

MARK-CODED SURFACE WITH FUNCTION FLAGS

FIELD OF INVENTION

The present invention relates generally to methods, systems and devices for interacting with computer systems, and in particular, to coded surfaces supporting such interaction.

CO-PENDING APPLICATIONS

Various methods, systems and apparatus relating to the present invention are disclosed in the following co-pending applications filed by the applicant or assignee of the present invention simultaneously with the present invention:

PCT/AU00/01442, PCT/AU00/01444, PCT/AU00/01446, PCT/AU00/01445, PCT/AU00/01450, PCT/AU00/01453, PCT/AU00/01448, PCT/AU00/01447, PCT/AU00/01459, PCT/AU00/01451, PCT/AU00/01454, PCT/AU00/01452, PCT/AU00/01443, PCT/AU00/01455, PCT/AU00/01456, PCT/AU00/01457, PCT/AU00/01458 and PCT/AU00/01449.

The disclosures of these co-pending applications are incorporated herein by cross-reference.

Various methods, systems and apparatus relating to the present invention are disclosed in the following co-pending applications filed by the applicant or assignee of the present invention on 20 October 2000:

PCT/AU00/01273, PCT/AU00/01279, PCT/AU00/01288, PCT/AU00/01282, PCT/AU00/01276, PCT/AU00/01280, PCT/AU00/01274, PCT/AU00/01289, PCT/AU00/01275, PCT/AU00/01277, PCT/AU00/01286, PCT/AU00/01281, PCT/AU00/01278, PCT/AU00/01287, PCT/AU00/01285, PCT/AU00/01284 and PCT/AU00/01283.

The disclosures of these co-pending applications are incorporated herein by cross-reference.

Various methods, systems and apparatus relating to the present invention are disclosed in the following co-pending applications filed by the applicant or assignee of the present invention on 15 September 2000:

PCT/AU00/01108, PCT/AU00/01110 and PCT/AU00/01111.

The disclosures of these co-pending applications are incorporated herein by cross-reference.

Various methods, systems and apparatus relating to the present invention are disclosed in the following co-pending applications filed by the applicant or assignee of the present invention on 30 June 2000:

PCT/AU00/00762, PCT/AU00/00763, PCT/AU00/00761, PCT/AU00/00760, PCT/AU00/00759, PCT/AU00/00758, PCT/AU00/00764, PCT/AU00/00765, PCT/AU00/00766, PCT/AU00/00767, PCT/AU00/00768, PCT/AU00/00773, PCT/AU00/00774, PCT/AU00/00775, PCT/AU00/00776, PCT/AU00/00777, PCT/AU00/00770, PCT/AU00/00769, PCT/AU00/00771, PCT/AU00/00772, PCT/AU00/00754, PCT/AU00/00755, PCT/AU00/00756 and PCT/AU00/00757.

The disclosures of these co-pending applications are incorporated herein by cross-reference.

Various methods, systems and apparatus relating to the present invention are disclosed in the following co-pending applications filed by the applicant or assignee of the present invention on 24 May 2000:

PCT/AU00/00518, PCT/AU00/00519, PCT/AU00/00520, PCT/AU00/00521, PCT/AU00/00522, PCT/AU00/00523, PCT/AU00/00524, PCT/AU00/00525, PCT/AU00/00526, PCT/AU00/00527, PCT/AU00/00528, PCT/AU00/00529, PCT/AU00/00530, PCT/AU00/00531, PCT/AU00/00532, PCT/AU00/00533, PCT/AU00/00534, PCT/AU00/00535, PCT/AU00/00536, PCT/AU00/00537, PCT/AU00/00538, PCT/AU00/00539, PCT/AU00/00540, PCT/AU00/00541, PCT/AU00/00542, PCT/AU00/00543, PCT/AU00/00544, PCT/AU00/00545, PCT/AU00/00547, PCT/AU00/00546, PCT/AU00/00554, PCT/AU00/00556, PCT/AU00/00557, PCT/AU00/00558, PCT/AU00/00559, PCT/AU00/00560, PCT/AU00/00561, PCT/AU00/00562, PCT/AU00/00563, PCT/AU00/00564, PCT/AU00/00565, PCT/AU00/00566, PCT/AU00/00567, PCT/AU00/00568, PCT/AU00/00569, PCT/AU00/00570, PCT/AU00/00571, PCT/AU00/00572, PCT/AU00/00573, PCT/AU00/00574, PCT/AU00/00575, PCT/AU00/00576, PCT/AU00/00577, PCT/AU00/00578, PCT/AU00/00579, PCT/AU00/00581, PCT/AU00/00580, PCT/AU00/00582, PCT/AU00/00587, PCT/AU00/00588, PCT/AU00/00589, PCT/AU00/00583, PCT/AU00/00593, PCT/AU00/00590, PCT/AU00/00591, PCT/AU00/00592, PCT/AU00/00594, PCT/AU00/00595, PCT/AU00/00596, PCT/AU00/00597, PCT/AU00/00598, PCT/AU00/00516, PCT/AU00/00517 and PCT/AU00/00511.

The disclosures of these co-pending applications are incorporated herein by cross-reference.

BACKGROUND

Systems have been described wherein pen- or stylus-like devices sense their own movement relative to a surface, for the purposes of providing input to a computer system, by detecting information which is invisibly but machine-readably encoded on the surface. See for example US Patent 5,477,012, US Patent 5,652,412 and PCT application WO99/50787. Visible information intended for human consumption may also be placed on the surface, although typically in a separate step which therefore may include recording an association between the two kinds of information. The encoded information doesn't include knowledge of the content of the visible information, and any content-specific feedback provided to a user during the user's interaction with the encoded information must be provided on the basis of the recorded association. If the recorded association is remote or unavailable, timely feedback may be impossible to provide.

SUMMARY OF INVENTION

According to a first aspect, the present invention provides a region, of a surface, having coded data, wherein at least some of the coded data is indicative of a plurality of locations of the region and includes first data indicative of at least one functional attribute of a part of the region, wherein the value of the first data is independent of the location of the respective part in the region.

