

In general, contact with the pressure-sensor array using multiple parts of the hand forfeits some degrees of freedom but introduces others. Consider the case of the whole hand pressed flat against the sensor array. The pressure of the fingers and thumb can be varied, yielding five parameters. The pressure of the ends of the fingers can be varied, yielding four parameters. The pressure of the palm can be varied, yielding one parameter. The tilt of the palm can be varied front-to-back and side-to-side, yielding two parameters. And the curvature of the thumb can be varied, yielding one parameter. Thus contact with the whole hand enables the user to vary 17 independent parameters.

Image Recognition

Contiguous regions of non-zero pressure resulting from contact with the pressure-sensor array will in general have a variety of shapes and sizes. The touchpad's image processing stage can include a capability for comparing these shapes and sizes with ranges of shapes and sizes corresponding to pre-defined categories. Contiguous regions of non-zero pressure can be classified as falling in one of these categories or, if they do not, can be treated as error conditions or subjected to special handling.

Once a contiguous region of non-zero pressure is classified as a member of a particular category, it can be processed in ways appropriate for the category. For instance, data acquisition and data compression can proceed in particular ways, the region can be represented using a data record with a particular structure, and the values of particular parameters can be derived for the region. In this way, data handling, parameter scaling, region interpretation and other functions, as well as user interface metaphors (discussed below), can be tailored to particular types of regions of contact. The classification of a region can also be used as another derived parameter. For example, different parts of the hand can be used to create contexts that determine how other parameters are interpreted. It is also possible to classify particular collections of regions of contact and subject them to special handling based on their classification. This can be done with particular sequences of images as well as with compound images. Compound images are created naturally by contact with the human hand (see, for example, Figures 5 and 6). The use of compound images and sequences of images will be considered in more detail below.

Context-Dependent Parameter Derivation and Interpretation

As just mentioned, it is possible to vary the parameters that are derived from an image and how they are interpreted as a function of the image's classification. This makes it possible to derive more parameters from images that are easier to manipulate (such as ones of the fingertip or finger end joint) and fewer parameters from images that are harder to manipulate (such as ones of the palm or wrist). In addition, the application interface can be used to interpret different types of images in different ways. For example, finger images can be used to control one set of operations, thumb images another and wrist images yet another. Further, the same type of image can be processed in different ways in different contexts. Consider the following example. If an image contains only one finger, then six parameters are derived. But if an image contains more than one finger, then three parameters are derived for each finger. The interpretations assigned to the same parameter can be different in the two cases. Note that in the second case, the three parameters each finger controls would be chosen because they are easy for each finger to manipulate independently of what the other fingers do.

Compound Images

An important special case of context-sensitive processing and interpretation involves the treatment of compound images. As mentioned earlier, compound images are images consisting of multiple, non-overlapping regions of non-zero pressure, though some images consisting of a single contiguous region of non-zero pressure may be treated as consisting of multiple, separate regions. A sample compound image, created by touching the sensor array with the left part of the left hand with the palm raised, is illustrated in Figure 7. This compound image consists of four contiguous regions of non-zero pressure against a background of pixels measuring zero pressure.

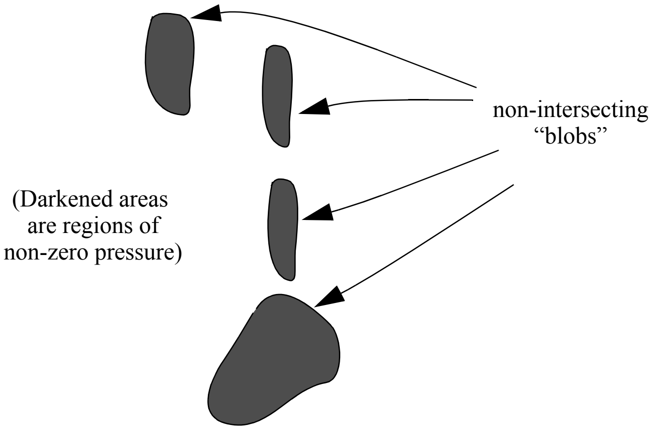

Figure 7. A Sample Compound Image

As discussed earlier, ink prints of the hand show that many other images, such as those of a flat finger, fist and open hand, are compound images. Further, ink prints of the hand also show that it is impossible to merge non-overlapping contiguous regions, provided the surface of the pressure-sensor array is not too elastic. For instance, no matter how tightly two fingers are pressed together and how hard they press on the surface of the pressure-sensor array, the images of the two fingers will not merge. These observations suggest that many compound images can be classified as belonging to a pre-defined category and subjected to processing appropriate for members of that category.

A compound image can be recognized as a member of a pre-defined type of compound image on the basis of several different sorts of criteria. These include the following:

- the presence of two or more images that would otherwise be separately processed whose proximity or relative orientation falls within a specified threshold
- the presence of one or more images that cannot occur individually
- specified combinations of images that would otherwise be processed separately
- combinations of these criteria

One or more of these criteria, as well as others, can be chosen as appropriate for a given application. Once a compound image is classified as an image of a particular type, there are a number of different ways in which it can be processed so that parameters are extracted and interpreted. Component images can be processed separately, or their processing can be combined. Examples of the latter include merging the calculation of particular parameters (such as average pressure) and treating component images as a single, merged image and calculating all parameters for the merged image.

