

METHOD OF DISPLAYING VISUAL INFORMATION ON A VIEWER

FIELD OF INVENTION

The present invention relates generally to devices for interacting with computer systems, and in particular, to a device which may be used to render otherwise passive surfaces audiovisually interactive via invisible surface codings.

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

Devices such as personal computers, personal digital assistants and even mobile phones may be used to interact with audiovisual information and with computer applications which have an audiovisual interface.

In general, these devices don't provide access to situated interactive information or application interfaces, e.g. to an online fan club chat group associated with a concert poster encountered at a train station, or a mortgage calculator associated with a flyer received through the mail. Each device must be used to seek out the desired information or application interface through a virtual space

accessible through the device, or the information or application must be brought to the device in a device-compatible physical format.

The present invention utilizes methods, systems and devices related to a system referred to as "netpage", described in our co-pending applications listed above, wherein invisible coded data is disposed on various surfaces to render the surfaces interactive in the manner of graphical user interfaces to computer systems.

SUMMARY OF INVENTION

In one broad form the invention provides a hand held viewing device with one or more sensors capable of sensing coded data. Images which include coded data are sensed by the viewer and decoded. The decoded information is transmitted to a computer system which associates the decoded data with one or more files stored on the system, using previously stored association data. The file or files are transmitted to the viewer and to the user via the viewer's display screen.

Accordingly, in one broad form, the invention provides a viewer for the download at least one file and the display of visual information based on the at least one file, the viewer including:

- at least one sensor for sensing coded data on or in a substrate and for generating first data;

- a transceiver for transmitting said first data or second data at least partially based on the first data to a computer system and for receiving at least one file associated with an identity derived from the first data from the computer system;

- at least one display device for outputting visual information based at least partially on said at least one file.

The viewer may include a decoder for receiving said first data from the sensor device and for outputting said second data to the transceiver;

The visual information may include text or images or both. Images may be still or moving images. Audio information may also be output by the device.

Where the display of the viewer overlays part or all of the substrate the visual information may include a component which substantially corresponds to the visual appearance of the part of the substrate overlain by the display, including any markings on the substrate visible to the average unaided human eye. The component may be aligned with markings on the substrate visible to the average unaided human eye or it may be aligned with the viewer. The component may be modified as the position of the viewer on the substrate changes, or under user control. The component may correspond to an area of the substrate larger or smaller than the area of the display.

After receiving one or more files, the viewer may sequentially displays all visual information based on the file(s).

The viewer may display first visual information corresponding to uncoded visible text or markings on the substrate and second visual information corresponding to coded visible text or blank areas on the substrate, so as to enable an authorized viewer or user to view secure or confidential information which is not available for view on the substrate.

In use the viewer may be moved across a substrate tiled with data encoded tags and the output is modified as the viewer is moved across the substrate so as to correspond to the interface on the substrate.

The visual output may correspond to the interface but be text in the same or a different language as the interface. The visual output may include the same, more or less information as the interface on the substrate. The interface may represent information, such as bank account information which is not displayed on the interface in full but is displayed in full via the viewer.

The device preferably has controls to enable the optical output to display information with a size the same, less than or more than the corresponding information on the interface.

The device preferably includes a touch sensitive overlay and the optical output includes interactive elements by which the user may modify the optical output by interaction with the touch screen.

The device may also display motion pictures, in which case the physical relationship between the viewer and the substrate once the coded data has been sensed and transmitted may be immaterial.

The viewer may also include memory into which a file or files, are downloaded for subsequent viewing.

When the viewer is lifted from the page the portion of the page with which the viewer was last in contact can be retained by the viewer and remain interactive. The viewer can do this by default, or alternatively only when the user 'freezes' or 'snaps' the current view before lifting the viewer from the page, as discussed in more detail below. The viewer may include controls for rotating and panning the view after the viewer has been lifted from the page, allowing the viewer to be used to navigate an entire page after only a single contact with it. The navigation controls may be in the form of a small joystick, or a pair of orthogonal thumb wheels, or may be provided via the touch-screen.

The viewer preferably senses the identity of the underlying page as well as its own position and orientation relative to the page using a netpage sensor embedded in the viewer. The viewer can contain multiple sensors (one in each corner, for example), so that it works when only partially overlapping a page or even when overlapping multiple pages. A single sensor located in the center of the viewer is sufficient for most purposes, however, and is most economical.

The viewer is capable of providing an enhanced view of the underlying page. In the dark, for example, it can provide a lit view of the page. When magnification is needed, it can provide a zoomed view of the page via its zoom controls. The viewer may also be able to provide a zoomed-out view of the page.

The viewer can provide an alternative (or additional) netpage data entry mechanism, i.e. the usual kinds of inputs, including drawing and handwriting and gestures, can be captured relative to the displayed page via the touch-screen instead of relative to the physical page via the netpage pen.

The viewer can provide access to dynamic content, such as audio and video, in the context of a physical page. Playback controls which provide access to interactive content are typically only shown when the page is viewed through the viewer. They may be shown on the printed page as well, although ideally in a form which clearly indicates that they are inactive. Status information such as playing time may also be shown on-screen. The underlying page layout should provide space for the display of interactive controls and status information, so that the printed page and on-screen display remain compatible. Interactive control is not limited to playback of streaming media such as audio and video. It can include arbitrary interaction with a computer application. It may, for example, include manipulation of an object in 3D.