Preferably, the functional attribute indicates at least one of the group comprising: a hyperlink, a hypertext link, a button, a drawing field, a text field and a signature field.

Some of the coded data may include data indicative of a location, data indicative of an identity, or data indicative of an identity and a location.

The identity may be indicative of a location in the region, or of an identity of the region.

According to a second aspect, the present invention provides a method of coding a region including applying coded data to at least a part of the region, at least some of the coded data being indicative of a plurality of locations and including first data which indicates at least one functional attribute of a part of the region, the value of the first data being independent of the location of the respective part in the region.

Preferably, the functional attribute indicates at least one of the group comprising: a hyperlink, a hypertext link, a button, a drawing field, a text field and a signature field.

According to a third aspect, the present invention provides a sensing device for use with a region having coded data thereon, at least some of the coded data being indicative of a plurality of locations of the region and including first data indicative of at least one functional attribute of a part of the region, wherein the value of the first data is independent of the location of the respective part in the region, the sensing device for sensing and reacting to said first data.

Preferably, the sensing device provides at least one indication to a user when it senses data indicative of a functional attribute of a part of the region.

Preferably, the indication is selected from the group comprising a visual, aural, haptic, tactile, vibratory and any other human sensory modality.

Preferably also, when the sensing device senses data indicative of a functional attribute of a part of the region, the sensing device transmits, to a computer system, said sensed data or data based at least partially on said sensed data.

According to a fourth aspect, the present invention provides a system including:
computer system, and

a sensing device for use with a region having coded data thereon, at least some of the coded data indicative of a plurality of locations of the region and including first data indicative of at least one functional attribute of a part of the region, wherein the value of the first data is independent of the location of the respective part in the region,

the sensing device for sensing said first data and for transmitting, to the computer system, said sensed first data or data based at least partially on said sensed first data,

wherein the computer system executes, or commences execution of, a function related to the functional attribute of the sensed first data.

Preferably, the computer system verifies that the functional attribute is consistent with an expected attribute of the respective part of the region.

BRIEF DESCRIPTION OF DRAWINGS

Preferred and other embodiments of the invention will now be described, by way of non-limiting example only, with reference to the accompanying drawings, in which:

Figure 1 is a schematic of a the relationship between a sample printed netpage and its online page description;

Figure 2 is a schematic view of an interaction between a pen with an attached netpage code sensor attachment, a netpage printer, a netpage page server, and a netpage application server;

Figure 3 is a schematic view of a high-level structure of a printed netpage and its online page description;

Figure 4a is a plan view showing a structure of a netpage tag;

Figure 4b is a plan view showing a relationship between a set of the tags shown in Figure 4a and a field of view of a netpage sensing device;

Figure 5a is a plan view showing an alternative structure of a netpage tag;

Figure 5b is a plan view showing a relationship between a set of the tags shown in Figure 5a and a field of view of a netpage sensing device;

Figure 5c is a plan view showing an arrangement of nine of the tags shown in Figure 5a where targets are shared between adjacent tags;

Figure 5d is a plan view showing the interleaving and rotation of the symbols of the four codewords of the tag shown in Figure 5a;

Figure 6a is a schematic view showing a tag diagonal relationship between a nib and a sensing device field of view at zero tilt;

Figure 6b is a schematic view showing a tag diagonal relationship between a nib and a sensing device field of view at maximum tilt;

Figure 6c is a schematic view showing a tag off-diagonal relationship between a nib and a sensing device field of view at zero tilt;

Figure 6d is a schematic view showing a tag off-diagonal relationship between a nib and a sensing device field of view at maximum tilt;

Figure 6e is a schematic view showing a tag parallel relationship between a nib and a sensing device field of view at zero tilt;

Figure 6f is a schematic view showing a tag parallel relationship between a nib and a sensing device field of view at maximum tilt;

Figure 7a is a schematic view showing a set of tags covered by a functional area bitmap; and

Figure 7b is a schematic view showing a set of subregions covered by a functional area bitmap.

DETAILED DESCRIPTION OF PREFERRED AND OTHER EMBODIMENTS

Note: Memjet™ is a trademark of Silverbrook Research Pty Ltd, Australia.

In the preferred embodiment, the invention is configured to work with the netpage networked computer system, a summary of which is given below and a detailed description of which is given in our co-pending applications, including in particular applications PCT/AU00/00565, PCT/AU00/00561, PCT/AU00/00569 and PCT/AU00/00578. It will be appreciated that not every implementation will necessarily embody all or even most of the specific details and extensions described in these applications in relation to the basic system. However, the system is described in its most complete form to assist in understanding the context in which the preferred embodiments and aspects of the present invention operate.

In brief summary, the preferred form of the netpage system employs a computer interface in the form of a mapped surface, that is, a physical surface which contains references to a map of the surface maintained in a computer system. The map references can be queried by an appropriate sensing device. Depending upon the specific implementation, the map references may be encoded visibly or invisibly, and defined in such a way that a local query on the mapped surface yields an unambiguous map reference both within the map and among different maps. The computer system can contain information about features on the mapped surface, and such information can be retrieved based on map references supplied by a sensing device used with the mapped surface. The information thus retrieved can take the form of actions which are initiated by the computer system on behalf of the operator in response to the operator's interaction with the surface features.

In its preferred form, the netpage system relies on the production of, and human interaction with, netpages. These are pages of text, graphics and images printed on ordinary paper or other media, but which work like interactive web pages. Information is encoded on each page using ink which is substantially invisible to the unaided human eye. The ink, however, and thereby the coded data, can be sensed by an optically imaging pen and transmitted to the netpage system.

In the preferred form, active buttons and hyperlinks on each page can be clicked with the pen to request information from the network or to signal preferences to a network server. In one embodiment, text written by hand on a netpage is automatically recognized and converted to computer text in the netpage system, allowing forms to be filled in. In other embodiments, signatures recorded on a netpage are automatically verified, allowing e-commerce transactions to be securely authorized.

