DREAM MACHINES

New Freedoms Through Computer Screens — a Minority Report

This is the flip side of Computer Lib.

Come Dream along with me:
The Best Is Yet To Be.

DREAM MACHINES

This is the flip side of Computer Lib.

(Feel free to begin here. The other side is just if you want to know more about computers, which are changeable devices for twiddling symbols. Otherwise skip it.)

(But if you change your mind it might be fun to browse.)

In a sense, the other side has been a come-on for this side. But it's an honest come-on: I figure the more you know, the readier you'll be for what I'm saying here. Not necessarily to agree or be "sold," but to think about it in the non-simple terms that are going to be necessary.

The material here has been chosen largely for its exhilarating and inspirational character. No matter what your background or technical knowledge, you'll be able to understand some of this, and not be able to understand some of the rest. That's partly from the hasty preparation of this book, and partly from the variety of interests I'm trying to comprise here. I want to present various dreams and their resulting dream machines, all legitimate.

If the computer is a projective system, or Rorschach inkblot, as alleged on the other side, the real projective systems-- the ones with projectors in them-- are all the more so. The things people try to do with movies, TV and the more glamorous uses of the computer, whereby it makes pictures on screens-- are strange inversions and foldovers of the rest of the mind and heart. That's the peculiar origami of the self.

Very well. This book-- this side, Dream Machines-- is meant to let you see the choice of dreams. Noting that every company and university seems to insist that its system is the wave of the future, I think it is more important than ever to have the alternatives spread out clearly.

But the "experts" are not going to be much help; they are part of the problem. On both sides, the academic and the industrial, they are being painfully pontifical and bombastic in the jarring new jargons (see "Babel's in Toyland," p. 4). Little clarity is spread by this. Few things are funnier than the pretensions of those who profess to dignity, sobriety and professionalism of their expert predictions-- especially when they, too are pouring out their own personal views under the guise of technicality. Most people don't dream of what's going to hit the fan. And the computer and electronics people are like generals preparing for the last war.

Frankly, I think it's an outrage making it look as if there's any kind of scientific basis to these things; there is an underlevel of technicality, but like the foundation of a cathedral, it serves only to support what rises from it. THE TECHNICALITIES MATTER A LOT, BUT THE UNIFYING VISION MATTERS MORE.

No more pencils, no more books.
No more teachers' dirty looks.
Kids' jingle with new meaning.

Everything is Deeply Intertwingled.

THE GESTALT, DEAR BRUTUS, IS NOT IN OUR STARS BUT IN OURSELVES.

DREAMS

Technology is an expression of man's dreams. If man did not indulge his fantasies, his thoughts alone would inhibit the development of technology itself. Ancient visionaries spoke of distant times and places, where men flew around and about, and some could see each other at great distance. The technological realities of today are already obsolete and the future of technology is bound only by the limits of our dreams. Modern communications media and in particular electronic media are outgrowths and extensions of those senses which have become dominant in our social development.

How Wachspress, "Hyper-Reality." © Auditac Ltd. 1973.

CHILDREN OF ALL AGES!

Ladies and gentlemen, the age of prestidigitative presentation and publishing is about to begin. Palpitating presentations, screen-scribbled, will dance to your desire, making manifest the many mysteries of winding wisdom. But if we are to rehumanize an increasingly brutal and disagreeable world, we must step up our efforts. And we must hurry. Hurry. Step right up.

Theodor H. Nelson, "Barnum-Tronics." Swarthmore College Alumni Bulletin, Dec 1970, 12-15.

"When you're dealing with media you're in show business, you know, whether you like it or not."

"Show business," he said. "Absolutely. We've gotta be in show business. We've gotta put together a team that will get us there."

I made a mental note to use the show business metaphor again, and continued. "IBM's real creative talent probably lies in other areas . . ."

Heywood Gould, Corporation Freak (Tower), 23.

The Great Robert Crumb. (From Zap Comix #0.)

This book has several simultaneous intentions: to orient the beginner in fields more complex and tied together than almost anybody realizes; nevertheless, to partially debunk several realms of expertise which I think deserve slightly less attention than they get; and to chart the right way, which I think uniquely continues the Western traditions of literature, scholarship and freedom. In this respect the book is much more old-fashoned than it may seem at the gee-whiz, very-now level.

The main ideas of this book I present not as my own, but as a curious species of revealed truth. It has all been obvious to me for some time, and I believe it should be obvious as well to anyone who has not been blinded by education. If you understand the problems of creative thinking and organizing ideas, if you have seen the bad things school so often does to people, if you understand the sociology of the intellectual world, and have ever loved a machine, then this book says nothing you do not know already.

For every dream, many details and intricacies have to be whittled and interlocked. Their joint ramifications must be deeply understood by the person who is trying to create whatever-it-is. Each confabulation of possibilities turns out to have the most intricate and exactly detailed results. (This is why I am so irritated by those who think "electronic media" are all alike.)

And each possible combination you choose has different precise structures implicit in it, arrangements and units which flow from these ramified details. Implicit in Radio lurk the Time Slot and the Program. But many of these possibilities remain unnoticed or unseen, for a variety of social or economic reasons.

Why does it matter?

It matters because we live in media, as fish live in water. (Many people are prisoners of the media, many are manipulators, and many want to use them to communicate artistic visions.)

But today, at this moment, we can and must design the media, design the molecules of our new water, and I believe the details of this design matter very deeply. They will be with us for a very long time, perhaps as long as man has left; perhaps if they are as good as they can be, man may even buy more time-- or the open-ended future most suppose remains. (See "Endgame," p. 68.)

So in these pages I hope to orient you somewhat to various of the proposed dreams. This is meant also to record the efforts of a few Brewster McClouds, each tinkering toward some new flight of fancy in his own sensoarium.

But bear in mind that hard-edged fantasy is the corner of tomorrow. The great American dream often becomes the great American novelty. After which it's a choice of style, size and financing plan.

The most exciting things here are those that involve computers: notably, because computers will embraced in every presentational medium and thoughtful medium very soon.

That's why this side is wedded to the other: if you want to understand computers, you can take the first step by turning the book over. I figure that the more you know about computers-- especially about minicomputers and the way on-line systems can respond to our slightest acts-- the better your imagination can flow between the technicalities, can slide the parts together, can discern the shapes of what you would have these things do. The computer is not a limitless partner, but it is deeply versatile; to work with it we must understand what it can do, the options and the costs.

My special concern, all too tightly framed here, is the use of computers to help people write, think and show. But I think presentation by computer is a branch of show biz and writing, not of psychology, engineering or pedagogy. This would be idle disputation if it did not have far-reaching consequences for the designs of the systems we are all going to have to live with. At worst, I fear these may lock us in; at best, I hope they can further the individualistic traditions of literature, film and scholarship. But we must create our brave new worlds with art, zest, intelligence, and the highest possible ideals.

I have not mentioned the emotions. Movies and books, music and even architecture have for all of us been part of important emotional moments. The same is going to happen with the new media. To work at a highly responsive computer display screen, for instance, can be deeply exciting, like flying an airplane through a canyon, or talking to somebody brilliant. This is as it should be. ("The reason is, and by rights ought to be, slave to the emotions." -- Bertrand Russell.)

In the design of our future media and systems, we should not shrink from this emotional aspect as a legitimate part of our fantic (see p. DM48) design. The substratum of technicalities and the mind-bending, gut-slamming effects they produce, are two sides of the same coin; and to understand the one is not necessarily to be alienated from the other.

Thus it is for the Wholiness of the human spirit, that we must design.

(COVER)
Our Caped Communicant faces his greatest challenge and adventure as... SUPERSTUDENT MEETS HYPERTEXT!
(see p.)

AUTHOR'S COUNTERCULTURE CREDENTIALS

Writer, showman, generalist. Gemini, moon in Libra, Gemini rising.
One-time seventh-grade dropout. I have relatively little interest in improving the educational system within the existing framework.
Author of what may have been world's first rock musical, "Anything & Everything," Swarthmore College, November 1957 (with Richard L. Caplan).
Photographer for a year at Dr. Lilly's dolphin lab (Communication Research Institute, Miami, Florida). Attendee of the Great Woodstock Festival (like many others), and it changed my life (as others have reported). What we are all looking for is not where we thought it was.
Lifelong media nut. Magazine collector; hung around TV studios as a child. Compulsive explainer. Gimmicist by disposition, computerman by accidestiny.

SPECIAL SUPPLEMENT TO THE THIRD PRINTING, August 1975.

Gee whiz, folks, here we are at another printing and already the big clock on the wall tells us that another year has gone by.

This supplement is mainly things that had to be mentioned, but it kind of assumes you've read the book itself or are generally familiar with computers. BOOKSTORE BROWSERS: avoid these four pages. NEW OWNER OF THE BOOK: Check that the pages are right, right.

SORRY THE TYPE STILL ISN'T BIGGER, but that will require thousands of bucks in new negatives-- meaning a lot more have to be sold as is.

—— o ——

YOUR UNDERGROUND COMPUTER MAGS

The redoubtable PCC is now six issues and six dollars a year. People's Computer Company, P.O.Box 310, Menlo Park CA 94025.

BYTE Magazine, $10/year if you hurry, $12 later, from Green Publishing Co., Peterborough, NH. Editorial: Carl Helmers, Box 378, Belmont MA 02178. Hardware-oriented.

Creative Computing: The Magazine of Recreational and Educational Computing. Ideametrics, P.O. Box 789-M, Morristown, NJ)7960. Weird variety of subscription rates: student $6, "individual" $8, "institutional" $15.

The Computer Hobbyist, $6/year, Box 295, Cary NC 27511. Hardware-oriented.

Computer Notes (for Altair users; from MITS).

Micro-8 Newsletter, for people really into the Intel. Hal Singer, Cabrillo High School, 4350 Constellation, Lompoc CA 93436.

and also

Simulation and Gaming News, Box 3039 University Station, Moscow, Idaho 83843.

Electronotes is the magazine for music synthesizer freaks. Bernie Hutchins, 60 Sheraton Drive, Ithaca NY 14850.

and something else entirely,

Privacy Journal, a monthly newsletter on problems of privacy, many or most of which involve computers. P.O. Box 8844, Washington, D.C. 20003; $15 a year.

(Note: it is of interest that a bill on computer privacy in this year's House of Representatives just happened to be HR 1984.)

—— ◇ ——

COPYRIGHT AND COPYWRONG

One individual I know, who relishes his counterculture image, told me with angry and shaking voice that he doesn't believe in copyright and that anything that gets near his computer belongs to him. Well, don't leave your manuscripts near such a person. (Why is it always the guys with cushy and secure jobs who tell you tweedle de dee, ideas should be free, and patents and copyrights are selfish?)

Actually, for the individual, one of the strongest forms of protection available is copyright. Far from obsolete, the copyright makes publishing, and the better computer software, possible. (It is not generally known that copyright violation is a felony.) (And ripping off a program you're supposed to pay for is not a brave guerrilla affirmation, like hitting Harold Geneen with a pie, but grand larceny.)

Now that Altairs and LSI-11s have got a lot of you guys dreaming about selling software, an important question is how to protect your work. Well, you have a champion.

Calvin Mooers (see pp. 18-21) is not only a genuine Computer Pioneer From The Forties, but, along with Herb Grosch, pioneered the Computer Counterculture. Grosch flaunted a beard in front of old man Watson, Mooers strove to make computers easy to use-- back when that was unheard of.

One of his current interests is in ways that small independent underground-type programmers can protect their developments. He and some associates are exploring the possible formation of a group for the legal protection of small software producers and owners.

Incidentally, when you think something you've written belongs to you-- a computer program, poem or whatever-- slap the following at the beginning, under the title:

© 1975 Irving Snerd

substituting, of course, your own name. And the year currently in effect. If computer printing is used, (C), using parentheses, is considered an acceptable substitute for c-in-a-circle.

This not only gives notice to potential Borrowers, but it has certain strong magical properties as a legal incantation. See your lawyer for details, but don't hesitate to apply it liberally to your own work; you may be glad you did.

CHECK YOUR BOOK NOW

A lot of copies of this book have not been put together correctly. We hope that's all over now, but if this book belongs to you please check it. Incorrectly-made books will be exchanged. within two weeks of purchase (address on p. 2). Otherwise you have a Collector's Item. CHECKED: ☐ Date:

ALL YOU NEED DO IS CHECK THE NUMBERS ON THE 'COMPUTER LIB' SIDE. They run straight through from cover to cover, even though the contents flip capriciously. If the letters "DM" appear anywhere amongst these plain numbers, you got a lemon.

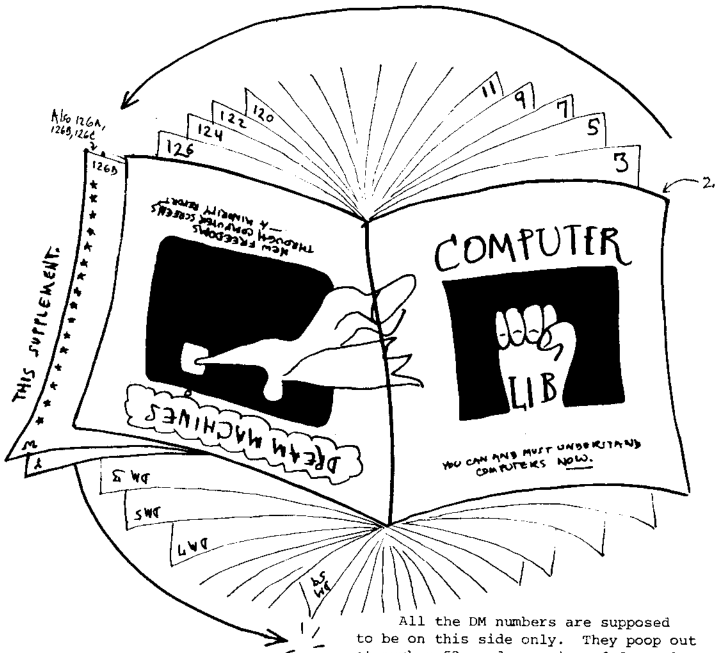

All the DM numbers are supposed to be on this side only. They poop out at number 59, and were intended merely for cross-reference.

THE ALTAIR STORY

It began with a bang last Christmas: the cover of Popular Electronics showed 'a computer you can build yourself for only $400'!

It was real. A young firm in Albuquerque called Micro Instrumentation and Telemetry Systems, or MITS, had finally done it: a computer for well under $1000. In a box not much bigger than a typewriter, a machine comparable to the Univac I. They called it the Altair 8800.

Of course, in a way this was an obvious step. The MITS computer was simply the packaging, as a computer, of a specific integrated circuit chip that had been on the market for some months. This chip, the Intel 8080, is a microprocessor, or two-level computer (see p. 44), generally employed for fixed purposes in cash registers, pinball machines, and the like. However, to make it a "general" computer-- with the engineering, hookups and accessories that entailed-- would be no small matter if taken seriously.

(cont. next page, last column)

MITS is at 6328 Linn N.E., Albuquerque NM 87108, 505/265-7553.

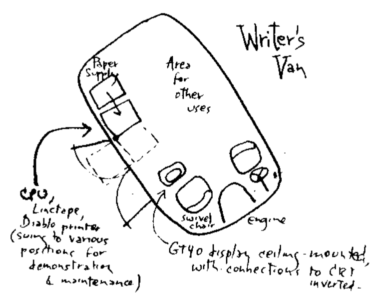

Next in computer hobbyism will obviously be the Computer Van. Already vans come with swivel thrones, four-track stereo, color TV; so this next step is obvious. But most important, recreational vehicles can be purchased on very long time-plans, sometimes seven years. (MITS has a demo van with Altair, floppy disk, lineprinter. It drives around showing off. But they'll sell you one like it for a trifling $29,000.) Now for mobile operation we redo the power supply...

UNDERGROUND PAGE

Dig It, All: THE GREENING OF COMPUTERDOM

1975 may be thought of as the year in which the computer underground suddenly appeared in full force. The Altair was probably the big crystallizing event.

Not that there wasn't a counterculture before. There were the games-players at every university, the prank programmers (see p. 48-9), and, wherever computers are the center of things, a shared experience of mischief and breakthrough. There was Computer People for Peace, a cliquey and unapproachable group with booths at the conferences (at least, their backs were always turned when you wanted to ask questions). There was the hobby fringe.

But now it's gone different. Instead of pretentious company names meant to appeal to obtuse businessmen, like Performance Measurement Systems Consultants Group and Bottom-Line-Tronics, the new companies have rock-group names like General Turtle, Inc., The Sphere and Loving Grace Cybernetics. In this new computer counterculture, the main computer companies are not IBM and Honeywell and Univac, but DEC and MITS and General Turtle; the standard computer is not the 370, but the 11 (or possibly the Altair or the 8). The standard language is not Fortran or Algol or PL/I, but BASIC. Instead of the big color TV that middle America wants, the underground computernik dreams of his own graphic setup forever running The Game of Life in color (see pp. 48-9 and pic p. DM26). (Of course that'll also require the color TV; see "Bit Maps," p. Z.)

In such a world, computers are not a tool but a way of life. The computer is toy, pet, checkerboard, music box and TV. Computers are for making music, computers are for getting people together via community memory, computers are for letter-writing, computers are for art and movie-making and the animated decoration of the home.

Computers are for games; a vast number of interactive game-programs are published and swapped around. Almost all are in the BASIC language. (Bob Albrecht's WHAT TO DO AFTER YOU HIT RETURN is said to be definitive-- $7.50 from People's Computer Company, 1919 Menalto Ave., Menlo Park CA 94025. See also their magazine PCC, as well as Simulation and Gaming News.) PLATO games, a somewhat different subspecies, are discussed on p. DM27.

The underground computer magazines have become a blizzard (see box). Albrecht's sprightly and successful PCC, originally oriented toward high and grade schools, has now branched into hobbyism as well. On the hardware side there is The Computer Hobbyist, and now a slick new hobby magazine, Byte, with a first printing of 50,000. On the educational side there is a swell new magazine called Creative Computing.

Then there is the Community Memory movement. The basic idea of Community Memory is to have a computer resource of information and ideas, commonly available. In its more glorified and mystical form, the idea seems to be to have a place, inside the computer, where information can be shared by The People, free of institutional obstruction or the profit motive.

This vision is perhaps unclear to others besides the author, but it attracts a variety of people interested in some form of grass roots revitalization of our society. Some of these are disillusioned sixties radicals who look to "community organization" as a building block for a new society; others are interested in more nuts-and-bolts applications, such as trying to make barter a viable economic form again, in an urban society with many nonstandard leftovers, skills and wants. (Presumably this would work by having the computer find pairs of people with matching wants and tradables; or even search out potential trades around multi-person rings.)

The first of these systems was Resource One, in San Francisco; I saw another Community Memory in Vancouver, which seemed to be in practice a sort of animated classified-ad system. A user sitting at the terminal can put in ads of his own, and can search through the entire file for key-words of interest. As there is no censorship, some rather surprising things get in there, for which I wish we had room.

(A newsletter of such projects, Community Communications, is being started by Lee Felsenstein, Loving Grace Cybernetics, 1807 Delaware St., Berkeley CA 94703.)

Even for those coming anew into the field-- the radio hams and amateur telescope makers who've laid their Master Charge cards on the line for the Altair-- computers represent a new social life. Amateur computer clubs have drawn startling numbers: for instance, the Los Angeles and San Francisco groups are currently pulling 100 members to their weekly meetings. (In San Francisco, contact Fred Moore, 558 Santa Cruz Avenue, Menlo Park CA 94025.)

This book and its surprise success probably rate mention of some sort in the world of underground computerdom, '74-75; although my underground status may be in jeopardy. I had intended to bypass the computer establishment, and certainly not expected to become assimilated therein; so the dozens of university class adoptions have come as a considerable shock, as have the acceptance and legitimation I had long since given up on. My heartfelt thanks for this response, and I'll try to live up to it. (How is discussed on p. Z, last column.)

But folks, this all is the merest beginning. As it says on diametrically the other side, p. 3,

COMPUTERS BELONG TO ALL MANKIND.

CORPORATE PAGE

• • • • •

BIG BROTHER AND AUNTIE TRUST

(All in the Family?)

Well, the anti-trust trial of IBM is underway. In an awkward start, opposing lead attorneys accused each other of professional misconduct, placing both men's careers under a cloud as the fight began.

The way it comes through in the trade press, the Government seems to be pulling punches and missing the point of what its own witnesses say. A large-scale botch may be in progress. (The Computer Industry Association, or IBM-hater's club, offers transcripts of the IBM trial, as well as daily summaries. The group's headquarters are now at 1911 N. Fort Meyer Drive, Rosslyn, Virginia.)

