Advances in Image Pickup and Display, Vol. 4, Ed. by Bernard Kazan

Book Review

Reviewed by Raymond L. Hallows, Jr.

This fourth in the Academic Press series of volumes on specialized imaging and display techniques comprises treatments of three major display subjects. Two of its chapters deal with passive imaging displays, namely, liquid-crystal matrices and cathodochromics. The third chapter presents a theoretical treatment of bilevel display image-processing that applies more appropriately to displays having relatively limited gray-scale capabilities.

This volume, like the previous volumes in the series, would probably be of greatest use to the general motion-picture or television technologist as a single-source reference on subjects of peripheral familiarity.

Edited by Dr. Bernard Kazan of Xerox Corp., it gives theoretical background and practical descriptions of the most successful examples of their kind, rather than attempting a comprehensive coverage of other developments that would be of interest to the specialist. I found this volume well suited to my interest level as a generalist. I found it necessary, however, to pursue a selected reference from the long reference list for the first chapter, in order both to answer my questions on points that had been given only cursory treatment, and also to assess the considerable additional depth and breadth of treatment that the references provide for the specialist.

Liquid-Crystal Matrix Displays

This first subject, "Liquid-Crystal Matrix Displays," written by Eiji Kaneko of the Hitachi Research Laboratory, treats salient features of rapidly developed technology that has resulted in the profusion of watch dial and hand-held calculator displays which are viewed by external light, with their attending low power and voltage requirements. With these attractive features conducive to portability, it is natural to expect that liquid-crystal displays would by now be competing for their share of the television and soft-copy display market. As Dr. Kazan puts it, "their good performance ... has stimulated extensive research efforts at extending their information capacity and at making them useful for displaying graphic and television images. Unfortunately, ... attempts to develop such displays using conventional addressing schemes have resulted in seriously reduced image contrast and limited viewing angles .... These problems are discussed in some detail, taking into account the electrical and optical properties of liquid-crystal materials. ... schemes [are also presented] for overcoming these limitations, followed by a discussion of these new devices."

The sections, ranging from operating principles of the reflective and transmissive versions of the twisted nematic field effect (TNFE) cell, preferred to cholesteric and smectic types, gave a practical, albeit partial, understanding of various properties affecting liquid-crystal cell operating characteristics. One of these is voltage threshold as a function of operating temperatures and frequency effects on the dielectric constant as measured in directions parallel and normal to the liquid crystal molecular alignment. This seems to be a less than ideal sequence in organizing the facts contained in this work.

For instance, the first time the reader is made aware that liquid-crystals are operated by ac voltage, even for steady-state conditions, is in a statement that "... their transmission is determined over a broad range of frequencies by the value of the rms voltage applied, irrespective of the shape of the voltage wave." This seems, even with the reference to supporting work, inadequately to answer the natural question, "Why is it apparently not feasible to operate liquid-crystal cells with steady-state voltages corresponding to the desired open and closed states?" That is, it would seem more instructive to provide a static characteristic of relative light transmission versus voltage for comparison with the presented dynamic plot, which shows alternating exponential build-up and decay imposed upon the switching step function.

It turned out that one of the reference articles, by Kmetz, entitled "Liquid-Crystal Display Prospects in Perspective," provided the apparent answer that static dc operation is accompanied by drastically life-shortening electrochemical reactions. It further turned out that this irreversible process was described for dynamic scattering displays which preceded the twisted nematic development. The present work mentions dynamic scattering by way of introduction only, giving extended treatment to the second-generation twisted nematic and induced birefringence devices.

Thus, it appears that the author and editor have tacitly assumed some of the background material, so well known to themselves, also to be common knowledge among their readers. Anticlimactically and almost parenthetically, about seven pages after the section is completed on Operating Principles (dealing with electro-optical response and frequency dependence) without mention of dc operation, the statement is made under a new section heading, "Amplitude-Selective X-Y Addressed Liquid-Crystal Matrix Displays," that "in order to avoid chemical etching of the electrodes, ac voltage pulses are usually used to drive liquid cells."

