

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
9 March 2006 (09.03.2006)

PCT

(10) International Publication Number
WO 2006/026013 A1

- (51) International Patent Classification:
G06T 15/10 (2006.01) **G06T 3/00** (2006.01)
- (21) International Application Number:
PCT/US2005/027152
- (22) International Filing Date: 29 July 2005 (29.07.2005)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
10/901,603 29 July 2004 (29.07.2004) US
- (71) Applicant: RAYTHEON COMPANY? [US/US]; 870
Winter Street, Waltham, MA 02451-1449 (US).

AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SM, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.

(84) Designated States (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

(72) Inventor: HOFF, Christopher, G.; 1200 South Catalina Avenue #201, Redondo Beach, CA 90277 (US).

(74) Agents: GUNTHER, John, E. et al.; c/o RAYTHEON COMPANY, (EO/E4/N119), 2000 East El Segundo Blvd, P.O. Box 902, El Segundo, CA 90245-0902 (US).


(81) Designated States (*unless otherwise indicated, for every kind of national protection available*): AE, AG, AL, AM,

Published:

- with international search report
- before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: MAPPING APPLICATION FOR RENDERING PIXEL IMAGERY



(57) Abstract: A mapping application relies on a modified affine transformation algorithm to provide near real-time navigation of imagery, track-up capability and integration of segment and picture data. The affine transformation is modified to map a "world space" (10) comprised of a plurality of pixel maps A and B onto a view port (16). The pixel data (24) for the maps (23) is preferably stored in vertical strips (26) of memory to reduce the likelihood of cache misses. The mapping application is most useful for rendering pixel imagery on a platform with limited processing power, limited memory and small display sizes, e.g., "hand held devices" such as a GPS handset, or Pocket PC.

MAPPING APPLICATION FOR RENDERING PIXEL IMAGERY

5 BACKGROUND OF THE INVENTION

Field of the Invention

This invention relates to mapping applications, and more specifically to a mapping application for rendering pixel imagery on platforms with limited processing power and memory and small display sizes.

10

Description of the Related Art

Hand held devices are rapidly evolving beyond their use as an address book, calendar, note pad, etc. They are merging with other technologies to provide cellular telephone service, email, wireless internet and other applications. Continuing
15 improvements in processor speed, memory, displays and battery power are making the hand held devices a powerful and ubiquitous tool for consumers, business, and the military.

Mapping applications allow users to display and navigate maps on a hand held display using GPS technology. Mapping application software should support the ability
20 to pan around a map, zoom in and out, manipulate navigation objects, and view a tracked position relative to displayed maps in near real-time, e.g. at least 10 frames per second. Military handsets usually have additional direction-finding hardware such as digital compass and accelerometers, which make the application feature of "track-up" (the ability to rotate the map relative to user orientation) desirable for that market. These
25 requirements have to be implemented on a battery-powered handheld device with limited processing power, memory and display size. Therefore, the application software must be very efficient.

Most consumer GPS mapping applications utilize line segment data to represent roads, land-masses, and other boundaries. Consumer applications such as those seen on
30 Garmin, Magellan, or standalone Pocket PC suites are capable of displaying line segment data. However, their frame-rate is limited by the need to redraw all of the line segments on-screen every time the view of the map is changed. These applications alleviate the

problem somewhat by preprocessing the line segment data into spatial data structures to skip the clipping step. They also employ several drawing tricks such as only drawing certain line segment data at particular zoom levels, e.g., only draw freeways at distant zoom levels, and only drawing all of the line segment data when the user “stays still”, i.e. doesn’t pan or zoom the screen rapidly. Such optimizations allow these applications to achieve a moderate performance level of a few frames per second at best. However, some of the drawing tricks employed can have the effect of temporarily disorienting the user when they are rapidly panning or zooming. Most of these applications do not support rotation.

10 Military applications and certain consumer applications use detailed imagery, like topographic maps or pictures taken from satellites or surveillance aircraft. These maps typically take the form of large picture imagery, composed of pixel data rather than line segment data. The application software that supports these types of maps must process a lot more data to perform the same display and navigation tasks. Applications such as
15 ESRI’s ArcPad and MapServer CE support displaying these types of maps. However, on a platform such as a PDA having limited memory and processing power, the supported frame rate is very slow, a slide show at best. These applications use rendering algorithms that are more efficient at panning than zooming, so they typically limit zoom capability to a “step” zoom between a few constant defined zoom levels so that the slow update rate is
20 less noticeable. This can have the effect of temporarily disorienting the user when “jumping” from one zoom level to the other. In addition, rotating picture maps is a feature that is completely unsupported by any existing PDA mapping tools.

A new approach to the PDA mapping application problem for pixel imagery is needed, especially in the military mapping space. The performance of panning and
25 zooming, particularly important because of the small PDA display size, needs to be increased by at least an order of magnitude to provide near real-time response. Ideally, this would be done without the limitations that the known approaches impose upon the user. A real-time “smooth-zoom” feature is desirable in situations where the user cannot afford to temporarily lose track of their location, such as when driving or operating within
30 a military environment. The desire for a “track-up” feature creates an unmet need to rotate maps, preferably with the same performance as zooming and panning. Finally,

there is a need to create an application that supports both line segment data and picture data, and still maintains good performance with both – particularly for the military whose theatres include both natural and urban environments.

