Andrew Glassner's Notebook

http://www.research.microsoft.com/research/graphics/glassner

Signs of Significance

Andrew Glassner

Microsoft Research

You're in an unfamiliar train station, and you're late. You know that you want the 22:15 train to Magritte, but as you look around, your tired arms holding heavy suitcases, you realize you're lost. Thousands of corridors with thousands of branches seem to stretch in every direction. So you look for the signs, and you see your destination in bright yellow letters on a big sign halfway down one of the tunnels. You take a breath, lift your bags, and walk to the train with confidence.

Signs are great. They provide visible labels of our geography. If you're a relativist, you never really know where you are, but you always know where you are with respect to some landmarks. Often signs serve as those landmarks, either simply as part of the remembered geography or as active sources of information.

Because they're so important, signs need to be very clear. Consider the special case of signs that change, like the varying destination signs in a train station. The information must be succinct and the presentation direct, contributing to the information without conflicting with it.

How do you avoid conflicting with something unpredictable? Be simple and bland. Simplicity means that you don't risk having the style of the presentation conflict with the underlying message. Blandness means you avoid contamination or distortion of the message. So a simple, bland sign should be perfectly legible and perfectly boring. If you're careful and clever, you can work in some design as well.

In my experience three types of varying sign are the most popular: mechanical, electromechancial, and electronic. Mechanical signs include a whiteboard and a movie marquee. Electromechanical signs often appear on the fronts of buses and in rotating billboards. Some are very clever indeed; I particularly like the grids of little disks, black on one side and yellow on the other, that flip in response to magnetic-field control.

In this column I'll focus on the purely electronic displays, typically implemented with light-emitting diodes (LEDs), liquid crystal displays (LCDs), or light-emitting panels. I think the evolution of the electronic digital display provides a fascinating example of graphic design cooperating with technology. Digital displays also demonstrate the practical importance of economy: Not only must they be easy to read, displays must be economical to make and buy—or no one will.

Early digital displays

Computer-controlled illuminated electronic displays probably began with the Nixie tube prominent in the 1950s and 60s. The Nixie tube was a clear glass envelope with a plastic base that looked like a regular vacuum tube, but with different insides. A loop of wire formed the digit "0," vertically centered in the tube. Behind it, another wire shaped the digit "1," and behind it a "2," and so on. Electrically, each wire acted as a cathode, with a single anode at the top. The tube itself contained neon gas. To display a particular number, you pulsed the appropriate pin at the base with around 200 volts (which you could then reduce to sustain the display). The neon gas around the selected wire would glow an orange-amber color, creating a clear and legible number. Because the digits were formed individually with wires, they looked quite nice, with curves and joins just the way you'd draw them. The glowing neon shone bright enough that even the numbers at the back were easily legible. Nixies were used for everything from large-scale electronic computers to desk calculators.

The Nixie numbers were attractive and very legible, but the tubes were expensive and ran at high voltages. A replacement tube consisted of a black sheet mounted vertically inside the envelope. Seven incandescent wires sat above the sheet, arranged in a rectangle with a horizontal bisector as in Figure 1a. This was the first "seven-

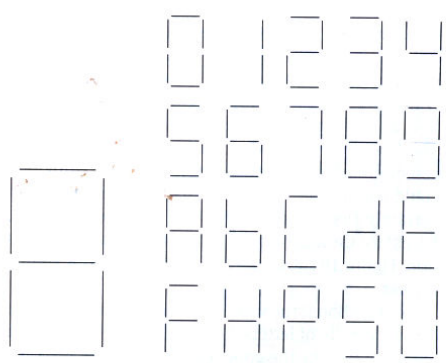

(a) (b)

1 (a) A simple seven-segment display made of wires. (b) The digits and a few recognizable letters.

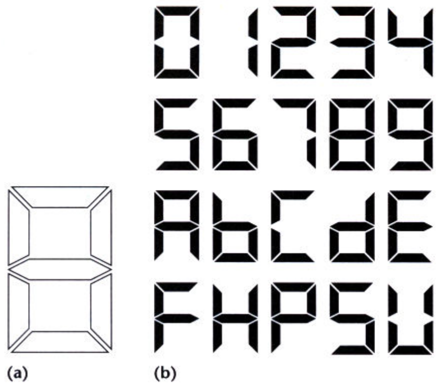

2 (a) The seven-segment display with bevels. (b) The digits and a few letters.