Other apparent editorial deficiencies, not obvious as such, at first impeded my grasp of the principles under discussion. For instance, Fig. 13 (a graph of the voltage margin ratio, α_{max}, applied to selected versus non-selected pixels as a function of the number of horizontal electrodes), shows a steep drop from $\alpha_{max} = 2$ for the minimum number (two) of electrodes, with a maximum change of slope for about ten electrodes ($\alpha_{max} = 1.35$) and a further asymptotic approach to $\alpha_{max} = 1.1$ for values above 100 to 150 electrodes. The text says that "... it is very difficult to produce satisfactory images if a matrix liquid-crystal panel has more than 100 horizontal electrodes." Since α_{max} is essentially constant for values between 100 and 150, I wondered why 150 electrodes could not be quoted as the upper practical limit, if 100 is feasible. The value of ten electrodes was seen to

---

Review of a book published 1981 by Academic Press, Inc., 111 Fifth Ave., New York, NY 10003. Reviewed by Raymond L. Hallows, Jr., Princeton, N.J. This review was received February 16, 1983.

be the point at which acceptable voltage margins diminished most rapidly and thus offered a more logical practical limit if other factors not shown on the graph did not influence the choice. Also, a 2:1 suppression of the α_{max} ordinate scale was noted, possibly to emphasize its critical nature as an influence on contrast. If so, good engineering practice would require that a break in the ordinate be used to call attention to the missing values between 0 and 1.0.

Possibly these apparent anomalies can be attributed to editorial or typographic errors. The text 13 pages later seems to clear up the first point by stating, "As previously noted, the contrast ratio . . . decreases greatly, particularly if the number of scanning electrodes exceeds 10 lines." As mentioned, the zero suppression should be noted by a zig-zag in the ordinal axis, but the violation is less serious if the zero at the abscissa level was intended to be shown as a 1.0.

These seemingly trivial apparent errors, because they are not immediately obvious as such, detracted considerably from my ease in grasping salient points in this complex subject; I hope that they will be corrected or clarified in a future edition along with the more evident errata, such as occasional grammatical disagreements in number between subject and verb.

I must also comment that although much attention is given to the complex interactions influencing the viewing contrast ratio, the only means I found of "getting calibrated" as to the contrast ratios actually achieved was in separate tables listing salient characteristics of various forms of liquid-crystal displays in a section entitled "Matrix Devices for Character Display." These notations vary from 5:1 for a reflective matrix character display to 28:1 for a thin-film transistor (TFT)-addressed display panel. The preceding "Operating Principles" section addresses various factors affecting contrast ratio of transmissive and reflective cells under such topics as crosstalk, amplitude-selective addressing, viewing angle dependence, multiplexing waveforms double-matrix panels and their driving waveforms and higher-order multiple-matrix panels. It speaks only qualitatively of contrast ratio effects, but gives quantitative curves and formulas on some of the above interrelated factors. The complexity of the interactions can readily be appreciated by a reading of the discussions, and I was left with a feeling of having "been there" rather than of having read from a neatly distilled design handbook from which the desired contrast ratio could be cranked out in a paper design.

Double Matrix Displays

The section on Double Matrix Displays for television would undoubtedly be of special interest to TV-oriented *Journal* readers. It gives an overview of well-known digital TV and data-handling techniques that are adapted to the special requirements of multiple-matrix panels to provide two kinds of waveforms, respectively, to the selected and nonselected scanning electrodes, and many pulse-width-modulated waveforms to the signal electrodes to establish the required gray levels. Photographic examples are given of the quality obtainable on double- and quad-matrix display panels, for which only about $^1/_{16}$ of the available NTSC pixels are displayed.