5

SUMMARY OF THE INVENTION

The present invention provides a mapping application for rendering pixel imagery on platform with limited processing power, memory and display size. The mapping tool provides the capability to navigate multiple images simultaneous, the ability to pan, scan,
10 zoom and rotate in near real-time, and enables “track-up” mode.

The mapping application first defines a world space, a coordinate system representing a “world”, typically 2D, in which the user can navigate. Bounding objects, typically rectangles, are defined in world space with each representing the extent and position of a particular map in world space. Although these objects contain no image
15 data of their own (they consist simply of bounding world space coordinates, specifically the four vertices in the case of a bounding rectangle), each object is associated with a particular map and its corresponding pixel data stored elsewhere in memory. In response to user input, these objects can be added, removed, associated with different maps, and resized. Also, map image data can be loaded, added or removed by the user, which in
20 turn can be used for association with world space objects.

A view port is defined within world space by bounding world space coordinates that represents what part of the “world” the user sees on the display. A typical rectangular view port is represented by its four rectangular vertices. The view port is translated, scaled, and rotated within world space in response to user input. The mapping
25 application then performs an affine transformation to assign world space coordinates to each pixel of the display represented by the view port. For every assigned world space coordinate that intersects with a world space map object, the mapping application converts the world space coordinate into a map space coordinate using a second affine transformation and copies the map’s pixel value to the display pixel. The mapping
30 application assigns a background value to all other non-intersecting pixels. Once all display pixels are assigned, an entire “frame” is suitably constructed, and the frame is copied to the actual hardware display buffer so that it is visible on screen.

In another embodiment, the execution time to perform the affine transformation is, on average over all rotations of the view port in world space, greatly improved by storing the map pixel data in vertical strips of memory. On average, this will greatly reduce the likelihood of a cache miss when rendering the maps. In this case pixel data is
5 sampled by determining the vertical strip in memory in which the pixel data resides, converting the map space coordinate into a vertical strip column coordinate, and sampling the pixel data from the vertical strip at the column coordinate.

In another embodiment, a hand held device includes non-volatile storage for storing a plurality of pixel maps and a mapping application. Once initiated, the mapping
10 application is loaded into main memory. User I/O is used to load maps from non-volatile storage into main memory in a world space coordinate system and to navigate world space to position a view port therein. The pixel data is stored in vertical strips in memory. A processor executes the mapping application to assign coordinates in world space to each pixel in the view port using an affine transformation, convert the world space
15 coordinates that intersect a map into map space coordinates using a second affine transformation, sample each pixel at its coordinates from the map's vertical strips stored in memory, assign the sampled pixel value to that pixel in the view port, assign a background value to the non-intersecting pixels and paint the display with the map and background pixels assigned to the view port.

20 If navigation hardware such as a GPS receiver is present and geo-referencing data is provided with the map image data, the user can navigate based on his/her location relative to the loaded maps. If direction-finding hardware is present, such as a digital compass or accelerometer, then the user's orientation can be displayed relative to the map's orientation. With direction-finding hardware, the mapping application will
25 automatically rotate the map according to the user's orientation, so that the direction the user is facing is always upward on the view port.

To further improve processing efficiency, the processor preferably executes the mapping application using integer fixed-point math whenever possible, writes and reads pixel data in double words, and constructs an entire graphics frame in an off-screen
30 double-buffer and then paints the entire frame to the view port at once.

These and other features and advantages of the invention will be apparent to those

skilled in the art from the following detailed description of preferred embodiments, taken together with the accompanying drawings, in which:

BRIEF DESCRIPTION OF THE DRAWINGS

5 FIG. 1 is a diagram providing an overview of the mapping application in accordance with the present invention;

 FIG. 2 is a flowchart of the display and navigation process from a user perspective;

 FIG. 3 is a flowchart of the affine mapping transform for painting the display with
10 the portion of world space defined by the view port;

 FIGs. 4a through 4i are illustrations of the steps of the affine transformation;

 FIGs. 5a and 5b are diagrams of a rotated view port in world space and the painted display;

 FIG. 6 is a plot of execution time versus rotation for rendering maps;

15 FIG. 7 is a 1-bit per pixel vector map;

 FIG. 8 is a perspective view of a hand held display implemented with the mapping tool; and

 FIG. 9 is a hardware block diagram of the hand held display.

20 DETAILED DESCRIPTION OF THE INVENTION

 The present invention relies on a modified affine transformation algorithm to provide near real-time navigation of imagery, track-up capability and integration of segment and picture data. The mapping application is most useful for rendering pixel imagery on a platform with limited processing power, limited memory and small display
25 sizes, e.g., "hand held devices" such as a GPS handset, a personal digital assistant (PDA) or cell phone. However, the techniques could be applied to other platforms such as car navigation systems, laptop computers or PCs.

 Affine transformations are algorithms most often used in 3D graphics for texturing polygons (triangles are most commonly used) with bitmaps. They are also used
30 in applications that allow the rotation of 2D images, such as Power Point®. Although the affine transformation algorithm is known, its use in and modification for a map rendering

application is new. The standard affine algorithm, although adequate for displaying a single image of limited size as is done in 3D graphics or Power Point[®], does not meet all of the requirements for mapping applications. Mapping applications are expected to have the ability to display multiple maps at the same time, handle very large maps and rotate the maps. Affine algorithms only paint a “map space” (referred to as “texture space” in 3d graphics) portion of a single bitmap onto a polygon, which prevents it from supporting multiple maps. Furthermore, the standard memory bitmaps are not optimized for rotation in very large maps.

