MICROSOFT AND OTHERS ARE EXPLOITING LCD SUBPIXEL CHARACTERISTICS IN AN ATTEMPT TO ECONOMICALLY IMPROVE THE DISPLAY QUALITY OF TYPEFACES AND OTHER FINE-DETAIL IMAGES. HERE'S THE THEORY BEHIND THE IMPLEMENTATION AND THE REALITY BEHIND THE HYPE, ALONG WITH A DOSE OF CONTROVERSY.

Display technology's results are compelling, but legacy is un"clear"

By Brian Dipert, Technical Editor

When, at Comdex 1998, Microsoft unveiled its ClearType program, the public response was predictable. Some self-proclaimed pundits claimed that the technology was overrated. Others pointed to work in the academic and business worlds that appeared to both predate and conceptually overlap Microsoft's "revolutionary technology." And a few other companies quietly but quickly ramped up their own display-enhancement programs.

Microsoft's actions since Comdex 1998 have been equally predictable. Typical of the "vaporware" tradition that's been a part of corporate lore since the earliest days, ClearType white papers and presentation slides have taken two years to transform into real-life, shipping products. In Microsoft's defense, its ClearType group's development time frames were not the sole defining factor of the products' implementation schedules. Those schedules also required synchronization with the roll-out of Windows CE 3.0 and supporting hardware and with the availability of e-book titles from publishing partners.

By late summer, Microsoft has released versions of its Reader software for the Pocket PC platform and for Windows 9x, ME, NT 4.0, and 2000. To the company's credit, the results look great on my notebook PC's LCD screen. And, with Adobe's competing CoolType under development, font-enhancement technology may soon be available for operating systems besides those that emerge from Redmond, WA.

SCREENS AND STRIPES

In today's most common LCD-screen layout, each pixel comprises one red, one green, and one blue subpixel (**Figure 1**). An 800×600-pixel screen, for example, actually contains 2400 subpixels per row. A viewer's eye and brain blend the three subpixel's information to create black (no color), white (all col-

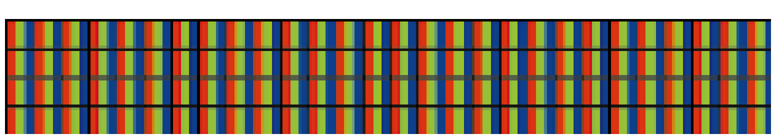

Figure 1 Each LCD pixel actually consists of red, blue, and green subpixels with same-color subpixels combining to form horizontal or vertical stripes (courtesy Gibson Research).

ors), or any color variation in between. **Figure 1's** pattern is RGB (red, green, blue), and successive rows' red pixels, for example, combine to form vertical red stripes. Other LCD variants, notably Apple's notebook PCs, employ a GBR pattern or arrange same-color subpixels in a horizontal-stripe configuration.

A brute-force-rendered font appears jagged when the display (or print, in the case of a laser or ink-jet printer) processor considers the LCD pixels in a binary fashion (**Figure 2a**). Edge antialiasing is a now-commonplace first-generation approach to improving image quality that first appeared, for example, in Microsoft's operating systems with the release of Windows 95's Plus Pack "font-smoothing" feature. Edge antialiasing goes beyond a simple "black" or "white" representation and instead gives each pixel a color tint or gray shade commensurate with the percentage of each color from the image contained within each pixel's boundaries (**Figure 2b**).

Antialiasing goes a long way toward improving perceived font quality, particularly when you view the display at a distance and at high-resolution settings. Look closely, or set the display to a low-resolution mode, however, and you notice that antialiasing has replaced the jagged edges with blurred edges. One seemingly obvious approach to eliminating jagged edges without introducing unacceptable blurring might be to simply increase the native LCD resolution (number of pixels per row and column) and correspondingly decrease the pixel pitch (the spacing between subpixels and therefore pixels) for a given-sized screen. Unfortunately, just as with other semiconductor-based devices, such as DRAM, cost exponentially increases, and yield exponentially decreases as the number of pixels grows (analogous to higher transistor counts) and as the pixel spacing narrows (analogous to smaller transistor dimensions).