As illustrated in Figure 1, a printed netpage 1 can represent a interactive form which can be filled in by the user both physically, on the printed page, and "electronically", via communication between the pen and the netpage system. The example shows a "Request" form containing name and address fields and a submit button. The netpage consists of graphic data 2 printed using visible ink, and coded data 3 printed as a collection of tags 4 using invisible ink. The corresponding page description 5, stored on the netpage network, describes the individual elements of the netpage. In particular it describes the type and spatial extent (zone) of each interactive element (i.e. text field or button in the example), to allow the netpage system to correctly interpret input via the netpage. The submit button 6, for example, has a zone 7 which corresponds to the

spatial extent of the corresponding graphic 8.

As illustrated in Figure 2, the netpage pen 101, a preferred form of which is described in our co-pending application PCT/AU00/00565, works in conjunction with a netpage printer 601, an Internet-connected printing appliance for home, office or mobile use. The pen is wireless and communicates securely with the netpage printer via a short-range radio link 9.

The netpage printer 601, preferred forms of which are described in our co-pending applications PCT/AU00/00561 and PCT/AU00/01281, is able to deliver, periodically or on demand, personalized newspapers, magazines, catalogs, brochures and other publications, all printed at high quality as interactive netpages. Unlike a personal computer, the netpage printer is an appliance which can be, for example, wall-mounted adjacent to an area where the morning news is first consumed, such as in a user's kitchen, near a breakfast table, or near the household's point of departure for the day. It also comes in tabletop, desktop, portable and miniature versions.

Netpages printed at their point of consumption combine the ease-of-use of paper with the timeliness and interactivity of an interactive medium.

As shown in Figure 2, the netpage pen 101 interacts with the coded data on a printed netpage 1 and communicates, via a short-range radio link 9, the interaction to a netpage printer. The printer 601 sends the interaction to the relevant netpage page server 10 for interpretation. In appropriate circumstances, the page server sends a corresponding message to application computer software running on a netpage application server 13. The application server may in turn send a response which is printed on the originating printer.

The netpage system is made considerably more convenient in the preferred embodiment by being used in conjunction with high-speed microelectromechanical system (MEMS) based inkjet (Memjet™) printers, for example as described in our co-pending application PCT/AU00/00578 (docket no. IJ52). In the preferred form of this technology, relatively high-speed and high-quality printing is made more affordable to consumers. In its preferred form, a netpage publication has the physical characteristics of a traditional newsmagazine, such as a set of letter-size glossy pages printed in full color on both sides, bound together for easy navigation and comfortable handling.

The netpage printer exploits the growing availability of broadband Internet access. The netpage printer can also operate with slower connections, but with longer delivery times and lower image quality. The netpage system can also be enabled using existing consumer inkjet and laser printers, although the system will operate more slowly and will therefore be less acceptable from a consumer's point of view. In other embodiments, the netpage system is hosted on a private intranet. In still other embodiments, the netpage system is hosted on a single computer or computer-enabled device, such as a printer.

Netpage publication servers 14 on the netpage network are configured to deliver print-quality publications to netpage printers. Periodical publications are delivered automatically to subscribing netpage printers via pointcasting and multicasting Internet protocols. Personalized publications are filtered and formatted according to individual user profiles.

A netpage printer can be configured to support any number of pens, and a pen can work with any number of netpage printers. In the preferred implementation, each netpage pen has a unique identifier. A household may have a collection of colored netpage pens, one assigned to each member of the family. This allows each user to maintain a distinct profile with respect to a netpage publication server or application server.

A netpage pen can also be registered with a netpage registration server 11 and linked to one or more payment card accounts. This allows e-commerce payments to be securely authorized using the netpage pen. The netpage registration server compares the signature captured by the netpage pen with a previously registered signature, allowing it to authenticate the user's identity to an e-commerce server. Other biometrics can also be used to verify identity. A version of the netpage pen includes fingerprint scanning, verified in a similar way by the netpage registration server.

Although a netpage printer may deliver periodicals such as the morning newspaper without user intervention, it can be configured never to deliver unsolicited junk mail. In its preferred form, it only delivers periodicals from subscribed or otherwise authorized sources. In this respect, the netpage printer is unlike a fax machine or e-mail account which is visible to any junk mailer who knows the telephone number or e-mail address.

Each object model in the system is described using a Unified Modeling Language (UML) class diagram. A class diagram consists of a set of object classes connected by relationships, and two kinds of relationships are of interest here: associations and generalizations. An association represents some kind of relationship between objects, i.e. between instances of classes. A generalization relates actual classes, and can be understood in the following way: if a class is thought of as the set of all objects of that class, and class A is a generalization of class B, then B is simply a subset of A. Each class is drawn as a rectangle labelled with the name of the class. It contains a list of the attributes of the class, separated from the name by a horizontal line, and a list of the operations of the class, separated from the attribute list by a horizontal line. In the class diagrams which follow, however, operations are never modelled. An association is drawn as a line joining two classes, optionally labelled at either end with the multiplicity of the association. The default multiplicity is one. An asterisk (*) indicates a multiplicity of "many", i.e. zero or more. Each association is optionally labelled with its name, and is also optionally labelled at either end with the role of the corresponding class. An open diamond indicates an aggregation association ("is-part-of"), and is drawn at the aggregator end of the association line. A generalization relationship ("is-a") is drawn as a solid line joining two classes, with an arrow (in the form of an open triangle) at the generalization end. When a class diagram is broken up into multiple diagrams, any class which is duplicated is shown with a dashed outline in all but the main diagram which defines it. It is shown with attributes only where it is defined.

Netpages are the foundation on which a netpage network is built. They provide a paper-based user interface to published information and interactive services. A netpage consists of a printed page (or other surface region) invisibly tagged with references to an online description of the page. The online page description is maintained persistently by a netpage page server. The page description describes the visible

layout and content of the page, including text, graphics and images. It also describes the input elements on the page, including buttons, hyperlinks, and input fields. A netpage allows markings made with a netpage pen on its surface to be simultaneously captured and processed by the netpage system.