An overview of the mapping application in accordance with the present invention is illustrated in Figure 1. World space **10** is a (u,v) coordinate system that can be populated with multiple rectangular regions **12** (more generally convex polygons) that signify the position and spatial extent of the pixel maps A and B. In fact, only the vertices of rectangles need to be stored in world space. A view port **16** is defined in an (x,y) coordinate system in display space **17** and represents what is seen on the display. The user selects and positions maps in world space or they are automatically positioned relative to one another according to geo-reference data. The user then positions the view port in world space by navigating via pan, zoom, and rotate, or the view port is automatically located using positional and orientation information. The view port's orientation in world space is changed by applying the requisite translation, scaling and rotation transforms to its four world space vertices.

The application maps world space coordinates (u',v') **18** to each pixel (x',y') **20** of the view port **16** using an affine transformation. For each pixel in the view port **16** that intersects a rectangle **12**, the application converts the world space coordinate (u',v') **18** using a second affine transformation into a map space coordinate (m,n) **21** in map space **22** for the map **23** (Map A) associated with that particular rectangle **12**. Image data **24** for Map A **12** is stored in vertical strips **26** in memory. Using the map space coordinate (m, n) **22**, the application determines the vertical strip **26** of map image data in which the pixel lies and converts the map space coordinate (m,n) **22** into a vertical strip column coordinate (r, s) **28**. Any overlapping maps are arbitrated based, for example, on priority or blending. The application samples the map's pixel data **24** and assigns a value to that pixel (x',y') **20** in the view port **16**. This process is repeated for each pixel in the view

port. Any non-intersecting pixels within the view port are assigned a background color value. The background value is typically one color such as black or white but is not so constrained. Once completed, all of the view port pixels (x',y') 20 are assigned a color value with which to "paint" the display. Any segment data is stored as a monochrome
5 bitmap and rendered like any of the other maps.

The mapping application incorporating the modified affine transformation is extremely fast and delivers very good visual results for surfaces that are parallel to the display's viewing plane, perfect for 2D maps. The concept of painting a "world space" onto the display allows the mapping application to display multiple maps at the same
10 time. The modified affine transformation provides the ability to pan, zoom and rotate an image in near real-time with any granularity and at any rotation angle the user chooses. Improvements in frame rates of 15x to 30x have been demonstrated with software that supports track-up and does not impose limitations on, for example, zoom level.

Mapping applications must deal with very large, variable image sizes (2000 or
15 more pixels per dimension is common). Standard affine algorithms typically define a constant image size (512x512 pixels is standard) to optimize for CPU cache size. If a standard affine algorithm were used to rotate a large image on a typical PDA cache-limited CPU to a vertical orientation (90 or 270 degrees), it would encounter memory inefficiencies so severe that it would incur an approximate 10x performance hit.
20 Because the mapping application does not have the luxury of choosing a small, constant pixel resolution for its images, its affine algorithm has been adapted to better cope with these memory inefficiencies. By presenting the image data to the affine algorithm in vertical strips of memory rather than in a standard bitmap memory orientation, the map software reduces the above situation's 10x performance penalty to 3x.

25 Because the map software is intended to operate on a PDA, other important, but more mundane optimizations to its affine algorithm include using integer fixed point over floating point math whenever possible, double word read/writes whenever possible, and double buffering.

The details of an embodiment of the mapping application for use in a hand held
30 device including the I/O interface with the user and a step-by-step presentation of the modified affine transformation are illustrated in Figures 2 through 6. The equations for

implementing the modified affine transformation assume a world space populated with a number of rectangular maps to be rendered on a rectangular display. The equations can be generalized to accommodate non-standard shapes as long as the display is a convex polygon. However, industry convention is to render rectangular maps onto rectangular displays, which allows for further optimization of the affine transform.

As shown in Figure 2, from a user's perspective the mapping application is very similar to existing applications except that the application allows the user to rotate maps on the display and provides for track-up mode in which the maps are automatically rotated to track user orientation in near real-time. More specifically, a user navigates an on-screen menu to start the mapping tool (step 50) and load map(s) from non-volatile storage into memory (step 52). Maps are positioned relative to one another according to their geo-location or user preference to define the current world space (step 54). Maps can also be removed from memory. Using I/O buttons and a touch screen, the user navigates the maps in world space (step 56) by creating and editing NAV objects, changing the orientation of the map view, and arbitrating map layering, blending, and visibility. If track-up mode is active, rotation is automatically controlled by direction-finding hardware. The mapping application performs the modified affine transformation as detailed below to paint the display with the portion of the map(s) in the view port at the selected translation, zoom, and rotation (step 58). The process iterates as the user continues to navigate through the existing maps or changes maps loaded into memory. When finished, the user uses the on-screen menu to exit the mapping tool (step 60).