State-Dependent Interpretation of Images

We have seen how the parameters derived and their interpretation can be varied depending on how the image is classified. More generally, the parameters derived and their interpretation can be varied as a function of particular states of the touchpad. The states can be controlled by the application interface or by the external system the touchpad controls. This permits such things as interpreting images in one way when one task is carried out and in a different way when a different task is carried out. For instance, if the touchpad is used to control a CAD/CAE workstation, images can be interpreted in one way when an icon is being selected and positioned, and in a different way when it is being edited, so that the same kinds of hand contact used to select and position the icon in the first mode are used to carry out such functions as sizing, rotating, warping and joining in the editing mode. In another example, the touchpad is used to control an electronic music interface, and the parameters derived and their interpretation vary with the type of sound created.

Sequences of Images

So far, we have considered only parameters derived from a single scan of the pressure-sensor array. It is also possible to derive parameters based on changes across multiple scans. One example is deriving the average velocity of a movement across the surface of the touchpad or a movement of striking the touchpad. In the case of an electronic musical instrument, the velocity of impact can be used to control the sound generated just as it is in modern keyboard and percussion music synthesizers. Velocity measurements can also be tested against thresholds to create logic signals. For example, a quick tilt or sudden downward movement can be used to trigger an event or a change in context.

Interpreting Sequences of Images as Sentences of a Language

In an earlier section it was pointed out that how an image is classified can itself serve as a parameter derived from the image. We can extend this idea to the case of temporal sequences of images. A natural way to do this is to think of sequences of images classified in particular ways as sentences in a language. A simple scheme for doing this is to treat each image as a word of the language. Parts of the image can be treated as letters of an alphabet, but a scheme that is likely to be more workable is to treat individual words as linguistic primitives, like characters in a pictographic language like Chinese.

Because movements of the hand (foot, etc.) can be fluid, it is necessary to demarcate temporal parts of the movement into individual words. On one approach, the end or beginning of a word is explicitly indicated. This can be done, for instance, by increasing the overall pressure beyond a specified threshold or holding a hand posture for a specified length of time. Alternatively, the demarcation between words can be implicit, as it is in speech and American sign language. However, such a scheme would be more difficult to implement, more prone to error and harder to learn. Fluid movements made within temporal constraints can themselves be treated as linguistic primitives, on analogy with words. Such fluid movements will be discussed briefly in the next section, in connection with the notion of a gesture.

Once the words of the language have been fixed, clauses or sentences can be constructed from the words according to syntactic rules. These rules will specify which combinations of words are admissible and which are not -- for instance <*subject, attribute*> might be admissible while <*attribute, subject*> is not, <*subject, verb, object*> might be admissible while <*object, verb, subject*> is not, and so on.

Gestures

Closely related to the idea of treating particular sequences of images as meaningful is the idea of a gesture. A gesture is a fluid movement over the surface of the pressure-sensor array, and can be represented as a temporally-varying set of parameters. The temporal variation of some or all of the parameters can be compared to a pre-specified range. A match results in the recognition of a gesture. Gestures, like static images, can be construed as words, though they can also be construed as sentences. When construed as words, syntactic rules can be used to specify admissible combinations of gestures. It is possible for meaning to be assigned to both gestures and static images. In this case, gestures can be treated as a different kind of linguistic entity from static images; for instance, gestures can be used to represent properties or relations while static images can be used to represent objects or things.

3 User Interface Metaphors

The touchpad naturally lends itself to metaphors useful in a wide range of user interface applications [1]. This will now be illustrated for the case, described earlier, of the tilt of a finger.

Simple Metaphors

The two parameters of tilt -- front-to-back and side-to-side -- can be interpreted in a number of ways. Four possible metaphors are illustrated in Figures 8a-d. In the metaphor depicted in Figure 8a, the two parameters of the finger's tilt determine coordinates on a "tilt" plane. The tilt plane is superimposed on a "location" plane, whose coordinates are determined by the finger's location on the pressure-sensor array. In Figure 8b, the two parameters are interpreted as determining two angles. In Figures 8c-d, one dimension of finger tilt is interpreted as an angle while the other is interpreted as a point on a line.

