

Lecture 5: Output Models

Fall 2004

6.831 UI Design and Implementation

1

Today's Topics

- Output models
- Automatic redraw
- Antialiasing & subpixel rendering
- Color models

Today's lecture continues our look into the mechanics of implementing user interfaces, by looking at **output** in more detail.

Our goal for these implementation lectures is not to teach any one particular GUI system or toolkit, but to give a survey of the issues involved in GUI programming and the range of solutions adopted by various systems. Presumably you've already encountered at least one GUI toolkit, probably Java Swing. These lectures should give you a sense for what's common and what's unusual in the toolkit you already know, and what you might expect to find when you pick up another GUI toolkit.

Three Output Models

- **Components**
 - Graphical objects arranged in a tree with automatic redraw
 - Example: Label object, Line object
 - Also called: views, interactors, widgets, controls, retained graphics
- **Strokes**
 - High-level drawing primitives: lines, shapes, curves, text
 - Example: drawText() method, drawLine() method
 - Also called: vector graphics, structured graphics
- **Pixels**
 - 2D array of pixels
 - Also called: raster, image, bitmap

Fall 2004

6.831 UI Design and Implementation

3

There are basically three ways to represent the output of a graphical user interface.

Components is the same as the view hierarchy we discussed last week. Parts of the display are represented by view objects arranged in a spatial hierarchy, with automatic redraw propagating down the hierarchy. There have been many names for this idea over the years; the GUI community hasn't managed to settle on a single preferred term.


Strokes draws output by making calls to high-level drawing primitives, like drawLine, drawRectangle, drawArc, and drawText.

Pixels regards the screen as an array of pixels and deals with the pixels directly.

All three output models appear in virtually every modern GUI application. The component model always appears at the very top level, for windows, and often for components within the windows as well. At some point, we reach the leaves of the view hierarchy, and the leaf views draw themselves with stroke calls. A graphics package then converts those strokes into pixels displayed on the screen. For performance reasons, a component may short-circuit the stroke package and draw pixels on the screen directly. On Windows, for example, video players do this using the DirectX interface to have direct control over a particular screen rectangle.

What model do each of the following representations use? HTML (component); Postscript laser printer (stroke input, pixel output); plotter (stroke input and output); PDF (stroke); LCD panel (pixel).

Example: Designing a Graph View



- Component model
 - Each node and edge is a component
 - A node might have two subcomponents: circle and label
- Stroke model
 - Graph view draws lines, rectangles and text
- Pixel model
 - Graph view has pixel images of the nodes

Fall 2004

6.831 UI Design and Implementation

4

Since every application uses all three models, the design question becomes: at which points in your application do you want to step down into a lower-level output model? Here's an example. Suppose you want to build a view that displays a graph of nodes and edges.

One approach would represent each node and edge in the graph by a component. Each node in turn might have two components, a rectangle and a label. Eventually, you'll get down to primitive components available in your GUI toolkit. Most GUI toolkits provide a label component; most don't provide a primitive circle component. One notable exception is Amulet, which has component equivalents for all the common drawing primitives. This would be a **pure component model**, at least from your application's point of view – stroke output and pixel output would still happen, but inside primitive components that you took from the library.

Alternatively, the top-level window might have *no* subcomponents. Instead, it would draw the entire graph by a sequence of stroke calls: `drawRectangle` for the node outlines, `drawText` for the labels, `drawLine` for the edges. This would be a **pure stroke model**.

Finally, your graph view might bypass stroke drawing and set pixels in the window directly. The text labels might be assembled by copying character images to the screen. This **pure pixel model** is rarely used nowadays, because it's the most work for the programmer, but it used to be the only way to program graphics.

Hybrid models for the graph view are certainly possible, in which some parts of the output use one model, and others use another model. The graph view might use components for nodes, but draw the edges itself as strokes. It might draw all the lines itself, but use label components for the text.

Issues in Choosing Output Models

- Layout
- Input
- Redraw
- Drawing order
- Heavyweight objects
- Device dependence

Fall 2004

6.831 UI Design and Implementation

5

Layout: Components remember where they were put, and draw themselves there. They also support automatic layout. With stroke or pixel models, you have to figure out (at drawing time) where each piece goes, and put it there.

Input: Components participate in event dispatch and propagation, and the system automatically does **hit-testing** (determining whether the mouse is over the component when an event occurs) for components, but not for strokes. If a graph node is a component, then it can receive its own click and drag events. If you stroked the node instead, then you have to write code to determine which node was clicked or dragged.