Other techniques are described for special applications such as enhancing contrast by using two driving frequencies, one of which serves to select the desired pixel and the other to apply an electric torque of negative sign to nonselected pixels, taking advantage of the frequency dependence of liquid-crystal materials. For long-term storage applications, phase change from the clear to the scattering mode can be obtained in certain mixtures of nematic and cholesteric materials for which the scattering state is retained for hours after the excitation voltage is removed.

Nonlinear element techniques are reported, using diodes, double-diodes for symmetrical thresholding, or varistors at the intersections of the scanning and signal electrodes. These techniques result in a sharper threshold for each pixel, thereby lessening the limitation to the number of scanning electrodes that may be used without crosstalk and consequent contrast limitation.

Field-effect-transistor and MOS-FET switching extend this concept to three-terminal devices, allowing switching-mode operation for storage between TV frames and charging within a TV line interval to approach a 100% duty ratio for each pixel. The problems of obtaining MOS transistor arrays on large-area silicon wafers are alleviated by a promising technique described, which uses insulated-gate, amorphous silicon FETs. The constructional details with diagrams and discussion are also given for those devices, with photos of the displays in most cases.

Finally, the chapter ventures into developmental devices for multicolor displays using the phenomenon of electronically controlled birefringence, the most promising of which, from a stability standpoint, are the deformation-of-vertical-aligned-phases (DAP) and the hybrid-aligned nematic (HAN). The DAP scheme, which exhibits transmission peaks in the red, green, and blue for white incident light for various applied voltages, unfortunately does not show sufficient color separation and has color variation with cell thickness, temperature, and viewing angle. The HAN scheme has low drive voltage and excellent color separation characteristics, but it suffers from having no voltage threshold to facilitate addressing an X-Y matrix system. The viewing angle dependency has been overcome by designing optical projection system versions of these matrix cells.

The TNFE cell has also been adapted for two-color display with selective absorption of white light based on crossed pleochroic polarizing filters. These organic dyes absorb more light at a given wavelength when the electric vector of the polarized light is parallel to the long axes of the dye molecules than when the alignment is perpendicular. This display has low brightness due to the large absorption of the polarizers and analyzers, but has potential for obtaining field sequential two-color displays from black-and-white CRTs.

Pleochroic dye displays, in which some pleochroic dye is dissolved in a liquid-crystal material, are reported as having the advantages of a wider viewing angle and higher brightness than the TNFE cell, since the polarizer is not used. Of course, the trade-off is that contrast ratio is limited, because only a small percentage of the dye can be dissolved in the liquid-crystal material.

The concluding remark is to the effect that despite the many difficulties in liquid-crystal technology, it is safe to predict that a palm-sized liquid-crystal TV display device will be produced in the near future. The statement is supported by the final reference of this comprehensive work on the mainstream of liquid-crystal matrix technology. It is entitled "Pocket-sized TV with Liquid-Crystal Display," and was published in the 1979 IEEE Tokyo Section.

Cathodochromics

The second chapter, on "Cathodochromics," written by Faughnan, Heyman, Gorog, and Shidlovsky of RCA Laboratories, deals with electron beam excitation of passively illuminated displays. The classes of alkali halide materials treated in this work form the basis for the CRT displays that are popularly known as "dark trace" tubes.

After a historical summary of cathodochromic device developments, the types and nature of color centers caused by impurities or vacancies in the crystal structure are discussed. To be useful as a visual display, the spectral absorption of a color center should peak near the most sensitive wavelength of human vision, 500 nm, and have an absorption-curve shape complementing the eye's luminosity-curve shape.

Two modes of cathodochromism are defined — optically and thermally reversible, the optically reversible mode being bleached by visible light or heat, whereas the thermally reversible-mode material usually requires a temperature higher than that required to erase the optically reversible material. The separate mechanisms for these phenomena are discussed in terms of detailed delvings into the crystal structures and energies involving electron-ion ballistics.