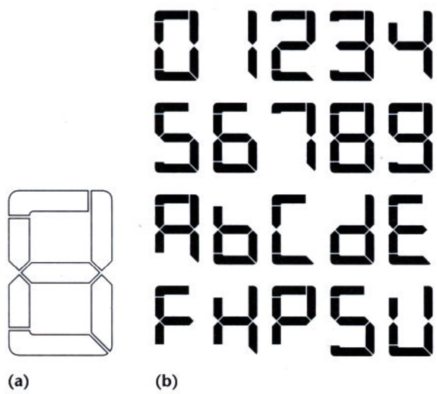

3 (a) The seven-segment display picks up some curves. (b) Notice that the numbers improve but the letters suffer.

segment display." By running electrical current through the different segments, you could cause them to light up, like the wires in your toaster. This incandescent design required plenty of juice, and eventually the filaments burned out.

The seven-segment displays had an economic advantage: They cost less than Nixies. But they looked much less attractive—the Nixie tube's nicely curved numbers became boxy in these displays. We lost our familiar pretty numbers and got something functional but bleak. Still, all 10 digits were clear and easily distinguishable. The design represents a fascinating compliance with the constraint of absolute economy, as I don't think you could pull this off with fewer than seven segments. It certainly fits our simplicity and blandness criteria. Except for the annihilation of beauty, it's a terrific solution. Figure 1b shows the digits formed by this simple format.

As a nice bonus the seven-segment display could also show a variety of letters. To my mind, the letters that worked best were the upper-case A, C, E, F, H, I, J, L, O, P, S, U and lower-case a, b, c, d, e, g, h, l, o, q, r, u. Not a bad list—it has all the vowels and a few consonants. Note also that it includes the first six letters of the alphabet, which was perfect for hexadecimal systems (where A-F stand for the numbers 10-16). You can also make some letters upside down. (To receive a calculator's greeting, punch in 07734 and turn the display upside-down.)

The seven-segment display easily made the leap to solid state. The first LEDs generally came as single blobs of light—a little red or green gumdrop with a couple of leads from the base. Soon after a single rectangular block housed seven independent LEDs, built to illuminate little polygons. The designers took the opportunity to bevel the edges, as shown in Figure 2.

I remember discovering the beauty of this display. One of my first homebrew electronics projects was an electronic deck of cards. I had used my paper-route money to buy an early RAM chip—the 7489 (remember TTL?)—with 64 bits of memory arranged as 16 words of 4 bits each. While wondering what to do with it, I realized that this organization corresponded naturally with a deck of cards, which has 13 cards in each of 4 suits. So I built an electronic card deck.

The user interface consisted of two fingertip-sized red pushbuttons (they were cheap) and two seven-segment displays. Pushing the "deal" button started a very high-frequency counter that ran through the addresses of the bits on the chip; releasing the button stopped the counter on the current bit. If that bit was 0 (indicating that the card had not yet been dealt), I set it to 1 and displayed the appropriate card on two seven-segment LEDs. If the bit was already 1, I stepped through the memory looking for the first 0 bit I could find (ignoring the last three words). If I came full circle, I shuffled the deck by resetting everything to 0, then dealt the next card. A "shuffle" button handled manual resets. I think I used about a half-dozen TTL chips and a 555 counter for this.

The magic seven-segment display made all four suits available: S for spades, C for clubs, H for hearts, and d for diamonds. The cards themselves could be shown by A for ace, 2-9 for number cards, J for jack, q for queen, and (here was a stretch) r for regal, for the king. It worked pretty well, except that the cheap versions of these memory chips kept burning out; they'd get stuck on a single address and never move. The three of hearts forever!

Many LCDs use curved corners, breaking up the symmetry as in Figure 3. Notice that although some of the numbers look better, the letters suffer. It's an interesting tweak—watch designers improved the design for what mattered to them (numbers) and let the other symbols suffer.

A different enhancement increased the number of segments to 16 by cutting the middle segment in half and adding a plus and an X to the display, as in Figure 4. I have only seen this display with little line segments, not with thickened segments like Figure 2. This display took a great step forward, making the entire alphabet available. Unfortunately, it hardly excelled in aesthetics; the B in particular relies pretty heavily on people guessing the letter from context. For some reason, I've seen this one a lot in elevators.

Moving to grids

Eventually people could stuff a whole bunch of little dot LEDs onto a chip. Most popular was the 5-by-7 grid, which let you represent all the digits and the alphabet with pretty good clarity, as in Figure 5. The results looked something like a Lite-Brite toy, but perfectly legible. Larger, premium grids usually measured 9-by-11 or 13-by-15.

I also remember meeting the grid.