ClearType and its peers exploit the higher subpixel resolution in one screen dimension (**Figure 2c**). The technique has an advantage if most of the chosen font features would benefit from finer detail in this same dimension. For example, as former IBM researcher Ron Feigenblatt explains in his ClearType Web site, Western sans serif fonts, such as Arial and Helvetica, contain many more vertical- than horizontal-line segment alternatives (see **sidebar** "Demos and documentation"). These fonts could harness the additional detail that increased horizontal resolution affords, such as that offered by vertically striped subpixel displays. Many serif fonts, such as Times Roman, contain additional decorative details that also benefit from greater horizontal resolution, as do bold and italic fonts.

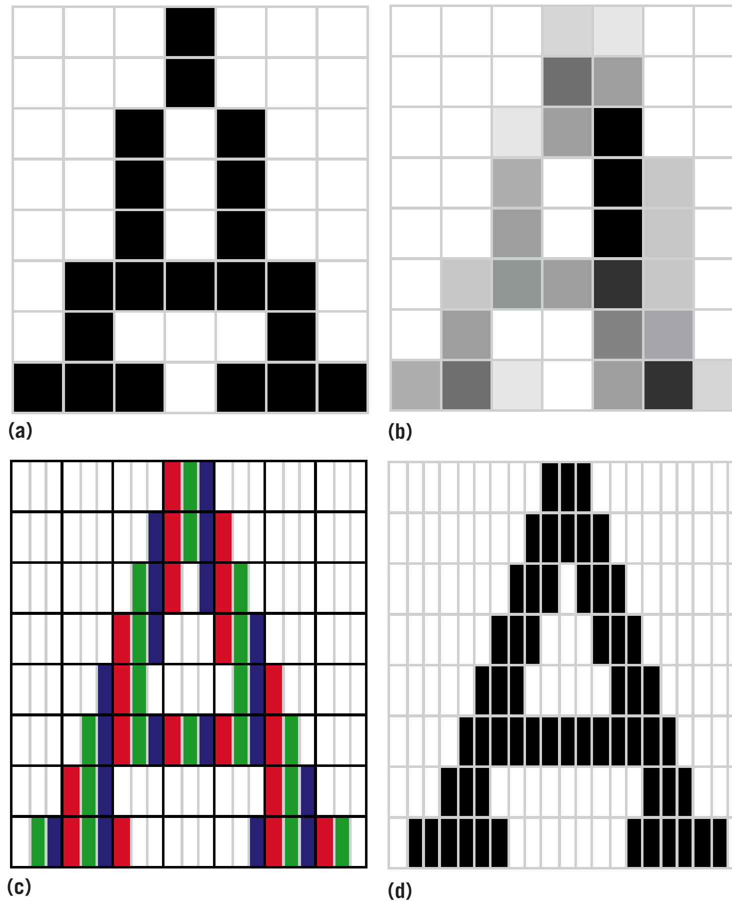

Figure 2 Pixel-based edge antialiasing (b) reduces the jaggedness of conventionally rendered fonts (a) and other sharp-edged objects but introduces sometimes-unacceptable edge softening in the process. Subpixel rendering (c) delivers a sharper edged approximation of a true higher resolution but far more expensive, fine-pitch display (d) (courtesy Gibson Research).

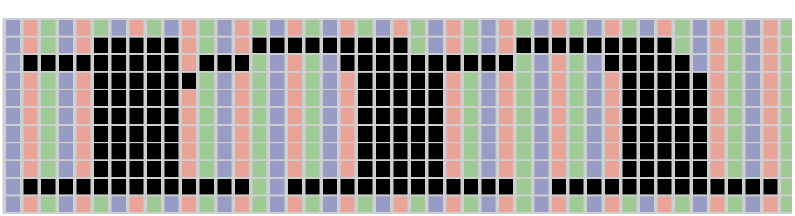

Figure 3 The sometimes-nontriplet subpixel combinations that result from ClearType and similar approaches can cause edge-color fringes (courtesy Gibson Research).

Looking beyond the font itself, the *kern*, the portion of a typeface that projects beyond the body or shank of a character, is another area that's amenable to the fine-tuning that increased subpixel resolution supports. Microsoft points out that the whole point of ClearType is to enable not only printlike character densities but also a paperlike, smooth, no-eyestrain reading style. Studies suggest that the eyes move four to five times when scanning across each line of text. Any more interruption than that amount makes reading tedious; shortly thereafter, a user typically gives up and sends the text to the printer for offline reading. For the computer display to become the future "paper," the user experience must be at least as good as that from a book or a 300-dpi printer output.