Preparation methods of sodalites — sintering and hydrothermal growth — are examined, and properties of three important cathodochromic materials — potassium chloride, calcium titanate, and sodalite — are discussed in detail. A method is given for measuring the optical absorption coefficient from which the CRT screen gain may be calculated from additional measurements of diffuse contrast and the screen weight per unit area depth of penetration, which is a function of electron-beam voltage. The optical absorption coefficient is an intrinsic property of the bulk material and is independent of device configuration and accelerating voltage. Thus it is useful in making objective comparisons of various materials.

The section on cathodochromic systems starts by discussing a major system-performance limitation, the addressing time, determined by the rate of depositing the required charge to a pixel to achieve the desired contrast ratio. The discussion brings out the observation that although cathodochromic devices can achieve contrast ratios of about 100:1, the practically useful limit is around 10:1, that of an ordinary newspaper. The explanation given was that the straight-line portions of contrast ratio (CR) versus flux (ϕ) plots when related to the slope-intercept form resulted in the equation $CR = d + b \log \phi$, and that for large exposures (flux), the contrast ratio increase is essentially logarithmic. Over the same range of practical interest, the Weber-Fechner law is cited, which states that the perceptible brightness change is proportional to the brightness itself, that is, the perceived contrast is also roughly a logarithmic function of the contrast stimulus. Thus, the perceived contrast for high values varies roughly as the log ($\log \phi$), so that once a practical coloration or contrast level is achieved, further exposure will give relatively small additional perceived contrast, and additional contrast is paid for in greatly increased writing time. The discussion gives further good examples and plots of practical cathodochromic device exposures to produce contrast ratios of 4:1 and 8:1, suitable for practical systems.

Experimental Systems

The experimental systems described in some detail include a cathodochromic display data terminal using a hybrid scan format, in which a one-line shift register was unloaded at constant rate during each line-writing period. The waiting time before the next written line, however, was dependent upon the information contained in the line and therefore on the time required to decode it. This non-redundant, aperiodic, scan method was found to be well suited to the storage properties of cathodochromic devices.

Another device, using a 10-in. potassium chloride storage tube, could store as many as 6 × 6 separately written images for simultaneous display, each containing up to 1024 6-bit, or 64 gray-level, pixels.

Plan-position-indicator (PPI) radar images using the dark trace tube have been found practical for the variable and sometimes high ambient illumination levels of aircraft panels. This PPI system uses an increased beam current modulation, roughly proportional to the radial distance from the tube center, to compensate for the decreased beam overlap with increased radius of scan.

A two-way video system using a single optically erased cathodochromic CRT for three imaging functions — pickup, storage, and display — would be of special interest to TV engineers. It comprises a sandwiched screen structure having separate layers of phosphor and cathodochromic material. In operation the pictorial image is entered by optical projection onto the cathodochromic surface after it has been darkened by scanning with a uniform dark current raster operated in the optical erase mode. The projected light image causes localized bleaching of the precolored screen, to produce a positive stored image. This image can be converted to video by scanning the phosphor with a low beam current and voltage, so that the resulting phosphor light output is modulated by the variably bleached cathodochromic layer and picked up by a multiplier phototube, à la flying-spot scanner.

Other applications include a random-access alphanumeric graphic terminal and a facsimile system making use of the compressed data techniques by reduction of redundancy that normal facsimile systems cannot utilize. Slow-scan TV and computer memory and buffer storage systems, with their use in hard-copy devices, are mentioned to round out the rather broad list of special purposes that cathodochromic devices can serve.

The further progress of these devices, it was concluded, are limited by electron-beam coloring sensitivities in both optically and thermally erased sodalite versions. Sodalite is, to date, considered the most generally promising material for further research. Fatigue effects and the relatively slow rate of writing and erasing, especially in selective erase applications, remain as limitations when they are compared with competing devices using conventional CRTs and semiconductor memories. This informative, well-written treatise predicts that future cathodochromic systems will be used for their combination of unique capabilities (simplicity, long storage time, and high resolution) which make them excellent for certain radar, narrow bandwidth image transmissions, and temporary very high resolution image recordings.