Multiple netpages can share the same page description. However, to allow input through otherwise identical pages to be distinguished, each netpage is assigned a unique page identifier. This page ID has sufficient precision to distinguish between a very large number of netpages.

Each reference to the page description is encoded in a printed tag. The tag identifies the unique page on which it appears, and thereby indirectly identifies the page description. The tag also identifies its own position on the page. Characteristics of the tags are described in more detail below.

Tags are printed in infrared-absorptive ink on any substrate which is infrared-reflective, such as ordinary paper. Near-infrared wavelengths are invisible to the human eye but are easily sensed by a solid-state image sensor with an appropriate filter.

A tag is sensed by an area image sensor in the netpage pen, and the tag data is transmitted to the netpage system via the nearest netpage printer. The pen is wireless and communicates with the netpage printer via a short-range radio link. Tags are sufficiently small and densely arranged that the pen can reliably image at least one tag even on a single click on the page. It is important that the pen recognize the page ID and position on every interaction with the page, since the interaction is stateless. Tags are error-correctably encoded to make them partially tolerant to surface damage.

The netpage page server maintains a unique page instance for each printed netpage, allowing it to maintain a distinct set of user-supplied values for input fields in the page description for each printed netpage.

The relationship between the page description, the page instance, and the printed netpage is shown in Figure 3. The printed netpage may be part of a printed netpage document 45. The page instance is associated with both the netpage printer which printed it and, if known, the netpage user who requested it.

In a preferred form, each tag identifies the region in which it appears, and the location of that tag within the region. A tag may also contain flags which relate to the region as a whole or to the tag. One or more flag bits may, for example, signal a tag sensing device to provide feedback indicative of a function associated with the immediate area of the tag, without the sensing device having to refer to a description of the region. A netpage pen may, for example, illuminate an "active area" LED when in the zone of a hyperlink.

In a preferred embodiment, each tag contains an easily recognized invariant structure which aids initial detection, and which assists in minimizing the effect of any warp induced by the surface or by the sensing process. The tags preferably tile the entire page, and are sufficiently small and densely arranged that the pen can reliably image at least one tag even on a single click on the page. It is important that the pen recognize the page ID and position on every interaction with the page, since the interaction is stateless.

In a preferred embodiment, the region to which a tag refers coincides with an entire page, and the region ID encoded in the tag is therefore synonymous with the page ID of the page on which the tag appears. In other embodiments, the region to which a tag refers can be an arbitrary subregion of a page or other surface.

For example, it can coincide with the zone of an interactive element, in which case the region ID can directly identify the interactive element.

Each tag contains typically contains 16 bits of tag ID, at least 90 bits of region ID, and a number of flag bits. Assuming a maximum tag density of 64 per square inch, a 16-bit tag ID supports a region size of up to 1024 square inches. Larger regions can be mapped continuously without increasing the tag ID precision simply by using abutting regions and maps. The distinction between a region ID and a tag ID is mostly one of convenience. For most purposes the concatenation of the two can be considered as a globally unique tag ID. Conversely, it may also be convenient to introduce structure into the tag ID, for example to define the x and y coordinates of the tag. A 90-bit region ID allows 2^{90} ($\sim 10^{27}$ or a thousand trillion trillion) different regions to be uniquely identified. Tags may also contain type information, and a region may be tagged with a mixture of tag types. For example, a region may be tagged with one set of tags encoding x coordinates and another set, interleaved with the first, encoding y coordinates.

In one embodiment, 120 bits of tag data are redundantly encoded using a (15, 5) Reed-Solomon code. This yields 360 encoded bits consisting of 6 codewords of 15 4-bit symbols each. The (15, 5) code allows up to 5 symbol errors to be corrected per codeword, i.e. it is tolerant of a symbol error rate of up to 33% per codeword. Each 4-bit symbol is represented in a spatially coherent way in the tag, and the symbols of the six codewords are interleaved spatially within the tag. This ensures that a burst error (an error affecting multiple spatially adjacent bits) damages a minimum number of symbols overall and a minimum number of symbols in any one codeword, thus maximising the likelihood that the burst error can be fully corrected.

Any suitable error-correcting code can be used in place of a (15, 5) Reed-Solomon code, for example a Reed-Solomon code with more or less redundancy, with the same or different symbol and codeword sizes; another block code; or a different kind of code, such as a convolutional code (see, for example, Stephen B. Wicker, *Error Control Systems for Digital Communication and Storage*, Prentice-Hall 1995, the contents of which are herein incorporated by cross-reference).

One embodiment of the physical representation of the tag, shown in Figure 4a and described in our co-pending application PCT/AU00/00569, includes fixed target structures 15, 16, 17 and variable data areas 18. The fixed target structures allow a sensing device such as the netpage pen to detect the tag and infer its three-dimensional orientation relative to the sensor. The data areas contain representations of the individual bits of the encoded tag data. To maximise its size, each data bit is represented by a radial wedge in the form of an area bounded by two radial lines and two concentric circular arcs. Each wedge has a minimum dimension of 8 dots at 1600 dpi and is designed so that its base (its inner arc), is at least equal to this minimum dimension. The height of the wedge in the radial direction is always equal to the minimum dimension. Each 4-bit data symbol is represented by an array of 2x2 wedges. The fifteen 4-bit data symbols of each of the six codewords are allocated to the four concentric symbol rings 18a to 18d in interleaved fashion. Symbols are allocated alternately in circular progression around the tag. The interleaving is designed to maximise the average spatial distance between any two symbols of the same codeword.

In order to support "single-click" interaction with a tagged region via a sensing device, the sensing device

must be able to see at least one entire tag in its field of view no matter where in the region or at what orientation it is positioned. The required diameter of the field of view of the sensing device is therefore a function of the size and spacing of the tags. Assuming a circular tag shape, the minimum diameter of the sensor field of view 193 is obtained when the tags are tiled on an equilateral triangular grid, as shown in Figure 4b.