Redraw: An automatic redraw algorithm means that components redraw themselves automatically when they have to. Furthermore, the redraw algorithm is efficient: it only redraws components whose extents intersect the damaged region. The stroke or pixel model would have to do this test by hand. In practice, most stroked components don't bother, simply redrawing everything whenever some part of the view needs to be redrawn.


Drawing order: It's easy for a parent to draw before (underneath) or after (on top of) all of its children. But it's not easy to interleave parent drawing with child drawing. So if you're using a hybrid model, with some parts of your view represented as components and others as strokes, then the components and strokes generally fall in two separate layers, and you can't have any complicated z-ordering relationships between strokes and components.

Heavyweight objects: Every component must be an object (and even an object with no fields costs about 20 bytes in Java). As we've seen, the view hierarchy is overloaded not just with drawing functions but also with event dispatch, automatic redraw, and automatic layout, so that further bulks up the class. The flyweight pattern used by InterView's Glyphs can reduce this cost somewhat. But views derived from large amounts of data – say, a 100,000-node graph – generally can't use a component model.

Device dependence: The stroke model is largely device independent. In fact, it's useful not just for displaying to screens, but also to printers, which have dramatically different resolution. The pixel model, on the other hand, is extremely device dependent. A directly-mapped pixel image won't look the same on a screen with a different resolution.

Drawing in the Component Model

- Drawing goes top down
 - Draw self (using strokes or pixels)
 - For each child component,
 - If child intersects clipping region then
 - intersect clipping region with child's bounding box
 - recursively draw child with clip region = intersection



Fall 2004

6.831 UI Design and Implementation


6

Here's how drawing works in the component model. Drawing is a top-down process: starting from the root of the component tree, each component draws itself, then draws each of its children recursively. The process is optimized by passing a **clipping region** to each component, indicating the area of the screen that needs to be drawn. Children that do not intersect the clipping region are simply skipped, not drawn. In the example above, nodes B and C would not need to be drawn. When a component partially intersects the clipping region, it must be drawn – but any strokes or pixels it draws when the clipping region is in effect will be masked against the clip region, so that only pixels falling inside the region actually make it onto the screen.

For the root component, the clipping region might be the entire screen. As drawing descends the component tree, however, the clipping region is intersected with each component's bounding box. So the clipping region for a component deep in the tree is the intersection of the bounding boxes of its ancestors.

For high performance, the clipping region is normally rectangular, using component **bounding boxes** rather than the components' actual shape. But it doesn't have to be that way. A clipping region can be an arbitrary shape on the screen. This can be very useful for visual effects: e.g., setting a string of text as your clipping region, and then painting an image through it like a stencil. Postscript was the first stroke model to allow this kind of nonrectangular clip region. Now many graphics toolkits support nonrectangular clip regions. For example, on Microsoft Windows and X Windows, you can create nonrectangular windows, which clip their children into a nonrectangular region.

Damage and Automatic Redraw



Fall 2004

6.831 UI Design and Implementation

7


When a component needs to change its appearance, it doesn't repaint itself directly. It *can't*, because the drawing process has to occur top-down through the component hierarchy: the component's ancestors and older siblings need to have a chance to paint themselves underneath it. (So, in Java, even though a component can call its `paint()` method directly, you shouldn't do it!)

Instead, the component asks the graphics system to repaint it at some time in the future. This request includes a **damaged region**, which is the part of the screen that needs to be repainted. Often, this is just the entire bounding box of the component; but complex components might figure out which part of the screen corresponds to the part of the model that changed, so that only that part is damaged.


The repaint request is then **queued** for later. Multiple pending repaint requests from different components are consolidated into a single damaged region, which is often represented just as a rectangle – the bounding box of all the damaged regions requested by individual components. That means that undamaged screen area is being considered damaged, but there's a tradeoff between the complexity of the damaged region representation and the cost of repainting.

Eventually – usually after the system has handled all the input events (mouse and keyboard) waiting on the queue -- the repaint request is finally satisfied, by setting the clipping region to the damaged region and redrawing the component tree from the root.

Naïve Redraw Causes Flashing Effects




Object moves




**Determine
damaged region**

Fall 2004



**Redraw parent
(children blink out!)**

6.831 UI Design and Implementation



Redraw children


8

There's an unfortunate side-effect of the automatic damage/redraw algorithm. If we draw a component tree directly to the screen, then moving a component can make the screen appear to flash – objects flickering while they move, and nearby objects flickering as well.