What is the point of it all? The Justice Department is seeking to break up IBM. (According to one theory, it needs points after the ITT-Hartford business.)

There is a lot of superstition about IBM in the land. The stock market took a huge dive when the Justice Department announced it would prosecute. But why? Hesh Wiener, editor of Computer Decisions, thinks IBM will be broken up: "The Justice Department wants it, and IBM wants it, and the stockholders will make more money. They've already drawn the dotted lines."

A key question is what difference it would make. (Remember what happened after they broke up Standard Oil? Not much.) A phony breakup would simply make the different divisions into different companies, leaving the product line and the cooperation intact; a more effective split would in some way foster competition among the daughter corporations. But what way?

One of IBM's more recent tricks is to overwhelm litigators by the quantity of documents supplied, many of which are stored on computers in full-text form. To give you an idea of the humungous magnitudes involved, some figures just came up in recent litigation with Sanders Associates. Sanders is suing IBM, and recently asked IBM how many documents IBM had that were 'pertinent' to the case. The reply: "Active files, approximately 906,054,000 pages; inactive files, approximately 421,660,000 pages." (Datamation, July 75, p. 129.) An unfortunate aftermath of a suit by Control Data, in which IBM settled, was the destruction of the great indexes which had been constructed to the vast file of IBM's records; the index is gone and unavailable for this case.

To a lot of people this just seems to have to do with the size of one corporation. In the author's opinion, however, the issue is the one big usual question, the issue of freedom in our time; and that is not a matter of bigness, but the style of IBM's control. Computers should make things easier in both our work and our private lives, and should help lighten our loads and enlighten our minds, clarifying the complexities of everything. Unfortunately, IBM's method of making money has a little too much to do with creating rigid and oppressive and pointlessly complex systems, fobbing them off as "scientific," and ensnaring its customers in complications by the techniques discussed on pp. 52-56.

People should be free to use computers as they ought to be used, each in his personal style regardless of his job title, amidst rushing menus of options and clarifying screen graphics, rather than each person and office worker being locked into his own "sternly allotted sandpile," as cummings put it. And that is the problem.

RECENT IBMOGRAPHY

Nancy Foy, The Sun Never Sets on IBM. Not reviewed at press time.
William Rodgers' Think is out in paperback, with an added chapter, from the New American Library.
Datamation devoted large sections of its February and March '75 issues to material on the IBM Problem.

THE BIG LIE
(Cybercrud '75, Antitrust Division)

The backbone of IBM's defense in antitrust is this whopper:

MAJOR PREMISE. "Computers are so complicated only a company as large as IBM can put together the technical teams necessary to make them work."

THE COROLLARY. "Computers are so complicated that there's just no way to make it possible for competitors to hook up their equipment to ours."

The truth: almost anybody can make sensible computers that work and tie together sensibly. Only IBM can do it wrong and make it stick.

THIBMK

Some useful words for discussing the IBM problem. (Thanks to Computer Decisions magazine, in which some of these were first published. ©1973, 1975 Theodor H. Nelson.)

ibmology
: the study of IBM.

ibmosophy
: the wisdom of IBM; ibmosoph, one wise to IBM.

ibmperceptible
: officially noticed by IBM.

enibmatic
: puzzlingly ibmish.

ibmbroglio
: IBM software.

ibmologism
: clumsy or inappropriate term, esp. one which misspeaks itself, such as "random access" for cyclical access, "direct access device"for indirectly accessible device, and "virtual system" for real system involving virtual huge memory.

hexadecimal
: trying to put a curse on the PDP-10.

EBCDORK
: rumored forthcoming code for the Future System (Extended Binarily Coded, Decimally Organized, Rbitrary Kludge).

ibmophily
: love of IBM.

ibmolatry
: worship of IBM.

ibmecile
: someone ibmersed in, or ibmued with, the ibmage.

ibmopoly
: 65% of the market, a great ibmposition.

ibmunity
: the safety and togetherness of IBM.

ibmpunity
: judgments against IBM, if any.

ibmoclasm
: the breaking up of IBM by the Justice Department.

On page 53 of this book I say: "I hope to be able to report in future editions of this book that IBM has moved firmly and credibly toward making its systems clear and simple to use, without requiring laborious attention to needless complications and oppressive rituals."

This has in fact occurred, and I so report.

In an earth-shaking announcement in January, IBM totally reversed the policy of its computer division for the last ten years. Yet so jaded is the press that this event was not, I think, properly recognized.

Astounding as it may be from the company that gave the world JCL and the MT/ST, in January IBM stepped into the world of easy computers, bringing out the System/32, a minicomputer for business. You can only rent it as an interactive terminal, with a program created by IBM which cannot be modified (called an Industry Application Package or IAP). But these little programs prompt users step-by-step through what they are supposed to be doing, and apparently are very clear and helpful for the naive.

This about-face is in many ways gratifying for those of us who have been advocating easy, screen-based systems for years and years. At long last it gives IBM's "legitimacy" to minicomputers for business, and it helps companies that already provide such services, such as Basic/Four.

It will be interesting to see if IBM knows how to make things simple, considering the experience they have lavished on the opposite policy. Anyway, with this move I would say that IBM has purged itself of at least 20% of its discernible evil, if this begins a real change.

A delicate problem will restrict the impact of the 32 itself, however. That is that IBM wants it used only as a gateway to its big computers; presumably, if users were allowed to program it, they'd find ways out of having to use the biggie.

WHAT WILL IBM DO NEXT?

As it happens, we know what IBM's biggest next move will be. It is something to be called the Future System (FS). FS will be a complete line of computers and communications techniques for them, but that's all we know; security is very tight. Supposedly FS exists and is running; but what is it? All we know is that its scheduled introduction has been pushed back from 1978 to sometime after 1980.

Anyway, I have asked a lot of savvy people what they thought FS was going to be, and here are some of the answers:

- A completely modular line of computers and terminals with a Unibus-type architecture. (RUMOR: this would eliminate SEs, CEs and Systems Analysts. Thus the postponement.)
- A micrprogrammed line of equipment, whose underware uses APL.
- A totally PL/I system.
- A line of equipment with ever-changing microprogrammed "fan-dance" interfaces, such that no competing manufacturer can ever find out what they are. (A charge by Herb Grosch and numerous peripheral manufacturers.)
- A complete and impregnable total system for all symbolic information, which can only be keyed into through IBM terminals, processed on IBM computers, transmitted through IBM satellites, and read out through IBM terminals. (FACT: IBM has applied for a satellite.)
- Totally compatible with existing 370 hardware.
- Totally incompatible with existing 370 hardware, but software-convertible. (IBM makes a lot of money on adapters and conversions.)
- A line of pocket-sized and portable equipment built around Magnetic Bubble Technology.
- A line of easy-to-use equipment with easy-to-use interactive software. (This would suddenly eliminate hundreds of thousands of programmers, but IBM doesn't owe them anything.)

"Man, whatever it is, it'll be sick."

CYBERCRUD '75
County Fair Division

At the Dutchess County Fair this year, there was a "computer handwriting analysis" booth. You wrote your name on a card (Hollerith, natch) and this was put through a slot. A typewriter (marked "IBM") printed out the "analysis."

I wasn't there, but it was almost certainly a brazen fake. Presumably the typewriter was an ordinary Mag Card Selectric, Memory Typewriter or the like. The flathouse operator could simply choose what he wanted the printout to say by the insertion of a card (on the former) or the twist of a dial (on the latter).

Incidentally, while IBM is probably the principal employer of Dutchess County, we should not assume direct complicity.

"THE OFFICE OF THE FUTURE"

A remarkable issue of Business Week (June 30, 1975) carried a 36-page section called an "executive briefing," whatever that is, on the Office of the Future, whatever that is.

The article was actually two things spliced together: "futuristic" gab around the title, and a report on the so-called "word processing industry." Word processing, a silly IBM term, means handling text by tricky office equipment (see "Type Righter," p. 14).

IBM controls the word processing market, with such machines as the Mag Card Selectric and the abominable (in my opinion) MT/ST. As reported by Business Week, IBM's basic strategy is to tell businessmen that they have to have a centralized typing pool of specially trained typists to use these things, so the office has to be reorganized. The secretaries hate the new organization because it makes them into keypunch operators-- the peon/executive dichotomy is a traditional aspect of IBM products, it would seem-- but the whole thing is put over as Modern.

Now Xerox has come up with a competitive machine, the 800 (see Diablo, p. Y), and Business Week intones that only these two firms have the savvy and capital to succeed in competing to create the Office of the Future. Well, this is hogwash.

The big mistake IBM's competitors always seem to make is to let IBM define the problem, and then go in to try to compete on the battleground, and in the terms, that IBM has laid out. But it is not sensible to play follow the leader on slippery logs through a boobytrapped swamp. Now Xerox has stepped onto the slippery log. But the right thing would be to unmask the absurdities of the IBM game with new initiatives which they cannot possibly emulate.

The office of the future, in the opinion of the author, will have nothing to do with the silly complexities of automatic typing. It will have screens, and keyboards, and possibly a printer for outgoing letters, but possibly not. All your business information will be callable to the screen instantly. An all-embracing data structure will hold every form of information -- numerical and textual-- in a cats'-cradle of linkages; and you, the user, whatever your job title, may quickly rove your screen through the entire information-space you are entitled to see. You will have to do no programming, and indeed "programs" will never be explicitly invoked at all; they will simply take effect as you get near, in the display space, something which needs update. A display-driven information complex.

X MARKS THE SPOT

When Xerox Corporation entered the computer business a few years ago, it announced that it was going to challenge IBM head-to-head for domination of the whole broad field of information, whatever that is. Xerox made copiers, but saw the handwriting on the drum: eventually the handling of written materials would cross over into the computing realm; few are sure in what way. (For three future directions that have been proposed, see Engelbart, pp. DM46-7, PLATO, pp. DM26-7, and Xanadu™, pp. DM56-7.)

The last July issue of Computerworld in '75, however, tolled of Xerox's abandonment of the computer field. Specifically, Xerox will stop making computers themselves, though they still will make accessories such as the hot Diablo printer (see p. Y). The news was presented in the framework of grand tragedy, the Promethean collapse of overextended ambitions. Evidently Xerox management pushed too hard in two incompatible directions-- building slowly for the eventual challenge of IBM, vs. showing profits quickly. The firm fell between the boat and the dock, joining RCA and General Electric and the other big companies that found they couldn't make it selling computers.

But Xerox is not as far out of the field as some might think.

In a secret mountain hideaway-- well, not too secret-- Xerox still has perhaps the sharpest bunch of computer rascals in the world. And they are planning way way ahead, to the time computers are practically free. If Xerox gives them their head, and doesn't cut back, the corporation will have little trouble in triumphantly returning to the field five or ten years from now, conceivably knocking IBM off its feet in the new markets of that day with a karate-like sweep.

This Place of Power is called Xerox Palo Alto Research Center, or Xerox PARC, and its atmosphere of California Mellow can mislead the unwary.

I spoke there a few years ago and found it an astonishing experience. First, there was a busy volleyball game outside when I arrived, and when I asked for the person I was going to see, the receptionperson said to pull up a beanbag and wait till he had finished playing volleyball. Later, when I addressed a group, it was in a room furnished only with a mountain of those beanbag sacks. As people came in, they would pull beanbags off the mountain and sit down on them.

So far so good: California Mellow. So I went into my rap, and everybody sat listening. I had no idea if I was getting through. Since what I try to tell people begins where technology stops-- moral precepts, as it were, for organizing ideas and systems in the world of the future (see this whole DM side)-- I'm used to people looking confused, or worried, or angry, or even walking out. There was none of that. Was I getting through? Or were they all just stoned?

I think I just sort of stopped and said, "Is everybody following this?"

There were smiles and I think someone said, "We're with you, Ted."

And they were. It was the only place I've ever spoken where the audience was on the same wavelength, going straight on into Systems Design for Future Man. Very moving.

This is obviously the place to tell you about Alan Kay and the Dynabook.

The hottest project at Xerox PARC is Alan Kay's Dynabook, formerly the Kiddy Computer. As lots of people will tell you, it's going to cost five hundred dollars, be small enough to carry around on a shoulder strap, have a built-in screen, run on batteries, and have all the books a kid wants to read from the screen stored on a cassette.

And the demos! They'll knock you out. On a color TV screen, they'll show you a wildly changing pageant of toy soldiers, photographs, beautiful patterns, all generated by the computer in real time (see "Bit Maps," p. Z). And if you're into computers, they'll show you how all this is run by the beautiful SMALLTALK language (it was previously called the Kiddy Computer, remember), which any bright child can learn and which has some awfully powerful features.

Now let's sort this all out.

There have been a lot of cons in the computer field, but this is not one of them. It's marvelously real.

So how come Xerox is leaving the computer field?

Answer: they're not exactly leaving; they're taking a break until they can sell this beauty for five hundred dollars.

What's the delay?

The Dynabook, or Kiddy Computer, is actually a PDP-10.

You're supposed to laugh. A PDP-10 is a big computer, the best. (See page 41.) A PDP-10 system costs hundreds of thousands of dollars.

But the last laugh will be Xerox's. The way computer prices are coming down, through integrated circuits ever more powerful and cheap, that PDP-10 can be sold for $500 in... (check your choice) __1978 __1979 __1980 __1981 __1982.

(Interesting anecdote: the guys at Xerox PARC asked to buy a PDP-10, but management bridled, seeing as how Xerox was in the computer business and made competitive machines. So the fellas, nothing daunted, built their own. They modestly say the parts only cost a few thousand.)

(Note: the above predictions are based, of course, on the assumption of Xerox management knowing what it's doing. Assumptions of this type in the computer field all too often turn out to be without basis. But we can hope.)

The TRAC® Language is now running time-shared, for general customers, on Computility (as mentioned on p. 21), and in a fancier version offered by Interactive Sciences Corp., 60 Brooks Drive, Braintree, Mass., 617/848-2600. Mooers has licensed the latter firm to run both his basic processor and "Advanced Developments" (rather secret) in file systems and computer control. Apparently he has some spectacular data-base stuff in there, but you won't be able to find out about it directly. Special packages are the specialty of Interactive Sciences, and with TRAC they can offer packages with both the data base stuff and other unusual capabilities. For instance, this time-sharing TRAC can itself call up other computers and sign into them, responding to messages as if it were a user.

The Computility version seems to run for about $12 an hour, the Interactive Sciences version for somewhat more-- but the latter firm is interested in selling whole packages, not user-diddling.

• • •

Mooers has recently received registration for his trademark **"THE HAPPY ROBOT."**

"DEC IS GETTING LIKE IBM"

is a complaint you hear everywhere. The resemblance is certainly not in salesmanship-- ha ha-- but in the way that the standard answer to questions has now become, "I don't know, that's not my department." People feel this with a certain bitterness because so many of DEC's fans loved it for not being like IBM. It's like when Jackie Kennedy married Onassis...

News of DEC

Despite its steadfastly insipid marketing, DEC has consolidated its position at the center of the small-computer maelstrom, and the PDP-11 has been consolidated as the small and medium-sized computer of choice among sophisticates. (The PDP-11 is also attracting considerable interest as a network computer. In one curious instance, First National City Bank of New York is creating a network of 11/45s.)

NEWS OF THE 11

The PDP-11 has now become the first computer to range in size, genuinely, from the tiny to the grand. During the last six months, DEC has brought out the smallest of the line, the LSI-11, all on a board the size of a sheet of typewriter paper, for SIX HUNDRED AND FIFTY DOLLARS. That includes the full computer and 4K of volatile fast memory, as well as built-in debugger.

However, as with many announcements, this is not quite the full story. This LSI-11 is without power supply and without Unibus. Indeed, it seems that the LSI-11 happens to be the very same thing as the 11/04, demurely announced last fall, which costs $2500 with power supply and Unibus, no front panel. The announcement of the LSI-11 then takes on the appearance of a reply to the grand MITS announcement of January (see p. W). Especially when it turns out that if you want one LSI-11, it costs a thousand. ("Buying clubs" are being formed with the idea of pooling resources for the quantity price; see "Cheap Computers," p. Y.)

(Sophisticates interested in putting the LSI-11 in other equipment have been quick to notice an unusual feature: it has an empty socket in which you may insert the ROM that gives you floating point (a very cheap option). For those of us who daydream about unusual functions, such as list processing or graphics or the like, this opening is very suggestive: with access to the microprogram instructions, a different ROM could be put in for fast implementation of whatever it was you wanted-- and your program would use for your nefarious purposes the binary commands ordinarily reserved for floating point.)

(While he may not be able to deal with that, a very savvy source person for the LSI-11 is Daniel L. Lewis at DEC in Rolling Meadows, Ill.)

At the high end of the line, a big PDP-11-- the model 70-- has been unveiled, revealing a full 32-bit machine, in the hundred-thousand-dollar class, with cache memory and time-sharing. (But what of the even bigger PDP-11 model 85, rumored to be whirring its thirty-six bits unseen in the Marlboro plant under yet another operating system? Will it mean that all the other 11s have had two more bits all this time? Ah, pity that nothing can be said about that here.)

Multiple operating systems are, indeed, the bane of the PDP-11 line. Not only are there DEC's own, like RSTS, RT-11, DOS and RSX, which suffer from a lack of file compatibility and sometimes won't even run the same object code; but now there has arisen a far grander operating system, UNIX.

UNIX-- the name's suggestiveness of harem guards is deceptive-- is really the son of MULTICS (see p. 45). But it was finished in much less time. Like Multics, it's a beauty. Like Multics, it was programmed in a higher language; the language it's programmed in, however, is called simply "C". The language was created by Brian Kernighan, author of a widely-praised book which he audaciously compiled out of incorrect programming examples from other people's books on programming. Unix itself was programmed in "C" by Ken Thompson and Dennis Ritchie.

Unix is a demon. Aside from all the usual features, it allows programs the magic property of splitting. Each program can throw off copies of itself, which run independently and themselves initiate further events. This sorcerer's-apprentice structure comes mainly from a Norwegian language called SIMULA, and also appears in Alan Kay's SMALLTALK language at Xerox PARC: regrettably, there is no room to discuss these here. (For Simula, see Ole-Johan Dahl and C.A.R. Hoare, "Hierarchical Program Structures," in Dahl, Dijkstra and Hoare, Structured Programming, Academic Press.) These features effectively change the character of programming completely. For instance, to simulate a number of objects interacting, the program can spin off a copy of itself for every object, and each copy (mimicking the real-world object), can then respond to its continually-changing environment as required. In other words, this type of language means that programs behave much more like the things being simulated than they ever did before.

SIMULA costs $20,000, and, as it happens, UNIX costs $20,000 (free to non-profit organizations). Unfortunately this raises certain grave questions, since the telephone company (of which Bell Labs is a branch) is not supposed to be in the computer programming business; and those who are in the business are dismayed by the idea of such a competitor.

DEC'S OTHER COMPUTERS

Rather than throw its corporate weight entirely behind the PDP-11, DEC has carved out certain areas in which it is trying to market its 12-bit and 18-bit machines, the PDP-8 and PDP-15. The PDP-8 is being pushed for business applications, with DEC's COBOL-like language; also a very nice version of the 8 has appeared, an excellent home computer, with 8K of core, two floppy disks, keyscope, and wet-printer option; this is the "Classic," at $12,000.

The 18-bit PDP-15 line is still being marketed. Perhaps in order to save it, it is being marketed as a "medium-sized" machine, with MUMPS (DEC's data-base system), with virtual huge memory, and with hot displays.

COMPETITIVE LOOKALIKES

Imitation of DEC computers is continuing.

One firm, Intersil, has put the PDP-8 on a chip for some $300. (However, as it usually turns out, by the time you get all the parts together it costs $3000 after all. But in quantity it's another story, and the individual price will drop soon enough.)

Intersil has also intimated that they are working on a chip to simulate the PDP-11. If so, this will of course bring them smack up against the patent that seems to have knocked out the Digital Computer Controls lookalike, the Lockheed Sue (at least its direct marketing), and engendered a lawsuit against Cal Data. But that remains to be seen. (Same for Godbout's 11 lookalike, mentioned on p. Y.)

FINANCING YOUR PDP

As you may know, you can't in general just rent a computer (except from IBM), but must commit for its full purchase price, since the falling prices of computers mean it will probably have no market value in a few years. (IBM's great power stems in large part from being the only computer company big enough to rent.)