As shown in Figures 3 through 5, the modified affine transformation works in response to the selection and positioning of maps (step 54) and the navigation of those maps by the user (step 54). The application defines world space 62 as shown in Figure 4a as a 2D coordinate system (u,v) occupied by rectangular regions 64 and 66 representing maps A and B (step 68). For the purposes of this description, world space will be considered horizontally increasing to the right and vertically increasing downward. View port 70 is defined in coordinate system (x,y) in display space as shown in Figure 4b and represents what is seen on the display in world space 62 (step 72). The rectangular view port 70 has display coordinates (x1, y1), (x2, y2), (x3, y3) and (x4, y4). For each of these display coordinates, a corresponding world space coordinate (u1, v1), (u2, v2), (u3, v3)

and (u4, v4) respectively, will be assigned.

The application maps world space coordinates to each pixel (x', y') 74 in the view port 70 (step 76) as shown in Figure 4c. More specifically, for each (x', y') display coordinate within the bounds of the view port 70, the application assigns a (u', v') world space coordinate using an affine transformation. To do this, the application computes these four slope values:

$$\begin{aligned} \text{dudx} &= (u3 - u1) / (x3 - x1) \\ \text{dvdx} &= (v3 - v1) / (x3 - x1) \\ \text{dudy} &= (u2 - u1) / (y2 - y1) \\ \text{dvdy} &= (v2 - v1) / (y2 - y1) \end{aligned}$$

The application solves for u' and v' using the formula:

$$\begin{aligned} u' &= u1 + (y' - y1) * \text{dudy} + (x' - x1) * \text{dudx} \\ v' &= v1 + (y' - y1) * \text{dvdy} + (x' - x1) * \text{dvdx} \end{aligned}$$

Using these sets of equations, the application can represent view port 70 in the (u,v) coordinate system in world space 62 for any translation, scale (zoom level) or rotation or combination thereof.

The mapping application is optimized for an affine transformation of a single rectangular polygon that represents the entire view port. The affine algorithm simplifies to performing two linear interpolations, one along the vertical axis of the view port and the second along the horizontal axis of the view port for each pixel of the display. The rectangle is horizontal in display space, so the algorithm only needs to linearly interpolate between the left and right edge of the rectangle. In addition, the left and right edge of the rectangle will always have the same slope in world space, which means that both the vertical and horizontal affine slopes of interpolation do not change. This allows the algorithm to pre-calculate constant slope values that apply to both edges. The algorithm then proceeds as normal for an affine transformation; filling the bottom horizontal scan line of the rectangle and up until it fills the entire view port.

The application then determines if mapped world space coordinates (u',v') 74 within view port 70 intersect with any of the rectangular map regions 64, 66 in world space 62 (step 78) as shown in Figure 4d. For each (u', v') coordinate, the application performs an intersection test with every rectangular map region in world space. If a (u',

v') coordinate falls within a map region, a pixel is sampled and assigned to the associated (x', y') pixel of the display. Given the view port shown in Figure 4d, all (u', v') coordinates falling within the shaded areas 79 will have to be sampled.

The application converts the world space coordinates (u', v') 74 in intersecting areas 79 to map space coordinates (m', n') 80 using a second affine transformation (step 82) as shown in Figure 4e. Assuming that a (u', v') coordinate falls within a rectangular map region defined by world space coordinates $(ua1, va1)$, $(ua2, va2)$, $(ua3, va3)$ and $(ua4, va4)$, and the map image data associated with that rectangular region has map space coordinates $(m1, n1)$, $(m2, n2)$, $(m3, n3)$ and $(m4, n4)$, then the application converts (u', v') to (m', n') as follows:

First, the application computes these four slope values:

$$\begin{aligned} dmd_u &= (m3 - m1) / (ua3 - ua1) \\ dnd_u &= (n3 - n1) / (ua3 - ua1) \\ dmd_v &= (m2 - m1) / (va2 - va1) \\ dnd_v &= (n2 - n1) / (va2 - va1) \end{aligned}$$

Then the application solves for m' and n' using the formula:

$$\begin{aligned} m' &= m1 + (v' - v1) * dmd_v + (u' - u1) * dmd_u \\ n' &= n1 + (v' - v1) * dnd_v + (u' - u1) * dnd_u \end{aligned}$$

Figure 4e shows the conversion of world space coordinate (u', v') , which falls within the Map A rectangular region 64, into an (m', n') map space coordinate to access the Map A image data 84. If all maps within world space are considered to be north oriented (which is true for most map data available), then the maps are not rotated in world space, which means that further optimization can be made. If the map is north oriented then $n3 - n1 = 0$ and $m2 - m1 = 0$. This makes the slope $dnd_u = 0$ and the slope $dmd_v = 0$. This simplifies the two equations to:

$$\begin{aligned} m' &= m1 + (u' - u1) * dmd_u \\ n' &= n1 + (v' - v1) * dnd_v \end{aligned}$$

Essentially, the equation becomes a linear interpolation along the horizontal axis of the map in map space that is only dependant on the horizontal coordinate to solve for m' and then a linear interpolation along the vertical axis of the map that is only dependant on the vertical coordinate to solve for n' .

Image data **84** is arranged in a series of vertical strips **86**, each of pixel width: StripWidth. Map coordinates (m', n') **80** can not be used to directly access the map image data. Instead the application uses (m', n') to first determine in which vertical strip the coordinate lies and then converts (m', n') into vertical strip column coordinate (r', s') **88** (step **90**). The RGB pixel value is sampled from the vertical memory strip and assigned to the corresponding view port display pixel (m', n') (step **92**).

