


JANUARY 1983

COMPUTER



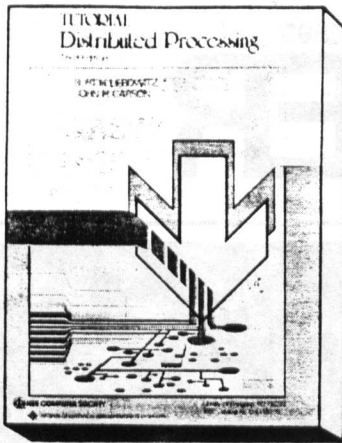
COMPUTER
ARCHITECTURES
FOR IMAGE PROCESSING



IEEE COMPUTER SOCIETY



THE INSTITUTE OF ELECTRICAL AND
ELECTRONICS ENGINEERS, INC.



The editors have gathered 64 of the finest papers on distributed processing for reprint in this volume. The first chapter is a general overview of the field; it is followed by chapters on technology, communications, intelligent terminals, networks, multiprocessor systems, and data bases. Each chapter has an introduction by either Liebowitz or Carson that gives an overview of the topic. 640 pp.

Order #363

TUTORIAL—DISTRIBUTED PROCESSING

Edited by **Burt H. Liebowitz**
and **John H. Carson**

3rd Edition, April 1981

Members—\$18.75 Non-members—\$25.00

Use order form on p. 128C

ABOUT THE COVER

The patterns in the cover design are the result of a biomedical image processing application in which white blood cells from a father, mother, and newborn are compared for evidence of genetic mutation. The 2-D electrophoresis process used in the comparison resolves the protein constituents of the human cells into a visible image, each protein being seen as a distinct spot whose position in the image corresponds to the protein's molecular weight (y-axis) and its molecular charge (x-axis).

In this study, supported by the National Institutes of Health and headed by noted geneticist James V. Neel of the University of Michigan, thousands of blood samples from family trios are examined for unusual protein variants in the blood of the child. The variants could be evidence of a parental germ cell mutation, which many believe is an early indicator of a genetically damaging radioactive or carcinogenic environment.

The study has brought together sophisticated biochemical and computer image processing techniques, which are being implemented on a device called the cytocomputer. Developed by Stanley Sternberg, whose article, "Biomedical Image Processing," appears in this issue, the cytocomputer is a pipelined neighborhood processor that successively processes the gel data, ultimately separating the gel into regions corresponding to a single spot on the pattern.

In the first stage of cytocomputer processing, background levels of the gel trios are detected and subtracted using a process based on mathematical morphology, an image-algebraic approach. The gels are then further processed to determine spot boundaries, and the boundaries are filtered smooth by a sequence of neighborhood processing steps. Finally, the cytocomputer searches for spots very close to background levels using a gel matching program. For more details on this process, see Sternberg's article on p. 22.

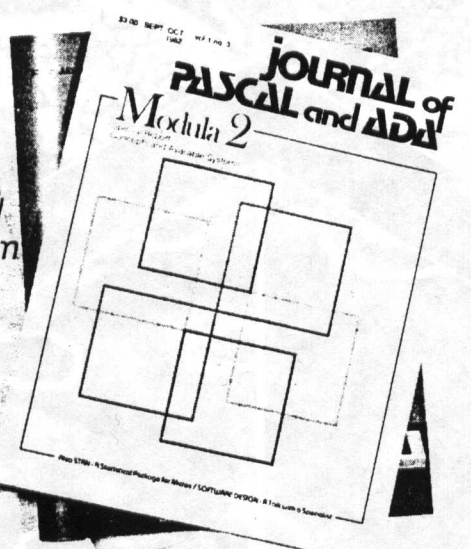
The most useful resource available!

JOURNAL of PASCAL and ADA

Subscribe to the journal that can help you program more efficiently.