The tag image processing and decoding performed by a sensing device such as the netpage pen is shown in Figure 6. While a captured image is being acquired from the image sensor, the dynamic range of the image is determined (at 20). The center of the range is then chosen as the binary threshold for the image 21. The image is then thresholded and segmented into connected pixel regions (i.e. shapes 23) (at 22). Shapes which are too small to represent tag target structures are discarded. The size and centroid of each shape is also computed.

Binary shape moments 25 are then computed (at 24) for each shape, and these provide the basis for subsequently locating target structures. Central shape moments are by their nature invariant of position, and can be easily made invariant of scale, aspect ratio and rotation.

The ring target structure 15 is the first to be located (at 26). A ring has the advantage of being very well behaved when perspective-distorted. Matching proceeds by aspect-normalizing and rotation-normalizing each shape's moments. Once its second-order moments are normalized the ring is easy to recognize even if the perspective distortion was significant. The ring's original aspect and rotation 27 together provide a useful approximation of the perspective transform.

The axis target structure 16 is the next to be located (at 28). Matching proceeds by applying the ring's normalizations to each shape's moments, and rotation-normalizing the resulting moments. Once its second-order moments are normalized the axis target is easily recognized. Note that one third order moment is required to disambiguate the two possible orientations of the axis. The shape is deliberately skewed to one side to make this possible. Note also that it is only possible to rotation-normalize the axis target after it has had the ring's normalizations applied, since the perspective distortion can hide the axis target's axis. The axis target's original rotation provides a useful approximation of the tag's rotation due to pen yaw 29.

The four perspective target structures 17 are the last to be located (at 30). Good estimates of their positions are computed based on their known spatial relationships to the ring and axis targets, the aspect and rotation of the ring, and the rotation of the axis. Matching proceeds by applying the ring's normalizations to each shape's moments. Once their second-order moments are normalized the circular perspective targets are easy to recognize, and the target closest to each estimated position is taken as a match. The original centroids of the four perspective targets are then taken to be the perspective-distorted corners 31 of a square of known size in tag space, and an eight-degree-of-freedom perspective transform 33 is inferred (at 32) based on solving the well-understood equations relating the four tag-space and image-space point pairs (see Heckbert, P., Fundamentals of Texture Mapping and Image Warping, Masters Thesis, Dept. of EECS, U. of California at Berkeley, Technical Report No. UCB/CSD 89/516, June 1989, the contents of which are herein incorporated by cross-reference).

The inferred tag-space to image-space perspective transform is used to project (at 36) each known data bit position in tag space into image space where the real-valued position is used to bilinearly interpolate (at 36) the four relevant adjacent pixels in the input image. The previously computed image threshold 21 is used to threshold the result to produce the final bit value 37.

Once all 360 data bits 37 have been obtained in this way, each of the six 60-bit Reed-Solomon codewords is decoded (at 38) to yield 20 decoded bits 39, or 120 decoded bits in total. Note that the codeword symbols are sampled in codeword order, so that codewords are implicitly de-interleaved during the sampling process.

The ring target 15 is only sought in a subarea of the image whose relationship to the image guarantees that the ring, if found, is part of a complete tag. If a complete tag is not found and successfully decoded, then no pen position is recorded for the current frame. Given adequate processing power and ideally a non-minimal field of view 193, an alternative strategy involves seeking another tag in the current image.

The obtained tag data indicates the identity of the region containing the tag and the position of the tag within the region. An accurate position 35 of the pen nib in the region, as well as the overall orientation 35 of the pen, is then inferred (at 34) from the perspective transform 33 observed on the tag and the known spatial relationship between the pen's physical axis and the pen's optical axis.

The tag structure just described is designed to allow both regular tilings of planar surfaces and irregular tilings of non-planar surfaces. Regular tilings are not, in general, possible on non-planar surfaces. In the more usual case of planar surfaces where regular tilings of tags are possible, i.e. surfaces such as sheets of paper and the like, more efficient tag structures can be used which exploit the regular nature of the tiling.

An alternative tag structure more suited to a regular tiling is shown in Figure 5a. The tag 4 is square and has four perspective targets 17. It is similar in structure to tags described by Bennett et al. in US Patent 5,051,746. The tag represents sixty 4-bit Reed-Solomon symbols 47, for a total of 240 bits. The tag represents each one bit as a dot 48, and each zero bit by the absence of the corresponding dot. The perspective targets are designed to be shared between adjacent tags, as shown in Figures 5b and 5c. Figure 5b shows a square tiling of 16 tags and the corresponding minimum field of view 193, which must span the diagonals of two tags. Figure 5c shows a square tiling of nine tags, containing all one bits for illustration purposes.

Using a (15, 7) Reed-Solomon code, 112 bits of tag data are redundantly encoded to produce 240 encoded bits. The four codewords are interleaved spatially within the tag to maximize resilience to burst errors. Assuming a 16-bit tag ID as before, this allows a region ID of up to 92 bits. The data-bearing dots 48 of the tag are designed to not overlap their neighbors, so that groups of tags cannot produce structures which resemble targets. This also saves ink. The perspective targets therefore allow detection of the tag, so further targets are not required. Tag image processing proceeds as described above, with the exception that steps 26 and 28 are omitted.

Although the tag may contain an orientation feature to allow disambiguation of the four possible orientations of the tag relative to the sensor, it is also possible to embed orientation data in the tag data. For example, the four codewords can be arranged so that each tag orientation contains one codeword placed at that

orientation, as shown in Figure 5d, where each symbol is labelled with the number of its codeword (1-4) and the position of the symbol within the codeword (A-O). Tag decoding then consists of decoding one codeword at each orientation. Each codeword can either contain a single bit indicating whether it is the first codeword, or two bits indicating which codeword it is. The latter approach has the advantage that if, say, the data content of only one codeword is required, then at most two codewords need to be decoded to obtain the desired data. This may be the case if the region ID is not expected to change within a stroke and is thus only decoded at the start of a stroke. Within a stroke only the codeword containing the tag ID is then desired. Furthermore, since the rotation of the sensing device changes slowly and predictably within a stroke, only one codeword typically needs to be decoded per frame.