Processing Images for Bilevel Digital Displays

The third, final chapter of this volume, "Processing Images for Bilevel Digital Displays," was written by C. N. Judice and R. D. Slusky of Bell Laboratories at Holmdel, N.J. It deals primarily with the requirement for gray-level appearance in hard- or soft-copy displays that are capable only of essentially bilevel operation. These devices are generally more specifically suited for text and graphics than conventional cathode-ray tubes, and frequently use elements which are addressed by rectangular matrixing techniques.

The possibility of changing the ratio of the bright-to-dark area within a given element or display site, which is exemplified by the subtractive halftone screening processes by which newspaper and magazine pictures are reproduced, cannot be used in most of the digital display devices of interest, because the element sizes are not variable.

Some rapid-response devices, such as ac gas plasma display panels, have been operated in a mode of rapidly switching a given element between "on" and "off" states with variable ratio of on- to off-time, or duty cycle, thereby controlling the average perceived brightness. Most of the devices of interest, however, do not have switching speeds that exceed the flicker threshold of our visual process, so the remaining variable to control is the distribution of the "on" and "off" display sites. The algorithms for obtaining gray levels comprise much of the substance of the authors' paper.

An overview is given, which establishes terminology applicable to both nongray-scale algorithms with one display site per picture element (pel) for single, or nonadaptive thresholding, in which the decision threshold is always a constant fraction (usually one-half) of the recordable brightness in the image being input for processing. Adaptive algorithms are then discussed, with examples given comparing the brightness of a given pel with the average brightness of surrounding pels. This strategy, called constrained averaging thresholding, allows the threshold of the processed pel to be modified to lie between the main "on" and "off" intensity clusters of the "intensity histogram," which plots the number of pels occurring at a given intensity level against the intensity range of the input image. The threshold established can thus be made greater than most of the dark, printed-text pels and smaller than most of the light, background pels. The local average can in effect "ride" the shading or other large-area non-uniformities in the image, reacting to differences between the text and background pel intensities which are statistically smaller than those possible if arbitrarily set absolute levels are compared.

Pseudogray-Scale Algorithms

The bilevel text example reviewed above is then extended to pseudogray-scale algorithms by providing the individual elements within a local pel region with random fractional thresholds between the extreme dark (0) and the extreme bright (1) limits of the input image. This algorithm by its very nature produces a picture having a grainy appearance, because the number of display sites per unit area is not great enough to insure a uniform threshold distribution.

The *periodic-random* algorithm is introduced to overcome that difficulty by spacing the thresholds uniformly across the quantization range but randomly within a fixed, normally square matrix, whereupon the boundaries of the local periodic random regions become visible. The review then progresses to the concept of *ordered dither thresholding*, whereby selected thresholds are assigned specific positions within a matrix, and the resulting dither matrix is repeated over the entire source image in the vertical and horizontal directions.

An additional provision is made to assure that the average distance between elements with successive threshold levels is as great as possible, so that the bright and dark patterns in the display are at the highest possible spatial frequency.

A digression is presented, with algorithms for coping with errors peculiar to specific reproducing processes: whereas the electronic ac plasma display panel can faithfully reproduce the foregoing dispersed-threshold algorithm, the various ink-on-paper processes have problems with ink spreading from printed areas into areas of intended white as a function of the ink dot spatial density, causing a darkening of the dithered image print. This can be remedied by arranging a *digital halftone matrix* in which the brightness thresholds of pels within each matrix area are arranged in positions so as to span the input image quantization range in ascending order, starting near the center of each repeated square matrix and progressing in spiral fashion around the square.

This strategy causes ink to spread into adjacent black dots rather than into white areas, thereby minimizing its excess darkening. This digital halftone-thresholding process produces a more uniform dot density than the ordered dither method but has an inherent loss of spatial resolution that, fortunately, can be eased by using the relatively high display-site density capability possible with the various printing methods.