To find which vertical strip contains (m', n') , the application applies the formula:

$$\text{VerticalStrip} = \text{floor}(m' / \text{StripWidth})$$

The application then applies the following formula to convert (m', n') to (r', s') :

$$r' = \text{mod}(m', \text{StripWidth})$$

$$s' = n'$$

Using the Vertical Strip index **94** and the (r', s') coordinate **88**, the application samples the RGB pixel value from the map image data **84** and assigns it to the (x', y') display coordinate.

By comparison an affine algorithm using a conventional row-by-row bitmap would perform poorly when rotating a large image because large images have a large horizontal width in pixels. If two pixels are in two different rows of an image, they will be (difference in image rows * image width in bytes) apart in memory. "Image width in bytes" will be large, in this case, because of the large horizontal width in pixels of the image. The second pixel is not likely to be in the same memory address region as the first pixel, which is currently loaded into cache memory on the CPU. If these two pixels are accessed in sequence, a CPU cache miss is likely to occur because of their distance apart in memory. A CPU cache miss forces the CPU to retrieve the second pixel from DRAM rather than cache, an operation an order of magnitude slower than if it was able to "cache-hit." If a large image is rotated to 90 or 270 degrees while zoomed out (the worst case), the affine algorithm will be accessing pixels in sequence that are multiple image rows apart. This means that the affine algorithm may cache-miss on every single pixel it samples from the image.

As just illustrated, the current mapping application splits the image into constant width vertical strips **86** whose pixel row width is a small fraction of the full pixel row

width of the image. Using the example above, this effectively makes “image width in bytes” smaller, increasing the chances that pixel 2 will be within the same memory region as pixel 1, making a cache-hit more likely. The tradeoff of the approach is that rotations of 0 degrees, which tend to access horizontal pixels in sequence, have the chance of
 5 accessing a pixel 1 in one vertical strip and a pixel 2 in a separate vertical strip, incurring a cache miss. However, the increased rate of cache missing for horizontal orientations will be orders of magnitude lower than the chance of cache missing for a vertical orientation with an unoptimized affine algorithm. The current approach provides balanced performance over all rotations as compared to the standard affine transform,
 10 which is optimum for 0 degrees but degrades rapidly with rotation. To implement this optimization, the mapping application as described above stores image data in a vertical strip format and its affine algorithm samples from a vertical strip format rather than a standard bitmap orientation. This optimization has been shown to reduce a 10x performance penalty in the worst case (rotated 90 or 270 degrees) to a 3x performance
 15 penalty.

The number of vertical strips allocated depends on the StripWidth in bytes, pixel width, and pixel color depth (bits per pixel). For example:

StripWidth = 128 bytes
 PixelWidth = 1200 pixels
 20 PixelColorDepth = 2 bytes (16 bits per pixel: 65,536 colors)
 #VerticalStrips Allocated = floor(PixelWidth * PixelColorDepth / StripWidth) +
 1

$$= (1200 * 2 / 128) + 1 = 19$$

It is not uncommon that maps will overlap in world space either because of their
 25 geo-reference or because the user places one map on top of another. In this case, the application must arbitrate how to display the overlapping portions 96 of the maps A and B associated with rectangular regions 64 and 66 (step 98) as shown in Figure 4g. The application may decide which map is “on top” and sample only that map or it may sample the multiple maps together and blend their color values in a weighted average using a
 30 linear interpolation. For example, the formula for blending MapA and MapB, where Map A is “on top” of Map B is:

$$(x', y') \text{ RGB vector} = (\text{sampler Map A RGB vector}) * \text{Alpha} + \\ (\text{sampler Map B RGB vector}) * (1 - \text{Alpha})$$

where Alpha is some fraction of 1 representing the transparency of Map A, 1 being opaque and 0 being completely invisible. Alpha is a settable property for each map in world space, representing how “transparent” the map is. Transparency allows the user to see two maps of the same region blended together without having to constantly flip between them. The application assigns all non-intersecting world space pixels **100** a background color value (step **102**) as shown in Figure 4h.

Once all view port pixels are assigned a color value, the display **104** is “painted” with the RGB values that have been assigned to the view port (step **106**). As shown in Figure 4i, map A **107** is considered to be on top of Map B **108**, and both are overlaid on background **109**.

The application continues to update world space **62** and the size and position of the view port **70** in world space in response to the user’s navigation commands and to perform the modified affine transformation and render the image data in near real-time. As shown in Figures 5a and 5b, the view port **70** has been rotated by 45° in world space **62**. The pixels (x',y') in the view port in display space are mapped to (u',v') coordinates in world space using the aforementioned modified affine transform. Because the image data for Maps A and B is stored in memory in vertical strips, the number of cache misses is minimized and the transformation is executed in near real-time. The world space coordinates inside the view port that intersect with a map are converted into map space coordinates (m',n'). The application determines the vertical strip in memory and the column coordinate (r',s') for each coordinate. The RGB pixel value is sampled from that location in the memory strip and assigned to the corresponding view port pixel (x',y') to paint the display.