Well, good old Digital Equipment Corporation has finally gotten into the leasing business! They have started a computer leasing company, Digital Leasing, in collaboration with U.S. Leasing. They will lease DEC equipment to individuals of good credit on terms up to seven years. Current rate on a 7-year lease is 2.3 percent a month.

DEC, the Halls

A wickedly funny description of DEC's home factory, fairly accurate, can be found in a neatsy belletristic book called Travels in Computerland by Ben Ross Schneider, Jr. (Addison-Wesley, paper, $6), esp. pp. 73-5.

THE ALTAIR STORY (cont.)

But MITS took it seriously, and offered with the Altair a small but complete line of terminals, disks, printers, interfaces, and, most important, service facilities.

The firm had innovated before, notably when they brought out the first hand-held calculator several years before. Just as they correctly anticipated that demand, they foresaw this one.

They also chose unerringly the right market to begin on: electronic hobbyists and kit-builders. The kit-maker enjoys the challenge of building a machine from only a diagram and a box of parts; and to be far from a repairman holds no terrors for him, for he is the repairman.

The price drop was not as dramatic as it might seem to the general public; nor is the computer quite as cheap as it seems at first glance. Contrary to a public impression, created by IBM and a muddled press over the years, that computers are huge and cost millions of dollars, very good computers have been available lately for a couple of thousand, not counting accessories.

But the accessories present a problem. On that score, the apparent rock-bottom price of the Altair may have been misleading, especially to kit-builders. A computer itself is a limp dishrag without memory, terminals and programs-- all of which pad the cost of the package. By the time you've added 8K memory, a terminal and BASIC software to your kit-built Altair, a thousand dollars has flown ($1400 if you buy it already assembled). Then if you want the disk (and who doesn't), that's at least fifteen hundred more.

Now kit-builders just starting may not see the point of all these fripperies; they aren't used to powers like that of a full computer, so coming to realize the immensity of it all may be a gradual awakening, with many happy soldering experiences on the way. Others may be brought up short as they sense what they're getting into.

This is partly a problem of MITS' trying to reach two consumer groups at once: the kit-builder, who may have thought a computer was a fancy switchbox, and now must enter a world he doesn't know, and the computer sophisticate, who looks at the bottom line for the cost of a complete package.

Indeed, MITS' low prices aren't that low. When it comes to price, they are about 50% ahead of the conventional competition. For instance, their $5000 setup (with terminal and disk) might be taken as roughly equivalent to the DEC Classic at around $10,000 (see p. Y).

But what you usually pay for in this field is service and fringe benefits. The fundamental test, it seems to me, is whether you can come back to the company with your problems. (They even answer correspondence about their customers' computer troubles.) MITS' principal contribution is really in the thought they have given to their market, and the depth with which they are serving it. They no doubt anticipated competitors who would supply accessories and undersell them (see p. Y). But they see the advantage in this: they even give out their mailing list to competitors who sell Altair memory boards cheaper! They were not out just for a quick buck; they appear to be thoroughly committed to full-spectrum computer service.

In eight months, MITS has gone from twenty-five to a hundred employees and sold OVER FOUR THOUSAND COMPUTERS, which is something like two or three percent of the computers in America. Today, the electronic nuts; tomorrow, the world.

Bob Albrecht, caliph of counterculture computerdom, highly endorses Altair Extended BASIC. Says it's terrific.

The main service center for Altairs has been the Albuquerque factory, but the first of their regional service centers has now opened in Nashville.

An Altair assembler is running on the PLATO system (see pp. DM26-7).

MITS prices are quite reasonable. If you buy a kit for anything in the Altair line, it's generally about 25% less than the assembled and fully-checked-out version.

The basic computer kit costs $439 ($621 assembled), but ignore that: it's like a car without an engine, seats or wheel. A complete package (their "Basic 1" set), with the computer, 8K of memory, terminal and 8K BASIC language is $1391. A more high-powered system with 12K of fast memory and double floppy disk is $6650, complete with their Extended Basic. There are many separate items, plans and options; it is possible, of course, to buy a packaged system from them for as much as you want to spend.

THE ALTAIR FACTORY OUTLET

Naturally it had to be in Los Angeles. The first "computer store," it seems, is at 11656 Pico (at Barrington), West L.A., ¼ mile west of the San Diego Freeway; 213/478-3168. Hours are 2 to 8 Wednesday to Friday, 10 to 6 on Saturday and Sunday. It's called the Arrowhead Computer Company, and they stock a line of Altairs.

The BIZ MINI BIZ

Andrew J. Singer (who says he is now "consultant to a small firm of astrologers"), announces to me that

"SOMEONE HAS DONE BUSINESS PROGRAMMING RIGHT."

He has only praise for the BASIC-oriented business systems offered by BASIC-FOUR Corp., to be found in major cities. Not only does the contract spell out what you get in sparkling detail, but the manual is written in English. And Andrew himself couldn't crash the system.

(The Basic-Four setup uses a mini from Microdata Corporation. Microdata itself sells a time-sharing business-type setup called REALITY, which is highly praised by John R. Levine, another young heavy.)

Very much in the BASIC game is Wang Labs; they offer a system with 4K, a BASIC interpreter (in firmware), display and cassette for under $6000. Wang has cleverly farmed the local programming problem out to a network of software houses, each responsible to its customers for their programs.

CUSTOM AUDIO WORK

Two bright guys in New York, Norman Schwartzman and Jerry Fischer, do good custom audio work. They are also an authorized TEAC repair station. NJ Electronics, 212/265-0116, 359 W. 45 (next to the Flying Saucer News).

CHEAP COMPUTERS

You already saw about MITS and the Altair on p. X.

MITS' new computer will be based on the Motorola 6800, and sell in kit form for around $300. But their main commitment is to the Altair, a line based on the Intel 8080, and the customers already into that machine will not be in any way let down, they say.

A computer kit based on the Motorola 6800, with 21K bytes of core, cassette recorder and TV display (32 chars. by 16 lines) is offered for $1745 by THE SPHERE, 96 East 500 South, Bountiful, Utah 84010.

Two computer kits, one built around the PACE and another a Nova lookalike, have eeen announced by Bill Godbout Electronics, Box 2355 Oakland Airport, Oakland CA 94614. He also plans an 11 lookalike.

Or you might get an LSI-11. An LSI-11 buying pool is being formed by Hal Lashley, Southern Cal Computer Society, P.O. Box 987, S. Pasadena CA 91030.

CHEAP TERMINALS

MITS has a 'very low-cost terminal' (the VLCT, yuk yuk) for $170 ($129 for kit).

Processor Technology, 2465 4th St., Berkeley 94710, makes a text display kit for the Altair for $160 (you supply the TV monitor and evidently the keyboard). 64 character per line, 16 lines.

A similar kit at a similar price is made by Southwest Technical Products.

Bootstrap Enterprises, Ann Arbor, are also working on a similar unit, called "The Dumb Terminal," with a color option.

MITS is committed now to building a video terminal, the CT-8096, that will provide both text and graphics. Following specs are not final.

PRICE IS TO BE ABOUT $1000.

It will have a keyboard and video monitor, plug straight into the Altair, and refresh from Altair memory modules-- which may double as regular memory, if you don't mind garbage on the screen.

It will have 24 lines of upper-case characters, 5x8 dots to the character, 80 characters to the line on a built-in monitor. In addition it will offer graphics from bit maps (see p. Z), either 120x120 or 240x240. (The resolution will be switch-selectable, if you have enough buffer memory; a screen of text takes 2K, so does a 120x120 picture, and 240x240 takes a whole 4K.) Buffer memory will also be dividable into separate "pages" of text or graphics; and two pages will be superimposable, interlacing alternate video fields (see pp. DM6-7). Note that refreshment is from random-access, rather than serial, memory, so that multiple fields cannot be overlaid.

OTHER ACCESSORIES

While none has been announced as yet, a music synthesizer that plugs into the Altair will almost certainly be available in 1976.

(Note that this could provide an entirely new form of interactive terminal if used with the Wachspress equipment; see nearby.)

A Selectric interface to the Altair is in the works at MITS.

Altair interfaces to the PDP-8, PDP-11 and Nova have not yet appeared. Why not?

DEC's own floppy disk, for the 8 and 11, finally came out. Price for 11: $3000 for one drive, $4000 for double.

LINCtape, which is virtually the same as DECtape but unpatented, has just come out at $2000 for one drive, including controller and interface to 11 or Nova (interrupt-driven). Note that the unit is compact and rugged, and may be more suitable than disk or cassette for those of us concerned about portable rigs and van-mounting. Computer Operations, Inc., 10774 Tucker St., Beltsville MD 20705. (The bad news: software costs $300 for the driver, plus $750 to DEC if you want operating system RT-11.)

Cambridge Memories, Inc., cleverly sells main memory banks for the 11 which can attach to two PDP-11s at once-- thus connecting the two machines without using DEC's expensive Unibus coupler.

Also for 11s: Formation, Inc., sells a curious programmer's console that traps and displays the last sixteen Unibus addresses referenced; and Fabri-Tek offers a cache memory for the PDP-11/45.

ROLL THE DRUMS...

IBM has announced a laser printer, the model 3800, which prints at 13,360 lines per minute. Like most of their printing machines, it's good. It is basically an electrostatic drum copier, like the original Xerox 914, on which the patents have run out. (Even Toshibafax now makes one.)

Anyway, this spectacular beast writes with a scanning laser on the electrostatic drum. But IBM has cleverly found the way around the problem of spoilage of the drum surface: instead of just a polished metal surface, it's a renewable surface, which is itself changed automatically when required for new images.

Moreover, you can have up to 18 type fonts, defined in 18x24 dot matrices. (THESE ARE KEPT ON A FLOPPY DISK, AND UP TO FOUR MAY BE CURRENT AT ANY ONE TIME.) Fonts are user-designable.

A flash-projector can put business forms on the drum.

Now for the bad news: base price is $310,000, plus extensive extra charges. Also it doesn't make carbon copies, and needs a new box of paper spliced on the end every 20 minutes.

(Canon, of Japan, has out a more modest laser printer that goes at only 4000 lines per minute. And it also plots. Burroughs is said to be trying to get the bugs out of a similar device.)

NOTE WELL: this musical staff was done in a few seconds on the Diablo, using its ordinary type-wheel.

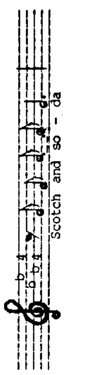

THE HIGHER-PRICED SPREAD: the Diablo can position the characters exactly, and repetitively, anywhere, by fine degrees. This means, for instance, that justified typesetting by computer becomes easy and cheap.

My spacing is now being incremented both horizontally and vertically.
My spacing is now being incremented both horizontally and vertically.
My spacing is now being incremented both horizontally and vertically.
My spacing is now being incremented both horizontally and vertically.
M y s p a c i n g i s n o w b e i n g i n c r e m e n t e d b o t h h o r i z o n t a l l y a n d v e r t i c a l l y.
M y s p a c i n g i s n o w b e i n g i n c r e m e n t e d b o t h h o r i z o n t a l l y a n d v e r t i c a l l y.
M y   s p a c i n g   i s   n o w   b e i n g   i n c r e m e n t e d   b o t h   h o r i z o n t a l l y   a n d   v e r t i c a l l y.

TRULY AMAZING NEWS

One of the buys of computer history is waiting up at American Used Computer, in Boston, 617/261-1100.

Memorex, for some unfathomable reason, built in the early 70s a computer intended to be upward-compatible from the 360/20. But it was not a 360.

Why did they do this? The kind of people who shop around would not buy 360/20s, and the kind of people who buy 360/20s would scarcely leave IBM's skirts at upgrade time.

Thus the Memorex 40 has, quite understandably, been discontinued. And all the ones they had left are waiting for you brand new up at American Used Computer for the heart-stopping price of

$3500.

That price includes 48k bytes of core.

Now for the bad news.

It comes bare-bones, with no software, and no hardware support. You get the wiring diagram with it, and a list of other owners, and you're on your own. AUC does have spare parts, however. And peripherals, mostly more expensive.

Mr. Monoson of AUC told me on the phone that it had 158 instructions, including 64-bit floating point, 32-bit binary. On studying the literature, however, it appears to me that the instruction-set he described is microprogrammed, with the micro-code intended to be read in at startup time. (There are 65 microinstructions.) Maybe you can get the microcode for those 158 instructions and maybe you can't. Maybe you don't care, if you're well enough fixed to handle one of these.

It comes in basic black, 2x5x4 feet, fits in a van, and supposedly does not need airconditioning. Supposedly plug-compatible with 370 peripherals! It's really a sixteen-bit machine, and it has eight sets of eight registers, having been designed to perform up to eight functions simultaneously.

So. 64 main registers, 4K dynamic microstore, 48K of memory, for about the price of a used PDP-11/10 with 4K. Smelling salts, anyone?

THE DIABLO & others

"Diabolo" was a game of the twenties that involved poking a spinning object. Oddly, that's what today's Diablo involves.

Redoubtable Max Palevsky, who brought you Scientific Data Systems (which Xerox bought and recently shut down), Rolling Stone and the movie "Marjoe"-- has another winner, which he's also sold to Xerox.

This is the Diablo company, which first made disks and now makes a sensational printing machine. It has a whirling plastic "daisy wheel" of type, interchangeable, and can type 30 characters per second in either direction, as well as draw pictures-- of a sort.

The basic difference between these printers and conventional typewriters, like the Selectric, is their use of servos rather than ratchets. This means their characters can be positioned in many intermediate positions, unlike the fixed positions available on an ordinary typewriter. For instance, the Diablo can position the type to 1/60 of an inch horizontally and 1/48 of an inch vertically. (Nice for justified typesetting.)

There are now a number of machines of this kind. First came the Diablo printer, officially the HyType I; then the engineers who built that went off and created a competitive printer called the QUME (pron.'kyoom'); now there's an improved Diablo HyType II; Interdata makes a competitive unit, the Carousel printer, with a little print cup; and to make things totally confused, there's a special model Diablo called the 800, which can't be connected to computers but is sold for office use as a "word processor."

A number of companies make terminals in the $5000 ballpark embracing one or the other of these printers. Gen-Com Systems makes one around the Diablo; Anderson-Jacobson makes one around the QUME. Xerox makes its own computer terminal, the 3010, around the Diablo I-- which, it should be noted, can be rented for as little as three months, at $190/month.

The one everybody wants for their computers is called the Xerox 800, but so far that is not available as a computer terminal. It goes faster than the other Diablos and offers typefaces that look beautiful for typesetting; much nicer, it seems, than the types currently available for the other Diablos.

For those interested in just hooking up the printer mechanism, for substantially less money than a whole terminal, interfaces for hooking the Diablo or QUME printers to PDP-8 or PDP-11 are available from Data Systems Design, Inc., 1122 University Avenue, Berkeley CA 94702.

SUGGESTIONS TO XEROX CONCERNING DIABLO PRINTERS. No charge.

1. *Sell the 800 as a terminal, for goodness sake.*
2. *Failing that, make those pretty typefaces available for the others.*
3. *Already you offer black and red ribbons; a blue and yellow ribbon would permit printing PICTURES IN FULL COLOR, a development of great interest to the many computer graphics freaks.*
4. *However, that would require somewhat finer positioning of the platen; say, 1/120 in both directions.*
5. *... failing which, you could put out a "graphic daisy wheel" with intermediate dot positions equivalent to dot positions between those now available.*
6. *Could the Diablo somehow be made to sound less like a dentist's drill?*
7. *How about a portable?*

LOGOMANCY

Dan Hillis and Radia Perlman, of the LOGO group at MIT, are working on a special "preliterate" terminal to allow non-readers (possibly including chimps and gorillas) to program in LOGO, especially on the General Turtle 2500 (see "Minsky's Computer," nearby). Plastic credit cards will have symbols for the various picture and music-box functions. To write a program, or create a movie on the scope, the user will insert function cards in slots. Color coding will be used for program transfer: a red card means "jump to the red subroutine." Since this is MIT, the full recursive power of the system will of course be available. (My hope is that chimpanzees and other little slotniks can be taught recursive program definition. Then will the public wake up to computers being easy?)

MINSKY'S COMPUTER also affectionately known as MARVIN'S COMPUTER, THE FLYING TURTLE

The great Marvin Minsky is renowned on five continents. Dean of the amorphous field of "artificial intelligence," and referred to without ambiguity as "Marvin" throughout computerland, he is a theoretician's theoretician.

But at the heart of every theoretician, I think, burns the dream that he will someday prove the outright, worldly importance of his thoughts. Like Destry, at last he will go to his suitcase and get out his guns, and the audience will cheer.

The great Marvin Minsky has come out blazing.

General Turtle, Incorporated, is a toy company that the team of Minsky and Papert put together to market their educational computer accessories. (See p. 57.)

They've sold a few, but the impact has been modest. And, as a member of the project puts it, "We wanted to get our ideas for education out to the world."

So they decided to build a terminal. But it grew, as terminal designs will. It is now the GTI 2500.

Remember the tortoise and the hare? This is the hairiest tortoise on four wheels. First deliveries this fall.

And here's what you get for your five thousand dollars--

THE KILLER CHELONIAN—

a 16-bit computer like none you ever saw.

- 8 working registers, in addition to PC.
- 32 scratchpad registers (70 nanosecond).
- 250-nanosecond I/O.
- 4K of main memory, 250 nanosecond. (Expandable, of course.)
- 1K OF MICROPROGRAM MEMORY, 40 NANOSECOND, DYNAMICALLY ALTERABLE. (Expandable to 4K.) Likewise 16 bits.

(YOU PROGRAM YOUR OWN INSTRUCTION-SET.)

- Cassette memory, 1 drive.
- Alphabetical display, standard video, with 8x16-dot character generator, 64 characters, DYNAMICALLY ALTERABLE. Also expandable.
- Vectoring graphic display with 2D rotation ("turtle geometry"-- lines are specified not by endpoints but by angles and length). 512x512 resolution, 1 million endpoints/sec refresh.
- Keyboard.

I asked Dan Hillis, a member of the group, about the possibility of installing the 2500 in a van. "Think of it as a recreational vehicle with the van optional," he said.

IN THE CHIPS

What makes possible the computer counter-culture and everything else is, of course, the spectacular development of electronic chip technology, the techniques of shrinking great electronic circuits to almost no size. Electronic rigs that were shoebox-size ten years ago are typically now etched on chips the size of your thumbnail and go for a few dollars, no matter what they contain.

A few years ago, the chips only contained building blocks, such as registers-- units for holding information temporarily. But now in the mid-seventies they have come to contain whole computers, or large sections of them. (The distinction between microprocessors and computers is taken up on p. 44.)

The first biggies were from Intel: the 8008 and then the 8080, a chip that has become the heart of the Altair (see p. X), as well as rival computers.

New computer chips keep coming out; people keep telling me to mention specific ones, but I can't keep track of them. The Motorola 6800 seems popular; it will soon be the heart of new computers from MITS and SPHERE (see p.W and Y). (An augmented and faster copy of the 6800 is reputedly being sold by MOS Technology for $20.) Another interesting computer chip is the PACE microprocessor from National Semiconductor, with four working registers and a ten-word stack; with 16K memory it costs $500. (The PACE is hidden in an automatic drink mixer and booze inventory controller from Electro Units Corp., San Jose, Calif. Adjusts prices to hours and can even water the drinks precisely. Claimed to make absentee ownership of bars practical.)

Because of chips, the price of computer main memory is collapsing apace. Something like a dollar a word in the sixties, it is something like like a dime a word now. But Intel now offers a storage chip holding 16K bits for $55, which is 3¢ a bit, and a friend of mine estimates that memory chips will cost 1/10 of a cent per word in 1976.

These cost collapses cause many to predict the end of disk and tape. But that's premature. While these zappier chips hold a lot for a little, their contents disappear when the lights go out. Until laser-punched tape comes along, disk and magnetic tape will be very much with us as long-term and backup storage devices.

Because of the action in chip technology, a potentially important movement in computer design may have been passed over: the "macromodules" developed at Washington University in St. Louis by Wes Clark (father of the original DEC modules), and associates.

The basic idea of the macromodule approach was to have computer subsections that were completely interpluggable. With them you can build any computer, to your own design, in a couple of days. The system exists now and it works just fine: counters, registers, memories can be attached quickly by cable.

Unfortunately, the cost is high and they haven't found a manufacturer. With chip prices falling, and chip know-how widespread, it's hard to justify charging ten or so times as much for components just because they can be plugged together faster. (Just as unfortunately, everything in the macromodule system is built on sections of twelve bits.) For this reason the St. Louis folk are having trouble getting commercial sponsorship. However, perhaps some bright hungry chip company, reading this, would like to get into the macromodule game. And presumably whittle the module down to the now-universal 8 bits.