Although exploiting subpixel techniques can make more display resolution available, these subpixels are single-colored (**Figure 2d**). Where less-than-complete RGB triplets exist, they can create annoying local-color imbalances, or "fringes," at borders between different-colored objects (**Figure 3**). These transitions, although less critical than excessively bright or dark luminance values, are particularly noticeable in portions of the color spectrum to which the eye is most sensitive.

Subpixel rendering overcomes the local color imbalance problem by proportionally spreading the "extra" subpixels' energy across neighboring subpixels, rebalancing any local discoloration while still keeping most of the luminance energy localized (**Figure 4a**). One downside of this rebalancing, aside from the processing time for correction, is that it creates a slight blurring of the object-to-object border, which is at odds with the additional resolution potential of the approach. However, this blurring is often less disagreeable than that seen with the full-pixel edge-antialiasing technique.

SYSTEM SUGGESTIONS

Several key system characteristics require your attention to maximize the rendering results. First, ClearType and subpixel alternatives work only with LCD screens, not with CRTs. The LCD's direct-addressing capability enables the system to precisely control the luminance and chrominance (hue and saturation) value of each subpixel, whereas, with a CRT, the image resolution and screen resolution are decoupled from each other. A CRT image's physical placement changes depending on the computer monitor or television's horizontal and vertical sizing and centering settings, as well as in response to other user-configured controls with which the rest of the system has little to no knowledge or influence.

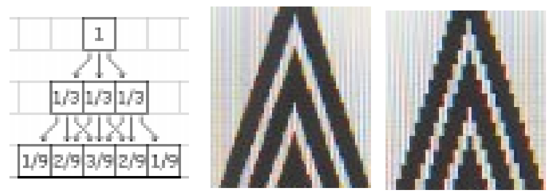

(a) (b) (c)

Figure 4 Compensating for color-fringe aberration is computationally complex (a), but the results (b) speak for themselves when you compare them with a display that doesn't harness subpixel-added resolution (c) (courtesy Gibson Research).

CRTs frequently exhibit uneven brightness, which is far more unpleasant to the eye than inaccurate color, across the screen. And, for each perceived image "pixel," the CRT's electron beam actually illuminates multiple "shadow-box" holes and corresponding phosphors. Pixel brightness and color are primarily a function of the center phosphor, but surrounding phosphors also have a secondary effect. *Misconvergence*, the inability of the monitor to correctly align red, green, and blue components of an image

DEMOS AND DOCUMENTATION

If you're interested in learning more about advanced font display techniques, Steve Gibson's *Sub-Pixel Font Rendering Technology* site (www.grc.com/cleartype.htm) should be your first stop. Gibson covers the topic in depth, and he includes a clever demo program, *Free and Clear*, that lets you investigate subpixel-rendering results across a range of typefaces and subpixel-energy-distribution-filter settings.

Gibson also provides a number of useful links to other information sources. One of the best is "Ron Feigenblatt's Remarks Microsoft ClearType" (www.geocities.com/SiliconValley/Ridge/6664/ClearType.html). Feigenblatt, while working at IBM, did much of the early research work on what became known as subpixel rendering. He, like Gibson, supports Microsoft's activities in this area, a refreshing stance.

You can also learn a lot from a visit to Ductus' Web site (www.ductus.com). The company includes lots of before-and-after images comparing nonantialiased, conventional antialiased, and Ductus' proprietary ClearView antialiased versions of fonts and graphics. Ductus also provides a Java-powered antialiased demo application with plenty of information and interaction opportunities.

And now for Microsoft. You can download your own version of the *Reader* program for notebook PCs from www.microsoft.com/reader and then visit Barnes and Noble's Web site for *Reader*-compatible e-books, some of which you can download for free. From the *Reader* site, you can also access a nearly hour-long, Windows Media-powered presentation on ClearType from Bill Hill, research scientist for e-books. And the Web site (www.microsoft.com/~jplatt/ClearType/default.htm) of Jim Platt, another Microsoft research scientist, provides a good jumping-off point to access two detailed ClearType published papers (from IEEE Signal Processing and the Society for Information Displays Symposium), image samples, and other information.

on the screen, is often visible as fringes of color at the edge of the screen or as color around text or graphics where it should be white. A subpixel-rendering-enabled application on a CRT, though, is still better than an uncorrected alternative, because the subpixel technique is analogous, at the pixel level, to edge antialiasing. Also, the types of portable systems that ClearType and alternatives target probably wouldn't normally use a CRT.