It is possible to dispense with perspective targets altogether and instead rely on the data representation being self-registering. In this case each bit value (or multi-bit value) is typically represented by an explicit glyph, i.e. no bit value is represented by the absence of a glyph. This ensures that the data grid is well-populated, and thus allows the grid to be reliably identified and its perspective distortion detected and subsequently corrected during data sampling. To allow tag boundaries to be detected, each tag data must contain a marker pattern, and these must be redundantly encoded to allow reliable detection. The overhead of such marker patterns is similar to the overhead of explicit perspective targets. One such scheme uses dots positioned at various points relative to grid vertices to represent different glyphs and hence different multi-bit values (see Anoto Technology Description, Anoto April 2000).

Decoding a tag results in a region ID, a tag ID, and a tag-relative pen transform. Before the tag ID and the tag-relative pen location can be translated into an absolute location within the tagged region, the location of the tag within the region must be known. This is given by a tag map, a function which maps each tag ID in a tagged region to a corresponding location. A tag map reflects the scheme used to tile the surface region with tags, and this can vary according to surface type. When multiple tagged regions share the same tiling scheme and the same tag numbering scheme, they can also share the same tag map. The tag map for a region must be retrievable via the region ID. Thus, given a region ID, a tag ID and a pen transform, the tag map can be retrieved, the tag ID can be translated into an absolute tag location within the region, and the tag-relative pen location can be added to the tag location to yield an absolute pen location within the region.

The tag ID may have a structure which assists translation through the tag map. It may, for example, encoded cartesian coordinates or polar coordinates, depending on the surface type on which it appears. The tag ID structure is dictated by and known to the tag map, and tag IDs associated with different tag maps may therefore have different structures.

Two distinct surface coding schemes are of interest, both of which use the tag structure described earlier in this section. The preferred coding scheme uses "location-indicating" tags as already discussed. An alternative coding scheme uses "object-indicating" (or "function-indicating") tags.

A location-indicating tag contains a tag ID which, when translated through the tag map associated with the tagged region, yields a unique tag location within the region. The tag-relative location of the pen is added to this tag location to yield the location of the pen within the region. This in turn is used to determine

the location of the pen relative to a user interface element in the page description associated with the region. Not only is the user interface element itself identified, but a location relative to the user interface element is identified. Location-indicating tags therefore trivially support the capture of an absolute pen path in the zone of a particular user interface element.

An object-indicating (or function-indicating) tag contains a tag ID which directly identifies a user interface element in the page description associated with the region (or equivalently, a function). All the tags in the zone of the user interface element identify the user interface element, making them all identical and therefore indistinguishable. Object-indicating tags do not, therefore, support the capture of an absolute pen path. They do, however, support the capture of a relative pen path. So long as the position sampling frequency exceeds twice the encountered tag frequency, the displacement from one sampled pen position to the next within a stroke can be unambiguously determined. As an alternative, the netpage pen 101 can contain a pair or motion-sensing accelerometers, as described in our co-pending application PCT/AU00/00565.

An embodiment of the present invention, in the form of tags which encode functional attributes, will now be described.

A tag may contain flags which relate to the region as a whole or to the tag.

One or more flag bits may, for example, signal a tag sensing device to provide feedback indicative of a function associated with the immediate area of the tag, without the sensing device having to refer to a description of the region.

A tag may contain one or more "active area" bits, indicating that a particular subregion in the vicinity of the tag is in the zone of a hyperlink (such as the zone 7 of the "Submit" button in Figure 1) or other active page element. The sensing device can provide feedback to the user, e.g. by illuminating an "active area" LED, whenever the sensing device is positioned in a subregion corresponding to a particular set bit. The sensing device can provide further feedback when the user clicks in the subregion, to indicate that the click was in an "active" subregion. This can provide confirmation to the user that a click on a hyperlink succeeded, even if the user is operating offline. The page server which eventually processes the corresponding digital ink input can know that the user received affirmative feedback at the time the click was captured, and can thus ensure that the click succeeds even if a hit test between the click and a precise definition of the hyperlink's zone would otherwise fail. The page server can deduce from the digital ink position samples, with reference to the page description, what feedback the user would have received. Alternatively, each position sample can include the corresponding "active area" bit value actually encountered.

A click in or other interaction with an "active area" can also signal a mobile sensing device and/or base station to connect to the network to allow the click or interaction to be handled in a timely fashion.

A tag may also contain one or more "signature area" bits, indicating that a particular subregion in the vicinity of the tag is in the zone of a signature field. The sensing device can provide feedback to the user, e.g. by illuminating a "signature area" LED, whenever the sensing device is positioned in a subregion corresponding to a particular set bit. This can provide feedback to the user that a signature is knowingly being captured by the system. On this basis, the user learns to associate the feedback with a genuine signature field,

and thereby learns to mistrust a field purporting to be a signature field but which fails to generate the proper feedback. Without such feedback a "Trojan horse" can be used to steal a user's signature via a drawing field masquerading as a signature field. To prevent the inclusion of fake "signature area" bits in the tags of a drawing field, the bits are automatically generated by the system. However, to prevent a forged page being used successfully, each position sample ideally includes the corresponding "signature area" bit value actually encountered. If a page server receives unexpected "signature area" bits then it can report the forgery, and at a minimum it can treat the input as a genuine signature, whose digital ink is never forwarded to an application anyway, but is instead used as the basis for verifying user identity and generating a digital signature of form content.

"Signature area" feedback can of course also be provided in the absence of functional area bits embedded in the tags, albeit with greater latency, simply on the basis of the page description.