BRIEF DESCRIPTION OF DRAWINGS

Figure 1 is a perspective view from above of an embodiment of the invention;

Figure 2 is a perspective view from below of the figure 1 device;

Figure 3 is an exploded perspective view from above of the Figure 1 device;

Figure 4 is a side view from the rear of the device of Figure 1;

Figure 5 is a plan view from above of the Figure 1 device;

Figure 6 is a side view from the right of the Figure 1 device;

Figure 7 is a cross-sectional view taken along line AA of Figure 5;

Figure 8 is a perspective view from above of the internal components of the Figure 1 device;

Figure 9 is a perspective view from below of the internal components of the Figure 1 device;

Figure 10 is an exploded perspective view of the underside of the PCB of the Figure 1 device;

Figure 11 is a plan view showing the Figure 1 device in use on a coded substrate in a first orientation;

Figure 12 is a view from the side of the Figure 11 arrangement;

Figure 13 is a plan view of the Figure 1 device in a second orientation on the substrate of Figure 11.

Figure 14 is a plan view of the Figure 1 device viewer showing a "lifted" view of a page;

Figure 15 is a plan view of the Figure 1 device showing a view of a page in darkness;

Figure 16 is a plan view of the Figure 1 device showing a "zoomed-in" view of a page;

Figure 17 is a plan view of the Figure 1 device showing a "zoomed-out" view of a page;

Figure 18 is an example of a movie library page;

Figure 19 is a plan view of a viewer showing a view of a playable video clip;

Figure 20 is a plan view of a viewer showing a snapped and lifted view of a playable video clip;

Figure 21 shows an example of a music library page;

Figure 22 shows a plan view of a viewer showing a view of playable music clip;

Figure 23 shows a plan view of a viewer showing a snapped and lifted view of a playing audio clip;

Figure 24 shows an example of PIN advice letter from bank;

Figure 25 shows a viewer showing a secure view of personal data;

Figure 26 shows an example of "fill-in" form;

Figure 27 shows a viewer allowing form fill-in by a user;

Figure 28 shows a schematic of viewer's electronic components;

Figure 29 shows a document element view; and

Figure 30 shows the viewer interaction with a netpage network

DETAILED DESCRIPTION OF PREFERRED AND OTHER EMBODIMENTS

Referring to Figs. 1 to 10, there is shown a viewer 100. The viewer has a color LCD screen 102, control buttons 186, 192 & 194, a speaker 106, volume control 108, an audio out jack 110 and a infrared sensing device 112. The LCD screen 102 has a touch sensitive overlay 132.

The viewer 100 is a netpage system enabled device and communicates with a netpage system in a similar manner to that disclosed in the co-pending applications referred to earlier and in particular to applications PCT/AU00/00561 and PCT/AU00/00565.

The viewer 100 has upper and lower moldings 114 and 116 respectively which encapsulate a PCB 118 and all of the electronic components are mounted on or connected to this PCB. Power for the device is supplied by a rechargeable 3 volt lithium ion battery 120. The battery is not user replaceable and is located within the casing.

The two moldings 114 and 116 clip together at the upper edge 122 and are held together by three screws 124 extending through apertures 126 in the lower molding to engage in threaded studs (not shown) in the upper molding 114. The screws pass through holes 128 in the PCB and sandwich the PCB between the lower bosses 130 and the studs in the upper molding.

The PCB includes four electrical switches which are acted on by the push buttons. The zoom buttons 186 allows the view to be zoomed in, zoomed out, and, when activated simultaneously, to be reset to normal scale. The snap button allows the nearest dynamic or interactive object to be snapped to the screen. It more generally allows the current view to be de-synchronized (or 'frozen' or 'lifted') from the underlying page. Optional pan controls allow the view to be panned independently of the underlying page. The power button allows the viewer to be switched on and off. The volume control, in the form of a potentiometer, allows the volume of the viewer's audio output to be controlled. The audio output is used to play audio content associated with a page, and optionally to provide operating feedback to the user.

A combined data and power connector 134 is connected to the PCB and is accessible via opening 136 in the casing. The connector 134 includes a power input socket 138 for recharging of the battery 120 and a data socket 140 for input/output of data to and from the device. The connector 134 is connected to the PCB via a flexible PCB 142 and connectors 144 and 146.

The speaker 106 is also connected to the PCB at its lower edge via connector 148 and is positioned between the PCB and the lower edge of the casing. The speaker may be hard wired to the PCB rather than being provided with a removable connector.

Digital audio generated by the viewer is converted to analog via a digital-to-analog converter (DAC) 173, is amplified by an amplifier 177 subject to the volume control 108, and is output to a speaker 106 or to an external audio device via an audio out jack 110. The speaker is disabled when the audio out jack 110 is in use.

The volume control 108 and audio out jack 110 are mounted directly on the left hand side of the PCB and are accessible via openings 154 and 156 respectively in the left hand side of the casing.