In high school, our school computer was a PDP-8/E, manufactured by Digital Equipment Corporation. It had an enormous 8-kilobyte core memory. Unfortunately, the Basic interpreter required about 4K, so that left just a few thousand bytes in which to store the text of your program; every character counted. Often several of us would tackle a project simultaneously but independently, trying to figure out the most elegant and compact solution, and learning tricks from each other along the way. I remember two of my tricks: starting my line numbers at 1 rather than 10 (thus saving one character per line), and saving words in string variables. For example, A$=" PHASER " requires 13 characters to make the assignment, but only four to use (for example, PRINT "YOUR"A$"EXPLODED"), so if you print the word "PHASER " more than twice you save characters and can write more code. We probably shared a lot of attitude with the guys who chiseled stone tablets.

We talked to the computer using two old Teletypes and saved our programs on paper tape. One of my favorite programs printed out banners on tape—you entered a message, and it punched holes in the tape to make the appropriate letters. It was basically a grid eight holes tall and as long as you liked. The nice thing about the paper tape is that it had a little ribbon of holes off-center, which engaged the teeth of a wheel in the reader to pull the tape through. This dotted line left five circular punches on one side and three on the other. This suggested three attractive designs: lower-case letters on the bottom five holes and upper-case filling all eight; using the upper three holes for ascenders (like the high part of "d"); and putting the five holes on top and leaving the lower three for descenders (such as the tail on "g"). All three designs had their charm, and I used to love to watch my banners chug out of the paper tape punch at 10 rows per second.

Several years ago in Vienna, I noticed an interesting mosaic display in the underground subway system, announcing train destinations. Figure 6 shows the display, along with the alphabet. Each segment could be individually illuminated, and glowed with a steady, flat yellow light. This display contains 66 elements and is almost symmetric. Note the addition of 45-degree edges, which provide little bevels to round out the sharp corners.

This display has a lot to recommend it: The letters look pretty nice (most of them) from a reasonable distance, where the little divots and bumps aren't so noticeable, and they're legible even from far away. The line weight is good, and there's no possible confusion among any of the letters or numbers. Some letters work great, the S and R particularly. However, the W and M aren't much to cheer about, and the V is a disaster. But still, the signs led me to my destination with utter clarity.

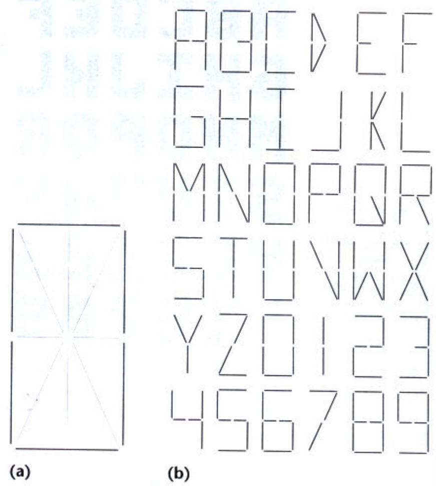

4 (a) The 16-segment display. (b) The alphabet and digits. Note the really lame B.

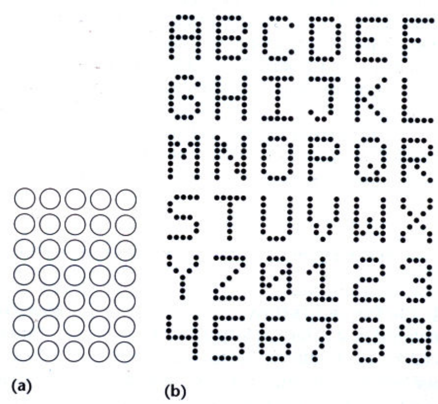

5 (a) The 5-by-7 grid. (b) The alphabet and digits.

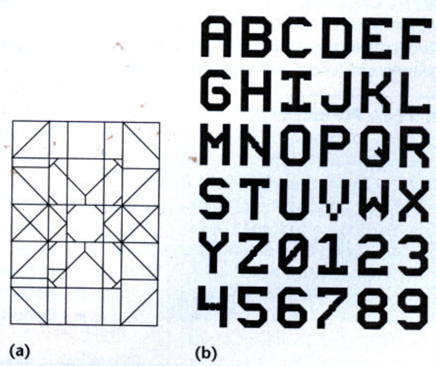

6 (a) The 66-element Vienna Underground mosaic. (b) The alphabet and digits. The V is quite unfortunate.