EQUIPMENT PAGE

WEIRD AND SEXY COMPUTERS

The most glamorous computers being built today all seem to be openly called by their developers' names: the Greenblatt computer, Minsky's computer, the Amdahl machine, the Cray computer.

THE AMDAHL COMPUTER

The Amdahl computer or System 470, a super-computer of the 360 series by one of the guys who designed them originally-- see p. 41-- is now available from Amdahl Corporation, 1250 East Arques Avenue, Sunnyvale CA 94086. (They are now advertising for systems people who know the insides of OS/MVT, VS, etc.). The first 470 is up and running at NASA's Institute for Space Studies, Columbia U. But IBM is said to be readying one of their famous "knockout" machines to do it in. (Datamation, July 75, 96.)

OCTOPUTSCH

Of course you've thought that hardwired setups were for sloppy analog types of thing. But here now we have THE CHESS MACHINE, under straightfaced construction at the MIT AI Lab, which will provide HARDWIRED THREAT ANALYSIS. Yes, its advanced perceptron architecture will supposedly be capable of analyzing threats to any given position in a GRAND PARALLEL FLASH. The impact of this astonishing development on the world of Electronic Chess, or anything else, for that matter, is totally impossible to predict.

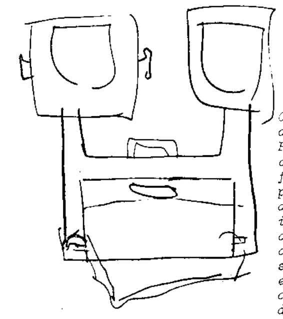

Over a very nice lunch at Rodity's in Chicago, Prof. Minsky and I discussed possible styling for his computer. He particularly liked the arrangement suggested in this sketch: a fold-down keyboard and the displays sort of on poles so they could be seen easily through a crowd of bystanders. The handle would only work, of course, with the scopes removed. We'll see later what it finally looks like.

THE GREENBLATT MACHINE

Unsatisfied with the structure of normal computers, they are building at MIT's AI Lab a computer whose native language is LISP. It will have 32 bits with virtual memory, and execute LISP like a bat out of hell.

In a refreshing reversal of trends, it will be for one user at a time. "Time-sharing is an idea whose time has gone," chuckles one participant. (Project MAC, where time-sharing grew up, was there.)

THE CRAY COMPUTER

Seymour Cray, master computer builder, created the 6600 system for Control Data. Indeed, he had the audacity to require CDC to build the computer factory on the property adjoining his own estate in Chippewa Falls, Minnesota. Now that he's broken off to start his own company (with money from CDC, among others), the new computer factory adjoins his estate on the other side. The Cray-1, another supercomputer, is nearing completion there.

AUDIO TRANSDOUCHER

Patent #3,875,932 has now been issued for How Wachspress' electronic sex machine or whatever it is (you saw it first on p. DM9). In the illustration we see it tickling a shmoo.

After you send Wachspress his fifty-buck royalty, you can either buy the kit or a pre-built model. Concave or convex, as the poet says. (Etchings are antediluvian and waterbeds are commonplace; as an invitation, what more incisive comeuppance could be profferéd?)

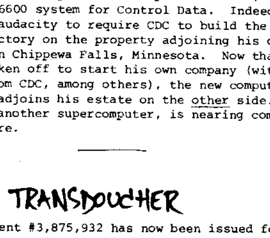

Speaking of Wachspress, it seems that the unusual I/O equipment offered by the Federal Screw Works (Troy, Mich.) is only a voice output device.

Surprisingly, a voice input device is now commercially available from Threshold Technology, Inc., Cinnaminson, NJ. For $10,500 you get a device that will recognize 32 spoken words, and microphones. (Each user has to train it on his 32 words, but separate vocabularies may be stored on the computer for different users or purposes. This is still some way from the fabled "talking computer"-- see pp. DM 13-14 for problems and objections-- but it's undeniably a useful step.)

GRAPHICS PAGE

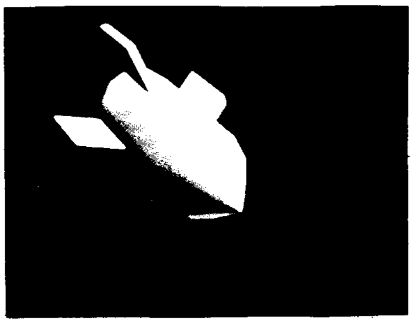

The halftone system of HUMRRO, rumored on p. DM38, is real. Clever indeed: it divides the half-tone problem into two parts, one the original picturing of the scene, the other its presentation in the terminal. That means that their system permits one central image generator to send out pictures to as many terminals as desired. Unlike the Watkins Box (see p. DM37), whose half-million-dollar opulence can be poured only on a single user at once, in this system the central resource can be distributed among various users, with each one's picture changed intermittently, or poured on a single user for full animation. Currently it runs in Fortran, transmitting encoded pictures to the unusual terminals required (built around Trinitrons). But a special central processor is foreseen.

The system is called CHARGE, and Ron Swallow, its developer, is indeed a hard charger. (Software: Bill Underhill and Roger Gunwaldsen.) Swallow's game isn't movies or engineering graphics; he wants CHARGE to compete head-to-head with PLATO (see pp. DM26-7). And at the prices he's talking about-- $5000 per terminal and $150,000 for the central processor-- who knows?

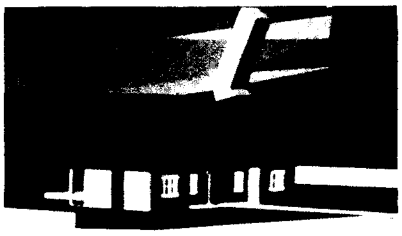

UNREAL ESTATE: for relaxation, Ron works on the "dream house" he keeps inside the system.

Millions of people saw computer graphics for the first time on the PBS "Ascent of Man" series, where a screen drawing of Early Man's skull was seen to rotate and gradually change in its features. This was startling even if you know about computer graphics, since it seemed to be proceeding from complex data concerning the entire skulls and their changes.

Not so. Actually what you saw was a series of skull drawings by Peter Foldes, a Parisian artist, with the computer generating transitional drawings between them. (Indeed, though you saw Prof. Bronowski next to the screen, you did not see him next to the screen at the same time the drawings were changing-- because that had to be filmed very slowly.)

The system was created by Nestor Burtnyk and Marcelli Wein, of the National Research Council of Canada. It currently runs only on an SEL 840A. (It was also used by the National Film Board of Canada for creating Foldes' splendid film "Hunger.") They can preview by rolling through bit-map video on a moving-head disk. (See Burtnyk and Wein, "Computer Generated Key-Frame Animation," J. SMPTE, March 71, 149-53.)

What about the animated figure that talks to Joe Gariagiola before baseball games? Haha. That's a rubber puppet matted in from a black box; the guy who does the voice works the mouth.

Many unlikely individuals have stormed that heartbreak town of Hollywood, leaving sadder but wiser-- but Ivan Sutherland, dean of computer graphics? Well, having found that the moviemakers are not ready for image synthesis-- the dreamsmiths unprepared, as it were, for the Total Forge-- he is sojourning at the Rand Corporation.

A fella named Charles McCarthy, of suburban Chicago, bought the "Computer Eye" from Spatial Data Systems, and will do mail-order picture conversions. He'll convert your favorite snapshot to a printout of the same subject made of light and dark letters. If you're interested in having the actual grey-scale data for processing in your own computer, inquire.

The Möbius Group, Inc., P.O. Box 306, Winfield IL 60190.

Want a computer-controlled videocassette recorder? The model to ask for is the Sony 2850, costing (gasp) some six thousand bucks. An interface to the PDP-11 is made by CMX Systems, 635 Vaqueros, Ave., Sunnyvale CA 94086.

Incidentally, scaled-down CMX editing setups are beginning to get around. For instance, they have a small setup in the pleasant offices of DJM Film & Tape, 4 East 46, NYC: three of the above Sonys and the CMX Model 50 control setup, using a PDP-11 and keyscope. Though prices are by the job, the basic charge is $75/hour. (Note that the big CMX setup, with a disk, is the model 300.)

Video Disks, Supposedly TURN, TURN, TURN

Since the forties, there have been continual announcements that video disks-- movies you play on your TV off a record-- were right around the corner. Earlier this year they were supposedly going to be available before Christmas. Now they might be on sale, "on a limited basis," in 1976. (TV Guide, 16 Aug 75, p. 7.) Because of the grave difficulties of engineering-- inaccuracies in punching the center hole mean the track can't help being off center, for instance-- some of us are skeptical.

Two systems have been confidently announced. Philips, the firm that gave us the audio cassette, has a system that will follow the spiral track on the disk from underneath with a laser. The disk turns at 30 revolutions per second, or one turn per TV frame, so it can supposedly freeze on one frame when desired.

The other system is from RCA, which has a long history of me-too announcements, but at least two of them made it big (the 45 record and the color TV system now used in the USA), so RCA should not be dismissed out of hand. Their disk system will supposedly go at 450 rpm (7.5 revolutions/second), but they still mean to track it with a needle. The man from TV Guide says he's seen it and it works perfectly, but I would personally look for hidden wires.

(MCA, an entertainment conglomerate, has hitched up with Philips and printed a catalog of all the movies they will supposedly make available on disk for the "MCA-Philips" system-- such as Destry Rides Again for around ten bucks. This is probably just a bluff; with the price of audio records what they are, no way is a movie going to cost ten bucks. But it makes RCA look weaker, which is probably the purpose.)

MOVIES FROM YOUR COMPUTER

The prospect surprised them, but MAGI (see p. DM36) allows as how they might let you make movies on their over-the-phone movie-making setup (sketched on p. DM36). Price to capable outders, if the software meshed, would be about $50 an hour. (Six hours makes one minute of film, not counting the phone bill. Cheap if you know movie economics.)

Meanwhile, John Lowry, at Digital Video Laboratories in Toronto, has been developing high-quality video suitable for transfer to theatrical film. He and they have developed a 655-line color system-- with heavy digital enhancement (see "Picture Processing," p. DM10). I scarcely believe my notes, but I saw it, and wrote down that it was comparable to 35mm studio color. The day of "electronic cameras"-- that is, film-quality video-- may be upon us soon.

About 1972, there was announced an electronically-controlled color filter that could change to any hue in nanoseconds. That would be just what we all need for color movies from COMs-- but what happened to it?

VECTOR DISPLAYS

At the high end of things, a firm called Three Rivers Company has come in with a 3D vectoring system (competitors discussed p. DM30). Supposedly they can pack a lot more lines on the screen.

The price of the GT40 display (see p. DM21), which all in all is one of the best displays on the market, has just dropped to $6500. To disguise this price drop, DEC gives you the smaller tube and no keyboard.

And at the low end, a firm called Megatek in San Diego offers line-drawing CRT controllers for $1000 to $3000. All permit animation. You have to supply the oscilloscope. Their equipment plugs into the PDP-11 or the Nova, or in one case connects in tandem to an ASCII time-sharing terminal (!).

The 11 and Nova models work directly from BASIC; your program in Basic puts line lists in the device's buffer memory. The time-sharing model converts incoming line lists from ASCII to binary and stores them internally. 256 lines with 8-bit resolution cost $1900, $1100 and $1600 for 11, Nova and t-s respectively; 1024 lines with 10-bit resolution cost $2800, $2000 and $2500 respectively. (Nova and 11 models can be completely updated in two refresh cycles, yielding as much animation as anyone can decently expect for the price. Software is supplied to provide display output from Nova, PDP-11 or time-sharing BASIC; also t-s Fortran.)

Meanwhile, for the hands-on electronics guy, Optical Electronics, Inc. makes all kinds of rotation modules. You can build your own 3D rotation setup out of their modules for a couple of thousand; but, of course, the fancy digital I/O for high-speed refreshment is not available. An interesting capability of the OEI equipment, though, is that you can build 4D- or even 5D-rotation systems out of their modules. Hmmm.

PLATO news

Excellent manuals on the PLATO system and TUTOR language are now available from CERL, University of Illinois, Urbana.

The next generation of PLATO terminals is coming down the line. The microfiche projector is withering away, as was easily foreseeable; meantime, steps are being taken toward a more high-performance terminal, by putting a computer in it. This is being done both by Jack Stifle, who has done it with the Intel chip, and Roger Johnson, who has the panel interfaced to an 11. (11 fans please note the implication: it is possible that the interface may be marketed.)

Meanwhile, PLATO-like terminals (the model AG-60) are about $5000 from Applications Group, Inc., P.O. Box 444B, Maumee, Ohio 43537. Note that these have standard non-PLATO interfaces and standard keyboards, but the Owens-Illinois plasma panel (erroneously called Corning elsewhere in the book) blazes in all its glory.

BIT MAPS

The main development in computer graphics in the last year has been the sudden upsurge of the bit-map approach to computer display. While the approach, and equipment for it-- like the Data Disk system-- have been around for some time, the falling price of electronics, especially in the memory area, have made it abruptly the cheapest and thus the most popular type of computer display for graphics.

A "bit map" is a series of dot positions, or bits, recorded in some form of fast memory and read out in sync to a conventional scanned video system (see pp. DM6-7). The one bits stand for dots or little squares, the zeroes for nothing, and the video system brightens the corresponding zones on the screen. This method has certain disadvantages-- parts of pictures cannot be automatically distinguished or separately animated, as with subroutining display (see "The Mind's Eye," esp. p. DM23)-- but for the money it's great. Sizes given refer to the number of squares in the rectangle of the picture.

BLACK-AND-WHITE

An off-the-shelf bit-map system for the PDP-11 or the Nova is available from Intermedia Systems, 20430 Town Center Lane, Cupertino CA 95014 ($2750 or $2500 respectively). May be ganged for grey-scale or color. It's 256x256.

For the Altair, the forthcoming 8096 display (see p. Y) will have 120x120 or 240x240 bit-map graphics, for prices starting around $1000.

COLOR

Extra bit maps, plus electronics, can get you color; if you double the number of bits you can double the number of available colors on your display, ad infinitum.

On the small side, 64x64 color wll shortly be available for the Altair from the Digital Group, Denver. A 128x128 color bit-map system for the 11 has just been announced by DEC (for "nuclear medicine" of all things-- but they will part with it to anybody for 8 or 10 thousand (not yet fixed)). They stress that this will be the first of a modular series of bit-map displays, with plugins for different degrees of resolution and different character generators.

Ramtek and Comtal both make 256x256 bit-map systems, priced in the $16,000 area.

Above this resolution special TV systems tend to be necessary. Both Ramtek and Comtal make very expensive systems for the purpose, using solid-state and disk respectively.

You may or may not have heard of the Advent TV projector, the most glorious TV thing there is. It costs $3500 and projects a four-foot picture in the best TV color you can find. A lot of guys are bit-mapping to it.

At MIT they've got bit-map color on the Advent at better than 400x500 resolution. (An option planned for the Flying Turtle (see p. Y) will allow its core memory to be used with the Advent as a bit-map display refresher.) At Comtal they're going for 1000x1000 on the Advent, rejiggering the electronics from scratch.

The most spectacular demonstration of bit-map color so far has no doubt been the film done by Dick Shoup et al. at Xerox PARC (see p. X), showing the super animation that's possible when big-computer resources are given over to bit-map animation. Their system is 600x800.

YOUR BIG DISPLAY PANELS

All those scoreboards and wisecracking light-grids, now that they are computer-controlled, raise all kinds of possibilities for non-frame animation. The big ones cost in the millions; a small one for shopping centers costs a hundred grand (Millenium Info Systems, Santa Clara CA).

Within a year or so, though, you ought to be able to get a nice animated display-panel of some sort for the side of your van, assuming you've got the computer inside.

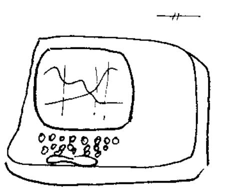

2D FROODY

A surprise something-or-other from DEC, the VT55, represents a breakthrough of some sort. But what were they thinking of?

"Graphic capability" has been added to an ordinary upper-case keyscope. Specifically, the ability to make two graphs, i.e., two wiggly lines (no more) somewhere between the left and right sides of the screen. You can also shade in under them, and add coordinate grids. It's $2500, and obviously great if you're bonkers for 2D graphs.

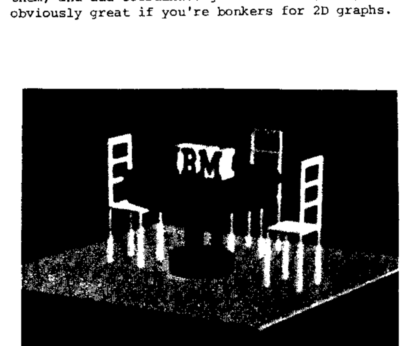

GUESS WHO'S COMING TO DINNER

IBM, which did not take part in its development, is sponsoring a $100,000 CHARGE installation at the University of Waterloo, in Canada.

NOTE FROM THE AUTHOR & PUBLISHER

Labor Day, 1975

MORE THANKS

In banging together this volume originally, I omitted thanking Hesh Wiener, brazen & brash young old-fashioned new editor of Computer Decisions, who has changed that publication from stolid to peppery.

Thanks also to my good friend Robert W. Fiddler, Esq., patent attorney and still an ex-philosophy professor at heart, for many delightful and witty conversations on problems of patent, copyright and the vagaries of intellectual property. Any harebrained ideas on these topics expressed here, however, are almost assuredly my own.

For much of the information in this supplement I am grateful to Bob Albrecht of PCC, mentioned here and there.

Finally, special thanks to Commander Hugo McCauley, better known to you as Hugo's Book Service, for his yeoman performance in shipping out the books-- not to mention carrying them up and down stairs, typing the mailing labels, checking for bad ones, and sending out all those notes of apology when we were out of books again and again and again. And to long-suffering Lois and Megan McCauley, my especial gratitude.

WHATEVER

The sea-to-shining-sea Nelson Empire now consists of a lot of unsold books, a 1K Altair and a second pair of shoes. My scheme for taking on Apprentice Generalists may have to wait awhile. So may Computer Lib, the film. But just wait.

Speaking of which, what about this book, hey, now?

Eventually there will be a new edition. Yes, the type is horrendously small, and that will have to be fixed. But that involves new negatives for every page, an expenditure of thousands of dollars, and some reconsideration of how this should all be set up.

There have been several interesting plans. One was to split the contents of this book into three books, add material, enlarge the type and have them each this size and price. Tentative titles were Computer Lib /Dream Machines, Computers Arise!/Computers Arouse!, and Guerrilla Computing/Electronic Monkeyshines. Sample cover, for Guerrilla Computing: King Kong climbing the front panel of a 370 holding Patty Hearst. (I also daydreamed about putting out a 10-volume encyclopedia in the same format, embracing psychology/sociology, biology/evolutionary strategy, history (as strategy)/more history (as mood and feeling), revolution versus continuity (a two-sided position paper)... the Gem-Maniacal Encyclopedia™. But reason has prevailed, and such forays have been postponed indefinitely.

The present plan is for Computer Lib to be rewritten and reset in bigger type, at least 256 pages, with at least 8 color pages and color cover. (We're talking about fall '76 or later.) Price will have to be $15. If you think that's a ripoff you can still get this one. (A number of people have complained to me about the $7 price tag of this volume. Have they ever bought other books?) Later I would like to put out an anthology of my favorite articles in the field, using the Computers Arise!/Computers Arouse! title and format, and with some good 3D if possible. In any case, I want to stay in the publishing game; I haven't had so much fun in years. Other projected volumes include The Inner Beyond, by Sheila McKenzie; Dirty Driving and the Strategy of Traffic by "Driver Ed;" and The Nelson Computer Glossary. Soon I hope to be able to typeset from my own computer, and possibly to share this facility.

This has been a most interesting year. I have been pleased to meet, and otherwise enjoy, the variety of clever, charming and/or lubricous persons who have sought me out since the book first appeared; as well as all the speaking engagements, soirees and whatnot.

I am delighted to receive relevant material and communications of any kind, although problems of time, disorganization and mood often preclude a Personal Type Reply.

It has been a real lift for my morale to share some of these ideas and enthusiasms with a wider public at last. It is you, finally, who have to care; and I am very glad you do.

As to the most important matters, there is a news blackout for the indefinite future. Please stand by.

Next year in Xanadu.