Tags typically have a diameter of a few millimeters. Defining functional areas (such as active areas and signature areas) at tag resolution may therefore not be ideal. To overcome this, a number of flag bits within a tag can be dedicated to defining a particular functional area, with each flag bit corresponding to a subregion in the vicinity of the tag. A group of such flag bits can define a bitmap whose organization corresponds to the spatial organization of the corresponding subregions. The subregions typically fully tile the area of the tag, so that any subregion of the area of the tag can be included in a functional area. Assuming a subdivision factor of n , i.e. a square tag is tiled with n by n subregions, a bitmap of n^2 bits is required.

The field of view of the sensing device is designed so that a complete tag is guaranteed to be captured at any position on a surface, subject only to the allowed tilt range of the sensing device. The center of the field of view therefore lies inside a complete tag. Furthermore, the field of view is not designed to guarantee that more than one complete tag is imaged at any particular position. Since the nib of the pen is offset from the optical axis by at least the diameter of a tag, i.e. to prevent it from obscuring the tag, it is possible for its position to lie outside the only complete tag imaged.

As a consequence of the nib being offset from the center of the field of view, any functional area bitmap within a tag must actually cover an area larger than the area of the tag.

If the nib offset is given by S and the pen tilt by θ , then the offset when the pen is tilted is given by:

$$S' = S / \cos \theta$$

As the tilt increases so does the nib offset, and with a maximum pitch-induced tilt of 48 degrees, as discussed in our co-pending application PCT/AU00/00569, the maximum nib offset is given by:

$$S' = 3S / 2$$

If the tag diameter is k and the nib offset S is minimized, i.e. it is equal to the tag diameter, then the maximum nib offset is approximated by:

$$S' = 3k / 2$$

Based on this assumption, the maximum number of tags (or tag subregions) spanned by the center

of the field of view (FOV) and the nib can be analyzed by considering various arrangements of the center-to-nib vector with respect to a square tag array.

In the diagrams the FOV is shown as a circle of appropriate diameter, with a small open circle representing its center. The active tag, which contains the FOV center, is shown as a square with a thick outline. Other relevant tags are shown as squares with thin outlines. The nib position is shown as a filled circle. In diagrams illustrating zero tilt, the center-to-nib separation is the same as the tag diameter, while in diagrams illustrating maximum tilt, the center-to-nib separation is as given above. The FOV center is always shown in a position which maximizes the number of tags spanned by the center-to-nib vector. The tag in which the nib lies is shown with four subregions.

When the center-to-nib vector is arranged diagonally to the tags and the pen is at zero tilt, as illustrated in Figure 6a, the nib lies in a tag diagonally adjacent to the active tag. When the pen is at maximum tilt, as illustrated in Figure 6b, the nib lies in a further diagonally adjacent tag. Because the center-to-nib separation increases rapidly with tilt while the size of the FOV increases only slowly, the tag which lies between the active tag and the tag in which the nib lies does not enter the FOV soon enough to provide an alternative active tag.

When the center-to-nib vector is arranged off-diagonally to the tags and the pen is at zero tilt, as illustrated in Figure 6c, the nib lies in a tag beyond a tag diagonally adjacent to the active tag. When the pen is at maximum tilt, as illustrated in Figure 6d, the nib lies in the same tag, but in a different subregion.

When the center-to-nib vector is arranged parallel to the tags and the pen is at zero tilt, as illustrated in Figure 6e, the nib lies in a tag two tags beyond the active tag. When the pen is at maximum tilt, as illustrated in Figure 6f, the nib lies in a tag three tags beyond the active tag. In this case, however, the tag adjacent to the active tag provides an alternative active tag, as illustrated in both Figure 6e and Figure 6f.

Figure 7a shows the consequent extent of a functional area bitmap when there is one subregion per tag. It has a storage requirement of 25 bits. Figure 7b shows the extent of a functional area bitmap when there are four subregions per tag. It has a storage requirement of 96 bits.

In general, a single bit in a functional area bitmap does not have to represent a subregion of a tag, nor do subregions necessarily have to non-overlapping. It depends somewhat on the meaning of the functional area and the behavior it is being used to control. For example, where an "active area" functional area is used to indicate to a mobile sensing device and/or base station that it should connect to the network to allow the input to be handled in a timely fashion, there is no need for the functional area to be precisely mapped, so long as the mapping is inclusive. A single bit per tag can be used to indicate whether, when a particular tag is active (i.e. contains the FOV center), it is possible for the nib to be located in an "active area". In this case the "subregion" would be a circle of radius $2S$ centered on the tag, in line with the foregoing discussion. Obviously more than one bit can still be used per tag to define multiple subregions.

The size of a functional area bitmap can obviously be minimized by minimizing the number of subregions per tag. It can also be minimized by limiting the allowed tilt range of the pen. Conversely, the size of the bitmap may increase if the nib offset is increased.

Feedback based on functional areas such as active areas and signature areas can of course also be provided by the pen based on information in the page description provided by the page server to the pen via the base station. Although this route has higher latency and doesn't function when the pen is offline, it has the advantage of not consuming potentially scarce tag data bits.

The present invention has been described with reference to a preferred embodiment and number of specific alternative embodiments. However, it will be appreciated by those skilled in the relevant fields that a number of other embodiments, differing from those specifically described, will also fall within the spirit and scope of the present invention. Accordingly, it will be understood that the invention is not intended to be limited to the specific embodiments described in the present specification, including documents incorporated by cross-reference as appropriate. The scope of the invention is only limited by the attached claims.