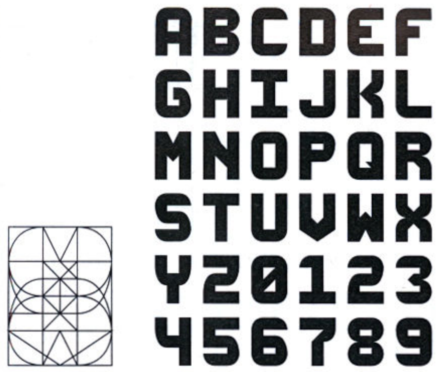

(a) (b)

7 **(a) My 55-element mosaic. (b) My alphabet and digits.**

Building my own

Just for fun, I tried cooking up my own display. I wanted to make letters with rounded forms, using fewer than the 66 digits in the Vienna display. The results of this amateur font design appear in Figure 7. On the up side, I managed to do it with only 55 segments. On the down side, my letters look much harsher than the Vienna letters, rather like a bold-style typeface. This might not be a bad thing, but I would have preferred a more delicate touch. Even so, most of my characters look pretty good to me.

This turned out to be a very interesting exercise. I started by making a list of the most important letters and letting them guide the design. Because my native tongue is English, I started with the English letter-frequency table: E is the most common letter in the alphabet, then the table continues with TAOINSHRDLU. These had to look really good, even if at the expense of infrequent letters. I tried to compromise carefully. For example, the two curved arms of the K aren't well distinguished on the right hand side of the letter, which could be a problem, but the sharp point where they join the left-hand vertical saves things and makes it a recognizable K.

It's fun to find various characteristics of letters that you can sacrifice or reinforce to keep them legible. As another example, I could have saved six segments if I'd left off the two diagonal lines at the bottom—they are used only for the V. It was tempting, but then I'd have ended up with the same embarrassing V as in the Vienna design. Why should V always be the one to suffer? (Historically, V has seniority; the letter forms U and W both derived from the V in the Middle Ages). I felt that the sharp point at the bottom was the essence of V in this design. So, I jumped from 49 to 55 segments—three on each side. Symmetry is both an ally and a nemesis in this problem.

I encourage you to come up with your own tiled character designs; you might enjoy the process as much as I did. If you make a really nice one, send me an Adobe Illustrator or Postscript version of the character format and the alphabet (with the digits). The best designs will appear in a future column. ■

Contact Glassner at Microsoft Research, One Microsoft Way, Redmond, WA 98052-6399, e-mail glassner@microsoft.com.

VRAIS 98 Call for Participation

The IEEE 1998 Virtual Reality Annual International Symposium will take place in Atlanta, Georgia on March 14-18, 1998. Paper submissions are due by **September 1, 1997**, as are panels, tutorials, workshop proposals, and posters. Larry Hodges of the Georgia Institute of Technology is conference chair (hodges@cc.gatech.edu). Program co-chairs are Grigore Burdea of Rutgers University (burdea@caip.rutgers.edu) and Susumu Tachi of the University of Tokyo (tachi@star.t.u-tokyo.ac.jp). Also go to http://www.eece.unm.edu/eece/conf/vrais on the Web.

Visualization 97 in Phoenix

IEEE Visualization 97, sponsored by the IEEE Computer Society Technical Committee on Computer Graphics in cooperation with ACM Siggraph, will be held **October 19-24, 1997** at the Sheraton Crescent Hotel in Phoenix, AZ. For conference information, consult the Web at http://www.erc.msstate.edu/vis97 or contact one of the conference co-chairs: Robert Moorhead (rjm@erc.msstate.edu) and Nancy Johnston (NEJohnston@lbl.gov).

Pacific Graphics 97 in Seoul

Pacific Graphics 97 will take place in Seoul, Korea, **October 13-16, 1997**. The yearly conference covers advances in computer graphics and applications. Conference co-chairs are C.M. Park of Postech and G.R. Oh of Seri. Program co-chairs are Yeong Gil Shin of Seoul National University (pg97@cglab.snu.ac.kr) and James Kwangjune Hahn of the George Washington University (hahn@seas.gwu.edu). Also consult the Web site http://cglab.snu.ac.kr/pg97.

Euro-American Visualization Workshop

The Codata Euro-American Workshop on Visualization of Information and Data asks, "Where are we, and where do we go from here?" To be held **June 24-25, 1997** in Paris, France, the workshop will focus on visualizing information and how users interact with digital information. Both English and French contributions are accepted, but presentations will mostly be in English because of international attendence. Contact workshop co-chairs Nahum Gershon at gershon@mitre.org and Jacques-Emile Dubois at dubois@paris7.jussieu.fr for more information.