TN

DM3

This book (both sides) is based in part on my talks at or before the American Chemical Society, the American Documentation Institute, the American Management Association, the Associated Press, the Association for Computing Machinery, the Central Intelligence Agency, the Institute of Electrical and Electronics Engineers, the Printing and Publishing Association, the Rand Corporation, the Society for Information Display, the Society of Motion Picture and Television Engineers, TIME Incorporated, Union Theological Seminary (the Auburn lectures), Xerox Palo Alto Research Center, and various art schools, colleges, universities and Joint Computer Conferences.

ACKNOWLEDGMENTS

Everybody at Chicago Circle Campus has been very sporting about this project. I am grateful not only for the encouragement and assistance of various individuals (especially Joseph I. Lipson, David C. Miller and Samuel Schrage), but for the atmosphere of support which has made this possible. My thanks to the Department of Art and the Office of Instructional Resources Development for freeing me from teaching duties, to the Computer Center and the Department of Chemistry for letting me use pictures of their equipment, and everybody for their encouragement.

I would like to thank the Walt Disney organization for their permission to depict their wonderful characters, and everyone else who furnished materials and permissions for the things herein.

Thanks also to those who looked over some of the material, especially Herbert Grosch of Computerworld, Dan McGurk of the Computer Industry Association, and William Rodgers.

I am particularly grateful to the many who have explained computers to me over the years, especially Dave Denniston, Robert Fenichel, Andrew J. Singer, John R. Levine.

My thanks to Tom Barnard for some of the early typing, and for the Porta-Xan.

I am grateful to Computer Decisions magazine for their good will, and help in researching computer image synthesis.

My roommate Tom DeFanti, mentioned elsewhere in this book, has been considerate beyond the call of duty in giving over all the first-floor rooms of our house to this project for six months.

My thanks finally to the many others whose good will has kept me going, in particular my former wife and eternal friend, Deborah Stone Nelson.

Special greetings to my friend and neighbor, Mrs. John R. Neill: I hope you enjoy the uses which your husband's illustrations of Tik-Tok the Machine Man have found here.

Lastly, for her contributions to morale (and for not footprinting the pasteups), let's have a warm hand for Pooky the Wonder Dog.

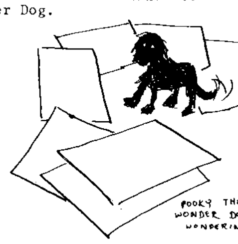

The occasional Oz illustrations are all by John R. Neill, from various out-of-copyright Oz books by L. Frank Baum, especially Ozma of Oz and Tik-Tok of Oz. Tik-Tok, the Machine Man, is the figure to whom occasional allegorical significance is attached here by juxtaposition.

The Oz picture in this spread is from The Patchwork Girl of Oz.

Thought you might wonder.

ARE WE MATERIALISTIC?

Persons of sagacity have been saying for some time that we are materialistic.

In an important sense this is not so.

The machines, and toys, and involvements we buy into, are in but a small proportion of cases owned simply as scores, for their cost as consumption symbols.

Rather, we buy things that REPRESENT IDEALS, hoping ourselves to partake of some abstraction or image-- the Playboy man, the Smart Businessman, the Clever Homemaker.

Each product tries to tell us it is the keystone of a way of life, and then, at least at that moment of purchase, we step into, we embrace that way of life, covering ourselves with the feeling, the aura, the magic we saw in the commercial.

This is not materialism. It is wishful grasping at miasma. (Following sentence optional.) It is communion, with the object seized simply the Objective Correlative of a hoped-for transsubstantiation. (Sorry.) It's a seeking, not to possess, but to belong.

GREAT AMERICAN MACHINE-DREAMERS

D.W. GRIFFITH-- took the movie-box and created the photoplay, no longer a twisted stage production.

WALT DISNEY-- created a hypnotic pantheon of kindly and innocent semi-animals, sentimentally universal, generally acceptable.

JOHN W. CAMPBELL-- as author and then editor of Astounding, turned American science-fiction from the Buck Rogers space opera to the human story, built around thought-out premises and structures.

IVAN SUTHERLAND-- programmed and systematized a computer setup for helping people think and work with deeply-structured pictorial information. (See p. DM 23.)

DOUG ENGELBART-- foresaw the use of computer screens as a way of expanding the mind, and over the last decade and a half has brought about just that.

And more, and on.

ANOTHER QUICKIE

Compare Alice, when she gets to Wonderland ("Deary me! Curioser and curioser!") with Dorothy Gale, transported to Oz ("How do I get back to Kansas?!!!") Fantasy ties in with everything, including American git-out-n-do-it.

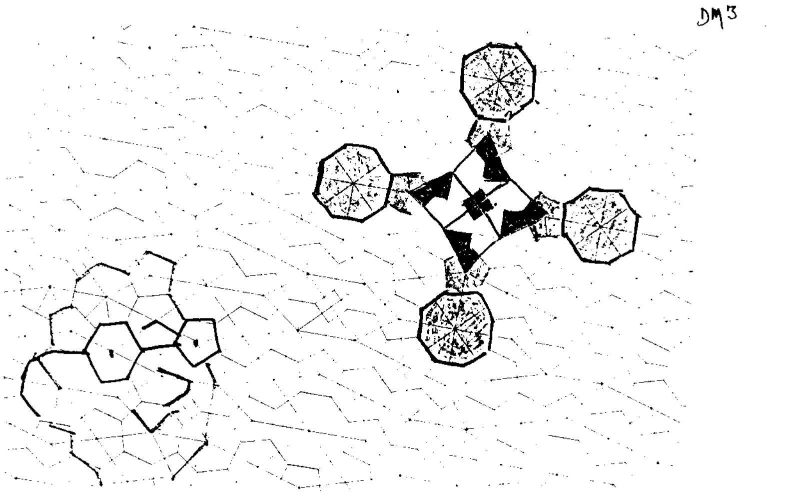

THE GREAT AMERICAN DREAMER

I already said on the other side that the computer is a Rorschach, and you make of it some wild reflection of what you are yourself. There is more to it than that.

America is the land where the machine is an intimate part of our fantasy life.

Germans are too literal, they can get off on well-oiled cogs. The French are too vague. (I've noticed that German science-fiction magazines had covers of machines and planets; French science-fiction magazines, of dragons and people with wings. Our science-fiction covers show people with machines. Intimately, emotionally.) German fantasy is icy and impersonal, French fantasy too personal, and American fantasy is splat in the middle, uniting both: man and machine, means and ends, emotion and details.

Men always longed to fly, but it was here that they first did. This is the land of the MOVIE, a fantasy fabricated with endless difficulty using various kinds of equipment.

The mad tinkerer is a fabled character in our fiction.

This is the land of the kandy kolor hot rod, the Hell's Angel chopper, the drive-in movie. And the wild hot-rod, in fact, is just the flip side of the deep-carpeted Cadillac: each is a fantasy, an extension of its owner's image of himself in the world.

Thus it was not an historical accident, but utterly predetermined, that in the hands of Americans the computer would become a way of realizing every conceivable wild fantasy that was dear to them.

This is perfectly all right. This is as it should be. This is the best part of our culture. Not "Let a hundred flowers bloom," but "Let a hundred gizmos clank." This has sped immeasurably the imaginative development of many different things we might want. I try here fairly to explain a few differences among them.

There is just one problem with all this. Now that all these things exist, or come nearer to existing, which ones will other people want? What will it be possible for everyone to have? And how can we tie all these things together?

(Note: this thesis is being advanced only half-seriously. There have been a number of exactly-dreamful Frenchmen, and for this three-nationality split to be really true, they would all have to have come from Alsace, next to Germany: Jules Verne, Daguerre, the brothers Montgolfier, the brothers Lumière, to name a few.)

CONTENTS of Dream Machines.

perhaps it is time

THE LEGEND OF HYPER-MAN

In the fantasies of their subjects, which they feel are the precursors of new artistic images that will in turn actualize themselves as another form of being, Masters and Houston see a new hero figure constantly recurring. This new hero is not the old "hero of a thousand faces," the individualist who suffers, dies, and is reborn, slaughtering and conquering along the way. Instead, he is Protean, capable of infinite changes in appearance and style, a magician, a Balthazar bringing gifts. He ruptures categories and confuses the senses, and in doing so he holds out the promise of fusion on a higher level.

If such a hero were to become the model for the approaching age, he would probably not be the founder of a mass movement or the god of a new religion. He would be more elusive, more changeful than his predecessors. He would be a sorcerer who treats the external world and the internal world on equal terms, giving spirit to the former and flesh to the latter. He would be a master of paradox and a player of games, speaking a new language. His one prayer might be the lines of Blake:

May God us keep
From single vision
And Newton's sleep.

-- *Kenneth Cavander, "Voyage of the Psychenauts." Harper's, Jan 74, p. 74.*

OUT THE DOOR IN '74

I have wanted to write an introduction to computers, and a separate book on Fantics, for years. But the idea of binding them back-to-back in a Whole Earth format, with lots of mischievous Enrichment material, didn't hit me till Jan 73. I have tried to add all the stimulating and exhilarating stuff I could find, especially personalizations, as on the other side; computers are deeply personal machines, contrary to legend, and so are showing-systems. I regret having to throw so many of my concerns into comic relief, but I hope that some readers will sense the seriousness below.

The final inspiration for this book came from something called the Domebook, that tells you straightforwardly how to make Geodesic Domes. And of course I'm blatantly imitating, in a way, the wonderful Whole Earth Catalog of Stewart Brand. As I think back, though, the tone also comes in part from Pete Seeger's wonderful banjo book, and Tom McCahill's automobile reviews in Mechanix Illustrated. As to the last aspect, that of taking my case to the public because the experts won't listen, the only precedent I can think of is Maj. Alexander de Seversky's Victory Through Air Power, telling the country how he thought we should win World War II.

This project, simple in principle, has been infinitely bothersome. Self-publication was necessary because no publisher could have comprehended the concept of this book; I heartily recommend Bill Henderson (ed.)'s The Publish-It-Yourself Handbook, $4 from The Pushcart Book Press, Box 845, Yonkers NY 10701.

The present product is not the book I had meant to write. Most is first-draft; how the sentences do run on. (Believe it or not, I do not like underlining things-- a first-draft expedient.) Fact-checking and bibliographies had to be largely abandoned. Better planning could have increased type size; and so on. Half the manuscript, and the glossary, had to be kicked aside; including sections on movies, "multi-media," microfilm, training simulators, augmented stage productions of the future, and goodness knows what. Sorry for all that.

... WITH A LITTLE HELP FROM MY FRIENDS

This project could never have been completed without the dedicated and extraordinary efforts of my wise and warm friends Sheila McKenzie and Wade Freeman, both faculty members at Circle, who have my deepest gratitude. They gave months and weeks of their good time to the tedious aspects of this project (which I continuously underestimated.) I hope it has been worth their work as well as my own. Ms. McKenzie, whose concern for intelligent change in education drove her to boundless efforts on this project, has also my deepest admiration.

The sad thing about it all is that 90% of these efforts are unnecessary. A decent computer text system (of which only a couple exist as yet) would have obviated all the finding-and-retyping problems. I feel deeply for everyone who has trouble writing by conventional means, and who wouldn't if only decent systems were available.

NT

APPARATUSES OF APPARITION

It seems different companies are all the time introducing wonderful new devices that will revolutionize, uh, whatever it is we do with, uh, information and stuff. Things you'll attach to your TV to get highbrow programs or dirty movies. Microfilm devices that will shrink the contents of the Vatican Library to a dot on your glasses. Goggles that show you holographic color movies. A pince-nez that lets you see the future. And so on.

Reading Popular Mechanics or the Saturday review of patents in the New York Times, you get the idea of Something Big, New and Wonderful About to Happen, so we'll all have access to anything, anytime, anywhere.

But it's been that way for decades, and with certain exceptions hasn't happened yet.

Here are some things that have caught on, and are mostly familiar to us all.

Book. Newspaper. Magazine. Radio (AM). Phonograph record (78). Tape recorder, ¼". Black-and-white television. Radio (FM). Phonograph record (33). Phonograph record (45). Color television. Tape cartridge (¼"). Tape cassette (Philips, ca. 1/8"). Stereo records and tapes. Oh yeah, and movies: 35mm, 16mm, 8mm, Super 8mm. Carousel projectors. View-master stereo viewers.

Here are some things in the process of catching on (and not assured of success): Quadrophonic sound. Dolby. Chromium dioxide tape emulsion. Super 16 movie format.

But for everything that did catch on, dozens didn't. Some examples: 12-inch 45 rpm records. 11.5 millimeter movies. RCA's ¼-inch tape cartridge, which became a model for the much smaller Philips. Wire recorders.

Then there are the things that caught on for awhile and went away. Stereopticons (and their beautiful descendant, the Tru-Vue, which I loved as a kid). Cylindrical recordings. Piano rolls. And so on.

Then there are the video recording systems. CBS' EVR died before it got anywhere. RCA's SelectaVision isn't out yet. 2-inch quad is standard in the studios, ½-inch Porta-Pak is standard among the Video Freaks, and it looks like Sony's 3/4" cartridge will win as the main sales and storage medium. (The Philips system here looks as though it won't make it, and 1-inch is dubious.) But what's this we hear about video disks (twenty-five years after they announced Phonevision. Ah, well.)?

The thing is, so many of these things seem to sound alike. They all mention "information retrieval," education, technology, possibly "the information explosion" and "the knowledge industry." Press releases or effusive newspaper articles may use phrases like "space-age," "futuristic," "McLuhanesque" or even "Orwellian" (though few people who use that word seem to know what Orwell stood for; see p. DM59).

And the intimidating company names! Outfits with names like General Learning, Inc., or Synergistic Cybernetics, Inc., or even Communications | Research | Machines, Inc. Surely such people must know what they are doing, to use such scientific-sounding phrases as these!

Then there are the business magazines. In the late sixties they were talking about "The Knowledge Industry" (a fiction, it turned out, of an economist's lumping a lot of things together oddly). Now they talk about the Cable TV outfits and the Video Cartridge outfits as though they're the cat's pajamas.

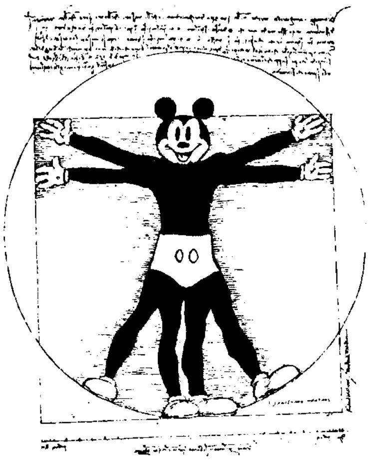

Emblem of 2d International Animation Film Festival in New York, Jan 74. © Walt Disney Productions.

THERE'S SNOW BUSINESS LIKE SHOW BUSINESS—

You Can't Tell the Experts Without They Program You
(Cf. "Calling a Spade a Spade, p. 12)

BABEL'S IN TOYLAND

Guy's Background	Tell-Tale Phrases & Jargumentation
Television:	"Media" (meaning television);
1. Video freaks	"Software" (meaning videotapes).
2. Network People	"Programming" (meaning competitive scheduling); "Software" (meaning fixed-length TV shows).
3. Cable Operators	Head end, upstream & downstream, back-channel, "interactive TV" (meaning any form of interactive computer system they can get in on).
Math/Engineering	Information theory, channel capacity, bandwidth, feedback, anything complex and irrelevant.
Display Engineering	Full duplex, echoplex, aspect ratio, scroll, cursor; "information transfer" (meaning telling or teaching); "data delivery" (act thereof).
Programmed Instruction, Computer-Assisted Instruction	"Software" (meaning sequential or branching tell-&-test materials); "Programming" (creating these); reinforcement schedules (meaning presentational order); "inputs" (meaning ideas and information); "feedback" (meaning replies); "simulations" (meaning pictures or events a user can influence).
Publishing	"Software" (meaning books).
Advertising, Public Relations, Marketing	"Demographics" (meaning factions); campaign strategy (meaning how you hit a market); "penetration" (meaning extent to which your stuff catches on); "Programming" (meaning anything whatever).
Artificial Intelligence	Anything mathematical; theorems, discriminators, neural nets; "programming" (meaning setting up anything very complicated and incomprehensible).
McLuhanatic	Global Village, mosaic, surround; "Programming" (meaning psychological indoctrination); anybody else's terms, dynamically infused with new senses.
Nelsonian	Medium (meaning stabilized presentational context); Writing and Creation (meaning thoughtful production of something presentable, whether sequential or not, in a medium); "Programming" (meaning giving exact instructions to a computer); media integrity, inventions & conventions; hypertext, thinkertoy, fantics.

Having spent some considerable time around and among these areas, I have developed considerable cynicism and a bad case of the giggles. Originally it all seemed to fit together and to be leading somewhere, but talking to people at all levels, and either giving advice or trying to interpret the advice of others, I am convinced that what we have here in this whole audio-visual-presentational whizbang field is nothing less than a very high order of collective insanity. The strange way companies adopt and drop various product lines, and verbalize what they think they are doing, seem to me a combination of lemmingism and a willingness to follow any Authority in an expensive suit. I have talked to enough vice-presidents and presidents of computer companies, publishing companies, networks, media outfits and so on, to be totally certain that they have no special knowledge or unusual basis of information; yet these people's remarks, as amplified through the business reporters, send the whole nation a-dithering. There are times I think everybody in Media is either deluded, misguided, lying or crazy.

THREE CRUCIAL POINTS.

1. SYSTEMS "IN THE HOME."

The emphasis has changed from trying to sell snazzy systems to the schools (which don't have the money) to the home. This in turn has convinced most people that the new systems have to be very limited, like jimmied-up TV sets. (We easily lose track of the fact that you can have anything "in the home" if you want to pay for it; and an economy in which Marantzes and snowmobiles have caught on big indicates that some people are going to be willing to pay for really hot stuff.)

2. CATCHING ON.

The key question is not how good a system is in the abstract, but whether it will catch on. (Obviously if we're public-spirited we want the best systems to catch on, of course.)

This matter of Catching On is a fickle and crucial business.

According to one anecdote, Mr. Bell couldn't interest anyone in his invention, which he was showing at some trade fair. Then who should come by but the Emperor of Brazil (!), who was about to leave with his retinue of advisers. "What is that?" asked the Emperor of Brazil. "Nothing to bother with," they said, and tried to rush him by, but he stopped and loved it, and ordered the first pair of telephones sold. This made the headlines, and the sale of telephones began.

Another anecdote. It is legendary that inventors overvalue their own work. Yet after Thomas Edison had invented the kinematograph, or "moving picture," a device you looked into turning a crank, he declined to build a projector for it, saying that the novelty would wear off. Obviously he did't quite see what "catching on" would mean here.

Wonderful Systems That Were Gonna Be

WHERE ARE THE SHOWS OF YESTERYEAR?

I once read a mind-blowing review article in Films in Review, early sixties I think, on schemes to make three-dimensional movies before 1930. There were dozens.

Then there was that multiscreen film Napoleon -- a legend-- done in the nineteen-twenties. (That one really existed.)

Phonevision, about 1947 or so, was going to store a half-hour movie on a 12-inch disk. Did they get the idea from the LP? Did they really think they could do it?

The German photo-gizmo, around 1950: a special camera that supposedly created a sculpture of what it was pointed at. (But how did it know what was behind things?)

A weird lens around 1950-- I think it was depicted as having a blue center and a red periphery, like a fifties hoodlum tail-light-- that was somehow going to find "residual traces" of color in black-and-white pictures, and make 'em into color, zowie, just by copying them.

Then there was the Panacolor Cartridge. During the Days of Madness-- 1968, I think it was -- a rather good little movie gadget was being pushed by a firm called Panacolor. It had ten parallel movie and audio tracks, I believe, on a 70mm strip. The prototypes were built by Zeiss.

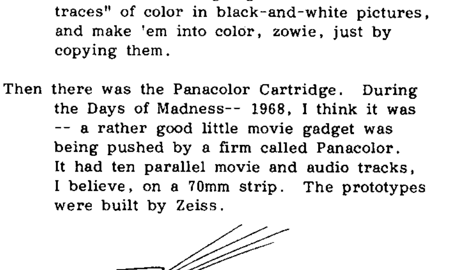

Their idea was that this was a compact movie projector. I kept trying to persuade the company's president that they had inadvertently designed a splendid device for branching movies (see "Hyperfilms," p. DM44).

Exercise for the reader: map out properties of the branching and expository structures implicit in such a device. (It's one-directional. Gotta rewind when you get to the end. But you can jump between tracks when it seems appropriate.)