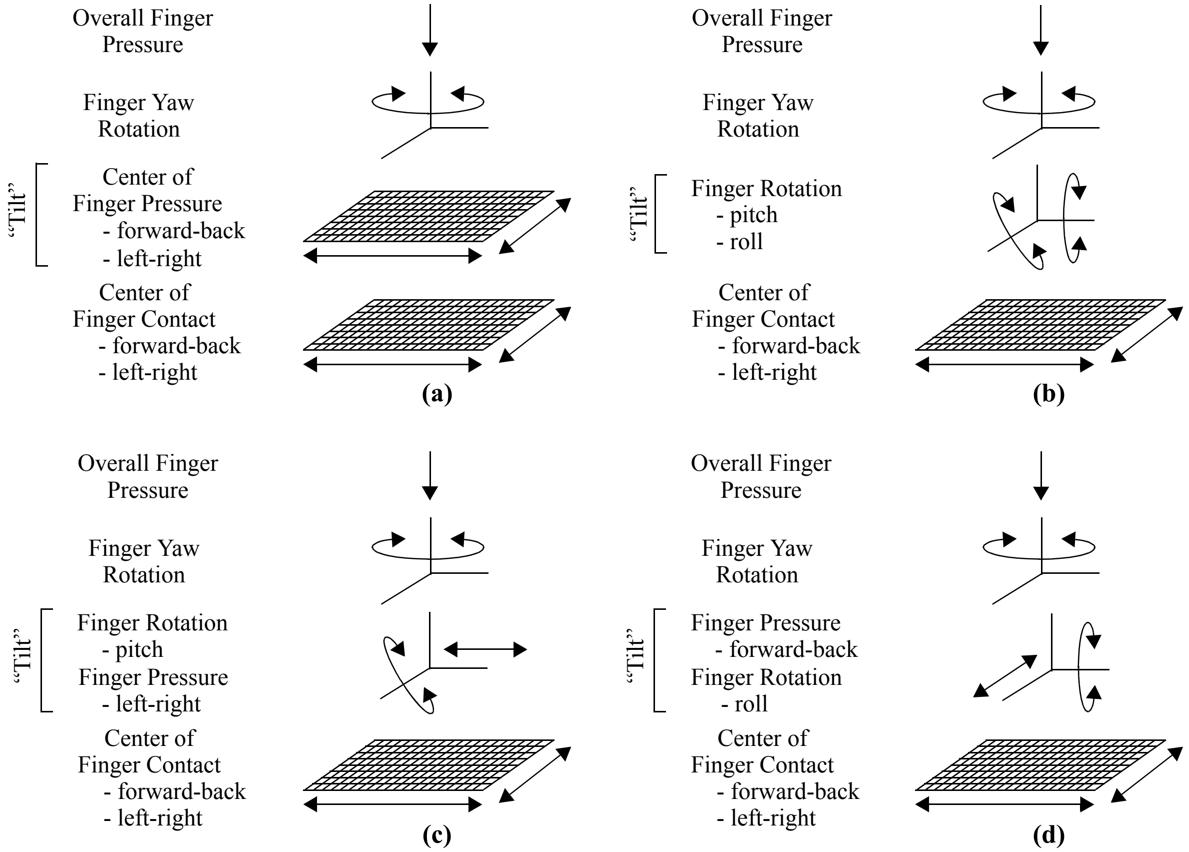

Figure 8. Some Possible Metaphors for the Tilt of a Finger

Compound Metaphors

As mentioned, in the metaphor depicted in Figure 8a, the two parameters of the finger's tilt are thought of as determining a location on a tilt plane. The location on the tilt plane can itself be interpreted in a number of ways. Some examples include:

- as a hierarchical or structured multi-parameter index

- as two separate input specifications which may be morphed or interpolated between according to the value of yet another parameter (such as overall pressure)
- as an input-plane/output-plane distinction for a two-input/two-output transformation
- as two parallel processes, each controlled by a two-dimensional control input akin to a mouse position, where:
  - the processes are separate, or
  - the processes may be caused to converge according to the value of yet another parameter (such as overall pressure) specifying the degree of correlation, angle of intersection, etc. between the two independent processes

Using User Interface Metaphors In a CAD/CAE Application

In a CAD/CAE workstation application, the tilt plane can be used to select icons from a two dimensional array, as shown in Figure 9. The finger's location on the sensor array can be used to specify the icon's destination, while other parameters can be used to select other attributes (size, color, orientation, etc.; these can be items on a menu). One parameter, such as the pressure of a finger, can be used to lock the icon so its attributes cannot be changed without performing an operation to unlock it. Thus a single, quick action of one finger performs functions that require several mouse or keyboard operations in conventional systems. In addition, since only one of the six parameters locks the icon so its attributes cannot be changed, the touchpad provides a novel capability: being able to rapidly compare variations in several different attributes of an object before settling on its final form. As a result, the touchpad eliminates many editing operations conventional systems require when a designer explores options.

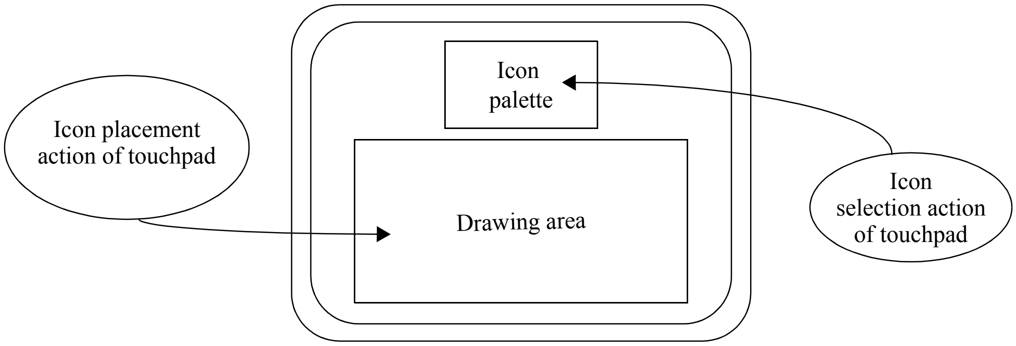

Figure 9. CAD/CAE Workstation Screen (schematic)