Next, the authors present three adaptive algorithms that do not use the periodic threshold matrixing common to the ordered dither and digital halftone algorithms: the *constrained average* assigns a threshold to each source-image picture element as a function of its own intensity and the average intensity of its eight surrounding pels. This can give significant edge enhancement, since the thresholds are a function of a local average and strongly depend on the distribution of the involved pel intensities. *Overflow counting* and its bi-directional processing counterpart, *error diffusion*, are presented as algorithms to correct for the amount by which a pel exceeds or falls short of its threshold in a locally averaged area.

This section on pseudogray-scale algorithms introduces a second strategy for reproducing gray-scale imaging, by which each pel is associated with a cluster of sites rather than a single display site, to provide improved intensity resolution. In contrast with the previously described spatially-ordered digital halftone input thresholding of pels, the multiple display site cluster, called a halftone cell in this context, represents an output brightness quantity roughly analogous to the halftone screening process used in printing. Each cell is arranged to display patterns varying from all-sites dark to all-bright in increments of one additional site so that an $n \times n$ site cell can display $n^2 + 1$ brightness levels. The shapes of the patterns are controlled by a computer algorithm. The cell pattern having the correct number of bright states is selected according to the brightness of its corresponding source-image pel.

The authors' comprehensively treated pseudogray-scale algorithms, for which I have attempted an inclusive overview in the preceding paragraphs, are capable of being modified to allow the control of brightness and contrast, desirable in any display system. It seems evident that changing the average of the thresholds of the bilevel display will change the overall brightness. Lowering the thresholds will cause more display sites to be available for a given source image, thereby increasing its brightness. Decreasing the average separation between numerically successive thresholds will increase the displayed contrast of the gray-scale portion for which the compression occurred, and vice versa.

Thresholding

A case of local gray-scale enhancement of special interest is known as histogram-balanced thresholding. Here the numerically successive thresholds are still located in the same places within the dither or digital halftone matrix, but the threshold values are now shifted so as to delineate equal-area regions under the histogram curve. Thus, for regions of the curve with the greatest number of pels, the thresholds will be relatively close-spaced, thereby providing improved intensity resolution. The necessary decrease in contrast for intensity regions not in the vicinity of the curve maximum is not likely to be detrimental, since it involves relatively few pels.

This technique requires either that the source image be scanned twice or that it be stored to obtain the pel distribution as a function of intensity level (histogram), from which the histogram-balanced thresholds can be determined. In general, it is shown, the histogram-balanced threshold technique improves tonal rendition only if the histogram curve has a single predominant maximum. A bimodal histogram distribution of pels, peaking at two separated intensity levels, will usually result in degraded intensity resolution in the processed image.

The authors describe another esoteric treatment of pseudogray-scale images — periodic intensity matching — by which the familiar aliasing errors (moiré patterns), generated by sampling a repeating spatial pattern below its Nyquist frequency, may be minimized by adding or subtracting a fixed increment at each threshold in the matrix, so as to force the display brightness to match the source intensity of the area represented in the image. This technique was also shown to result in some cases where no improvement is obtained. Furthermore, a varying processing time is needed for each threshold matrix, thereby complicating the implementation of this asynchronous scheme.

Further modifications of the basic algorithms are explored, to provide improved pseudogray-scale algorithms for the dual purpose of reproducing bilevel images, such as text, as well as continuous-tone images.

The first of these, called adaptive threshold compression, computes the average intensity of the pels surrounding the pel to be processed. The input brightness thresholds are then compressed to span only the intensity range of the local group in a way analogous to the constrained average thresholding, for purely bilevel imaging, described earlier in this review. This technique is now seen to affect the average brightness of the image as well as the desired local contrast and edge enhancement. A correction is made that involves an equivalent modification of the source-image pel intensities and use of the thresholds obtained before compression.