The battery 120 is also positioned between the PCB and the lower edge of the casing. A removable connector 150 is provided to connect the battery to the PCB. However, since the battery is not intended to be user replaceable, a permanent connection may be used. The battery is preferably shaped to fit the casing and accordingly has a beveled corner 152.

Mounted on the lower surface of the PCB is the optical sensor device 112 capable of detecting infrared markings on a substrate. The sensing device 112 comprises an infrared LED 160 and an image sensor 162. In use infrared light is emitted from the LED 160 and passed through an optical guide 164 and then through an aperture 166 in the lower molding. Reflected light passes through the aperture 166, the optical molding 164 and is focused onto the CCD 162. The optical molding preferably includes a lens 168 and may include a beam splitter/combiner to allow light from and to the LED and the CCD to follow the same path. The LED 160 may be strobed in synchrony with image capture to prevent motion-blurring of captured tag images. The image sensor 162 typically consists of a 200x200 pixel CCD or CMOS image sensor with a near-infrared bandpass filter.

The PCB also includes a processor chip 170, DRAM 172, flash ROM 174, display controller 250 for controlling the LCD, a transceiver chip 178 and an aerial 180.

In order to display video, the color display 102 usefully has a 4:3 aspect ratio, although with the advent of DVD and digital broadcasting, digital video content with wide screen aspect ratios of 16:9 and wider are becoming more prevalent. The aspect ratio used is not critical. The minimum useful resolution is preferably SIF resolution, which has a square-pixel equivalent of 320x240 pixels. In order to display text, the display preferably has a pitch of at least 100 pixels per inch, giving a maximum diagonal display size of four inches. Larger displays with correspondingly more pixels provide obvious benefits, although with the added expense of a larger form factor and greater processing and communications bandwidth requirements.

The display ideally mimics the optical properties of the underlying paper, i.e. the paper's high reflectivity, high contrast, and Lambertian reflectance function. Although the display may be a conventional back-lit color LCD, such as an active-matrix twisted-nematic (TN) color LCD, it is more usefully a reflective display, such as a passive-matrix cholesteric color LCD from Kent Displays Inc, an electrophoretic display from E Ink Corporation, or a rotating ball display from Xerox Corporation, Inc. The display may also be a low cost organic LED (OLED) display.

Display output generated by the viewer is routed to the display 102 via the display controller 250. Assuming a 320x240 RGB pixel display, the display controller 250 has an associated or embedded 0.25 Mbyte single-buffered or 0.5 Mbyte double-buffered display memory 181.

A dedicated compressed video and audio decoder 171 which produces square-pixel progressive-scan digital video and digital audio output is also provided. To handle MPEG-1 encoded video and audio, a video and audio decoder similar to a C-Cube CL680 decoder may be used. To handle MPEG-2 encoded video and audio, a video and audio decoder similar to C-Cube's ZiVA-3 decoder may be used. An MPEG-1 decoder typically uses a 4 Mbit DRAM during decoding, while an MPEG-2

decoder typically uses a 16 Mbit SRAM during decoding. The decoder memory 179 may be dedicated to the decoder, or may be part of a memory 172 shared with the processor.

The processor unit 175 controls and coordinates the various electronic components of the viewer. The processor executes software which monitors, via the sensor(s) 112, the identity of the underlying page and the position of the viewer relative to the page; communicates the identity and position data to a netpage base station via the wireless transceiver 178; receives identity- and position-related page data from the base station via the transceiver; renders user output to the color display 102 and audio output; and interprets user input captured via the user interface buttons 104 and the screen's touch sensor 132. The embedded software executed by the processor is stored in the non-volatile memory 174, which is typically a ROM and/or flash memory. Identity information unique to the viewer, as well as communications encryption keys, are also stored in non-volatile memory. During execution the processor utilizes faster volatile memory, typically in the form of a 256 Mbit (32 Mbyte) dynamic RAM (DRAM) 172.

The processor unit 175 communicates with the other components via a shared bus 183. The processor 175, the bus 183, and any number of other components may be integrated into a single chip. As indicated in Fig. 28 the integrated components may include the digital transceiver controller 183, the video decoder interface 187 and the tag image sensor interface 185. If a fingerprint sensor 189 is provided a fingerprint sensor interface 191 may be incorporated in the integrated components. In a more highly integrated chip, the integrated components may also include the display controller 250, the image sensor 112, the compressed video and audio decoder 171, the audio digital-to-analog converter (DAC) 173 and the memory 172. The analog radio transceiver 178 is unlikely to be integrated in the same chip, but may be integrated in the same package. A parallel interface 193 links the buttons 104, touch sensor 132 and the LED 160 to the bus 183.

The processor 170 is sufficiently powerful to render page content at interactive rates, i.e. at least 10 Hz, and to transform video decompressed by the video decoder and merge it with other page content. If it is sufficiently powerful, then it may also perform video and audio decoding, obviating the need for a video and audio decoder.