When an object moves, it needs to be erased from its original position and drawn in its new position. The erasure is done by redrawing all the objects in the view hierarchy that intersect this damaged region. If the drawing is done directly on the screen, this means that all the objects in the damaged region temporarily *disappear*, before being redrawn. Depending on how screen refreshes are timed with respect to the drawing, and how long it takes to draw a complicated object or multiple layers of the hierarchy, these partial redraws may be briefly visible on the monitor, causing a perceptible flicker.

Double-Buffering

- Double-buffering solves the flashing problem



Fall 2004

6.831 UI Design and Implementation

9

Double-buffering solves this flickering problem. An identical copy of the screen contents is kept in a memory buffer. (In practice, this may be only the part of the screen belonging to some subtree of the view hierarchy that cares about double-buffering.) This memory buffer is used as the drawing surface for the automatic damage/redraw algorithm. After drawing is complete, the damaged region is just copied to screen as a block of pixels. Double-buffering reduces flickering for two reasons: first, because the pixel copy is generally faster than redrawing the view hierarchy, so there's less chance that a screen refresh will catch it half-done; and second, because unmoving objects that happen to be caught, as innocent victims, in the damaged region are never erased from the screen, only from the memory buffer.

It's a waste for every individual view to double-buffer itself. If any of your ancestors is double-buffered, then you'll derive the benefit of it. So double-buffering is usually applied to top-level windows.

Why is it called double-buffering? Because it used to be implemented by two interchangeable buffers in video memory. While one buffer was showing, you'd draw the next frame of animation into the other buffer. Then you'd just tell the video hardware to switch which buffer it was showing, a very fast operation that required no copying and was done during the CRT's vertical refresh interval so it produced no flicker at all.

Stroke Model

- Drawing surface
 - Also called drawable (X Windows), GDI (MS Win)
 - Screen, memory buffer, print driver, file, remote screen
- Graphics context
 - Encapsulates drawing parameters so they don't have to be passed with each call to a drawing primitive
 - Font, color, line width, fill pattern, etc.
- Coordinate system
 - Origin, scale, rotation
- Clipping region
- Drawing primitives
 - Line, circle, ellipse, arc, rectangle, text, polyline, shapes

Fall 2004

6.831 UI Design and Implementation

10

We've already considered the component model in some detail. So now, let's look at the stroke model.

Every stroke model has some notion of a **drawing surface**. The screen is only one place where drawing might go. Another common drawing surface is a memory buffer, which is an array of pixels just like the screen. Unlike the screen, however, a memory buffer can have arbitrary dimensions. The ability to draw to a memory buffer is essential for double-buffering. Another target is a printer driver, which forwards the drawing instructions on to a printer. Although most printers have a pixel model internally (when the ink actually hits the paper), the driver often uses a stroke model to communicate with the printer, for compact transmission. Postscript, for example, is a stroke model.

Most stroke models also include some kind of a **graphics context**, an object that bundles up drawing parameters like color, line properties (width, end cap, join style), fill properties (pattern), and font.

The stroke model may also provide a current **coordinate system**, which can be translated, scaled, and rotated around the drawing surface. We've already discussed the **clipping region**, which acts like a stencil for the drawing. Finally, a stroke model must provide a set of **drawing primitives**, function calls that actually produce graphical output.

Many systems combine all these responsibilities into a single object. Java's Graphics object is a good example of this approach. In other toolkits, the drawing surface and graphics context are independent objects that are passed along with drawing calls.

When state like graphics context, coordinate system, and clipping region are embedded in the drawing surface, the surface must provide some way to save and restore the context. A key reason for this is so that parent views can pass the drawing surface down to a child's draw method without fear that the child will change the graphics context. In Java, for example, the context can be saved by Graphics.create(), which makes a copy of the Graphics object. Notice that this only duplicates the graphics context; it doesn't duplicate the drawing surface, which is still the same.