The combining of this compressed dither threshold technique with the previously described periodic intensity matching can achieve a rendition that was stated to have fewer bilevel aliasing errors, greater text and graphics legibility, and pronounced edges. The degree of edge enhancement was explained as dependent upon the ratio selected for threshold compression, which must take place before intensity matching. Indeed, the authors' much-used test image, showing a building comprising glass window-panel sections reflecting a cumulus-clouded sky, and with autos parked in the foreground behind a row of shrubs, provided a combination of high-contrast edges, subtle gray shadings, and fine detail to show this process to good advantage with respect to virtually all of the previous algorithms of the series.

Algorithm Switching

The final dual-purpose technique, called algorithm switching, struck me, however, as somewhat of an anticlimax. Its purpose is to provide the best of both the gray-level and bilevel worlds. Alas, the glass building source-image appears crisscrossed with a coarse diagonal pattern at all but the lightest of gray levels in sky, windows, and automobiles; the window partitions have lost their sharp edges; and the shrubbery detail is obscured.

Algorithm switching was explained as the result of monitoring the image signal in the region surrounding the processed pel by using a Laplace operator technique. This forms a summation of second-order partial derivitives that is generally greater for a manifestly gray-scale image than for a bilevel image. The switching example combined dither thresholds for about 60% of the sites, with the remainder using digital halftone thresholds.

Let us recall that the authors had acknowledged a loss of spatial resolution inherent in digital halftone thresholding, and that halftone thresholding produces a more uniform dot texture than does ordered dither. At the display site density of 40/cm shown for nearly all of the algorithms that were illustrated, the halftone threshold disadvantage far outweighs its uniformity advantage. Perhaps the authors could have added a few specific words of explanation and not-otherwise-evident justification to help balance the unspoken 1000 words, negative in this example, that their switching illustration is proverbially said to be worth. Certainly they are to be commended for showing the various processes under the same general conditions of processing parameters and should not short-change the value of an interesting concept for the sake of objectivity.

While on the subject of the photographic examples, I must also comment on the glaring spurious patterns in the source-images of a woman's portrait processed with the constrained average algorithm (Figs. 15 and 16). These defects show as nearly vertical, regularly-spaced striations running at slightly different angles, mainly in the upper two-thirds of each photograph.

From my own participation in experiments with electronic halftones by CRT techniques, I recalled a very minute deflection superimposed upon the desired deflection to individual halftone-dot positions. This modulated the white spaces between large dots, printed as black, to a degree that "hum bars," similar in appearance to the periodic disturbance in the authors' two examples, appeared in the darker regions of the output copy. My negative "white dot" experimental process, however, was essentially immune to the same degree of spurious deflection modulation by virtue of small deflections merely moving the small white dots, used to produce the near-black tones, within the black surround without changing their area.

Possibly some such phenomenon was not considered worthy of explanation in the authors' examples. I would agree only if, unlike the example cited from my experience, the spurious modulation susceptibility is not inherent in some particular aspect of the afflicted display technique. In any case, I would suggest that such defects, always a problem and tolerable in developmental systems only if they do not obscure the main object of the demonstration, should be explained. Then, the knowledgeable but not expert reader in this rapidly developing specialized field could gain insight and appreciation for the practical process constraints.

The chapter on modifications to basic algorithms concludes with a short section on processing images for multilevel displays having an intrinsic gray-scale limitation so that contouring is still evident. Such technology might include liquid-crystal or thin-film electroluminescent devices. The discussion is extended to include color displays, such as developmental color ink jet plotters; the authors state that their basic algorithms can likewise be applied properly to encode luminance and chrominance signals. They consider this to be an area worthy of further research.

Bilevel Image Representation and Data Reduction

Judice and Slusky's next section, Bilevel Image Representation and Data Reduction, treats techniques that take advantage of the redundancy and correlation in bilevel images to allow the image to be described with a substantially reduced number of data bits. These compression techniques involve two steps — image coding and message coding. Image coding uses the unequal probability of occurrence of ones and zeros in many subjects, such as a page of printed text, to maximize entropy, or stated differently, make the probability that a given bit will be a one or a zero very close to one-half. This process will be recognized as involving *run length coding* by those who are acquainted with the rudiments of information theory.