The transceiver 178 is typically a short-range radio transceiver. It may support any of a number of wireless transmission standards, including Bluetooth/IEEE 802.15, IEEE 802.11, HomeRF/SWAP, HIPERLAN, and OpenAir. Bluetooth/IEEE 802.15, IEEE 802.11-1997, HIPERLAN, OpenAir, and HomeRF/SWAP all support transmission rates in the range of 1 to 2 Mbit/s. IEEE 802.11b supports transmission rates of 5.5 Mbit/s and 11Mbit/s. HIPERLAN also supports a transmission rate of 24Mbit/s in an alternative mode. Beyond these currently-supported wireless LAN (WLAN) standards, next-generation WLAN standards promise to support transmission rates of 100 Mbit/s and higher.

The viewer may alternatively be connected to the base station by cable, or may utilize a non-radio-frequency wireless transport, such as infrared. IEEE 802.11, for example, optionally utilizes an infrared transport. IrDA also utilizes an infrared transport.

The lower surface of the lower casing 16 is provided with four outer feet 182 and four inner feet 184. The outer feet are located near the corners of the device whilst the inner feet 184 are located at the corners of a square centered on the lens 168.

All eight feet extend the same distance and enable the device to be placed on a planar surface with a small gap between the surface and the general plane of the lower molding. The feet have a rounded contact surface and this aids in sliding the device across the surface. The inner feet 184 ensure that if one or more of the outer feet 182 are not supported the device does not tip relative to the surface. They also aid in keeping the surface flat near the sensing device 112.

The sensor device 112 is infrared sensitive. The image sensor 162 is sensitive to infrared light, either inherently or by use of filters and the LED 160 emits infrared light, again inherently or by use of filters. The lens 168 is focused on the plane of the inner and outer feet 182 and 184, as this is where a substrate to be sensed will be located. The sensor device is capable of detecting infrared absorptive tags, such as netpage tags. For a full description of the processes involved, reference is made to co-pending application PCT/AU00/00565 referred to earlier. The CCD 162, the LED 160 and processing functions incorporated in the processor chip 170 are similar to those disclosed in the co-pending application.

The device is thus capable of sensing and decoding netpage tags on a substrate. Image data captured by the CCD 162 is sent to the processor 175 and decoded to a region ID (or page ID) and tag ID. The region ID and tag ID, together with other necessary information, is transmitted to the netpage system via the transceiver chip 178 and aerial 180. The netpage system resolves the region ID and tag ID to corresponding document data which it transmits back to the device. The processor 180 receives the data via the aerial 180 and transceiver. The processor renders the data for display on the color display 102 via the display controller 250.

The LCD screen 102 is overlaid with a touch sensitive overlay 132. When viewing information via the LCD screen, the screen will include controls to enable the user to control the device by touching the screen with a stylus.

USAGE

Referring to figures 11 and 12, in use the device 100 is placed on a substrate 200 having netpage tags 202 tiled over its surface. The substrate 200 may be paper, electronic paper such as used by E Ink Corporation, a plastic sheet or any other suitable substrate. As seen in Figure 12 the sensor device 112 senses one or more of the tags 202, decodes the coded information and transmits this decoded

information to the netpage system. In Figure 11 the substrate carries human discernable text or other markings 204 and so is readable without the device 100.

The device may be used with a surface which only carries netpage tags and so appears as a blank document to a user without an authorized viewer. As discussed in our earlier applications, each tag incorporates data which identifies the page it is on and its location within the page. The netpage system is thus capable of determining the location of the viewer 100 and so can extract information corresponding to that position. When displaying static images or text the displayed information usually corresponds to the human discernable text or other markings 204 on the substrate. Additionally the tags include information which enables the device and/or system to derive an orientation of the device relative to the page from the tags. This enables the information displayed to be rotated relative to the device to match the orientation of the text 204. Thus information 206 displayed by the viewer appears to be what is on the page in the corresponding position under the viewer.

When the viewer first comes into contact with a new page, it downloads the corresponding page description from the relevant netpage page server. It then renders the viewer's view according to the current view transform, i.e. according to the viewer's current zoom setting and its position and rotation relative to the underlying page. Whenever the view transform changes, i.e. because the user moves the viewer or changes the viewer's zoom setting, the viewer re-renders the view according to the new view transform. For the purposes of the following discussion, the size of the viewer's window onto the page is also assumed to be part of the view transform.

Whenever the view transform changes, the viewer transmits the view transform to the netpage page server responsible for the underlying page. This allows the page server to commence streaming dynamic objects which have come into view and to cease streaming dynamic objects which are no longer in view. It also allows the page server to provide the viewer with static objects, such as images, at a suitable resolution.

As the device is moved the sensor device 112 images the same or different tags, which enables the device 100 and/or system to update the device's relative position on the page and to scroll the LCD display as the device moves. The position of the viewer relative to the page can easily be determined from the image of a single tag; as the viewer moves the image of the tag changes and from this change in image the position relative to the tag can be determined. The system "knows" the absolute position of the tag on the page from its tag ID and so the absolute position of the viewer on the page may be derived. The viewer senses tags, and thereby its time-varying position relative to the substrate, and sufficiently frequently that movement of the viewer results in a smooth, as opposed to a jerky, scrolling of the display.