Additional operations can be performed through recognizing and extracting parameters from contact by additional fingers, the wrist or other parts of the hand (see Figures 5-6 and the associated discussion). For example, to select a symbol from a two-dimensional palette and position it in a two-dimensional CAD drawing the following assignment can be used:

Symbol selection:

- left-right tilt: left-right movement through the palette
- front-back tilt: up-down movement through the palette

Symbol placement and attribute selection:

- left-right position: left-right position of the symbol
- front-back position: up-down position of the symbol
- rotation: rotation of the symbol in the drawing
- overall pressure: sizing of the symbol by steps
- tap of additional finger: toggle between locking the symbol in the drawing to prevent changes or unlocking it for changes
- tap of thumb: undo changes to the symbol
- wrist pressure tap: toggle between adding a new object and selecting a previously placed object

This is just one possible assignment of hand movements to CAD operations; many variations are possible.

4 Extending the Touchpad's Capabilities

Though the discussion so far provides an idea of the basic capabilities of the touchpad, it only hints at the full range. Some extensions of the basic capabilities will be discussed below. Further extensions are discussed in [1-3].

Partitioned Operation

So far, we have considered the touchpad only in its "unpartitioned" modes -- i.e. the entire pressure-sensor array is treated as a single, functionally uniform region that responds to contact in the same way regardless of the part of the array that is touched. However, the touchpad can also be configured, under software control, to operate in "partitioned" modes -- i.e. the surface of the touchpad is divided into multiple, functionally-distinct regions. Each partition operates independently of the others, and different partitions can respond to the same kind of contact in different ways. There are many possible ways the touchpad can be partitioned, and which is most suitable is determined by such considerations as the application and user preferences. A simple example of partitioned operation involves configuring the surface of the pressure-sensor array to emulate a control panel. Configured in this way, the touchpad can allow the user to perform simultaneously both discrete operations (e.g. via "buttons" and "selectors") and continuous operations (e.g. via "sliders"). The partition boundaries and how the partitions operate can be stored in a file, and different files can be associated with different users, applications or tasks.

Figures 10a-b illustrate two possible partitions of a pressure-sensor array. Figure 10a shows a partition that is symmetric with respect to both the array's left-right axis and its front-back axis. The surface can also be partitioned so it is asymmetric with respect to either or both axes. Figure 10b illustrates the second case. It is also possible to operate in a partitioned mode and an unpartitioned mode concurrently. For example, images of fingertips can be processed using a partitioned mode while images of other parts of the hand are processed using an unpartitioned mode. As another example, the rotation angle and tilt of the user's fingers can be assigned to different parameters in different partitions while the pressure of the fingers is assigned to the same parameter in all partitions.

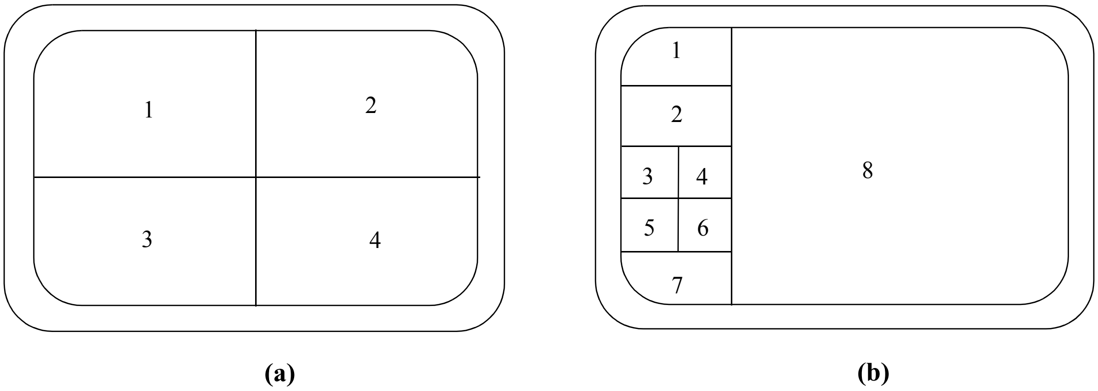

Figure 10. Examples of Partitions of a Pressure-Sensor Array

Visual Displays

The touchpad can incorporate a visual display. One important use for a display is to indicate the borders of partitions when the touchpad is operated in a partitioned mode. The display can mark the partitions with lines showing their boundaries, or with color, text or other visual effects. For more information about visual displays, see [2].

Configurable Operation

The touchpad provides great flexibility in the kinds of images that can be recognized, the parameters for which values are derived, how the parameters are assigned to control signals and other features of its operation. These can be fixed, and it can be advantageous to do so if the functionality of an application is limited or well defined. However, the operation of the touchpad can be configurable, so that a user, administrator or manufacturer can tailor the touchpad's operation to meet particular requirements. This can be done with a graphical user interface. Alternatively, or in conjunction with a graphical user interface, the user can train the touchpad by example to interpret specific patterns of hand contact in specific ways. The resulting configuration can be stored in a file, and the touchpad can load a specific file depending on the user, the application, or the mode or task within an application.