Other techniques such as *blob encoding* and *relative address coding* extend the process to two dimensions and gain a data-reduction factor of two or three. *Chain link encoding* outlines areas having constant brightness, encoding each link in the chain. This method is described as being especially adaptable to shape analysis, image rotation, size scaling, smoothing, and editing.

The pseudogray-scale image technique, as opposed to imaging of bilevel subjects, will not work directly with these redundancy reduction techniques, because the representation of a given gray level depends upon the statistical spacing or microstructure of ones and zeros within a display site. The authors point out that the bit stream of a pseudogray-scale image comprises two separate components. One is based on the level of the image, known only from its general statistical makeup, and the other is based on the defined, and therefore exactly known, thresholding algorithm.

Bit regrouping coding is used to take advantage of the known nonadaptive pseudogray-scale algorithm and achieve data compression. Here the display bits, corresponding to pels with equal or nearly equal assigned thresholds, are grouped together to form long rows of ones or zeros, capable of being encoded for considerable data reduction. This process, likewise, may take place on a one-dimensional level by grouping bits having like thresholds with others from the same line, or in two dimensions, by grouping bits throughout the image according to their respective equal thresholds.

Block encoding, using a look-up table to generate the correct bit-code pattern, and *predictive encoding*, which uses some known statistical property of the processed pel such as the states of previously coded pels and transmits only the differences from the predicted values, are additional compression techniques treated by the authors.

Message coding, the second step in data reduction, is a broadly defined part of the vast field of data communication technology. Thus it is not as directly inherent in the process of representing images in bilevel form as the image coding methods treated in some depth by the authors and summarized above.

Here Judice and Slusky give an example of a variable-length code, using short code words to represent source-image values that occur frequently and longer code words to represent less frequent values. The efficiency of this strategem obviously depends upon how well the coding lengths and source-image value recurrence frequencies are matched. The most familiar example that I would interpret as being analogous to this requirement is the international Morse code, in which the most frequently occurring letters in English usage are represented by the shortest groups of dots and dashes, and the less frequent characters, at least to some rough extent, are formed by longer code groups.

A further section on the animation of pseudogray-scale images deals primarily with the problems of twinkling, or scintillation, that occur in devices having limited gray-scale capability. An example is the ac plasma panel, used for presenting animations in TV fashion, updated frame by frame. The scintillation results from noise-triggering of pels that happen to operate close to their assigned dither threshold levels, so that they randomly change state from one frame to the next.

Various forms of hysteresis are described. One is in the time domain, by which a change is not permitted in a given pel until it has been called for in two successive frames. Hysteresis of the dither thresholds is effected by requiring that a site in a given state must exceed its dither threshold by some predetermined value to allow switching to the opposite state to take place in the following frame.

An alternate approach is to provide cyclically varied threshold values within the dither matrix, so that over a period of, say, four frames the thresholds at any position will span the quantization range, and the time-averaged light output over the display site is proportional to the intensity of the associated source-image pel. This time-varying dither also minimizes pattern and edge-crawl effects in animations and can be adapted for reducing or eliminating flicker effects attending the use of interlaced fields.

The final section on Bilevel Imaging System Design considers four main areas: algorithm selection, display medium selection, system configurations, and animation circuitry. This review will not present details, many of which are undoubtedly well known to engineers in the image display and video processing fields where digital methods are gaining prominence.

Conclusion

My aim in presenting the somewhat labored summary-review of this scholarly treatise by eminent Bell Labs authorities has been to give readers of the *Journal* an overview and nodding acquaintance with aspects of video processing with which most of us have not dealt, and which apply to systems that, as yet, are largely outside the purview of conventional television display technology. This review has served its purpose if the reader's appetite has been whetted for the more-than-casual reading of this work that is required if one is to approach a full appreciation of its depth and precision and of the sophistication that can now underlie the presentation of imagery on display devices with less than optimum gray-scale capability.