The information 206 displayed in the viewer is preferably aligned with the markings 204 on the page 200, so that it appears as if the viewer 100 is transparent. Thus the orientation of the viewer 100 does not alter the orientation of the displayed information 206, as seen in Fig. 13.

However, as seen in Fig. 14, the text 206 displayed on the LCD screen 102 need not remain aligned with the markings 204 on the page 200 but may remain aligned with the viewer 100. The user may use the 'snap' button 192 to instruct the viewer to display information aligned with the viewer, not the page text 204.

The viewer 100, if provided with backlighting of the LCD 132, is capable of displaying information in the dark. Fig. 15 shows a page 200 with a viewer 100 displaying information 206 corresponding to information on the page which is not visible due to lack of ambient light.

Whilst the information 206 shown on the LCD 102 is similar to the printed information 204, it need not be identical. The left hand control button 186 is a rocker type device. When pressed on the left, the button activates switch 188 and when pushed on the right the button activates switch 190. Pushing down centrally or equally on the left and right activates both switches 188 and 190. Pressing on the left switch 186 causes the display to zoom in whilst pressing on the right causes the display to zoom out. Pressing centrally resets the display to the device default, which may be a 1 to 1 ratio compared to the text 204. Fig. 16 shows a viewer 100 displaying "zoomed-in" text 206 whilst Fig. 17 shows "zoomed-out" text 206, both being of text 204 on a page 200.

The device is not limited to simply reproducing the text printed on a page. Because the printed text is associated with one or more electronic documents in the netpage system, the device can provide more than just the text to the user. Where a document has references, such as footnotes or hyperlinks, these references may be shown on screen. The user may then view those references by touching the reference on screen. This is picked up by the touch sensitive overlay 132 and the appropriate information displayed, either in a new window or as an overlay on the existing text. This information may be displayed for a set period, after which the display reverts to its original display. The system may also cause additional information to be shown on the LCD screen automatically. Where the page is mainly text the device 100 may be set to a read mode whereby as the user reads the page the text is recited and output via speaker 106. This would be of benefit to people learning to read, such as children or those learning a foreign language. Using a stylus and the touch sensitive overlay the user indicates where they are reading and the corresponding text is spoken by the machine. Pronunciation of a word or phrase may be practiced by repeatedly drawing the stylus over the relevant text.

The text displayed need not be in the same language or script as the text on the paper. The device may be set to display a single language, when possible. Thus if a netpage document exists in both English and Italian, for example, the device could be set to display the Italian text even though the printed

document is in English. Where a document only exists in one language, the system may be configured to either display the document in that language or nothing except a message to the effect that document is not available in the desired language.

If desired the device may be set to an "automatic recital" mode in which the screen automatically scrolls through the entire document, with or without a corresponding aural recital. An aural recital may also occur without a corresponding display of text. The speed of playback may be controlled by the user. Since each tag 202 on the page identifies the document the device merely needs to sense one tag on the page for those features to be activated.

Depending on the document size and device memory, an entire document may be downloaded to the device or "chunks" of data may be downloaded as required or as expected to be required. A simple text document, even of many pages, is relatively small and may be buffered in its entirety without needing extra memory; the device is capable of playing video, as will be discussed later, and the buffer needed to accommodate variations in data reception for video display is more than sufficient to accommodate most text documents. Where an entire document is downloaded, preferably the download is configured to initially send data corresponding to the location of the viewer on the page.

When viewing text or static images the device may have two modes. The first mode, described above, maintains a correlation between the information displayed on the LCD screen and the position of the device to the substrate. In the second mode, the device may be moved without affecting the display; the scrolling and autoplay feature discussed above being a subject of this feature. To set the device to the second mode the user may press the 'snap' button 192. Once in the second mode movement of the device across the paper does not automatically affect the information displayed. The information displayed continues to correspond to the location when the viewer was placed in this mode. Movement of information across the display in this mode is controlled by the user. A document will continue to be displayed or be accessible until the device is placed on a different netpage encoded substrate, at which point the device may normally revert to its first mode. Alternatively, the user must explicitly revert to the first mode by pressing the 'snap' button again.

Figs. 18, 19 and 20 provide an example of the interaction of the viewer with a video clip. Fig. 18 shows a sample 'movie library' page 210 containing multiple video clip icons 212. Again the page is tiled with netpage tags, not shown. Fig. 19 shows the view the viewer provides of a video clip icon, which includes playback controls 214, which may be activated via the viewer's touch-screen, and the playing time. The screen-displayed information 214,216 remains fixed relative to the displayed text 218 as the viewer moves. Fig. 20 shows the view the viewer shows once the nearest video clip is 'snapped' to the screen, i.e. with the video clip filling the screen, ready to be played. As seen, all of the playback controls and playing time are displayed when the view is snapped. The user snaps the nearest interactive or

dynamic object to the centre of the screen by pressing the viewer's snap button. Alternatively or additionally, an interactive object may snap to the screen automatically when activated, e.g. when played in the case of a video clip. In the example, status information and controls are superimposed on the video clip to maximize the video clip's use of the screen. The viewer 100 typically suppresses the display of the status information 216 and playback controls 214 once the clip is playing, either until the clip terminates or until the user initiates further interaction by touching the touch-screen 132. The playing time is preferably displayed in units of hours, minutes and seconds, and is preceded by a current chapter number if the clip has multiple chapters. The playback controls 214 include stop, play, pause, skip backwards and skip forwards. When the clip is stopped or paused, a play button is shown. When the clip is playing, a pause button is shown in place of the play button. The skip controls skip to the start/end of the current chapter, or, if already at the start/end, to the start of the previous/next chapter. If held, rather than pressed momentarily, the skip controls scan backwards/forwards through the clip, i.e. they play the clip at a greater than normal rate backwards/forwards.

