

and client-side processing, users are enabled to view the entire content of billions of existing Web pages using handheld devices in a simple and reasonable way.

In one embodiment, the client software may be a plug-in to a Web browser, such as Netscape Navigator or Microsoft Internet Explorer. Such a plug-in might have the browser download the data and display it in a sub-window of the browser. Alternatively, the client software may be a Java applet running in a browser. As another option, the client software may be a stand-alone program that interfaces with the proxy server or proxy software directly. The client software may bypass the proxy when requesting information that won't be translated to vectors, such as bitmaps.

With reference to FIG. 6, client-side processing proceeds in the following manner. In a block 160, the vector representation data (i.e., vectorized HTML content and compressed bitmap content) for the web page is gathered at the client. Typically, this data will be stored in a cache at the client as it is being received, and the client simply retrieved the data from the cache. In a block 162, a display list of vectors is built. This process is well known in the CAD arts, and is enabling rapid zooming of vector-based objects. In a block 164, user selectable scale and offset (pan) values are determined. Based on various user interactions with the user-interface of the client, the user is enabled to control the zoom (size) and offset of the rendered page. For example, suppose the user provides zoom and offset inputs to produce a rendered page 210E, as shown in FIG. 4E. In this rendered page, the original origin is now off of the screen (the page image is shifted upward and toward the left—see FIG. 4F), and the view has been scaled approximately 1.3 times.

Next, in a block 166, the vectors and boundary boxes are processed based on the scale and offset, and a bounding box defining the limits of the display content is determined. The results of this step are shown in FIG. 4F, while FIG. 4G shows specific details on how the vectors and bounding boxes corresponding to image objects 250B and 250B' (now 250B' and 252B', respectively) are processed. Logically, there are generally two ways to scale and offset the rendered content. In one embodiment, vectors and bounding boxes are mapped to a virtual display area in memory that has much greater resolution (e.g., 100,000×100,000 pixels) than any real display, and a virtual display limit bounding box is scaled and moved around over the virtual display area. Accordingly, during subsequent processing described below, objects falling within the display bounding box are rendered by reducing the scaling of those objects in the virtual display to how the objects will appear on the client device display relative to the virtual display bounding box. In the alternate, a fixed reference frame corresponding to the display resolution of the client device screen is maintained, wherein all vectors and bounding boxes are scaled and offset relative to the fixed reference frame. Each scheme has its advantages and disadvantages. One advantage of the second method is that the display bounding box is always maintained to have a size that matches the resolution of the content display area on the client device.

As shown in FIG. 4G, respective offsets in X and Y, ($-\Delta X$ and $-\Delta Y$ in the Figure) are applied to the starting point of each of the vectors. The vectors are then scaled by a scale factor "SF." The results of the new vectors are depicted by vectors 250D" and 252D". This produces a new datum for each object's bounding box that is relative to rendered page datum 262, which remains fixed. As discussed above, only a portion of the display screen will actually be used to display content (as defined by a display limit bounding box 266 in this embodiment), while other portions of the screen, including box 264, will comprise a generally fixed-size user interface.

Accordingly, rendered page datum 262 is not located at the upper left hand corner of the display area, although it possibly could be located at this point when either the current user interface is inactive (i.e., the display portion of the user interface is temporary disabled) or the user interface is contained in other portions of the display.

This foregoing process establishes a starting point (the new datum) for where the content in each object's bounding box will be rendered. At this point, each object's bounding box is then drawn from its new datum using the scaling factor. For example, in the original web page 210D (FIG. 4D), bounding box 250B had an X-axis datum of 150 pixels, a Y-axis datum of 225 pixels, and a height and width of 180×350 pixels. In contrast, after being offset and scaled, bounding box 250B' has an X-axis datum of $150 * SF - \Delta X$, a Y-axis datum of $225 * SF - \Delta Y$, and a height and width of $180 * SF \times 350 * SF$.

Returning to the flowchart of FIG. 6, once the vectors and bounding boxes are offset and scaled, content corresponding to objects having at least a portion of their bounding boxes falling within the display limit bounding box is retrieved from the client device's display list in a block 168. For examples, as shown in FIG. 4F, content corresponding to all of the objects except for those falling entirely outside of display limit bounding box 266 (objects 216, 238, 240, 242 and 244) is retrieved from the display list. That content is then scaled in a block 170. For image content, this comprises decompressing and scaling the compressed bitmaps corresponding to those images. For text content, this comprises scaling the font (i.e., typeface) that the text content portions of the web page are written in the parent HTML document and any referenced documents. There are various techniques for typeface scaling that may be implemented here, depending on the available resources provided by the operating system of the client device. For example, for WINDOWS™ operating systems, many TRUETYPE™ fonts are available, which use a common scalable definition for each font, enabling those fonts to be scaled to just about any size. In other cases, such as current PDA (e.g., Palm Pilots) operating systems, there is no existing feature that supports scaling fonts. As a result, bitmapped fonts of different font sizes and styles may be used. In addition to scaling image and text content, other types of content, such as separator lines and borders may also be scaled by block 170.

The process is completed in a block 172, wherein those portions of the scaled content falling within the display limit bounding box are rendered on the client device's display.

As discussed above, it is foreseen that the invention will be used with client devices having small, low resolution displays, such as PDAs and pocket PCs. Examples of various views of an exemplary web pages obtained from the YAHOO™ web site are shown in FIGS. 7A-B, 8A-B and 9A-B. For instance, FIG. 7A represents how the YAHOO™ home page might appear on a Palm IIIc color PDA.

In addition to directly scaling and offsetting content, the client user-interface software for PDA's provides additional functionality. For instance, a user may select to view a column (results represented in FIG. 7B by tapping that column with a stylus, as shown in FIG. 7A). Similarly, the user may select to zoom in on an image by tapping the image with the stylus, as shown in FIGS. 8A and 8B, or select to view a paragraph in an article by tapping on the paragraph, as shown in FIGS. 9A and 9B. It is noted that in some instances, the display of the paragraph may be reformatted to fit the characteristics of the display, rather than following the original format in the zoom-out view.