CLAIMS

1. A region having coded data thereon, wherein at least some of the coded data includes data indicative of a functional attribute of a part of the region and wherein the coded data comprises a plurality of substantially undifferentiated marks positioned relative to a set of predetermined nominal mark positions.
2. The region according to claim 1, wherein each mark represents one of a plurality of coded values.
3. The region according to claim 2, wherein the coded value represented by a mark is at least partially determined by the position of the mark relative to the set of predetermined nominal mark positions.
4. The region according to claim 1, wherein the predetermined nominal mark positions are arranged relative to a plurality of nominal lines.
5. The region according to claim 4 wherein at least one pair of the lines are spaced apart from one another by a distance of less than 100 millimeters.
6. The region according to claim 1, wherein the marks have centers and wherein the relative spacing of the centers of at least one pair of adjacent marks is less than 12 millimeters.
7. The region according to claim 1, wherein each mark has a minimum radius and a maximum radius and wherein the length of a mark's minimum radius is greater than 40% of the length of that mark's maximum radius.
8. The region according to claim 1, wherein any 10 millimeter diameter area of the coded data indicative of the functional attribute of a part of the region includes sufficient information to identify the functional attribute.
9. The region of claim 1 wherein the functional attribute indicates at least one of the group comprising: a hyperlink, a hypertext link, a button, a drawing field, a text field and a signature field.
10. The region of claim 1 wherein some of the coded data also includes data indicative of a location.
11. The region of claim 1 wherein some of the coded data also includes data indicative of an identity.
12. The region of claim 1 wherein some of the coded data also includes data indicative of an identity and a location.
13. The region of claim 11 wherein the identity is indicative of a location in the region.

14. The region of claim 11 wherein the identity is indicative of an identity of the region.
15. A method of coding a region including applying coded data to a least a part of the region, the coded data being indicative of a functional attribute of a part of the region and comprising a plurality of substantially undifferentiated marks positioned relative to a set of predetermined nominal mark positions.
16. The method according to claim 15, wherein each mark represents one of a plurality of coded values.
17. The method according to claim 16, wherein the coded value represented by a mark is at least partially determined by the position of the mark relative to the set of predetermined nominal mark positions.
18. The method according to claim 15, wherein the predetermined nominal mark positions are arranged relative to a plurality of nominal lines.
19. The method according to claim 18, wherein at least one pair of the lines are spaced apart from one another by a distance of less than 100 millimeters.
20. The method according to claim 15, wherein the marks have centers and wherein the relative spacing of the centers of at least one pair of adjacent marks is less than 12 millimeters.
21. The method according to claim 15, wherein each mark has a minimum radius and a maximum radius and wherein the length of a mark's minimum radius is greater than 40% of the length of that mark's maximum radius.
22. The method according to claim 15, wherein any 10 millimeter diameter area of the coded data indicative of the functional attribute of a part of the region includes sufficient information to identify the functional attribute.
23. The method of claim 15 wherein the functional attribute indicates at least one of the group comprising: a hyperlink, a hypertext link, a button, a drawing field, a text field and a signature field.
24. The method of claim 15 wherein some of the coded data also includes data indicative of a location.
25. The method of claim 15 wherein some of the coded data also includes data indicative of an identity.
26. The method of claim 15 wherein some of the coded data also includes data indicative of an identity and a location.

27. The method of claim 15 wherein the identity is indicative of a location in the region.
28. The method of claim 15 wherein the identity is indicative of an identity of the region.
29. A sensing device for use with a region having coded data thereon, the coded data comprising a plurality of substantially undifferentiated marks positioned relative to a set of predetermined nominal mark positions and at least some of the coded data including data indicative of a functional attribute of a part of the region, the sensing device being adapted to sense and react to data indicative of a functional attribute of a part of the region.
30. The sensing device according to claim 29, wherein each mark represents one of a plurality of coded values.
31. The sensing device according to claim 30, wherein the coded value represented by a mark is at least partially determined by the position of the mark relative to the set of predetermined nominal mark positions.
32. The sensing device according to claim 29, wherein the predetermined nominal mark positions are arranged relative to a plurality of nominal lines.
33. The sensing device according to claim 32 wherein at least one pair of the lines are spaced apart from one another by a distance of less than 100 millimeters.
34. The sensing device according to claim 29, wherein the marks have centers and wherein the relative spacing of the centers of at least one pair of adjacent marks is less than 12 millimeters.
35. The sensing device according to claim 29, wherein each mark has a minimum radius and a maximum radius and wherein the length of a mark's minimum radius is greater than 40% of the length of that mark's maximum radius.
36. The sensing device according to claim 29, wherein any 10 millimeter diameter area of the coded data indicative of the functional attribute of a part of the region includes sufficient information to identify the functional attribute.
37. The sensing device of claim 29 wherein the sensing device provides at least one indication to a user when it senses data indicative of a functional attribute of a part of the region.
38. The sensing device of claim 37 wherein the indication is selected from the group comprising a visual, aural, haptic, tactile, vibratory and any other human sensory modality.

39. The sensing device of claim 29, wherein when the sensing device senses data indicative of a functional attribute of a part of the region, the sensing device transmits, to a computer system, said sensed data or data based at least partially on said sensed data.
40. A system including:
computer system, and
a sensing device for use with a region having coded data thereon, the coded data comprising a plurality of substantially undifferentiated marks positioned relative to a set of predetermined nominal mark positions and at least some of the coded data including data indicative of a functional attribute of a part of the region,
the sensing device for sensing coded data indicative of a functional attribute of a part of the region and for transmitting, to the computer system, said sensed data or data based at least partially on said sensed data,
wherein the computer system executes, or commences execution of, a function related to the functional attribute of the sensed data.
41. The system of claim 40, wherein the computer system verifies that the functional attribute is consistent with an expected attribute of the respective part of the region.
42. The system according to claim 40, wherein each mark represents one of a plurality of coded values.
43. The system according to claim 42, wherein the coded value represented by a mark is at least partially determined by the position of the mark relative to the set of predetermined nominal mark positions.
44. The system according to claim 40, wherein the predetermined nominal mark positions are arranged relative to a plurality of nominal lines.
45. The system according to claim 44 wherein at least one pair of the lines are spaced apart from one another by a distance of less than 100 millimeters.
46. The system according to claim 40, wherein the marks have centers and wherein the relative spacing of the centers of at least one pair of adjacent marks is less than 12 millimeters.
47. The system according to claim 40, wherein each mark has a minimum radius and a maximum radius and wherein the length of a mark's minimum radius is greater than 40% of the length of that mark's maximum radius.
48. The system according to claim 40, wherein any 10 millimeter diameter area of the coded data indicative of the functional attribute of a part of the region includes sufficient information to identify the functional attribute.