As an alternative to on-screen controls, the viewer may provide dedicated playback control buttons which control whichever dynamic object is current. It may also provide a dedicated display, such as a monochrome segment LCD, for the display of chapter (or track) information and playing time.

The viewer streams dynamic objects such as video clips from any number of servers via the current netpage base station. The viewer is optimized to only stream objects which are currently visible on the viewer's screen. When an object ceases to be displayed, it may be configured to automatically pause, or it may continue to 'play' on the server, without consuming communications bandwidth, so that when it is once again displayed its playback has progressed in time as expected by the user.

Figs. 21, 22 and 23 provide an example of the interaction of the viewer 100 with an audio clip. Fig. 21 shows a sample 'music library' page 220 containing multiple audio clip entries 224. The entries may be grouped under a title 222. Fig. 22 shows the view the viewer provides of a set of audio clip entries 224, which includes an indication of a 'current' audio clip entry with an associated play button 226 which may be activated via the viewer's touch-screen. The current entry is preferably simply the entry closest to the center of the screen and so changes as the viewer is moved, unless a clip is playing. If no entry is sufficiently close, then no entry is current. Fig. 23 shows the view the viewer provides once the nearest audio clip is snapped to the screen or selected, i.e. with clip-related information filling the screen. Note that, as illustrated in the example, a snapped view can be quite different from any view directly derivable from the page itself.

As described above, when the viewer is lifted from the page the portion of the page with which the viewer was last in contact remains displayed and interactive. It may also be advantageous for the user to be able to 'freeze' the current view so that the viewer no longer synchronizes the view with the

underlying page. When the user snaps the nearest interactive or dynamic object to the screen the view is effectively frozen. When the user presses the snap button again, the viewer once again synchronizes the view with the underlying page. When there is no interactive or dynamic object to be snapped, the viewer may still interpret the snap as a freeze command, but may do no more than rotate the current view so that it is oriented the right way up, i.e. so that the axes of the screen and of the displayed page are aligned. The snap button may therefore be used as a freeze button, but as an alternative the viewer may provide a freeze button separate from the snap button.

The viewer may provide a mode wherein the axes of the screen and of the displayed page are always kept aligned even though the viewer is otherwise synchronizing the view with the underlying page, i.e. the viewer synchronizes the view according to the position and identity of the underlying page, but not according to the rotation of the underlying page relative to the viewer. This mode has the advantage that the displayed view always looks sensible to the user operating the viewer. It has the further advantage that rotation is never allowed to degrade the quality of the displayed image of the page. If rotation is never respected by the viewer, then the viewer is also freed from ever having to compute a rotated page image. As a compromise, the viewer may respect the nearest multiple of ninety degrees to its actual rotation relative to the page.

Figs. 24 and 25 provide an example of the interaction of the viewer with 'secure' data requiring authorized access. Fig. 24 shows an example personal identification number (PIN) advice letter 230 from a bank. Both the account number 232 and PIN 234 are suppressed on the printed page, i.e. printed as a series of X's. Figure 14 shows the view the viewer provides of the PIN advice letter. The view shows the actual account number 236 and PIN 238, on the basis that the netpage system has verified that the user of the viewer is the owner of the data. The advice letter is recorded on a netpage registration server as being owned by a particular user, i.e. by the user to whom the letter was originally addressed. The viewer is recorded as owned by the same user. The identity of the actual user is optionally verified with respect to the identity of the owner of the viewer by verifying the fingerprint of the actual viewer against the fingerprint of the owner. The fingerprint, or some other biometric, is recorded by the registration server for the owner. The viewer may incorporate a fingerprint sensor 189 for the purposes of fingerprint capture. Providing a biometric such as a fingerprint or signature to the viewer may place the viewer in a 'secure' mode wherein secure information is accessible for a period of time. An indicator, such as a lit LED, may indicate this mode to the user. Alternatively, secure access is only granted while the user's biometric is being sensed. In the example, the actual account number and PIN is then only shown while the user's finger is in contact with the fingerprint sensor.

Figs. 26 and 27 provide an example of the interaction of the viewer with a fill-in form 240. The form has a number of fill-in fields 242. Fig. 26 shows an example form. Fig. 27 shows the view the viewer provides of the form. The stylus 244 can be used to fill in the fields of the form via the touch-

screen. The example shows the form 240 being filled in by hand, pending handwriting recognition according to normal netpage mechanisms.