The relative performance of an affine transformation that uses standard row-by-row bitmaps versus the modified affine transformation that uses vertical strips of memory is shown in Figure 6, which plots execution time versus rotation. The execution time **110** for a row-by-row bitmap is optimum for rotations of 0° and 180° but increases rapidly and experiences worst case at 90° and 270° where the probability of experiencing a cache miss is the highest. By comparison the execution time **112** for a vertical strip bitmap is

slightly larger for rotations of 0° and 180° but only increases slightly with rotation. The modified affine transformation effectively smoothes out the performance curve for all rotations of large bitmaps. In hand held mapping devices, and particularly those that operate in track-up mode, on average the execution time using the modified affine transformation will be orders of magnitude less than the standard affine transform, providing near real-time performance. Furthermore, the execution time **112** can be optimized for a particular platform by modifying the StripWidth; a wider strip reduces the chance of cache miss for rotations of 0° and 180° and a narrower strip reduces the chance of cache miss for rotations of 90° and 270° .

Finally, map software is typically required to render line segment data. As mentioned before, the performance limiter for consumer navigation software is how many line segments (roads, borders, etc...) are displayed on screen at once, since every line has to be drawn per frame. It would seem that an affine algorithm would be useless in this respect since it is intended to render image data, not line segment data. However, the mapping application uses the affine algorithm in an innovative way that eliminates the need to draw line segments every frame. The mapping application draws all of the line segment data **118** to a large off-screen monochrome bitmap **120** as illustrated in Figure 7, and then feeds the bitmap into the affine engine like any other picture image. The color white (binary "1") of the monochrome bitmap is treated as transparent so that it appears to the user that only the black pixels (the rendered line segments **118**) are drawn on screen. The big advantage of the approach is that the number of line segments rendered becomes irrelevant to performance since they are all pre-rendered. The disadvantages are the memory hit from creating a large bitmap and that line thickness will scale with all of the other picture imagery (for most types of line segment data, it is desirable for line segments to stay the same thickness regardless of zoom level).

As mentioned previously, the mapping application is particularly well suited for rendering pixel imagery on hand held devices with limited processing power and memory and small display sizes. The capability to navigate multiple maps in near real-time and operate in track-up mode will be critical to many military and commercial applications.

As shown in Figures 8 and 9, the mapping application is embodied in a hand held device **130** such as a military GPS handset, Pocket PC or a Palm Pilot, hereafter referred to as a

"PDA". These PDAs will typically have more processing power and memory than typical consumer devices but far less capability than PCs or even laptop computers.

To start the mapping application, the user employs an I/O device **132** e.g., a stylus and touch screen or buttons, to inform the operating system (OS) running on the processor **134** to load the mapping application software **136** from non-volatile storage **138** into main memory **140**. For a Windows OS, the user double clicks on the executable icon **139** representing the mapping application. Non-volatile storage **138** in a PDA is typically flash memory. Main memory **140** is almost always some form of dynamic RAM, usually SDRAM or DDR SDRAM. The PDA is typically used in a wireless, mobile environment and thus requires a mobile power source **141** such as a lithium-ion or alkaline battery.

The PDA processor **134** typically has a small amount of cache memory **142**, 16 or 32 kb compared to 512 kb for typical desktop processors by today's standards. Cache memory **142** is useful for speeding up the processor's interaction with main memory **140** by adding a super-fast portion of memory that can "cache" areas of main memory so that the processor spends most of its time interacting with fast cache instead of slow main memory. Cache **142** "mirrors" portions of main memory **140** that the processor **134** has recognized as memory address ranges of frequent use. When the processor **134** interacts with a memory address that happens to be within the range "mirrored" by cache, the processor interacts with the cache instead of main memory, speeding things up (cache hit). If the processor accesses a memory address range that is not mirrored by cache, a cache miss occurs, the processor gets the value from main memory instead, and a portion of cache is flushed and loaded to mirror the address range that contains the newly accessed address.

Once in main memory **140**, the OS allows the processor **134** to execute the mapping application. The mapping application establishes its own internal data structures in main memory and then displays a view port **144** on a display **146** such as a small touch screen LCD. View port **144** is an area of the display **146** where map imagery **148** can be seen by the user. Using the on-screen interface, the user loads image data **150** from non-volatile storage **138** into main memory **140** where it is preferably stored in vertical strips. Once in main memory **140**, the mapping application uses the map image

data 150 to render maps 148 onto the view port 144. Map Image data is usually large, so only a subset of maps is typically loaded into memory, usually selected according to the immediate location of the user. The mapping application uses this information to define the current instantiation of world space.

5 Given a loaded set of maps 148 displayed on display 146, the user uses I/O 132 to create and edit navigation objects such as waypoint, routes, and areas of interest on the maps. The user can also pan, zoom, and rotate the maps according to his or her desire. The mapping application responds to changes in the view port translation, scale and orientation in world space and performs the modified affine transformation to paint the
10 view port on the display with the appropriate portion of world space. Because image data is preferably stored in vertical strips of main memory, incidence of cache misses averaged over all rotations will be dramatically reduced, hence execution time will be much faster.