The interface between the system and the screen should ideally be all-digital without any color-accuracy-degrading digital-to-analog and analog-to-digital conversions, such as those that occur with a conventional VGA cable. Microsoft points out that this recommendation is less critical than the CRT-versus-LCD differentiation, because analog interfaces are delivering increasingly higher quality. In a system with an integrated display, such as a notebook PC or palmtop computer, an all-digital connection is almost always present (such as a LVDS interface). And with the emergence of DVI, a pure digital interface to external LCD monitor is also available.

If the system incorrectly guesses the subpixel proximity ordering, RGB or BGR, the results can be unpleasant (**Figure 5**). Similarly, the system must know whether the striping is horizontal or vertical so that it can correctly determine in which dimension the additional subpixel resolution exists. For a system with an integrated display, such as a single-function e-book reader, you can hard-code these settings into the system software. A more general-purpose product, such as Microsoft's PC-targeted Reader, which must be compatible with a range of both integrated and external displays, would typically provide a one-time user-configuration utility.

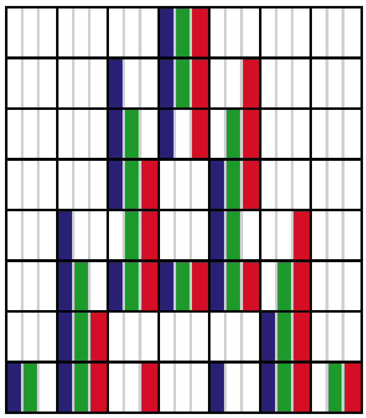

Figure 5 If the subpixel algorithm engine misidentifies the RGB or BGR subpixel striping pattern, the results can be quite disagreeable (courtesy Gibson Research).

Subpixel rendering might at first glance seem ideal for the display of detailed images on a television, such as with a WebTV or another advanced set-top box. Unfortunately, Microsoft doesn't expect the technology to provide as compelling results in this area as the results achievable with LCDs. The main limitations are neither the analog-composite, S-video, or component-video interconnect nor the low-resolution interlaced CRT-display technique that TVs use. The primary culprit is the notoriously color-inaccurate TV signal. As an indication of this constraint, you may have heard people translate "NTSC" as "never the same color twice." Ironically, though, Apple more than two decades ago used subpixel techniques with its TV-connected Apple II computer (see **sidebar "**Creation controversy").

WHAT'S THE CATCH?

Microsoft's documentation asserts that "ClearType delivers as much as 300% higher resolution than conventional font rendering, the same quality as antialiasing at 25% smaller font size and 1.8 times more text in a given resolution screen." If these claims have your marketing-hype alarm ringing, I can only commend you on your prescience. Subpixel rendering's results can significantly vary, depending on a number of factors, and detractors' grumbling that Microsoft presented ClearType in its best possible light during the '98 Comdex unveiling are in some respects valid.

Black objects, representing fully off pixels and, therefore, subpixels, on a white background or vice versa, represents the best color mixture for subpixel rendering. Fortunately, the black-on-white scheme is also the most common print-on-paper combination. The chosen font; its size; its vertical-versus-horizontal line-segment ratio; and the presence of serif, bold, italic, and other details have significant influence on the perceived improvement that subpixel rendering delivers. For these and other reasons, Microsoft has developed special ClearType-optimized fonts and is encouraging e-book publishers to use them.

Eastern fonts, such as kanji, for example, contain more horizontal lines than do Western fonts and would gain proportionally less benefit from a vertically striped display. Another important Eastern versus Western distinction is that, when successive character viewing occurs in a left-to-right or right-to-left fashion, horizontal kerning accuracy is most important. What, though, of some Asian languages, for example, which read from the top to the bottom of a page? In this case, increased horizontal resolution would be of little benefit.

The relationship is critical between where font and kerning characteristics can exploit additional resolution if available and where the display's subpixel striping provides it. Match Microsoft's *Reader* with a vertically striped and landscape-oriented notebook PC, and you'll take advantage of the available horizontal subpixel resolution. However, you'll be viewing the portrait-oriented "pages" only in the center of the landscape-oriented screen, thereby wasting overall resolution by not using all the pixels that would be available if you turned the computer 90° (**Figure 6**).