Training for Individual Users

The operation of the touchpad can be improved, and errors in its interpretation of images reduced, if it is trained to recognize the patterns of contact of individual users. This is akin to training a speech-recognition system to understand the speech of a particular user. A natural approach is to have the user make a specified sequence of hand postures and to calibrate the system to them. This procedure can also be used to train the user, for instance, by indicating whether postures are ill-formed, or whether the contact with the pressure-sensor array is too light or too hard.

Foot Operation

Although virtually all discussion so far has assumed the touchpad is operated with the user's hand, it is also possible to operate the touchpad with the foot. Images of the foot can be created using a bare foot or a foot wearing a sock or a shoe; using a bare foot or a foot wearing a sock allows more variation in the images produced, and so finer control, than using a foot wearing a shoe.

Figures 11a-d shows images created using a foot. Figures 11a-b show images created using a foot wearing a shoe, with Figure 11a showing an image of a heel and Figure 11b showing an image of a toe-box. As can be seen by comparing Figures 11a and 11b, images of the heel and the toe-box have significant differences and can easily be distinguished. Rocking the foot back and forth will produce a sequence of images alternating between ones of the heel and ones of the toe-box. For each kind of image, the following parameters can be independently controlled:

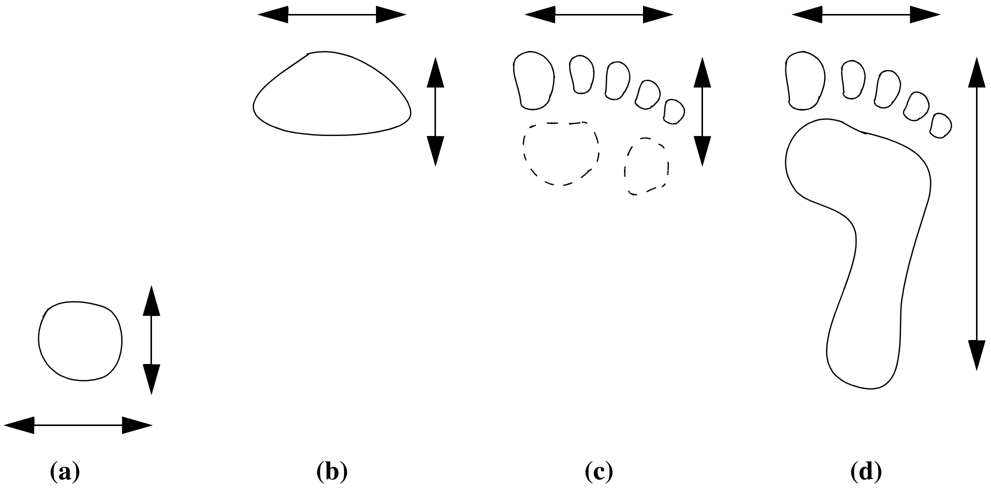

Figure 11. Outlines of Images Created with a Foot

- the left-right location of the geometric center of the image
- the front-back location of the geometric center
- the left-right location of the center of pressure, varied by tilting the foot left-to-right
- the front-back location of the center of pressure, varied by tilting the foot forward and back
- the total pressure

Most forms of contact between the pressure-sensor array and the toe-box will produce oblong images, as in Figure 11b. The angle of rotation of these images can be varied, thus providing control over an additional parameter. Similarly, contact with the entire foot, resulting in images that combine those shown in Figures 11a-b, also produces images whose angle of rotation can be measured and varied. For some types of shoes, it is also possible to measure an angle of rotation for the heel.

Figures 11c-d show images created by the front of a foot; images created by the heel resemble that shown in Figure 11a of the heel of a shoe. As the foot is rocked forward, the two regions with dotted outlines become increasingly small and eventually vanish. As the foot is rocked backward, those regions reappear and grow larger. As in the case of a foot wearing a shoe, rocking a bare foot forward and back will result in a sequence of images alternating between ones of the front of the foot and ones of the heel. Under significant pressure, the images of the front and the heel merge, as shown in Figure 11d. For all the images -- images of the front of the foot and the back of the foot pressed with moderate force, and images of the entire foot pressed with significant force -- it is possible to independently control the left-right and front-back locations of the geometric center of the image, the left-right and front-back locations of the center of pressure, and the total pressure. In addition, in the case of the front of the foot pressed with moderate force and the entire foot pressed with significant force, it will be possible to measure and vary an angle of rotation.

More information about using the feet to operate floor-controllers for music, lighting and industrial applications is provided in [4]

Modular Implementations

The software/firmware techniques used to partition a pressure-sensor array into functionally distinct regions (discussed above) can be adapted to link multiple pressure-sensor arrays so they behave collectively as one large array. The techniques can also be adapted to make it possible for a pressure-sensor array to be divided into several smaller, functional pressure-sensor arrays.