Anyway, it's gone now.

The Great Robert Crumb. (From Zap Comix #0.)

"In the news* there is no truth**; and in the truth**, there is no news*."

— Modern Russian proverb.

* Izvestia.
** Pravda.

HARDWARE, SOFTWARE AND WHATNOT (reprise)

Among the many odd things that have resulted from the collision of computer people with educators, publishers and others has been the respectful imitation of computer ways by those who didn't quite understand them. Again, the cargo cult.*

The most dismal of these practices has been the adoption of the term "software" for any intellectual or artistic property.** This wholly loses the distinction, made on the other side of the book, between:

hardware (programmable equipment)

software (programs, detailed plans of operation that the hardware carries out)

contents or data (material which is worked on by, moved in or presented by the hardware under control of the software)

In other words, hardware and software together make an environment; data or contents move and appear in that environment.

The publishing-and-picturefolk have missed this distinction entirely. Not realizing that their productions are the contents (material, matter, data, stuff, message...) that come and go in the prefabricated hardware-software entironments, they have mushed this together into a state of self-feeding confusion.

(The matter has not been helped by the computer-assisted instruction people-- see p. DM 15 -- whose branching productions seemed to them enough like computer programs to be called "software.")

* Primitives exposed to "civilized" man imitate his ways ridiculously in religious rituals, hoping for the shipments of canned goods, etc. that his behavior seems to bring down from parts unknown.

ON EX-SPURT-TEASE

** "Mere corroborative detail, to enhance an otherwise uninteresting narrative..."

Pooh-Bah,
Lord High
Everything
Else

3. STANDARDIZATION

In order for something to Catch On, it has to be standardized. Unfortunately, there is motivation for different companies to make their own little changes in order to restrict users to its own products. The best example of how to avoid this: Philips patented its audio cartridge to the teeth, but then granted everybody free use of the patent provided they adhered to the exact standardization. The result has been the system's spectacular success, and Philips, rather than dominating a small market, has a share of a far larger market, and hence makes more money. That's a virtue-rewarded kind of story.

The other problem with standardization, though, is that we tend to standardize too soon. We standardized on AM radio, even though FM would probably have been better. (One Major Armstrong, a great figure in the development of radio, committed suicide when nobody would accept FM. If he could only have heard our FM of today, he might have said "Oh, nuts," and lived.)

Another example. When they designed the Touch-Tone phone pad, the Bell people evidently saw no reason to have it match the adding machine panel, so they put "1" in the upper left rather than the lower left. Now there are lots of people who use both arrangements, every day, and at least one of them curses the designers' lack of consideration.

Another interesting example of Catching On: during the early sixties, it was fun being at places where they were just getting Xerox copiers for the first time. Everyone would argue that nobody needed a copier. Then, grudgingly, one would be ordered. The first month's use invariably would exceed the estimate for the first year, and go up and up from there.

The worst aspect of the confusion among the corporations is that certain deficiencies and crudities of vision slip into the mix. Unless our new media and their exact ramifications and concomitants are planned with the greatest care, everybody stands to lose. We must understand the detailed properties of media. (The first question to ask, when somebody is showing you the Latest and Greatest, is: "What are the properties and qualities of the medium?" The followup questions come easily with experience: How often do you have to change it, what are the branching options, what part could somebody accidentally put in backwards, are there distracting complications? etc.)

I am unpersuaded by McLuhan. His insights are remarkable, yet suspicious: he supposes that electronic media are all the same. How can this be? Here we may now decide what electronic media we want in the future-- and this decision, I would say, is one of the most important we have to face.

The engineers seem to be quite the opposite of McLuhan: somehow to them it's always a multiple-choice, multi-engineering problem, different every time; "this technique is good for A, that technique is good for B." But the net effect is the same: "electronic media are generally the same." I would claim that the're all different, all ten million of them (TV being only one electronic medium out of the lot), and the differences matter very very much, and only a few can catch on. So it matters very much which. Some are great, some are lousy, some are subtly bad, having a locked-in information structure, built deep-down into the system. (Example: the fixed "query modes" built into some systems.)

One last point. Everybody only has a 24-hour day. Most people, if they increase consumption of one medium (like magazines or books) will cut down on another (like TV). This drastically reduces the sorts of growth some people have been expecting. Except, now, if we can begin to replace some of the inane paper-shuffling and paper-losing of the business world, and replace the creepy activities of the school (as now generally constituted) with a more golden use of time and mind. Read on.

THANATOPSYS

A self-employed repairman of mobile homes named Donald Wells has invented a solar-powered tombstone that can show movies and still pictures of the departed, along with appropriate organ music and any last words or eulogies selected by the deceased.

The device is activated by a remote control device carried by a visitor to the gravesite. The movies would be shown on a twelve-inch screen mounted next to the epitaph.

"You could also have pictures of Christ ascending to heaven or Christ on the cross, whatever you want," says Wells. "It adds a whole new dimension to going to the cemetery...."

Cleveland Plain Dealer (Quoted in National Lampoon True Facts, May 74, 10.)

"The Emperor has no clothes on!"

Small Boy
(name withheld)

Last year I actually heard a phone company lecturer say that in the future we will have "Instant Access to Anything, Anytime, Anywhere."

What they're pushing is Picturephone, which it seems to me is unnecessary, wasteful and generally unfeasible.

(See: Robert J. Robinson, "Picturephone-- Who Needs It?", Datamation 15 Nov 71, 152.)

ON USING MEDIA

In any medium-- written, visual, filmic or whatever-- you generate instantaneously an atmosphere, a patina, a miasma of style, involvement, personality (perhaps implicit), outlook, portent. Consider--

The complacency of the Sulzbergers' New York Times--
The cynicism and mischief of Krassner's Realist--
The perkiness and sense of freedom of "Sesame Street"--
The personalized, focussed foreboding of Orson Welles films; as distinct from the impersonalized, focussed foreboding of Hitchcock--

Next to this matter of mood, all else pales: the actual constraints and structures of media, the expositions and complications of particular cognitive works and presentations within media, are as nothing.

"MEDIA" IN THE CLASSROOM

Time after time, the educational establishment has thought some great revolution would come through getting new kinds of equipment into the classroom.

First it was movies. More recently it's been "audio-visual" stuff, teaching machines, film loops and computer-assisted instruction.

In no cases have the enthusiasts for these systems seen how the equipment would fit into conventional education-- or, more likely, screw the teacher up. Teachers are embarrassed and flustered when they have to monkey with equipment in addition to everything else, and fitting the available canned materials into their lesson plans doesn't work out well, either.

The only real possibilities for change lie in systems that will change the instructor's position from a manager to a helper. Many teachers will like this, many will not.

PAY CAREFUL ATTENTION

when somebody shows you an electronic or other presentational system, device or whatever.

A certain kind of slight-of-hand goes on. It's very easy to get fooled. They may show you one thing and persuade you you've seen another.

And if you're canny enough to ask about a feature you haven't seen they'll always say,

"WE'RE WORKING ON IT."

It's only dishonest if they say, "It'll be ready next month."

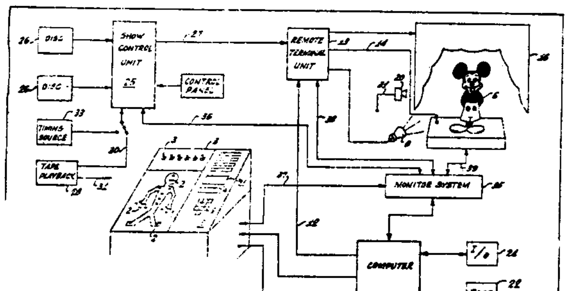

Patent 3,767,901 is for the Disney Audio-Animatronics system, which now basically consists of the manipulation of rubber puppets by minicomputer, through cables and puffs of air.

VIDEO. The happy medium? some mutterings

Would you believe there was television broadcasting over the airwaves in the nineteen-twenties? The thing is, it used bizarre spinning equipment because there were no CRTs (see "Lightning in a Bottle," nearby.) Only with the development of radar in World War II did there also come a practicable Cathode Ray Tube, making home television feasible.

But the big companies were at first very conservative in their marketing, figuring television would be a luxury item only. It took a man named Madman Muntz, who caricatured himself in a Napoleon hat, to see that millions would buy television if the price was right. So he came out with Muntz TV in the late forties. As I recall, the Muntz TV cost $100 and had one tuning knob. (This was less intimidating than the row of knobs on more expensive sets.) I don't know how Muntz came out on it all, but his opening of the mass market made the bigger corporations realize it was there. (This same thing may yet happen again in newer media.)

Originally all there was was Krazy Kat and Farmer Brown cartoons. But behold, sooner than you could say "vertical hold," there were Sid Caesar and Imogene Coca on the Admiral Show, and we were off.

A quarter of a century later, the best of television is no better and the bulk of television is about as bad as it ever was.

We "understand" television. That is, we know what a TV show is, how it fits together and so on.

ICECUBES

But what people don't realize about TV is that the governing feature is the time-slot. In any medium with time-slots, whether TV, radio or classroom education, the time-slot rules behavior. Whatever can happen is as constrained as icecubes in a tray.

This is the limiting factor when optimists try to use TV for teaching. If it's coming over a cable, everything has to be scheduled around it, and the contents are clipped and constrained to fit the time-slot. It may be better with videotape.

CABLES

In the last dozen years, Cable TV, or CATV, has become big business. A Video Cable is a high-capacity electrical carrier that runs through a given neighborhood or region. Business and individuals may "subscribe" and get their own sets hooked onto the cable.

What this does first of all is improve reception. The fouled-up video picture caused by such extraneous objects as the World Trade Center in New York can be corrected by hooking into the video cable: you get a nice, sharp picture.

In addition, though, the cable offers extra channels.

Now, the businessmen who have been throwing together these video cable outfits are aiming for something. They have been thinking that these extra channels would net them a lot of money: by showing things on them that can't be offered on the air-- highbrow drama, or perhaps X-rated stuff-- they could get extra revenue. (You'd pay extra to watch it by buying an unscrambler, or whatever.)

This is turning into somewhat of a disappointment.

The cable people had foreseen, evidently, that people would stay home in droves to see the new offerings on the cable. In Show Business it's easy to forget, though, that everybody has only twentyfour hours in a day, and far less than 24 hours to dispose of freely; so every leisure occupation is competing with every other leisure occupation. Moreover, the residual leisure occupation, when there's nothing else to do, is TV. It would seem that few people would watch more television if it were better, but many would watch less if they could afford to go out.

EXTRA CHANNELS

In recent years, a number of extra channels have been made available by law. These are the UHF, or Ultra High Frequency channels. These, like cables, represent a consumer breakthrough but will have only negligible impact.

THE PROBLEM OF ORGANIZATION

Whatever else you may say about them, the networks and TV stations are at least organized as going concerns within the institutional structures of the country. Ideas of "community television" and other such schemes which call for some new form of social organization to spring forth are about as plausible as "community control" of schools and police-- or at best likely to be as influential as "community social centers."

INTERACTIVE TV?

Some people, I won't say who, have gotten a lot of money for something they call "interactive television." What this turns out to mean is any form of computer time-sharing that will use home TV terminals and video cables. The questions are why use home TV terminals and video cables, insofar as they would seem to promise only comparatively low-grade performance; and whether these people have thought out anything about the potential characteristics of the various media they propose with such abandon. Nothing I have seen or heard about this is reassuring.

"ALTERNATE" TELEVISION, or VIDEO FREAKS

In recent years, many young folks have taken to video as a way of life. In the most extreme cases they say things like "the written word is dead," prompted perhaps by McLuhan. I have found it rather difficult to talk to video freaks. (It may be that some of them are against spoken words as well.) I really just don't know what they're about.

The work of these people is as exuberant as it is strange. I haven't seen much of it or understood much of what I have seen.

In some cases, "alternative television" simply means documentaries outside the normal framework of ownership and reporting. In one example cited by Shamberg (see bibliography), video freaks did excellent coverage of the 1968 Republican convention. People were allowed to speak for themselves, unlike "normal" TV journalism where "commentators" tell you what they see.

Now, this is hardly revolutionary; it is just good documentary-making that shucks dumb traditions artistically, much like the Pennebaker films. However, video enthusiasts claim it is somehow different, and indeed claim that video is different in principle from films. I have been unable to get a satisfactory clarification of this idea.

Video is being used in other ways, harder to understand, by artists (best defined as persons called "artists" within the art world today). Very odd "video pieces" have been shown at art shows, where the object seems to be to confuse the viewer-- or knock him into a condition of Enlarged Perspective, shall we say. And a variety of non-objective videotapes are now being created. (A gallery show in 1969 was called "Video as a Creative Medium" -- implying sarcastically that it had not been before, on the airwaves.)

Some video freaks think of video as intrinsically radical or Revolutionary. In this respect they differ interestingly from, say, the editors of the National Lampoon. The editors of the National Lampoon appear to be political radicals, but do not suggest that the very media of cartoon and joke-piece are themselves revolutionary. Some video freaks appear to be persuaded that the medium of television itself is inherently a vehicle for change.

I can understand one interesting sense in which this may be true: Shamberg talks about video as a method of self-discovery. Seeing yourself on TV does, of course, confer certain insights. But Shamberg suggests it may expand people's consciousness in larger ways-- allowing people to see the bleakness of certain pursuits (he uses the example of Shopping), for instance. But if this does hit home to people, it doesn't seem to me to be the medium that's doing it but the selected content-- as in all previous media. Maybe I've missed the point in some way.

These developments are all very interesting. It can be hoped that those trying to develop new forms of communication will make an effort to communicate better with those who, like the author, often cannot comprehend what they are doing.

" But decentralized transmission of information should be dominant, not fugitive. Each citizen of Media-America should guaranteed as a birthright access to the means of distribution of information."

(Shamberg, p. 67)

" Well, we went down there with our Porta-Pak and tried to take it inside. A guard came over and said we couldn't and even threw one of us out of the booth while the other was inside. A guard telling you what to do in a cybernetic environment?"

(Shamberg, p. 53)

("Cybernetic" is evidently a code word here for what they think is good, true, beautiful and inevitable. Cf. p. DM 13.)

" About the only generalization to be made is that community video will be subversive to any group, bureaucracy, or individual which feels threatened by a coalescing of grassroots consciousness. Because not only does decentralized TV serve as an early warning system, it puts people in touch with one another about common grievances."

(Shamberg, p. 57)

BIBLIOGRAPHY

Michael Shamberg and Raindance Corporation, Guerrilla Television. (Holt, $4.)

TUBE, an underground TV magazine. $8/yr. TUBE, 1826 Spaight St., Madison, WI 53704.

Cable Report, $7/yr. 192 N. Clark St., Room 607, Chicago. Samples $1. "SCANDAL IS RAMPANT in the cable television industry. Only Cable Report follows cable TV developments from the citizen's perspective and tells you what's happening and what's going wrong." Ad in Chicago READER.

Nicholas Johnson, How to Talk Back to Your Television Set. Bantam, 95¢

LIGHTNING IN A BOTTLE:

THE CATHODE-RAY TUBE

A cathode-ray tube is actually a bottle filled with a vacuum and some funny electrical equipment. The equipment in the neck of the bottle shoots a beam of electrons toward the bottom of the bottle.

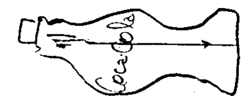

This beam of electrons is called, more or less for historical reasons, a cathode ray. Think of it as a straw that can be wiggled in the bottle.

Actually the bottle is shaped so as to have a large viewing area at the bottom (the screen), and this screen is coated with something that glows when electrons hit it. Such a chemical is called a phosphor.

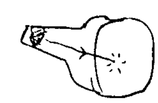

Now, two useful things can be done with this beam.

1) It can be made brighter by increasing the voltage, which increases the number of electrons in the beam.

2) The beam can be moved! That is, it can be made to play around the face of the tube the way you can slosh the stream of a garden hose back and forth on the lawn; or wiggle a straw in a coke bottle. The beam can be moved with either magnetism or static electricity. This is applied in the neck of the bottle-- or even from outside the neck-- by deflection plates, whose electrical pulsations determine the pattern the beam traces on the screen. (Note that the beam can be moved on the screen at great speed.)

The vertical deflection plates can pull the beam up or down on the screen, controlled by a signal to them;

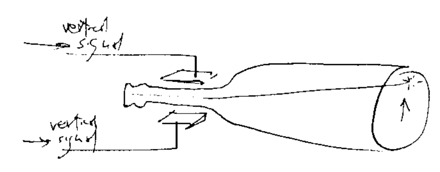

the horizontal deflection plates can pull the beam sideways on the screen, controlled by a signal to them.

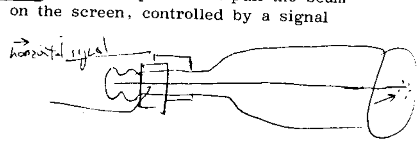

By sending combined signals to both horizontal and vertical deflection plates, we can make the end of the beam-- a bright dot on the screen, sometimes called a flying spot-- jump around in any pattern on the screen. A repeated pattern of the beam on the face of the CRT is called a raster.

From these two capabilities-- brightening and moving the beam-- a number of very special technologies emerge:

TELEVISION uses a zig-zag scanning pattern which repeats over and over. This zigzag pattern is always the same, night and day.

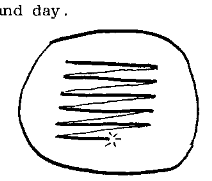

You can usually see the lines clearly on a black-and-white set. The picture consists of the changing pattern of brightness of this beam, which comes in over the airwaves as the television signal.

IF WE HADN'T STANDARDIZED TV WHEN WE DID, WE'D HAVE A BETTER SYSTEM NOW.

Let this be a lesson: standardize on the best system, not necessarily the first.

RADAR DISPLAY uses a CRT to show reflected images around where the radar antenna is standing. This uses a scanning raster of a star shape, brightening the beam when reflected images are received.

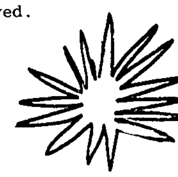

COMPUTER CRT GRAPHICS generally use the CRT in still another way: the beam is moved around the screen in straight lines from point to point. (Between different parts of the picture the beam is darkened, turned very low so you don't see it.)

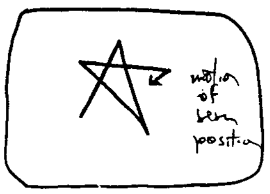

Because the image on a normal CRT fades quickly, the computer must ordinarily draw the picture again and again and again. (Methods for this are discussed on p. DM 22-3.)

SPECIAL KINDS OF CATHODE-RAY TUBES

The CRT is not merely a single invention, but an entire family of inventions. The ordinary CRT, which we have discussed, is viewed at one end by a human being, has an image which fades quickly, and can have its flying spot driven in any kind of raster or pattern.

Here are some other kinds of CRT:

The picture transmitter, which has different versions and names: Vidicon, Image Orthicon, Plumbicon, etc. THIS IS THE MAGICAL DEVICE THAT MAKES THE TELEVISION CAMERA WORK, AND YET, BY GOSH, IT'S JUST ANOTHER CRT. Except instead of the picture coming into it as an electrical signal and out of it as an optical image, the picture comes into it as an optical image and goes out of it as an electrical signal.

How can this be?

The tube sits inside the television camera, which is an ordinary camera, like, with a lens projecting a picture through a dark chamber onto a sensitive surface. But instead of the surface being a film, the surface is the faceplate of a CRT with some kind of a special pickup phosphor:

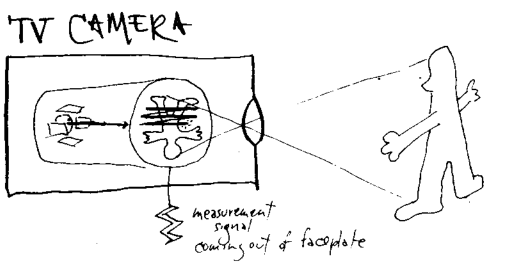

The electron beam, which is just like any other electron beam, is made to zigzag across the faceplate in a standard television raster. And the special phosphor of the tube measures the brightness of the picture at the spot the beam is hitting. I have no idea how this happens, but it's chemical and electronical and mysterious, and is based on the way the phosphor interacts with the light from one side and the electrons from the other side at the same time. Anyhow, a measurement signal comes out of the faceplate, indicating how bright the projected picture is in the very spot the electron beam is now hitting.

As the beam criss-crosses the faceplate in the zig-zag television raster, then, a continuously changing output signal from the faceplate shows the brightnesses all across the successive lines of the scan.