No other publication consistently gives you this kind of applicable information on Pascal and Ada:

- **New developments, book reviews, and software reviews** that keep you informed about the field of Pascal and Ada computing.
- **Articles that put you in touch with other Pascal and Ada programmers.** You'll find reports on how people from many different fields are solving their programming problems. Someone's software tool, procedure, or algorithm may be just the trouble-saving solution you need. A Q/A column also allows you to get our readers' help with a sticky problem, or share a clever solution.
- **Detailed reports and code from applications** in statistics, engineering, simulations, graphics, computer aided instruction, text processing, analog to digital conversions, and more. You'll also find concise, advanced, informative articles on designing, implementing, and maintaining programs.
- **Articles that keep you up to date on Ada's expanding role.** You'll find substantive articles on the latest developments, compiler reviews, and applications software.



- **The means to use its resources efficiently.** Programs and procedures from each issue are available over telephone line, saving you costly hours of hand typing and debugging.

Get applicable information. Subscribe to the **JOURNAL OF PASCAL AND ADA.**

JOURNAL of PASCAL and ADA

P.O. Box 384
Orem, UT 84057

1 year (6 issues) \$14.00
 Outside U.S. \$21.00
 One trial issue \$3.00

Name _____
 Street _____
 City _____ State _____
 Country _____ Zip _____


Payment enclosed U.S. dollars drawn on U.S. bank
 VISA MasterCard American Express

Card # _____
 Exp. date _____
 Signature _____

The burden on image analysts in medical fields has led to the automated processing of pictorial data. Here, a device called the cytocomputer searches for genetic mutations.

Biomedical Image Processing

Stanley R. Sternberg, CytoSystems Corporation



A computer revolution has occurred not only in technical fields but also in medicine, where vast amounts of information must be processed quickly and accurately. Nowhere is the need for image processing techniques more apparent than in clinical diagnosis or mass screening applications where data take the form of digital images. New high-resolution scanning techniques such as computed tomography, nuclear magnetic resonance, positron emission tomography, and digital radiography produce images containing immense amounts of relevant information for medical analysis. But as these scanning techniques become more vital to clinical diagnosis, the work for specialists who must visually examine the resultant images increases. In many cases, quantitative data in the form of measurements and counts are needed to supplement nonimage patient data, and the manual extraction of these data is a time-consuming and costly step in an otherwise automated process. Furthermore, subtle variants of shade and shape can be the earliest clues to a diagnosis, placing the additional burden of complete thoroughness on the examining specialist.

For the last five years, the University of Michigan and the Environmental Research Institute of Michigan have conducted a unique series of studies that involve the processing of biomedical imagery on a highly parallel computer specifically designed for image processing. System designers have incorporated the requirements of extracting a verifiable answer from an image in a reasonable time into an integrated approach to hardware and software design. The system includes a parallel pipelined image processor, called a *cytocomputer*, and a high-level language specifically created for image processing,

C-3PL, the cytocomputer parallel picture processing language.

These studies have involved a great many people from both the medical and engineering communities and have highlighted the interdisciplinary aspects of biomedical image processing. The methods have been tested in anatomy, developmental biology, nuclear medicine, cardiology, and transplant rejection. The general consensus is that quantification by automated image analysis not only increases diagnostic accuracy but also provides significant data not obtainable from qualitative analysis alone.

One study in particular, on which descriptions in this article are based, involves a joint effort by the University of Michigan's human genetics and electrical and computer engineering departments and is supported by a grant from the National Cancer Institute. Basically, automated image analysis is being applied via sophisticated biochemical and computer techniques to derive an accurate estimate of the mutation rate for the human species.

Study overview

We are becoming increasingly concerned about human exposure to environmental elements, particularly those agents that may be carcinogens or mutagens. The long-term effects of exposure to chemical poisons and low-level radiation must first be quantitatively understood before their ultimate costs can be taken into account in

the political process of their control. As a step in quantifying the impact of mutagenic agents on human populations, techniques are being developed to detect mutation and to estimate the mutation rate in sample populations.