Rotate the notebook-PC display, though, and you create a horizontally striped LCD whose increased subpixel resolution Western typefaces and kerning can't fully exploit. Hewlett-Packard made some interesting trade-offs in this area when developing its latest generation Jornado Pocket PCs. The company chose a graphics controller that supports only 2^8 (256) colors, adequate for displaying simple graphical items, such as icons, but inferior to 2^{16}-color competitors' displays for showcasing still and video images. However, HP, unlike its competitors, chose to use a 320×240-pixel horizontally striped LCD. In the 240×320-pixel portrait orientation common to all Pocket PCs, the now vertically striped display delivers outstanding ClearType text quality. The HP results also indicate that ClearType doesn't require a 16- or 24-bit color display to work its subpixel magic.

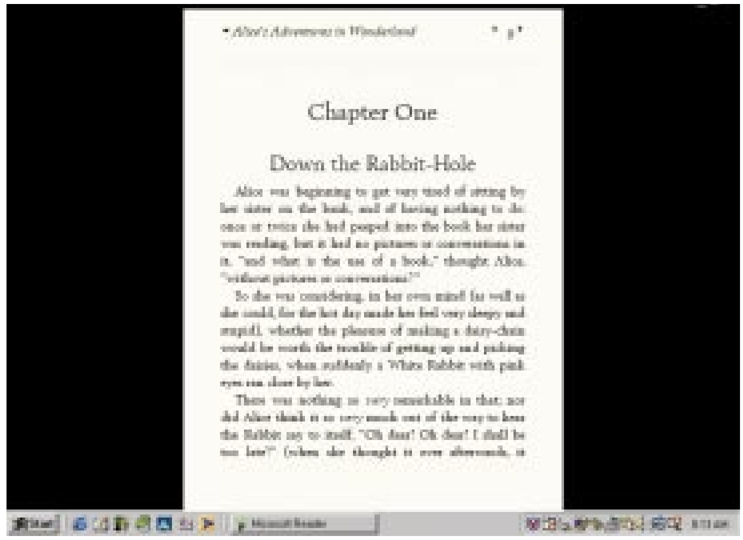

Figure 6 **Couple Microsoft's *Reader* software with a conventional vertically striped, landscape-oriented LCD, and you can harness subpixel resolution with Western fonts and left-to-right or right-to-left reading patterns, but displaying portrait-oriented pages leaves a lot of pixel landscape unused.**

You may be familiar with 3dfx Interactive's aggressive marketing of its latest graphics chips' hardware-assisted full-scene-antialiasing capability. The company's competitors claim that full-scene antialiasing unacceptably degrades frame-rate performance and that you should instead just run the display at a higher resolution where the "jaggies" won't be visible. You might remember this trade-off tug of war as you evaluate the similarly contentious system impact of subpixel rendering.

Ductus, a provider of more traditional pixel-based edge antialiasing hardware and software, is one of ClearType's most vocal detractors. The company claims that Microsoft should be putting all the energy it's currently expending on ClearType into making its operating systems' icons, buttons, dialogue boxes, and other graphical elements resolution-independent. Ductus asserts that the industry would be better served not by Microsoft's "Band-Aid" ClearType resolution enhancements but by us-

CREATION CONTROVERSY

Microsoft has a long-debated reputation as a company that relies more on adopting and mass-marketing techniques developed elsewhere than in creating its own innovations. Because of this reputation, skeptics immediately dissect any Microsoft announcement of a breakthrough technology. Such was the case with the ClearType project, particularly after patent and documentation searches uncovered subpixel-like programming that the Apple II used more than two decades ago. Apple, of course, is the company from whom Microsoft supposedly stole ideas for the graphical user interface, the mouse, and other features, further feeding the ClearType critics' flames.

Apple's technique took advantage of the fact that the Apple II hooked up to a television and harnessed a characteristic of the NTSC color subcarrier, which created a left-to-right allotment of available colors. The Apple II's highest resolution mode, 2803192 pixels, comprised 140 horizontal elements per row, each divided into green and purple subpixels. One reason for this approach was to extract the highest possible quality color out of the NTSC video signal. Designers could also use the approach to combine adjacent nonwhite subpixels and increase the effective display resolution.