And that is the television signal. Together with synchronizing information, it's what goes out over the airwaves, down your antenna and into your set. Your set, obeying the synchronizing information, brightens and darkens its own beam in proportion to the brightness of the individual teeny regions of the faceplate in the television camera. And this produces the scintillating surface we call television.

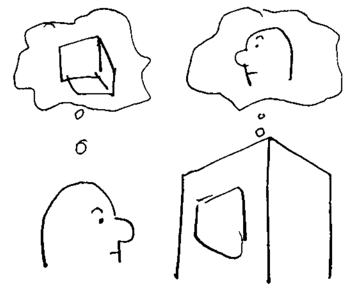

The color tube is a weird beast indeed. There are several types, but we'll only talk about the simplest (and many think the best), Sony's Trinitron(TM) tube.

This is an ordinary CRT which has, instead of a uniform coating on the faceplate, tiny vertical stripes of three primary colors-- red, blue and green. (You thought the primary colors were red, blue and yellow, didn't you. If you're mixing pigments that happens to be true. For some ungodly reason, however, if you're mixing lights, the colors that yield all others turn out to be red, green and blue; it turns out that yellow light can be made out of red and green. If you don't believe me go to a chintzy hardware store, get a red and a green bulb, turn 'em on and see what happens in a white-walled room.)

At any rate, color television uses additional color signals, and in the Trinitron these control the response of the faceplate. If the color signal says "green" as the electron dot crosses a certain part of the screen, the color signal tells the green stripes that they're free to light up when hit. If it's Yellow Time, the signal tells both the red stripes and the green, and so side by side they light up red and green, as the beam crosses them, but the total effect from more than a few inches is Yellow.

Most American color TV sets, however, at least up till this year, used something very different, something entirely weird called the Shadow Mask Tube. I'll spare you the picture, but there were several different electron beams -- often referred to jokingly as the "red electron beam," "blue electron beam" and "green electron beam," though of course they were identical in character. These hit a perforated sieve, up near the screen, called the shadow mask, and the color signal tweaked the unwanted beams so they did not hit different-colored phosphor dots that were intricately arranged on the screen. I'm sorry I started to explain this.

Multigun tubes have more than one electron gun and more than one electron beam. They can be used in different ways (aside from the old shadow-mask TV tube, mentioned above).

For instance, one gun can be driven in a video raster, to show television, while another gun can be used as a computer display, drawing individual lines with no regard to the TV pattern.

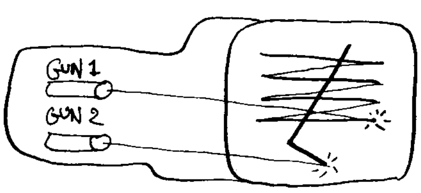

The storage CRT comes in two flavors: viewable and non-viewable. But what it does is very neat: it holds the picture on the screen. The mechanisms for this are of various types, and it's all weird and electronic, but the idea is that once something is put on the screen by the electron beam, it stays and stays. Up to several minutes, usually. The main manufacturers are Tektronix, Princeton Electronic Products, and Hughes Aircraft; each of these three has a product that works by a different method.

Note: Tektronix' tube is built into a number of different computer displays, and is recognizable by its Kelly green surface. They themselves make complete computer terminals around this scope for $4000 and up, but lots of other people put it in their products also. It shows whatever has already been put on the screen, and the electron beam does not have to repeat the action. However, it usually only stays lit for about a minute.

Princeton Electronic Products (guess where) is a much smaller outfit, so perhaps it is appropriate that they make a much smaller storage tube. It is about one inch square at its storage end, and you don't look at it directly. Instead, an image can be stored on it either wth a TV raster or by computer-driven line drawing. After the image is stored on it, though, it functions as a TV camera: the picture stored on the plate can be read out with a scanning raster, exactly as if it were a picture transmitter in a television camera. The Princeton folks have built a quite expensive, but quite splendid, complete terminal around this device: it can hold both video and computer-drawn pictures, superimposed or combined, and sends them back out in standard black-and-white TV. $12000.

CRTS which bring in a picture one way (such as a video raster) and send it back out another way (such as by letting a computer search out individual points) are called scan converters.

A word about this last method. It is often desired by computer people to turn a picture into some form of data (see p. DM 10). Scan converters, usually by the three manufacturers named above, can be hooked up to let the computer program poke around in the picture and measure the brightness of the picture in arbitrary places. A device which examines the brightness of something in arbitrary places is called a flying spot scanner.) Here are some different kinds of flying-spot scanners:

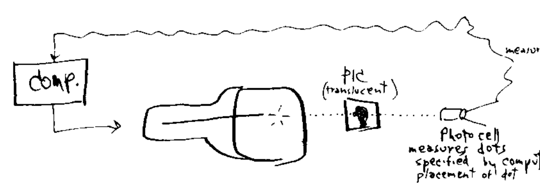

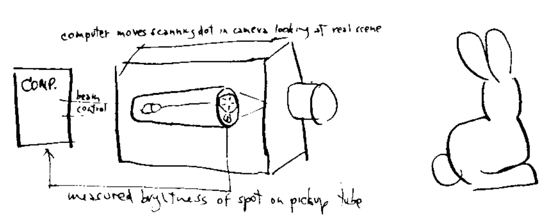

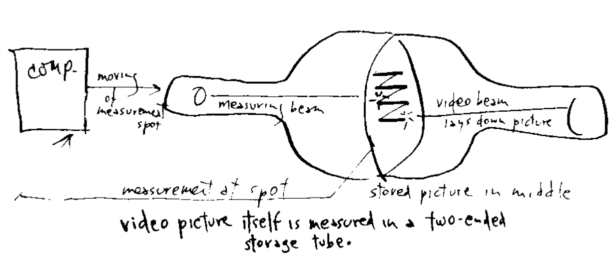

I have heard it said that it might be possible to build a CRT with a changeable mirror surface: that is, the screen becomes mirrored temporarily where it is being hit with the electron beam. Interesting. This would mean that you could make computer displays (and TV) bright and projectable to any degree, say, by pouring a super-intensity laser beam on it. "Be great for writing 'Coca-Cola' on the moon," says a friend of mine. If you believe in astral projection.

BIBLIOGRAPHY: Color TV Training Manual, Sams & Co./Bobbs-Merrill ($7), is a well-illustrated and intelligent introduction to the TV use of CRTs.

SANDIN'S IMAGE PROCESSOR

Dan Sandin, professor of Art at U. of Illinois, Chicago Circle, says very wise things (having been a physicist), and we were going to have a whole section on that, but as you can see there wasn't room.

Daniel J. Sandin (pronounced san-DEEN) has spent the last several years putting together a device he currently calls the IP (Image Processor). It's a system of circuits for changing and colorizing TV. What follows is the first published description of it.

I regret that the following is probably one of the most difficult sections of this book. (If you know nothing about video, read ~~the opposite page~~ DM 6-7 first.)

The idea is basically to create a completely generalized system for altering the color and brightness of video images. (I.e., the system does not move them on the screen. Thus it differs from the Computer Image line of video-twisting graphics systems, which alter positions of objects; see p. DM 39 . Note also that rather similar facilities exist as part of, e.g., the Scanimate system, p. DM 39 .)

This means that basically Sandin's system plays with the part of the TV signal called z, or brightness (as distinct from x or y, the signals for horizontal and vertical movement of the dot. See ~~opposite page~~ DM 6-7).

Now, as a physicist and field-theoretician, Sandin approached this as a problem in generality; and indeed, the style of generalization should be appreciated. Sandin repeatedly chose flexibility and power rather than obviousness in the parts he created. The resulting system is both parsimonious and productive.

His first important decision was that all parts of the system should be compatible and idiot-proof, so that any user could frivolously plug it together any way at all without burning out the circuits.

Indeed, Sandin decided to build it like a music synthesizer: by making all systems electrically compatible (as they are on the Moog and its progeny), any signal can be used to alter or influence any other signal. This is a very profound decision, whose far-flung results have not yet been fully explored even among Sandin's rather fanatical students.

Basically, the incoming video image is "stripped" of its synchronizing information, so that all signals turning up in the guts of the machine may be freely modified. Only at the final output stage are the jots and tittles of the video signal put back on.

Thus the first and last blocks of the Image Processor act like bookends, between which the other modules have their fun. The first block makes the incoming signal into "naked" video, the last block dresses it up respectably again.

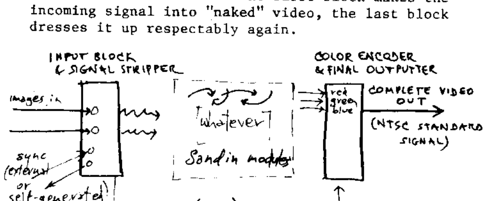

For the sake of clarity we will refer to the outputs as pictures, or as black, white or grey, which they would be if they went straight out to a screen; but they may be turned back into the system and function as inputs as well. "White" means +.5 volts, "black" means -.5 volts.

Let us consider, then, Sandin's modules and what they do individually to the brightness signal z. Combinations are beyond the scope of this article

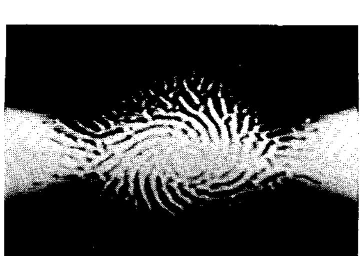

What Dan's processor can do to television is not to be believed.

Savage colors or delicate off-whites, solarizations and pictures on top of pictures. Then through "video feedback" (pointing a TV camera at a TV screen), the system can generate throbbing animated cobwebs and spirals of its own. Shown.

1. ADDER-MULTIPLIER. This combines two input channels, either directly or as specified by a third.

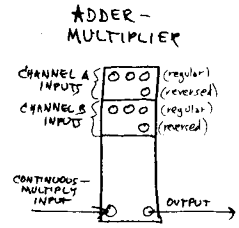

The channel A inputs are added together and multiplied by C; the channel B inputs are added together and multiplied by the reverse of C; both results are added to make the output. (NOTE: this unit is used among other things, for fades and keying.)

2. COMPARATOR. This is like Kodalith film, making an image into stark black and white. Its output is pure black or white. One input signal (the video) is compared with another input signal (reference level, other video, whatever).

While one is greater the output goes all black, and while the other is greater it goes all white.

3. VALUE SCRAMBLER. This is a single module dividing the picture into eight levels. It may be thought of as eight of the above comparators, dividing the brightness spectrum by quantum jumps. The floor and ceiling of the signal to be divided are specified by the two control channels, but the dividing lines between them are then automatically determined. Each corresponding output level may be controlled by a knob.

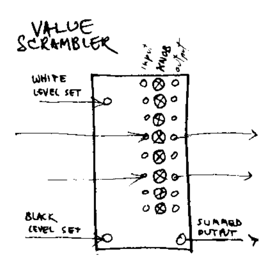

Thus from a range of input values, we get an output step-function each of whose brightnesses is individually adjustable.

Note that these devices may be arranged in parallel, thus dividing the brightness spectrum into as many levels as desired.

4. OSCILLATOR MODULE (very unusual). Sandin's oscillators are voltage controlled, just like the ones in music synthesizers. However, if given any kind of a sync signal, they lock into the nearest multiple (or submultiple) within the specified range. (But then the control signal, if any, tweaks it higher or lower.) Standardized output comes in sine, square and sawtooth.

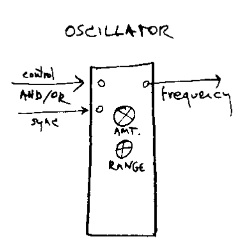

The two planned uses were A) with a sync, to generate fixed patterns, and B) without a sync, to generate movable patterns. If both inputs are used, it becomes a stubborn lock-on voltage-controlled oscillator, which tends to grab at passing submultiples.

5. DIFFERENTIATOR. Basically this sees edges in the picture, or any other part of a scan-line whose color is changing. Its output is proportional to change occurring in the brightness of a scan-line, As the input goes from black to white its output is light; as the input goes from white to black its output is dark. (The input hole selected determines the amount of multiplication.)

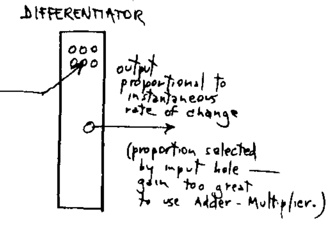

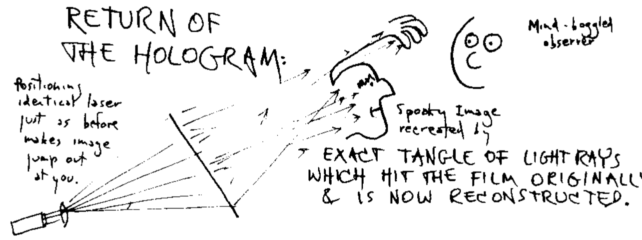

Diagram of how hologram is made, p. DM 20.

HOLOGRAPHY
puts you in the picture
and you, and you...

Holography is one of those Modern Miracles that we really can't get into. It is mind-blowing, influential, and of unclear importance.

Theoretically predicted by Dennis Gabor, the hologram (Greek "whole picture") was finally made to work in the late fifties by Leith and Upatnieks. Since then dozens of other types of holograms have been experimented with, including color holograms, movie holograms, video holograms, audio holograms and gracious know what.

Basically a hologram is an all-around picture. It doesn't look like a picture, but looks like a smudged fingerprint or other mistake of some kind.

Yet it is a marvel.

A basic hologram (-- actually it should be called a laser hologram or Leith-Upatnieks hologram, but we've no time for such distinctions--) is one of these smudgy pictures which, when viewed under a proper laser setup, shows you a three-dimensional picture. Worse than that: as you move your head, the picture changes correspondingly. It looks, not like the flat surface it is, but like a lit-up box with a model in it.

What does the hologram do? Actually it recreates, not a single view, but the entire tangle of light rays that are reflected from the real object. Even down to bright reflections, which scintillate in the usual way, as from chromium.

The only problem: ordinarily they have to be used with laser light, which is spookily one-colored.

Notes from all over: art stylist Salvador Dali presided at an unveiling of "the world's first 360° hologram" at a New York gallery not long ago. The subject was song stylist Alice Cooper.

The Haunted House at Disney World in Florida will ride you through a building full of holograms. That's one way to move through ghosts, all right.

There is a New York School of Holography.

6. FUNCTION GENERATOR. This device is hardest to explain. Let's do it in terms of that first module, the Adder-Multiplier. Know how the Adder-Multiplier puts out either a positive or a negative picture, depending on which input you select?

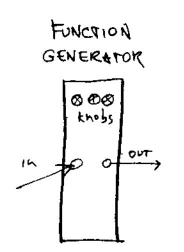

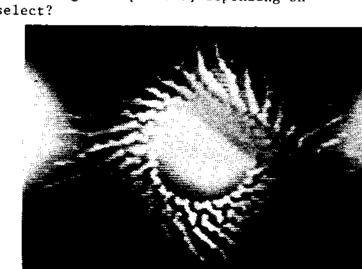

Well, the Function Generator divides the input brightnesses into three ranges, and multiplies each range positive or negative, in proportion to its own knob setting.

Thus the combined setting of the three knobs generates a "function," or curve, from the slopes of the individual settings. See graph. What in photography is called "solarization" represents just one of these combined settings. The others are nameless.

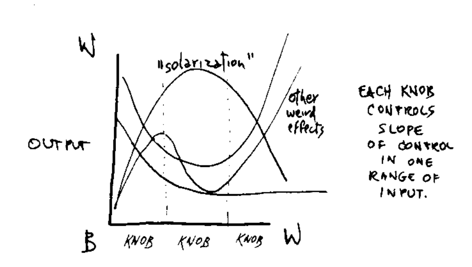

7. COLOR ENCODER MODULE. This is the last block. Into it go three signals, the desired red, blue and green; and out comes standard NTSC video.

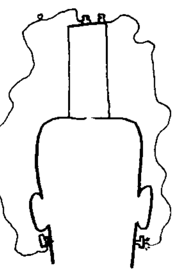

BODY ELECTRONICS

"I sing the body electric..." -- Walt Whitman

There are various people who want to attach electronics to people's bodies and brains.

There are basically two starting points for this ambition. One is authoritarian, the other is altruistic. I am not sure both schools are not equally dangerous, however.

Let's consider first the authoritarians. Prof. Delgado of Yale has demonstrated that any creature's behavior can be controlled by jolts to the brain. Delgado has dealt especially with the negative circuits of the brain, that is, places where an electrical impulse causes pain (or "negative reinforcement"). In Delgado's most stunning demonstration, he stopped a charging bull with just a teeny radio signal. Enthusiastically Delgado tells us how fine this sort of thing would be for controlling Undesirable Human Behavior, too.

Now, let's consider just what we're talking about. In these experiments, needles are implanted in the creature's brain. This can involve removing a section of the skull, or it can be done merely by hammering a long hollow needle straight into the skull and thus the brain.

The researcher, or whatever we want to call him, had better know what he is doing. But due to the remarkable mass action of the brain, the destruction caused by such needles will have not observable effects if done properly.

The hollow needle, once in place, becomes a tube for shielded electrical wires, whose bare metallic tips may then be used to carry little electrical jolts, to whatever brain tissue is reached by the tip of the needle, whenever tiny signals are applied.

Now there are regions of the brain, distributed irregularly through its mysterious contents, which are loosely called the "pleasure" and "pain" systems. They are called that because of what the organism does when you jolt it in those places. (We do not know whether jolts to these areas really cause pleasure or pain, because these things haven't been done to human beings. Yet. The creatures it has been done to can't tell us just how it feels; thus "pleasure" and "pain" are in quotation marks. For now.)

Anyway, what happens is this. If you stimulate a creature in the "pain" system it tends to stop what it is doing-- this is called negative reinforcement-- and if you stimulate it in the pleasure system, it tends to do more of what it was doing. Positive reinforcement.

Now, to some people this suggests wonderful possibilities.

Delgado, for instance, believes that this technology gives us everything we need for the control of Anti-Social Tendencies. Criminals, psychopaths and Bad Guys in general-- all can be effectively "cured" (i.e., put on their best behavior) by these techniques. All we have to do, heh heh, is get into their heads, heh heh, habits of proper behavior. And with these new techniques of reinforcement, we can really teach 'em.

Unfortunately Delgado is probably right.

In principle this is just a drastic form of behavior control on the B.F. Skinner model (depicted also in Nineteen Eighty-Four and A Clockwork Orange). The new system is more stark and startling because of its violation of the individual's body interior, but not in principle different.

Skinner has the same naive, simpleminded solutions for everything. All "we" have to do-- using "we" to mean society, the good guys, good guys acting on behalf of society, etc.-- is control the behavior of the bad guys, and everything will be better, and "we" can accomplish anything "we" desire.

The reader may see several problems with this.

In the first place (and the last), there is the obvious question of who we are, and if we are going to control other people, who is going to control us.

At a time when our "highest" leaders show themselves preoccupied with low retaliations and lower initiatives, we can wonder indeed if it is not more important to prevent anyone from ever getting this kind of control over humans than to facilitate it.

Even if that weren't a problem, there is the more simpleminded question of who in the existing system would use such techniques. It turns out, of course, that they would be added to what is laughably called the Correctional System, or even more laughably called the Justice System. All the sadists you could possibly want work there. (And no doubt some very nice guys-- but experiments have demonstrated horrifically that decent people, turned into "guards" even for a short time, adopt the patterns of brutality we have known from time immemorial.)

So, like truncheons and electric shock therapy and solitary confinement and everything else, these techniques-- if they are used-- will enter the realm of Available Punishments, not to be used with clinical precision but with gratuitously brutalizing intent, new tools for punitivity and sadism. The "correctional" system would have to be magically corrected itself before such tools could be employed without simply making things worse. And the prospect is not good.

Such schemes grow, of course, from a caricature of the malefactor-- thinking him to be some sort of miswired circuit, rather than a human being caught up in anger, pain, humiliation and unemployment.

(There are also a lot of canards about Free Will, but these do nothing for either side in this controversy.)

NEW FACULTIES

Starting from an entirely different outlook, various designers and bio-engineers are trying to add things to the human body and nervous system, for the voluntary benefit of the recipient.

A number of research and development efforts are aimed at helping those with sensory impairments, and electronics obviously is going to involved.

An example: a firm called Listening, Inc. in Boston, founded by Wayne Batteau (whom John W. Campbell considered one of the Great Men of Our Time), devised a system for helping the totally deaf to hear. Supposedly this could transmit the actual sensation of hearing into the nervous system by some scarcely-understood form of electrical induction. The machine was sold off; whether it ever got a safety rating I don't know.