The techniques used to link multiple pressure-sensor arrays to form one large one can be exploited to implement a small pressure-sensor array as a chip. Such chips can be laid as tiles in an array. The chips can also incorporate data acquisition and image processing functions. Another possible implementation uses "flexible electronics" fabrication techniques to implement a pressure-sensor array together with its supporting electronics as a fabric. The fabric can be designed so that cutting it at modular boundaries allows unimpeded operation of the resulting pieces of fabric, and pieces of fabric can be joined to create larger, functional pieces. Some possible applications of chip and fabric implementations of the touchpad are described below. For more information about modular implementations of the touchpad, see [1,2].

5 Further Examples of Applications

The touchpad has many possible applications, some of which have already been described. This section describes a few more. Depending on the scale of production, the features supported and whether custom integrated circuits such as gate arrays are used, it may be possible to considerably reduce the cost of the pressure-sensor array, electronics and packaging.

CAD/CAE and Computer Graphics

As some of the examples given above suggest, one possible use of the touchpad is as an interface for CAD/CAE applications in particular and computer graphics applications in general. Because conventional computer pointing devices provide control of only two continuous parameters at once while the touchpad provides control of many more, it provides a significantly more powerful and efficient way to control such applications, and it does so in an ergonomically felicitous way.

Real-Time Machine Control and Robotics Control

The capability the touchpad provides to easily and rapidly adjust large numbers of parameters makes it attractive for many real-time machine control applications. The touchpad is particularly well suited for applications that involve controlling the orientation, position and movement of physical objects. One example of this is controlling a robot arm, such as that shown in Figure 12. A robot arm like this might be used in manufacturing or handling hazardous materials.

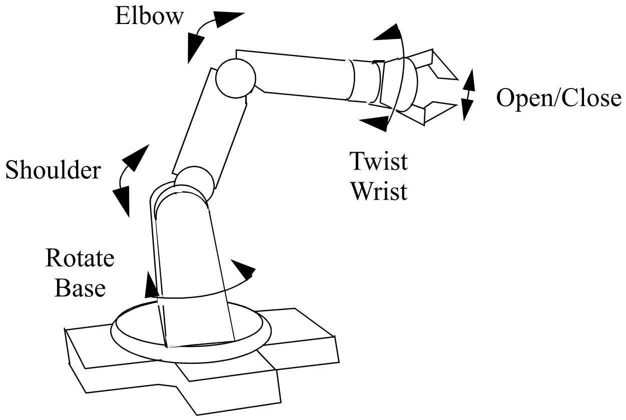

Figure 12. Primitive Robot Arm Movements

Here we consider how the touchpad can be used to control the robot arm, though the basic idea can be applied in many other contexts where similar mechanisms are manipulated, as in positioning a remote camera, performing remote surgery or manipulating devices like the robot arm that are part of a truck, ship or space vehicle. We assume that the robot arm permits five kinds of bidirectional movements, as shown in Figure 12:

- rotating the arm with respect to its base
- pivoting at the shoulder joint
- pivoting at the elbow joint
- twisting the wrist
- opening and closing the pincers

To illustrate how the touchpad can be used to control the robot arm, we will consider a simple example. Movements of a hand can be assigned to movements of the robot arm as follows:

- left-right finger translation: left-right translation of the pincers

- forward-back finger translation: forward-back translation of the pincers
- decreased-increased finger pressure: up-down translation of the pincers
- left-right finger rotation: wrist rotation
- thumb tap: close pincers
- finger tap: open pincers

This assignment is illustrated in Figure 13. The six components of the assignment can specify an absolute or relative value of the specified parameter. Note that there is a reference point shown in the figure, located between the pincers. Because the pincers can move in three spatial dimensions, the reference point has three coordinates, (x, y, z), where x is the left-right position of the pincers, y is the forward-back position and z is the vertical position. These coordinates correspond to, respectively, the left-right position of the finger on the touchpad, the forward-back position and the amount of downward pressure. Because the location of the reference point is affected by rotating the arm with respect to its base, pivoting at the shoulder joint and pivoting at the elbow joint (but not by twisting the wrist or opening and closing the pincers), this correspondence is maintained by means of computations, performed by the touchpad's application interface, that convert changes in the location of the reference point to movements of the base, shoulder joint and elbow joint.