Netpage input, in the form of digital ink, is labeled to indicate how it was captured, i.e. through a netpage pen interacting with a printed netpage, or through a stylus interacting with the touch-sensitive screen of a viewer. The digital ink is also labeled with any change in the authorization state of the viewer, e.g. with respect to the presence of the user's finger on the viewer's fingerprint sensor. This allows the page server, when interpreting netpage input in relation to the corresponding page description, to ignore document elements not visible to the user when the input was captured.

Digital ink generated through a viewer is otherwise suitably transformed to look like it came from a netpage pen interacting with the underlying page, i.e. the digital ink is labeled with the page ID of the underlying page, and the coordinates of position samples in the digital ink are transformed from the coordinate system of the viewer into the coordinate system of the underlying page, according to the view transform current when the position sample was captured. As an alternative, the viewer may provide an insertion point specifiable via the stylus, an associated text cursor, and text entry via an on-screen keyboard or via character-wise character entry and recognition, in the manner of current personal digital assistants (PDAs).

The viewer may alternatively or additionally contain a mobile telephone transceiver for longer-range communication with a netpage server via a mobile telephone network. If the transceiver supports a third-generation 'always-on' packet-switched connection, then the viewer may download or stream page content at will. If the transceiver only supports a circuit-switched connection, then the viewer may choose to connect and download a full page description whenever it encounters a new page, provide local navigation of the page via the downloaded page description, and only make further connections when dynamic content needs to be displayed or when hyperlinks are activated.

If the viewer incorporates a longer-range transceiver, then it may act as a netpage base station for wireless netpage pens and other netpage sensing devices.

When used to interrogate secure data, the viewer may incorporate a fingerprint sensor, such as the Vericom FPS200, to allow it to continuously monitor and verify the identity of the user. The processor typically interrogates the fingerprint sensor via a serial interface.

The viewer may incorporate several levels of power management. After a period of inactivity the viewer may inactivate the display. After a longer period of inactivity the processor may enter a power-conserving quiescent state. Power management may be coupled with the image sensor and/or the touch sensor, allowing wake-up on screen interaction or movement. The viewer may also incorporate an accelerometer for this purpose.

The viewer may be configured with software allowing it to function as a personal digital assistant (PDA). The components and functions of the viewer may also be incorporated into a mobile phone.

The viewer handles a variety of page content, including styled text, outline graphics, bitmapped images, audio, and video. While audio and video are by their nature dynamic (i.e. time-varying), text, graphics and images are normally static. They may, however, be part of a dynamic stream representing the output of an interactive application, executing remotely or within the viewer. A local application may be defined by a script object which is interpreted by the viewer, e.g. coded in Java or similar.

Page content falls into three categories: (1) static elements which appear on the printed page; (2) static elements which only appear through the viewer; and (3) dynamic elements which only appear through the viewer. The first category includes all the visual elements of the page. The second category includes, for example, elements associated with controls for dynamic content, e.g. playback controls for a video clip, and elements associated with secure information. The third category includes anything of a dynamic nature, e.g. a video clip. A dynamic element may have a static counterpart in the page description so that it has a meaningful visual appearance on the printed page. A video clip, for example, may be represented by a rectangle of the appropriate size showing the title of the clip.

Document elements in the second and third categories have associated view filters which restrict their appearance. The 'viewer view' filter restricts the element to appearing through a viewer. The 'secure view' filter restricts the element to appearing through an authorized viewer. The view filter class diagram is shown in Fig. 29.

A digital ink stream is generated with a maximum data rate of about 5 Kbit/s [13]. The viewer generates a view transform stream with a similar maximum data rate. The viewer therefore generates wireless communications traffic to the base station with a maximum data rate of about 10 Kbit/s, which is negligible.

To minimize wireless communications traffic from the base station to the viewer, it may be mandated that only one dynamic object can be active at a time. Even though more than one object may be visible through the viewer, the page server may deactivate all but one such object. It may, for example, only leave the dynamic object closest to the center of the viewer's display active. The user may be able to explicitly activate an alternative dynamic object by clicking on it with the stylus, or the user may be required to explicitly activate any dynamic object, thus implicitly deactivating any currently active object.

Where a dynamic object represents an application window, the page server typically relays to the application any changes to the viewer's view transform which relate to the application window. This allows the application to generate output when the window is visible, and to tailor the output to the actual

overlap between the window and the viewer's view. The output from the application may be in the form of a video stream, or it may be in the form of a stream of graphics commands operating on text, graphics and images.

It is possible to display output from a remote application through the viewer, and capture stylus input through the viewer and route it to the remote application, without the application having to be explicitly designed for interaction with the viewer. Applications typically interact with user interface devices such as bitmapped displays, keyboards and pointing devices in a device-independent manner by utilizing the services of a windowing system, and this allows the user interface devices to be remote from the applications without changing how the applications interact with them.

As an alternative to rendering the view on demand, the viewer may render the entire page into a memory buffer and then simply display a suitably transformed window into the buffer. However, since the rotation of relatively low-resolution text produces unacceptable text quality, the page must either be rendered at a higher resolution than the display resolution and appropriately lowpass filtered before display, or the display resolution must be higher than assumed so far.