 If navigation hardware 152 such as a GPS receiver is present on the PDA and geo-referencing data is provided with the map image data 148, the user can navigate based on
15 his/her location relative to the loaded maps. If direction-finding hardware 154 is present, such as a digital compass or accelerometer, then the user's orientation can be displayed relative to the map's orientation. Direction can also be computed with GPS when the user moves to a new location at a minimum velocity with a constant direction. With direction-finding hardware 154 and the mapping application's ability to rotate maps in
20 near real-time, an additional feature of "track-up" is enabled. In "track-up" mode, the mapping application will automatically rotate the map according to the user's orientation, so that the direction the user is facing is always upward on the view port 144.

 Other important optimizations have been added to the affine algorithm to make it operate efficiently on a hand held device including using integer fixed point over floating
25 point math whenever possible, double-word read/writes to write out the map data and double-buffering to paint the entire view port at once.

 Embedded processors developed for PDA's such as Intel's XSCALE and Bulverde typically lack floating point units. Embedded processors normally do not run applications that favor intensive floating point calculations, so the floating point unit is removed to
30 save power and space for integrated components like memory controllers, display controllers, etc... Whenever an embedded processor needs to execute a floating point

operation, it must emulate it with a series of integer instructions. The performance hit when doing a large number of real number calculations is high.

Because the mapping application in performing the affine transformation deals with slopes of lines during its linear interpolation, the processor must do several real
5 number calculations per pixel. This can significantly degrade performance in an affine transformation.

The technique of fixed point math is used to solve this performance issue. Fixed point math is a way of using integers to emulate real numbers. To produce a fraction, an integer is multiplied by a fixed constant. This size of the fixed constant represents the
10 range of the "fractional" part of the fixed point number. So, if you multiplied a number by 1000, you would have a fractional range of 0-999, or in decimal 0.000 to 0.999. This fractional part can be used in normal integer calculations while still maintaining a fractional value of the number, much like doing calculations with a real number. After carrying out real number-like calculations with fixed numbers, the result can be returned
15 to the actual integer value by dividing by the fixed constant.

In the mapping application, the current fixed constant used is 8096, giving it an accuracy of almost 4 digits after the decimal point. This leaves a range of +/- 265,252 for the integer portion of the number using signed 32 bit integers. The mapping application uses a fixed constant that is a power of 2 because the processor can use fast shift
20 operations instead of the slower integer multiply and divide operations. A fixed point implementation of the mapping application is approximately 2.5 times faster than a floating point implementation.

Memory performance is much more efficient when writing/reading double words (32 bit values) as opposed to writing/reading individual bytes (8 bit values) when dealing
25 with memory in sequential orientation (it is faster to read a single double word than read 4 bytes in sequence). An affine transformation must access map data at an unknown location, so there is no easy way to optimize for word writes within this portion of the algorithm. However, the view port's memory is in sequential order, so when writing out the map data, it is advantageous to do so in double word writes. For this reason, the
30 mapping application works on as many pixels at a time that can fill a double word, i.e. for a 4 bit per pixel display, it will work on 8 pixels at a time, because 1 double word = 32

bits / 4 bits per pixel = 8 pixels. It then composes the resulting values for each of those pixels into a full double word that is then written to the view port. This optimization produces performance gains of about 5-10%.

5 Embedded processors are typically inefficient when it comes to updating the display. The processor has to access un-cached display memory in DRAM or go out on the bus to update an external LCD controller's embedded memory. For this reason, the mapping application's drawing algorithm constructs an entire graphics frame in an off-screen double-buffer, and then copies the entire frame to the view port using a single Windows API BitBlt call. This has the advantage of utilizing memory caching during
10 frame construction, limiting the number of (slow) Windows API calls to one, and removing "flicker" from the display, because the entire display will be updated all at once.

While several illustrative embodiments of the invention have been shown and described, numerous variations and alternate embodiments will occur to those skilled in
15 the art. Such variations and alternate embodiments are contemplated, and can be made without departing from the spirit and scope of the invention as defined in the appended claims.

I CLAIM:

1. A mapping application for use with a hand-held device in which a plurality of maps are stored as pixel data in memory for presentation on a display, a processor executing the mapping application comprising the steps of:
 - selecting one or more objects **64,66** assigned to different said maps and
 - 5 arranging them in a world space coordinate system **62** (steps **54,56**), each said object defining the extent of its map in world space;
 - positioning a view port **70** in world space (step **72**);
 - using an affine transformation to assign coordinates (u',v') in the world space to pixels **74** (x',y') in the view port (step **76**);
 - 10 converting the world space coordinates (u',v') that intersect an object into map space coordinates (m', n') **80** and sampling the map pixel data **84** for the object to assign values to the pixels in the view port (steps **78,82,90,92,98**);
 - assigning a background value to the non-intersecting pixels **100** in the view port (step **102**); and
 - 15 painting the display **104** with the map **107,109** and background pixels assigned to the view port (step **106**).
2. The mapping application of claim 1, wherein the first affine transformation performs a first linear interpolation along a vertical axis of the view port **70** and a second linear interpolation along a horizontal axis of the view port **70** (step **76**), and wherein the second affine transformation performs a linear interpolation along a
 - 5 vertical axis of the object **64** in map space and a second linear interpolation along a horizontal axis of the object **64** (step **82**).
3. The mapping application of claim 1, wherein said pixel data **84** for each said map is stored in a plurality of vertical strips **86** in memory, said pixel data being sampled (step **92**) by,
 - determining the vertical strip **86** in memory in which the pixel data resides;
 - 5 converting the map space coordinate into a vertical strip column coordinate **88**; and
 - sampling the pixel data from the vertical strip at the column coordinate **88**.