Converting Strokes to Pixels

- Is (0,0) the center of the top-left pixel, or is it the upper left corner of the pixel?
 - MS Win: center of pixel
 - Java: upper left corner
- Where is line (0,0) – (10,0) actually drawn?
 - MS Win: endpoint pixel excluded
 - Java Graphics: pen hangs down and right
 - Java Graphics2D: antialiased pen, optional 1/2 pixel adjustments made for compatibility
- Where is empty rectangle (0,0) – (10,10) drawn?
 - MSWin: connecting those pixels
 - Java: extends one row below and one column right
- Where is filled rectangle (0,0) – (10,10) drawn?
 - MSWin: 121 pixels
 - Java: 100 pixels

Fall 2004

6.831 UI Design and Implementation

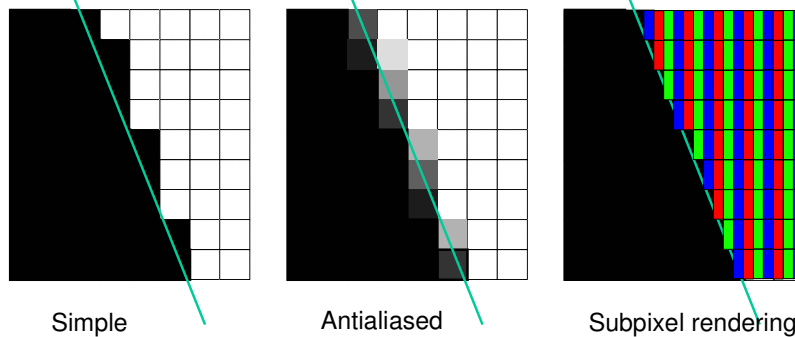
11

When you're using a stroke model, it's important to understand how the strokes are actually converted into pixels. Different platforms make different choices.

One question concerns how stroke coordinates, which represent zero-dimensional points, are translated into pixel coordinates, which are 2-dimensional squares. Microsoft Windows places the stroke coordinate at the center of the corresponding pixel; Java's stroke model places the stroke coordinates *between* pixels.

The other questions concern which pixels are actually drawn when you request a line or a rectangle.

Antialiasing and Subpixel Rendering



Fall 2004

6.831 UI Design and Implementation

12

It's beyond the scope of this lecture to talk about algorithms for converting a stroke into pixels. But you should be aware of some important techniques for making strokes look good.

One of these techniques is **antialiasing**, which is a way to make an edge look smoother. Instead of making a binary decision between whether to color a pixel black or white, antialiasing uses a shade of gray whose value varies depending on how much of the pixel is covered by the edge. In practice, the edge is between two arbitrary colors, not just black and white, so antialiasing chooses a point on the gradient between those two colors. The overall effect is a fuzzier but smoother edge.

Subpixel rendering takes this a step further. Every pixel on an LCD screen consists of three discrete pixels side-by-side: red, green, and blue. So we can get a horizontal resolution which is three times the nominal pixel resolution of the screen, simply by choosing the colors of the pixels along the edge so that the appropriate subpixels are light or dark. It only works on LCD screens, not CRTs, because CRT pixels are often arranged in triangles, and because CRTs are analog, so the blue in a single "pixel" usually consists of a bunch of blue phosphor dots (interspersed with green and red phosphor dots). You also have to be careful to smooth out the edge to avoid color fringing effects on perfectly vertical edges. And it works best for high-contrast edges, like this edge between black and white. Subpixel rendering is ideal for text rendering, since text is usually small, high-contrast, and benefits the most from a boost in horizontal resolution. Windows XP includes ClearType, an implementation of subpixel rendering for Windows fonts. For more about subpixel rendering, see Steve Gibson, "Sub-Pixel Font Rendering Technology", <http://grc.com/cleartype.htm>

Color Models

- RGB: cube
 - Red, green, blue
- HSV: hexagonal cone
 - Hue: kind of color
 - Angle around cone
 - Saturation: amount of pure color
 - 0% = gray, 100% = pure color
 - Value: brightness
 - 0% = dark, 100% = bright
- HLS: double-hexagonal cone
 - Hue, lightness, saturation
 - Pulls up center of HSV model, so that only white has lightness 1.0 and pure colors have lightness 0.5
- Cyan-Magenta-Yellow(-Black)
 - Used for printing, where pigments absorb wavelengths instead of generating them

Fall 2004

6.831 UI Design and Implementation

13

We learned a bit about how humans perceive color when we talked about human capabilities. Now let's look at how colors are represented in GUI software.