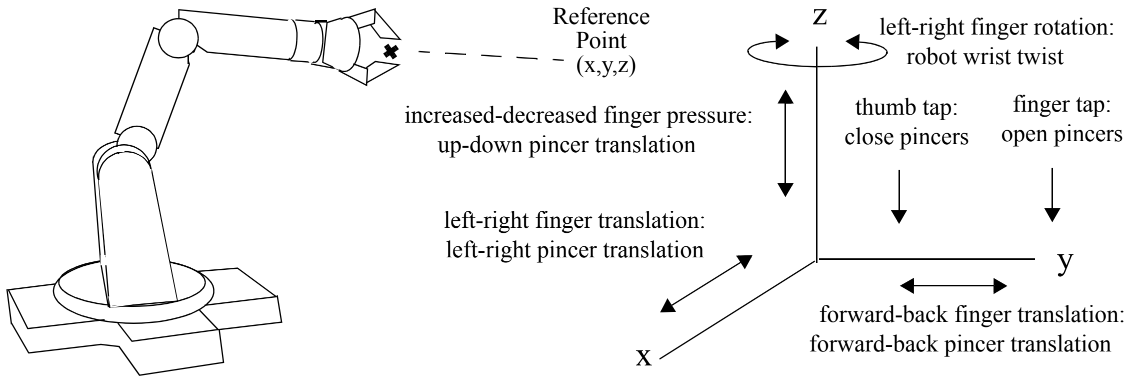

Figure 13. Robot Arm Control Using Fingers

Many other assignments of hand movements to robot arm movements are possible. For instance, in addition to the assignments given above, the left-right finger tilt can be used to control the amount of force exerted by the pincers. Or, as an alternative to the above assignment, the rotation of the arm with respect to its base, the pivoting of the shoulder and the pivoting of the elbow can be directly controlled by, respectively, the first three movements in the above assignment.

Assistive Technology for the Physically Disabled

Because the touchpad is well suited for real-time machine control, can be used to control many parameters simultaneously, and is easily customized and reconfigured, it has tremendous potential as a human-machine interface for the physically disabled. As such, it can have a positive impact on many lives. Since physical disabilities take many forms -- a characteristic that tends to limit the market and utility of most human-machine interfaces for the physically disabled -- the ease with which the touchpad can be customized and reconfigured makes it possible to create a single product that can accommodate a wide range of individual needs. For instance, in cases in which a person has lost or has limited use of a limb, the touchpad can be configured so it can be operated with limbs that the person can use, the nose, the lips, an object held in the mouth and so on. Further, the particular types of images and particular combinations of images that are recognized, and how they are interpreted, can be tailored to the abilities and needs of the individual user. In cases in which disabilities are particularly severe, a single touchpad can serve as a human-machine interface for a wide range of other systems, such as computers, appliances, locomotion systems and communication devices.

Electronic Musical Instruments

The original inspiration for the touchpad was for applications in electronic music. The idea of the touchpad came from study of the Indian table and baya drums. The drums provide a remarkable degree of expressive richness because there are many different modes of vibration that can be induced and modulated by simple forms of finger, palm and wrist contact. Further, because of the nature of the contact used to play them, the sounds they make can be varied precisely and rapidly. Thus, in the case of these drums, there is a collection of simple hand movements that can be rapidly made that provide a high degree of control over the sounds produced.

One possible application of the touchpad involves extending and generalizing the idea of the drums so that more forms of hand contact can be used, and they can be used to produce a wider range of sounds. This can be done, for instance, by controlling sampled or electronically synthesized sounds. And the touchpad can be used in many other ways besides as a percussion instrument. One example is as a device for “timbral finger painting,” in which variations in hand contact are used to produce corresponding variations along many different dimensions of the timbre of the sound created. A touchpad like this can be implemented as a stand-alone device, but can also be affixed to instrument keys, instrument bodies, mallet grips and so on. A variety of metaphors can be used, ranging from ones that are very natural (such as associating a finger’s left-right position, left-right tilt or left-right rotation with left-right stereo panning), to ones that are more abstract (such as associating increases in finger pressure with increases in volume or spectral complexity), to ones that are highly abstract (such as associating a finger’s left-right position with the position along a dimension of timbral variation).

Many types of self-contained products with internal sound generation mechanisms can be created. A particularly advantageous way to implement them is as devices that generate signals in MIDI (Musical Instrument Digital Interface) format, a format, used by hundreds of manufacturers, for real-time communication of control signals between electronic music devices. Products supporting MIDI input include almost all keyboard synthesizers, percussion synthesizers, digital recording playback (“sampling”) equipment and the most advanced electronic music synthesis equipment available. Thus a single touchpad product that provides MIDI output can be used to control almost all commercially-available electronic music devices and would therefore have the largest possible market. A comparatively small development effort can subsequently be used to create specialized implementations optimized for particular uses. For more details about music applications of the touchpad, see [1,3].

Lighting Control for the Performing Arts

The capability provided by the touchpad to control large numbers of parameters in real time also makes it attractive for lighting control in performances, particularly where multi-channel or motor-controlled lighting is used. The touchpad naturally lends itself to geometric metaphors, and the surface of the touchpad can be partitioned into separate cells, each associated with a particular light, which can be used to control pan, tilt, brightness, position, zoom-level, gel, gel-pattern orientation and so on. In addition, the touchpad opens the door to many new real-time lighting control and lighting design possibilities. For instance, it can be used to create new types of effects by enabling the operator to continuously vary the location, size and color of lighting regions created by over-stage lighting arrays, and to gradually merge in blending effects or modulations. Such a capability has a reasonable chance of having a notable impact on the design of lighting systems and on the stage lighting industry.

It is also possible to use the touchpad to control both musical processes and lighting in a performance. In such applications, some partitions or hand movements can be used to control either musical processes or lighting exclusively, while others can be used for simultaneous control of both so that the musical processes and lighting vary in corresponding ways. For more information about lighting applications, see [3].