Another frequently touted example of subpixel "prior art" (a term doubtlessly familiar to those of you who've ever been involved in patent applications) is the work that Ron Feigenblatt did while working for IBM. He accurately points out the various examples of subpixel-rendering research and implementation that preceded Microsoft, a list that also included Honeywell, Xerox and the big-screen Mitsubishi Diamond Vision systems popular in athletic event stadiums.

Feigenblatt is careful, though, to distinguish between an "idea" and an "implementation": Legal protection cannot cover just an idea; the idea must also include a practical implementation. Microsoft's ClearType probably to at least some extent stands on the shoulders of earlier academic and commercial subpixel pioneers. However, the fact that the company targeted LCDs, whereas earlier work focused on CRTs, may alone be sufficient cause to approve its in-process patent applications. But that's for the lawyers to decide.

ing 150- or 200-dpi rather than today's 72-dpi LCDs. Conveniently for Ductus, this approach requires a rendering technology no more advanced than the company's edge antialiasing.

I believe that Ductus' arguments are fundamentally self-serving. Despite its limitations, subpixel rendering provides a practical and effective bridge from today's LCD technology to a future when large, fine-pitch LCDs will be affordable to the masses. Ductus also ignores the earlier-mentioned economic reality (particularly in low-volume initial production) that LCDs share with other mass-fabricated semiconductor devices.

Some criticisms of subpixel rendering are, however, valid. Current implementations run exclusively in software on the host CPU before handing off the pixels' color information to the LCD controller. Users expect that the display will quickly respond to their inputs, in spite of the additional subpixel-rendering overhead. This requirement places additional processing cost and power consumption impacts on the system. Chip vendors, however, have already moved 2- and 3-D graphics functions that systems originally implemented in software to dedicated hardware in the graphics subsystem. Similarly, graphics chips that now accelerate conventional edge antialiasing in hardware may in the future add subpixel rendering to their repertoires.

Keep in mind, too, that Microsoft will most likely mate ClearType (just as with other Microsoft applications, such as Internet Explorer and Windows Media Technologies) to a limited set of operating systems, most of which also will come from Microsoft. Diverse operating-system portability is a compelling selling point of Adobe's upcoming CoolType subpixel-rendering technology. Adobe plans to build CoolType support into both its Acrobat and its recently acquired Glassbook Readers. This strategy assumes that the company can survive any legal challenge that might come from its Redmond nemesis. The Free and Clear subpixel-rendering program from Steve Gibson of Gibson Research demonstrates that not only large companies, but also individuals working under the open-source software model can implement the theory behind subpixel rendering—assuming that those individuals can navigate the patent land mine.□

ACRONYMS

DVI: Digital Visual Interface
LVDS: low-voltage differential signaling
NTSC: National Television Systems Committee
VGA: video graphics adapter

You can reach Technical Editor Brian Dipert at 1-916-454-5242, fax 1-916-454-5101, e-mail bdipert@pacbell.net.

ACKNOWLEDGMENTS

Particular thanks go to Steve Gibson from Gibson Research and to Bill Hill and John Platt from Microsoft Research for their extensive assistance and for the graphics that Gibson provided.

FOR MORE INFORMATION...

For more information on products such as those discussed in this article, go to our information-request page at www.rscahners.ims.ca/ednmag/. When you contact any of the following manufacturers directly, please let them know you read about their products in *EDN*.

Adobe Systems
1-408-536-6000
www.adobe.com
Enter No. 335

Ductus
1-650-968-7974
www.ductus.com
Enter No. 336

Microsoft
1-425-882-8080
www.microsoft.com
Enter No. 337

OTHER COMPANIES MENTIONED IN THIS ARTICLE

3dfx Interactive
www.3dfx.com
Apple
www.apple.com
Barnes & Noble
www.bn.com
Casio
www.casio.com
Compaq
www.compaq.com
Gibson Research
www.grc.com
Honeywell
www.honeywell.com
Hewlett-Packard
www.hp.com
IBM
www.ibm.com
Mitsubishi
www.amasis.com/diamondvision
Xerox
www.xerox.com

SUPER INFO NUMBER
For more information on the products available from all of the vendors listed in this box, enter No. 338 at www.rscahners.ims.ca/ednmag/.