In the study of human mutation rate at the University of Michigan, placental cord blood samples are collected from newborns as are venous samples from their mothers and fathers. Samples of lymphocytes, red blood cell membranes, red blood cell contents, and blood plasma are being subjected to 2-D gel electrophoresis. A mutation is indicated when a variant of an identified protein appears in a child that is not present in either parent. The challenge is to develop that combination of specimens and techniques that will yield the greatest amount of information about mutation per unit of effort.

For the last 36 years, geneticist J. V. Neel of the University of Michigan School of Medicine has studied germ cell mutations in radiation-exposed Japanese populations.^{1,2} Using biochemical methods to detect protein variants at a cellular level, the approach was aimed at detecting single amino acid substitutions in a child's proteins that have been caused by mutations of parent's ova or sperm germ cells. Recently, 2-D electrophoresis systems in which cellular proteins are automatically mapped according to both their molecular weight and molecular charge have been developed. The 2-D gel electrophoresis system is sufficiently sensitive to detect protein charge changes caused by single amino acid substitutions. The significant advantage of the 2-D gel system over other biochemical methods is that many proteins are treated in parallel on a single gel, each protein a potential candidate for mutation. By contrast, traditional methods of estimating mutation rate have relied on the appearance of congenital defects, stillbirths, and infant survival rates, all of which severely limit the size of sample populations.

The method of separating proteins in two dimensions on a polyacrilimide gel medium is due mainly to O'Farrell,³ with substantial improvements to the method by Anderson and Anderson.⁴ First, the cellular material, usually blood, is fractionated into different cell types by centrifuging. Separation in the first gel dimension is by isoelectric focusing in which the proteins are separated on the basis of their molecular charge, or their isoelectric point. Separation in the second dimension is done by electrophoresis after the proteins have been treated by a detergent that masks the proteins' molecular charge and permits electrophoresis to resolve proteins on the basis of their molecular weights. At the end of the second step, each protein has migrated to a position (x,y) on the 2-D gel, where x reflects the molecular charge (pH) of the protein and y reflects its molecular weight. To see the proteins in the gel and perform qualitative and quantitative analysis, the proteins are either radioactively labeled during their synthesis and detected by autoradiography or stained at the end of electrophoresis. The individual proteins appear on the stained gel or autoradiogram as spots of different size and intensity. The integrated density of each spot is proportional to the amount of a given protein in the sample. Figure 1 illustrates a silver-stained 2-D electrophoretic gel.

The cytocomputer: a biomedical image processing tool

Development and use. In the University of Michigan mutation rate study, my colleagues and I are seeking rare events—only one mutation will occur in 100,000 samples—and thousands of mother, father, and child gel trios must be examined spot by spot for variations in protein patterns. This tedious task is clearly best suited to computerized visual comparison rather than human vision alone. (For more on computer analysis of 2-D electrophoretic gels in general, see Skolnick et al.⁵)

Searching for a rare protein mutation in a large number of gels requires that processing be very fast and that operator interaction be required only if a mutation is suspected. For this reason, we have directed efforts to the implementation of highly automated gel processing on a special-purpose computer characterized by a large capacity for parallel processing. The *cytocomputer*, which was developed at the Environmental Research Institute of Michigan, runs programs that apply sequences of neighborhood transformations to digitized gel images. Cytocomputer image processing operations are based on the concepts of cellular automata. Every cell or picture element (pixel) of an image is subjected to an identical sequence of time-discrete transformations, the transformed value of a pixel being determined by the values of a finite group of cells that make up its "neighborhood." Each neighborhood transformation is performed in an individual cytocomputer processing element referred to as a processing stage.

A cytocomputer is a serial pipeline of programmable processing stages, in which each stage performs a single




Figure 1. Two-dimensional silver-stained electrophoretic gel of red blood cell contents. The spots are individual proteins, spatially separated according to their molecular weight vertically and molecular charge horizontally. Prepared by Barnett Rosenblum of the University of Michigan.

transformation of the processing sequence on an entire image. Images are entered into a cytocomputer in a line-scanned format and progress through the pipeline of processing stages at a real-time rate. Following an initial delay to fill the pipeline, images can be processed at the same rate they are scanned.