As a further alternative, the viewer may rely on the page server or the base station to render the page and supply it to the viewer, typically in compressed form. In this case the viewer is still responsible for integrating dynamic streams with the static page view.

As yet another alternative, the viewer may rely on the page server or the base station to render the current view and supply it to the viewer, typically in compressed form. In this case the viewer need contain no support whatsoever for interpreting and rendering the variety of elements which may comprise a page description. Assuming a 320x240 RGB pixel display, a 10 Hz update rate, and 10:1 compression, this requires a data rate of about 1.8 Mbit/s, which is within the realms of practicality using currently-available personal-area, local-area and wide-area wireless networking technologies as discussed above.

Fig. 30 illustrates the interaction between the netpage viewer and the netpage network. The viewer's time-varying view transform, as well as digital ink captured via the viewer, are transmitted (at A) to a netpage base station. The base station relays (at B) the view transform and digital ink to the page server which holds the electronic description of the physical page with which the viewer is interacting, as identified by the page ID embedded in the tags of the page. The page server interprets the view transform and digital ink with respect to the page, and initially responds (at C) with the page description. It subsequently transmits (at C) dynamic object updates and streaming dynamic content. The base station relays (at D) data from the page server to the viewer.

The page server may notify (at E) an application that the viewer's view onto a window associated with the application has changed, and the application may respond (at F) with commands

instructing the viewer to update parts of its current view. The page server relays (at C) these commands to the base station, which relays (at D) the commands on to the viewer.

Handwritten text can be transferred to the underlying page before or after it is recognized by the netpage system and converted to computer text. If transferred after conversion, the viewer can transfer the computer text to the underlying page in place of the handwritten text. Similarly, a signature can be transferred to the page before or after it is used by the netpage system to verify the identity of the user. If transferred after verification, the viewer can transfer a verification icon and/or a serial number in place of the signature.

Whilst the display of the preferred embodiment is visual, a Braille viewer may be made according to the invention. A Braille viewer is analogous to the optical viewer. It displays in registered fashion and via a tactile Braille display the text of a netpage. It also displays text labels corresponding to graphics and images, where available in the page description. Because Braille characters are typically rendered at larger sizes than visually-recognized characters, the Braille viewer typically presents a zoomed view of the underlying page. This means that as the viewer is moved over the page, the display moves correspondingly faster.

In other embodiments the viewer can be transparent when it is not in contact with a page, so that while it is being placed on a page, the page is seen through the viewer. This can be implemented using a LCD shutter which is reflective when opaque, positioned behind a transmissive color display. However, the netpage viewer design is simplified if the viewer is always opaque, and only ever virtually transparent.

Whilst the invention has been described with reference to the netpage system which uses invisible tags, the invention is not limited to the netpage system or the use of invisible tags. If desired, the invention may utilize tags or codes which are visible to the average unaided human eye, such as bar codes. The tags need not merely encode an identity which is then used to look up the relevant files. The tags may encode instructions at a higher level. For example a tag may encode an instruction of "play chapter 99". If invisible tags are used they need not be limited to the tags disclosed in relation to the netpage system. Other tagging systems are available and any suitable tagging system may be used. The invention is not limited to the use of inks which absorb certain wavelengths or fluoresce certain wavelengths. Magnetic inks, surface modification, including apertures, modification of the structure of the substrate itself all fall within the scope of the invention. The systems and methods to link the viewer of the present invention and the source of the document data are not limited to netpage systems. A viewer may be linked by a cable to a single computer, rather than a network of computers.

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 method of creating a display in a viewer having a sensor for sensing coded data on or in a substrate, wherein the coded data is indicative of at least one reference point the method comprising
 - overlying the substrate with said viewer;
 - sensing using the sensor of the viewer at least some of the coded data overlain by the viewer;
 - generating first data indicative of the sensed coded data;
 - transmitting the first data to a computer system;
 - receiving display data from the computer system, the display data being at least partially derived from the sensed coded data;
 - generating visual information using said received display data and a position of the viewer relative to the substrate indicated by at least some of the sensed coded data; and
 - outputting the visual information to a display device of the viewer.
2. The method of claim 1 including generating a component of the visual information which substantially corresponds to the visual appearance of the portion of the substrate overlain by the display, including any markings on the substrate visible to the average unaided human eye.
3. The method of claim 2 wherein the component is aligned with markings on the substrate visible to the average unaided human eye.
4. The method of claim 2 wherein the component is aligned with the viewer.
5. The method of claim 2 wherein the component is modified as the position of the viewer relative to the substrate changes.
6. The method of claim 2 wherein the component corresponds to an area of

the substrate larger than the area of the display.

7. The method of claim 2 wherein the component corresponds to an area of the substrate smaller than the area of the display.

8. The method of claim 1 including sequentially displaying all visual information based on said display data.

9. The method of claim 1 including displaying visual information which is not present on the substrate.

10. The method of claim 1 further including authenticating a user, wherein the viewer displays visual information which is not present on the substrate only to an authenticated user.

11. The method of claim 10 wherein the authenticating step includes using at least one biometric sensor.

12. The method of claim 10 wherein the authenticating step includes using at least one fingerprint sensor.