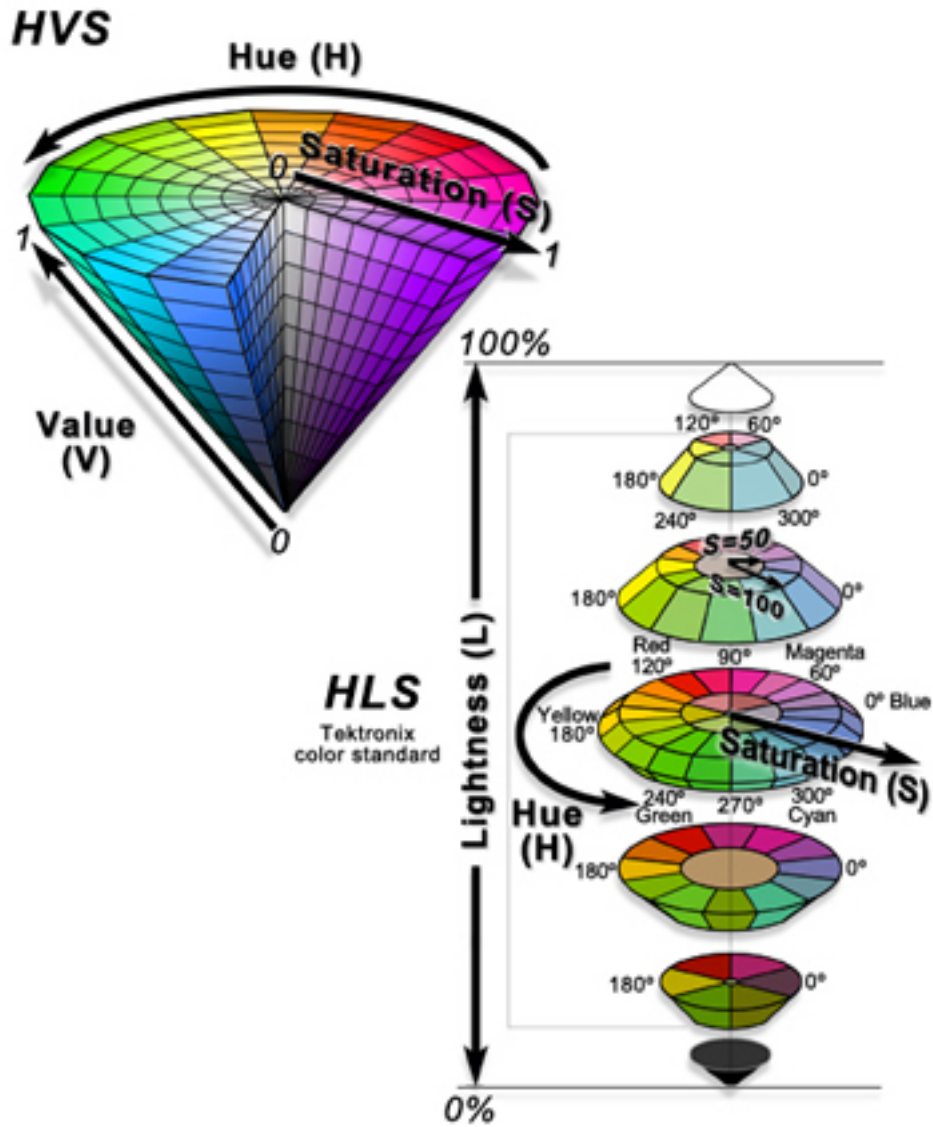
At the lowest level, the RGB model rules. The RGB model is a unit cube, with (0,0,0) corresponding to black, (1, 1, 1) corresponding to white, and the three dimensions measuring levels of red, green, and blue. The RGB model is used directly by CRT and LCD monitors for display, since each pixel in a monitor has separate red, green, and blue components.

HSV (hue, saturation value) is a better model for how humans perceive color, and more useful for building usable interfaces. HSV is a cone. We've already encountered hue and value in our discussion of visual variables. Saturation is the degree of color, as opposed to grayness. Colors with zero saturation are shades of gray; colors with 100% saturation are pure colors.

HLS (hue, lightness, saturation) is a symmetrical relative of the HSV model, which is elegant. See the pictures on the next page.

Finally, the CMYK (cyan, magenta, yellow, and sometimes black) is similar to the RGB model, but used for print colors.

HSV & HLS



Here are some pictures of the HSV and HLS models.

Transparency

- **Alpha** is a pixel's transparency
 - from 0.0 (transparent) to 1.0 (opaque)
 - so each pixel has red, green, blue, and alpha values
- Uses for alpha
 - Antialiasing
 - Nonrectangular images
 - Translucent components
 - Clipping regions with antialiased edges

Fall 2004

6.831 UI Design and Implementation

15

Modern color models add a fourth channel: the pixel's **alpha** value, which is its transparency. Simple image formats like GIF support only two values of alpha: 1 (opaque) or 0 (transparent). PNG has better support, allowing image pixels to be translucent, with alpha values between 0 and 1.