Digital images and cellular automata have a common conceptual framework. Each pixel of a digital image can be thought of as a cell in a given state. If we define a neighborhood and a cell transition function on a digital image, then we can apply the transition function to modify or transform the configuration of cell states into new configurations. Of critical importance then is whether neighborhoods and transition functions exist that will cause images to be transformed in a predictable and useful manner when subjected to long sequences of neighborhood operations.

The image processing language we are investigating differs from conventional approaches in that the basic manipulative unit is pictorial and operations deal with images as wholes. Image processing is treated as a computation involving images as variables in algebraic expressions. These expressions may combine several images through both logical and geometrical relationships.

Architecture. Cytocomputer operations are implemented in highly efficient cellular computer architectures, and the computations are very fast. The Cyto I cytocomputer executes 140 million neighborhood operations per second in an 88-stage pipeline. Image processing algorithms are constructed as well-formed strings of primitives that are either variables representing images or image operations. The image being processed is referred to as the *active image*. Other images referred to in an image-algebraic expression are called *structuring elements*. In an image processing algorithm, we can modify the active image by probing it with structuring elements or combining it with other active images.

The neighborhood in a cellular space determines the set of structuring elements that can be employed in a single neighborhood transformation. Each pixel of a digital image belongs to a window of pixels composed of the given pixel and its neighbors. All structuring elements used in a neighborhood transformation must be sub-images of the window.

Consider a two-dimensional cellular array, where each cell of the lattice has connections with a finite collection of other cells that make up its input. The geometric pattern of the cells input to a given cell is the same as the pattern of the points in the neighborhood. Figure 2 illustrates a cellular array with the connection pattern for a 3×3 window configuration. In Figure 3, each cell of the array consists of a register for storing the state of the cell and a transition module that computes the new value of the cell state as a function of the states of the cells in the window. When a common clock pulse is applied to each cell in the array, all cell state registers pass from their previous state to a new state as determined by the programming of the transition logic module.

Although digitized biomedical image dimensions can often exceed 1000×1000 pixels, the largest arrays yet produced are only on the order of 100×100 . (See Potter's article on the MPP in this issue.) Large images must be partitioned into image segments, and each segment processed in turn. However, segment border effects propagate into the segment when multiple neighborhood transformations are applied, necessitating extremely costly and time-consuming I/O hardware and software subsystems for rapid segment swapping. The problem is remedied in the pipeline architecture where parallelism of image operations is used instead of pixel parallelism. (Danielsson and Leviardi give a good review of image processor architectures.⁶)

A cytocomputer consists of a serial pipeline of commonly clocked neighborhood processing stages (Figure 4).⁷ Shift registers within each stage store two contiguous




Figure 2. Array similar to cellular automata of identical cells connected to their nearest neighbor for iterative neighborhood processing of digital images.




Figure 3. Cellular array block diagram. Inputs from neighborhood cells form the address to a RAM lookup tabulation of the neighborhood transition function.