Automated Assembly and Handling

The touchpad can be used for machine sensing of inanimate objects. As a “workpad,” it can be used to sense the location and orientation of objects placed on it. This would be useful for automated assembly and handling applications. A system like this can be quite inexpensive and available long before comparably priced computer-vision systems that perform the same functions. As another example, small pressure-sensor arrays can be put on the fingers of a robot’s hand to improve its dexterity, since information about the pressure distribution on the fingers can enhance the sensitivity of the hand and improve its ability to handle delicate or awkward objects

Generalized Controller Product

Because the same basic features of the touchpad are useful for a wide variety of applications, it is also possible to create a general-purpose controller product. A product like this can incorporate a simple (perhaps menu-driven) configuration mechanism so consumers are free to develop their own applications. This would provide an opportunity to identify applications that would otherwise be overlooked or unforeseen. Once such applications have been identified, those with the greatest potential for high-volume markets can be addressed by the creation of specialized products optimized for those particular applications.

Chip Applications

As mentioned earlier, it is possible to implement the touchpad as a dedicated chip, which can be laid as a tile in an array of such chips. Such chips have a number of potential uses. One possible use is for an application described earlier, in which small pressure-sensor arrays are put on the fingers of a robot’s hand to improve its sensitivity and dexterity. There are also many possible uses in music applications. For instance, chips can be affixed to keyboard keys, as shown in Figure 14, and to the handles of mallets. For more information about music applications, see [3].

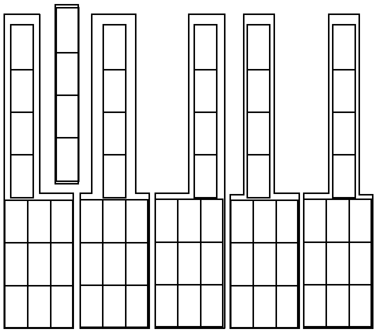

Figure 14. Arrays of Touchpad Chips on Instrument Keys

Fabric Applications

As mentioned earlier, it is possible to implement the touchpad as a flexible fabric. There are many possible applications for such a fabric, particularly if it can be stretched taut. These applications include:

- surface-sensing chairs and walls
- surface-sensing drum heads
- pressure-sensing rugs
- pressure-sensing clothing
- fitness-workout floor mats

In a fabric implementation, the touchpad could also be installed inside or on the surface of a glove or a sock. A glove or sock touchpad can be used in the performing arts, and a sock touchpad (or a touchpad installed on the bottom of a shoe) can be used for gait analysis of humans or animals.

6 Further Information

This document has described the user-level operation and some potential applications of a novel tactile user interface, a touchpad that incorporates a pressure-sensor array and exploits image processing techniques to extract the values of a large number of parameters from the resulting images. This document is based on U.S. Patent 6,570,078 [1] and related issued and pending patents [4-7], all licensable from NRI®. Detailed hardware and software reference designs can be discussed under negotiable terms. All financial or in-kind proceeds from such arrangements will be used to fund pure academic research at NRI®. Contact *inquiries@newrenaissanceinstitute.com* for more information.

References

[1] NRI®, U.S. Patent 6,570,078, “Tactile, Visual, and Array Controllers for Real-Time Control of Music Signal Processing, Mixing, Video, and Lighting,” May 27, 2003 (PDF available at www.newrenaissanceinstitute.com/Issued.PatentsPDF/Tactile_Visual_Array.pdf).

[2] NRI® Whitepaper, “The NRI® Rich Touchpad: Technology and Implementations,” 2004 (PDF available at www.newrenaissanceinstitute.com/WP_list.htm).

[3] NRI® Whitepaper, “The NRI® Rich Touchpad: Music and Performance Applications,” 2004 (PDF available at www.newrenaissanceinstitute.com/WP_list.htm).

[4] NRI®, U.S. Patent 6,689,947, “Real-Time Floor Controller for Control of Music, Signal Processing, Mixing, Video, Lighting, and Other Systems,” February 10, 2004 (PDF available at www.newrenaissanceinstitute.com/PreGrant-PDF/FloorController.pdf).

[5] NRI®, U.S. Pre-Grant Patent Publication 2004/0099131, “Transcending Extensions of Classical South Asian Musical Instruments,” May 27, 2004. (PDF available at www.newrenaissanceinstitute.com/PreGrant-PDF/SouthAsian.pdf).

[6] NRI®, U.S. Pre-Grant Patent Publication 2004/0074379, “Functional Extensions of Traditional Music Keyboards,” April 22, 2004. (PDF available at www.newrenaissanceinstitute.com/PreGrant-PDF/Keyboards.pdf).

[7] NRI®, U.S. Pre-Grant Patent Publication 2004/0069125, “Performance Environments Supporting Interactions Among Performers and Self-Organizing Processes,” April 15, 2004 (PDF available at www.newrenaissanceinstitute.com/PreGrant-PDF/SelfOrganized.pdf).

[8] Westerman, W. and Elias, J., U.S. Patent 6,570,557, “Multitouch System and Method for Emulating Modifier Keys via Fingertip Chords,” November 27, 2001.

[9] Westerman, W. and Elias, J., U.S. Patent 6,323,846, “Method and Apparatus for Integrating Manual Input,” November 27, 2001.