4. The mapping application of claim 3, whereby storing the pixel data 84 in vertical strips 86 of memory, on average, increases the execution time of the first affine transformation (step 76) for view port rotations of 0° and 180° and reduces the execution time for view port rotations of 90° and 270°.
5. The mapping application of claim 1, wherein the first affine transformation automatically rotates the view port 70 in world space 62 according to the display's orientation so that the direction a user is facing is always upward on the view port.
6. A hand held device 130, comprising:
a display 146;
non-volatile storage 138 for storing a plurality of pixel maps 150 and a mapping application 136;
5 main memory 140 into which the mapping application is loaded once initiated and segments of said pixel maps are stored in vertical strips in response to user input;
a processor 134 having cache memory 142 that stores the pixel data from the most recently accessed vertical strips of main memory; and
user I/O 132 for loading maps from non-volatile storage into main memory in
10 a world space coordinate system and for navigating world space to position a view port 144 therein,
said processor 134 executing the mapping application 136 to position one or more objects in world space that define the position and spatial extent of the maps 148, assign coordinates in the world space to each pixel in the view port 144 using a
15 first affine transformation, convert the world space coordinates that intersect an object into map space coordinates using a second affine transformation, sample each pixel at its coordinates from the map's vertical strips stored in memory 140, 142, assign the sampled pixel value to that pixel in the view port, assign a background value to the non-intersecting pixels and paint the display 146 with the map and background pixels
20 assigned to the view port.
7. The hand held device of claim 6, wherein the processor 134 executes the first affine transformation by performing a first linear interpolation along a vertical axis of

the view port **144** and a second linear interpolation along a horizontal axis of the view port **144**, and executes the second affine transformation by performing a linear interpolation along a vertical axis of the object in map space and a second linear interpolation along a horizontal axis of the object.

8. The hand held device of claim 6, wherein the processor **134** samples said pixel data **84** from the vertical strips **86** by determining the vertical strip in memory in which the pixel data resides, converting the map space coordinate into a vertical strip column coordinate **88**, and sampling the pixel data from the vertical strip at the column coordinate.

9. The hand held device of claim 6, wherein the processor **134** executes the mapping application using integer fixed-point math whenever possible, writes out the pixel data to the view port in double words and constructs an entire graphics frame in an off-screen double-buffer and then paints the entire frame to the view port at once.

10. The hand held device of claim 6, further comprising direction-finding hardware **154** whereby the first affine transformation automatically rotates the view port in world space according to the display's orientation so that the direction a user is facing is always upward on the view port.




Fig. 1

Fig. 2





Fig. 4a




Fig. 3




Fig. 4b




Fig. 4c




Fig. 4d




Fig. 4e




Fig. 4f




Fig. 4g




Fig. 4h




Fig. 4i




Fig. 5a




Fig. 5b




Fig. 6




Fig. 7a

Fig. 7b




Fig. 8




Fig. 9

INTERNATIONAL SEARCH REPORT

International Application No

PCT/US2005/027152

A. CLASSIFICATION OF SUBJECT MATTER

G06T17/50 G06T3/00

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

G06T

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, WPI Data, INSPEC

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	EP 0 510 939 A (PIONEER ELECTRONIC CORPORATION) 28 October 1992 (1992-10-28) abstract page 3, line 43 - line 51 page 4, line 5 - line 17	1-10
X	US 2001/038718 A1 (KUMAR RAKESH ET AL) 8 November 2001 (2001-11-08) abstract paragraph '0054! - paragraph '0060! paragraph '0104! - paragraph '0107!	1
A	EP 0 841 537 A (XANAVI INFORMATICS CORPORATION) 13 May 1998 (1998-05-13) abstract ----- -/--	1-10



Further documents are listed in the continuation of box C.



Patent family members are listed in annex.

* Special categories of cited documents:

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

- *T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- *X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- *Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- *G* document member of the same patent family

Date of the actual completion of the international search

23 January 2006

Date of mailing of the international search report

31/01/2006

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

González Arias, P

INTERNATIONAL SEARCH REPORT

International Application No

PCT/US2005/027152

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	EP 0 738 874 A (HITACHI, LTD; NISSAN MOTOR CO., LTD; XANAVI INFORMATICS CORPORATION). 23 October 1996 (1996-10-23) abstract -----	1-10

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/US2005/027152

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
EP 0510939	A	28-10-1992	DE 69206672 D1	25-01-1996
			DE 69206672 T2	02-05-1996
			JP 3108122 B2	13-11-2000
			JP 6088731 A	29-03-1994
			US 5396430 A	07-03-1995
US 2001038718	A1	08-11-2001	NONE	
EP 0841537	A	13-05-1998	JP 3375258 B2	10-02-2003
			JP 10143066 A	29-05-1998
			KR 268071 B1	01-12-2000
			US 6175802 B1	16-01-2001
EP 0738874	A	23-10-1996	DE 69633421 D1	28-10-2004
			DE 69633421 T2	29-09-2005
			JP 3474022 B2	08-12-2003
			JP 8292720 A	05-11-1996
			US 6141014 A	31-10-2000