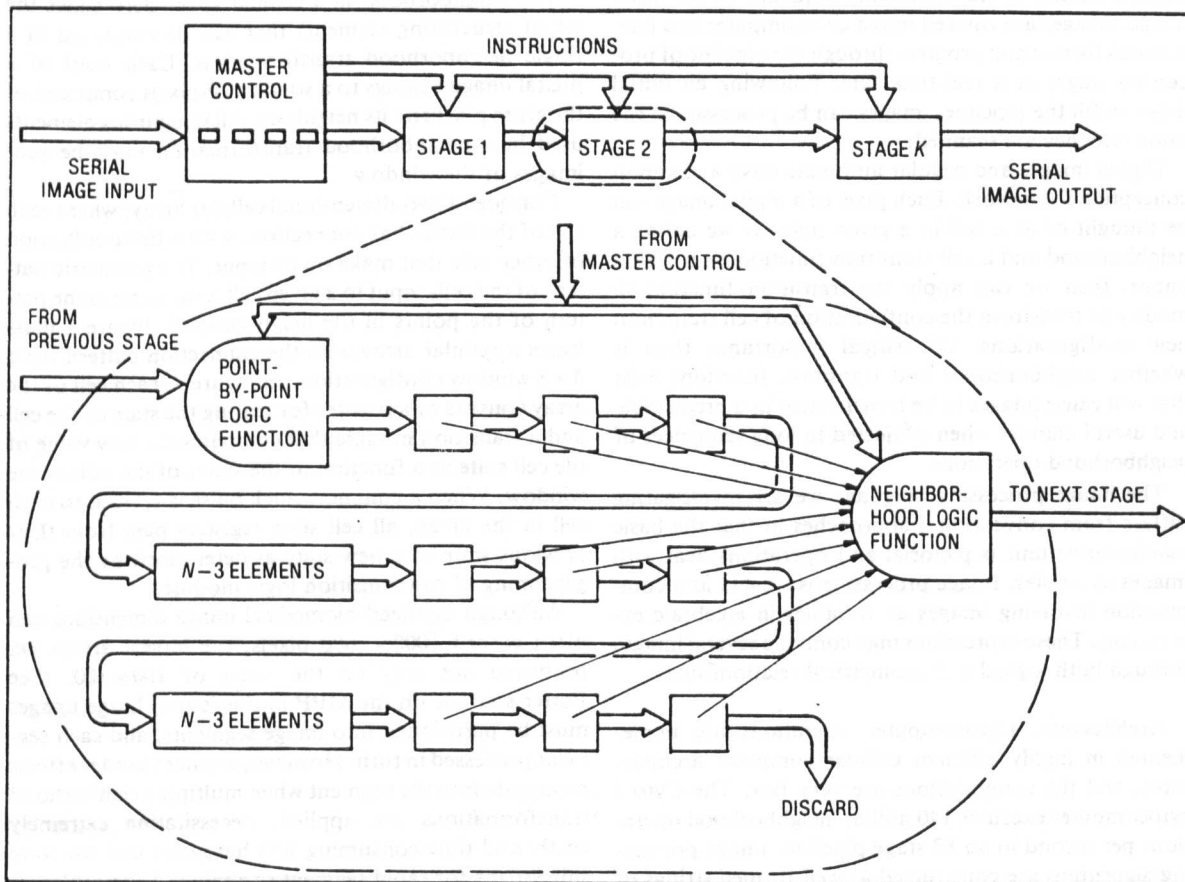


Figure 4. Cytocomputer block diagram.⁷ Serpentine shift register delays in each stage serially configure neighborhood inputs to the neighborhood logic module. A pipeline of K stages executes a K step neighborhood processing algorithm in real time. (Copyright IEEE, 1981.)

n pixel scan lines, and window registers hold the nine neighborhood pixels that constitute the 3×3 input to the neighborhood transition logic module. All neighborhood transformations and data transfers are computed within a single clock period. In each discrete clock period, a new pixel is clocked into the stage, and simultaneously the contents of all shift register delays are shifted by one element. In addition, operations that do not involve the states of the pixel's neighbors, such as gray-value scaling and bit setting, are performed in a separate point transition logic module to simplify the neighborhood transition logic circuit. Because stage input and output occur at the same rate, stages can be cascaded, with all stages operating in parallel. We can visualize a series of 3×3 windows following each other across the image, each processing the previous stage's output as shown in Figure 5.

Transformations implemented in a cytocomputer stage fall into two main categories. In the first, called the silhouette, or 2-D, transformation, image algebra operations are applied to planar binary images. In the second, called the umbra, or 3-D, transformation the same set of operations are performed on gray-scale images, where the gray value represents the brightness of a picture point or its height above an arbitrary reference plane. In umbra transformations the structuring elements are umbras of subimages in a $3 \times 3 \times 3$ window.




Figure 5. Cytocomputer stage windows sequentially scanning across a digital image. That portion of the image processed by the i th stage is immediately available for processing by the $(i + 1)$ th stage.




Figure 6. The morphological operations, erosion, dilation, and opening, illustrated on a binary example of electrophoretic gel spots. The process of opening separates the two connected spots and cleans them.