This is the sort of thing people would like to do for the blind, as well.

Now, in principle, it might be possible to transmit an image in some way to the actual visual area of the cerebral cortex. (This might or might not involve opening the skull.) Somebody's working on it.

In a related trend, numerous design groups are attempting to extend the capabilities of the human body, by means of things variously called possums, waldoes and telefactors.

"Possums" (from Latin "I can") are devices to aid the handicapped in moving, grasping and controlling. Whatever motions the person can make are electronically transposed to whatever realm of control is needed, such as typewriting or guiding a wheelchair. ("Waldo" is Heinlein's term for a possum that can be operated at a distance.)

In the space program, though, they call them telefactors. A telefactor is a device which converts or adapts body movements by magnification or remote mimicking. Unlike possums, they are meant to be operated by people with normal faculties, but to provide, for example, superhuman strength: cradled in a larger telefactor body, a man can pick up immense loads, as the movements of his arms are converted to the movements of the greater robot arms.

Telefactors can also work from far, far away. Thus a man sitting in a booth can control, with the movements of his own arms, the artificial arms of a robot vehicle on another planet.

(This whole realm of sensory and motor mechanics and transposition is an important aspect of what I call "Fantics," discussed on pp. DM 78-51).

Then there are those who, like How Wachspress (see nearby), want to expand man's senses beyond the ordinary, into new sensory realms, by hooking him to various electronics.

THOUGHTS

There are two problems in all of this. The first and worst, of course, is who controls and what will hold them back from the most evil doings. Recent history, both at home and abroad, suggests the answers are discouraging.

The second problem, wispish and theoretical next to that other, is whether in turning toward bizarre new pleasures and involvements, we will not lose track of all that is human. (Of course this is a question that is asked by somebody whenever anything at all changes. But that doesn't mean it is always inappropriate.)

In the face both of potential evil and dehumanization, though, we can wish there were some boundary, some good and conspicuous stopping place at which to say: no further, like the three-mile limit in international law of old. I personally think it should be the human skin. Perhaps that's old-fashioned, being long breached by the Pacemaker. But what other lines can we draw?

The prospects are horrorshow, me droogies.

BIBLIOGRAPHY

T.D. Sterling, E.A. Bering, Jr., S.V. Pollack and H. Vaughan, Jr., Visual Prosthesis: The Interdisciplinary Dialog. ACM Monograph. $21.

PSYCHO-ACOUSTIC DILDONICS

I originally hadn't intended to include anything like this in the book, wanting it to be a family-style access catalog and all that, but this particular item seems fairly important.

Remember how we laughed at the Orgasmotron in Woody Allen's Sleeper? Well, it turns out not to be a joke.

An individual named How (not Howard) Wachspress, electronicker-in-residence at a San Francisco radio station, has been developing just that, except that he has more elevated purposes in mind. The secret was broken to the world in Oui magazine earlier this year; but Hefner, the publisher, evidently held back the more startling photographs of a model in electronically-induced ecstasy.

Wachspress' devices transpose sound (as audio signals) into feelings; you touch your body with an open-ended tube or other soft fixture attached to his device-- which in turn is attached to a hi-fi.

The sensations, it is claimed, are profound and moving. You may take them anywhere on your body; the effect is deeply relaxing and emotionally engrossing. Wachspress thinks he has reached an entire neurological system that wasn't known before, much like Olds' discovery of the "pleasure center" in the brain; he sees it as a new modality of experience and a generalization of music and touch. That is the main point. "Hyper-reality" is where he says it gets you: a point curiously congruent with the author's own notions of hypertext and hypermedia as extensions of the mental life.

This said, we can consider the prurient aspects of Wachspress' Auditac and Teletac devices (which he intends to market in a couple of years as hi-fi accessories, b'gosh). When played with the right audio, in the right places, and a good operator at the controls, they provide a sexual experience said to be of a high order.

Wachspress' work ties in interestingly with today's "awareness" movement, of which Esalen is the spiritual center, which holds that we have gotten out of touch with our bodies, our feelings, our native perceptions. As such, the Wachspress machines may be an unfolding-mechanism for the unfeeling tightness of Modern Man-- as well as a less profound treatment for "marital difficulties" and Why-Can't-Johnny-Come-Lately.

Inscrutable San Francisco! Wachspress gave a number of demonstrations of his devices in Bay Area churches, until he became disturbed at immodest uses of the probe by female communicants who had stood in line to try the machine.

(Auditac, Ltd., Dept. CLB,
1940 Washington St.,
San Francisco CA 94109.)

Harry Mendell, a good friend of mine, rigged an interesting experiment while he was still in high school.

He used a little Hewlett-Packard minicomputer, which the manufacturer had generously loaned to his Knights of Columbus Computer Club of Haddonfield, N.J.

Harry hooked the Hewlett-Packard up to a CRT display (see pp. DM 6-7, DM 22-3). At the top of the CRT, following his program, the computer continuously displayed the letters of the alphabet. A little marker (called a cursor) would skip along underneath the letters, acting as a marker for each of them in turn.

Harry rigged one more external device: a set of electrodes. These would be strapped, harmlessly, to the head of a subject. Harry's computer program used these electrodes to measure alpha rhythm, one of the mysterious pulses in the brain that come and go.

Every time the subject flashed alpha, Harry's program would copy the letter above the cursor to the bottom of the screen.

Sitting in this rig, subjects were able to learn, rather quickly, TO TYPE WORDS AND SENTENCES. Just by flashing alpha rhythm when the cursor was under the right letters.

Jubilant, Harry showed this setup to an eminent neurophysiologist from a great university nearby, a man specializing in electrode hookups. Harry was a highschool student and did not understand about Professionalism.

"What's so great about that?" sniffed the eminent professional. "I can type faster."

So Harry dropped that and went on to other stuff.

PICTURE PROCESSING

"Picture processing" is an important technology, largely separate from the rest of computer graphics. It means taking an incoming picture, usually a photograph, and doing something to it. (Some now call this area "computer pictorics.")

First of all, there is image enhancement. This means taking pictures, dividing them into points whose brightness is separately measured, and then using special techniques for making the picture better. To people familiar with photography, this may seem impossible; to photographers it is a maxim that photographs always lose quality at each step. Nevertheless, various mathematical techniques such as Fourier Analysis (mentioned elsewhere) do just that, producing a new data structure improving on the original data. Surfaces appear smoother, edges sharper.

(These techniques have been extensively used to clean up photographs sent back from our unmanned space vehicles-- both those used exploring other planets and those spying on our own-- see Secret Sentries in Space, Bibliography.)

Then there are recognizers-- programs that look at the data structure from an input picture, and try to discern the lines, corners and other features of the picture. (While your eye instantly sees these things, computers do not, and must look at the dots of a picture one-by-one. How to analyze pictures in such tedious sequences is no simple matter.)

For recognizing more complex objects in pictures-- boxes, spheres, faces or whatever-- more complex structure-analyzing programs are necessary. As the possibilities of what might be in a picture increase, these increasingly become guessing programs. (This becomes a branch of artificial intelligence, a misleading term for a curious field, discussed on pp. 12-14.)

Numerous computer people think it is important to match up our computer graphic display systems (described variously on this side of the book) to image input systems. This is a matter of taste.

These are all basically techniques for making a data structure. Any data stored in computers must have, of course, a data structure-- which basically means any arrangement of information you choose. (see p. 26-9.)

These various techniques are intended to create reduced data structures, recording only the "most important" data of the picture-- from which new and varying pictures may be created, reflecting the "true" structures originally shown in the initial picture. How much it's going to be possible to create these data structures from input pictures remains to be seen; some of us think it's not going to be generally worthwhile.

BIBLIOGRAPHY

Azriel Rosenfeld, "Progress in Picture Processing 1969-71." ACM Computing Surveys June 73, 81-108.

Ken Knowlton and Leon Harmon, "Computer-Produced Grey Scales." Computer Graphics and Image Processing, April 72, 1-20.

Philip J. Klass, Secret Sentries in Space. Random, 1971, $8. Interesting general book on geopolitical strategy and orbital photoreconnaissance. "Now-it-can-be-told" approach.

SATELLITE PICTURES OF YOUR OWN HOME COUNTY, OR WHATEVER

You can get pictures of any area you want from ERTS (Earth Resources Observation Systems) satellites, from EROS Data Center (no, not a dating service, see p. 64), Sioux Falls SD 57198, or call 605/594-6511 bet. 7 AM & PM central time.

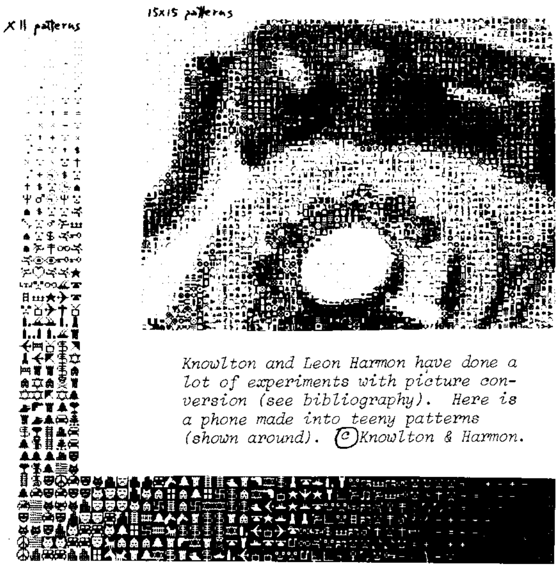

Knowlton and Leon Harmon have done a lot of experiments with picture conversion (see bibliography). Here is a phone made into teeny patterns (shown around). ©Knowlton & Harmon.

LIZZIE OF THE LINEPRINTER

A famous converted picture. The painting was divided into 100,000 brightness-measured spots by H. Philip Peterson of Control Data Corporation; then each dot was made into a square of overprinted letters on the printing device. The program allowed 100 levels of grey. Above: Control Data's version, reprinted by permission. Below: a cut-down version that often turns up. (From original flat 2D artwork by Len DaVinci of Medici Associates.)

NOTE: this is not a "computer picture." There is no such thing. It's a quantization put out on a lineprinter.

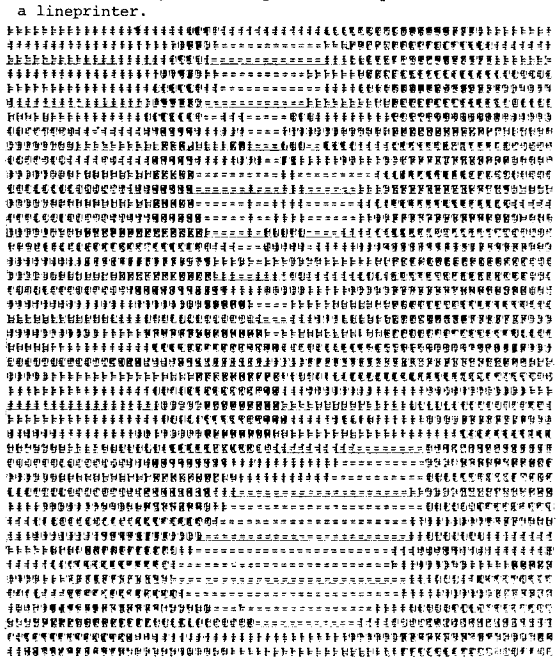

KEN KNOWLTON

Kenneth Knowlton is a Bell Labs lifer. Tall, patrician and gracious, his work, like Sutherland's, shows the inner light of unifying intelligence. He works in Max Mathews' section of Bell Labs at Murray Hill, where they do all that interesting stuff with music and perceptual psychology and so on. During the last decade, Knowlton has turned out vast quantities of articles, processed pictures, movies, and actual computer languages; while any ordinary man would be satisfied to be so productive, apparently he does a lot of other things in his work that he doesn't talk about.

Some of Knowlton's best-known work has been in picture processing, where he has converted photographs into mosaics of tiny patterns-- which nevertheless show the original.

His first widely-known language was BEFLIX (BEll Labs movie-making system); this was programmed for the 7094 in the early sixties. BEFLIX allowed the user to create motion pictures by a clever mosaic process that used the output camera more efficiently. (Actually, the lens was thrown out of focus manually and the entire frame created as a mosaic of alphabetical characters; this did the whole thing much more quickly and inexpensively.)

(Some of the clever data-handling techniques of BEFLIX Knowlton then turned around and used in L^6, a language which made these techniques available to other computer people. This may sound like only a computer technicality, but it's the sort of thing that's widely appreciated. (L^6 stands for "belL Labs' Lower-Level List Language."))

Wanting to get outside artists interested in BEFLIX and related media, he worked for a time with film-maker Stan Vanderbeek; from this Knowlton saw that artists' needs were more intricate than he had anticipated. Augmenting BEFLIX with some of the things Vanderbeek asked for, Knowlton came up with a new language called TARPS (Two-Dimensional Alpha-Numeric Raster Picture System). This in turn led to EXPLOR (EXPlicit(ly provided 2D Patterns,) Local (neighborhood) Operations, and Randomness). EXPLOR is fascinating because of its originality and generality-- not only does it modify pictures and serve as an artist's tool, but it has fascinating properties as a computer language and may even have applications in complex simulations for technical purposes.

Since Vanderbeek, Knowlton has entered into a long and fruitful collaboration with Lillian Schwartz, a talented artist. Their many films have been clever, startling and powerful. I must say that they grow on you: I liked them at first, but when I saw five or six in a row this January, I found them just incredible. Because they are abstract, and full of fast-changing patterns and reversals, they take some adjusting to; but they're worth seeing over and over.

EXPLOR may be thought of as a highly generalized version of Conway's game of Life (see p. 48). You start with two-dimensional patterns as your data structure; these can be abstractions or even converted photographs, as in a recent Knowlton-Schwartz film showing Muybridge's Running Man. In your EXPLOR program, you may then cause the pattern to change by degrees, each cell of the pattern reacting to the cells around it or to random events as specified by the programmer.

EXPLOR, running without external data, comes up with some extraordinary snakeskin and Jack Frost patterns. But its uses in traffic simulation and various other studies of populations in space could be very interesting.

EXPLOR has obvious artistic applications. Lillian Schwartz is using it extensively in filmmaking. It's now running on a minicomputer feeding to a modified Sony Trinitron color TV. (This color setup was created by Mike Noll and is described in a recent issue of the CACM, though only for black-and-white TV; the color is more recent. It stores the color picture as a list of sequential colors represented in the computer's core memory, each dot being represented. Cf. "Boyell's Terrarium," p. DM38.)

Knowlton has used EXPLOR for teaching computer art at the University of California; the language is available programmed in "medium size" Fortran from Harry Huskey, Dept. of Information and Computer Science, U. of Cal. at Santa Cruz, Santa Cruz, California.

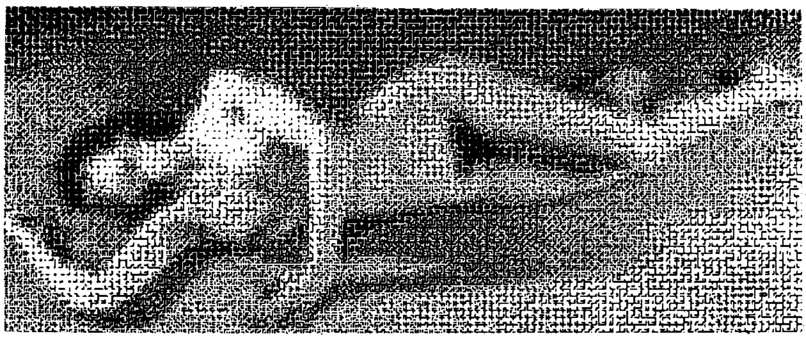

This is a non-simple picture conversion. The original photograph was converted into measured points; but these were in turn made into grow-together patterns by a program in the EXPLOR language. ⓒ *Knowlton & Harmon.*

AUDIO

Wish there were room to talk about plain regular audio here-- matters like "binaural" recording, and Why don't they make hi-fi systems based on a Grand Bus (see p. 42)? But there's no room here.

AUDIO AND COMPUTERS

People are occasionally still startled to hear that computers can make sound and music. They can indeed.

First of all, note that an incoming sound is a fluctuating voltage and can thus be turned into a data structure, i.e., a string of measurements.

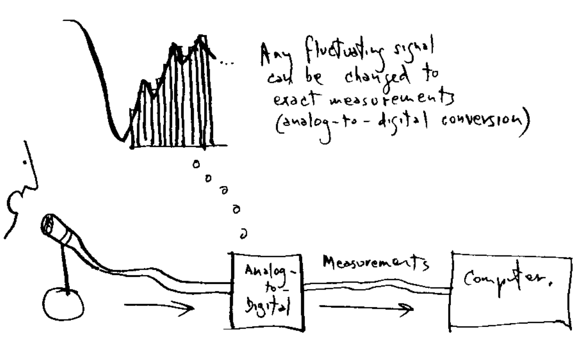

To make sound by computer is the obverse. If the computer can be set up to send out a string of measurements, these can be turned back into a fluctuating voltage, and thus make sounds.

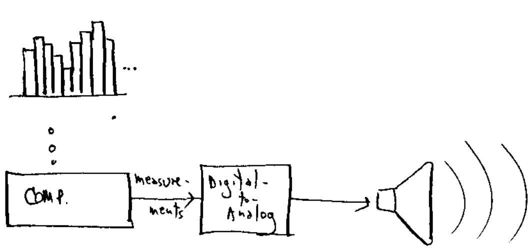

In the easiest case, the computer can just send back out the voltages it originally got in. This is rather ridiculous-- using the computer just as a recording device-- but it's a clear and simple example.

The question after that is what next: how to have the computer make interesting streams of output measurements, i.e., sounds and tones.

There are numerous methods we can't go into. Max Mathews, at Bell Labs, has for years been doing music by computer; his current system is called GROOVE. Heinz von Foerster, at the University of Illinois (Urbana), has been doing the same. Another lab at MIT has just gotten a PDP-11/45 (see p. 42) for the same purpose.

(The problem is: can the computer keep up with the output rate needed to make music in real time? maybe the 11/45 can.)

Another approach is to relieve the computer itself from making the tones, and use other devices-- music synthesizers-- for this, controlled by the computer. This is essentially the approach taken with General Turtle's Music Box (see p. 57), and at the Columbia-Princeton Electronic Music Center, where their RCA Mark II music synthesizer-- an immense one-of-a-kind jobbie-- is under more general computer control.

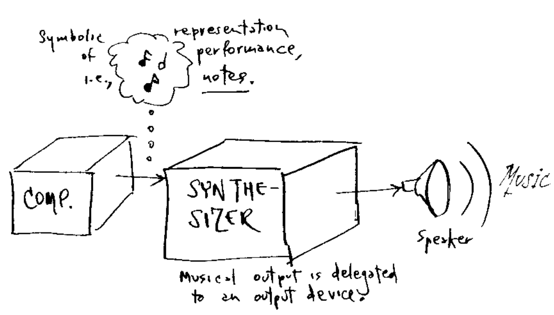

MUSICAL NOTATION

Note that the computer handling of musical notes, as symbols, is another task entirely, closely resembling computer text handling (mentioned variously in the book). A high-power structured-text system or Thinkertoy (see p. DM 52-55) is fine for storing and presenting written music.

And, of course, such stored musical notation (a data structure) can obviously be played by the hookups mentioned.

SPEECH BY COMPUTER

You may have heard about various kinds of "talking computer." This deserves some explanation.

Computers may be made to "talk" by various means. One is through an output device that simply stores recordings of separate words or syllables, which the computer selects with appropriate timing. (Machines of this type have been sold by both IBM and Cognitronics for a long time.)

A deeper approach is to have the computer synthesize speech from phonemes, or actually make the tones and noises of which speech is composed. These are very tricky matters. Bell Labs, and others, have been working on many of these approaches.

The real problem, of course, is how to decide what to say. (This was discussed under Artificial Intelligence, p. DM 12-17.)

AUDIO ANALYSIS AND ENHANCEMENT

The problem of analyzing audio is very like the problem of analyzing pictures (see p. DM 10), and indeed some of the same techniques are used. The audio goes into the computer as a stream of measurements, and the selfsame technique of Fourier Analysis is employed. This reduces the audio to a series of frequency measurements over time-- but, paradoxically, loses little of the fidelity.

Once audio is reduced to Fourier patterns, it can be reconstituted in various ways: changed in timing and pitch independently, or enhanced by polishing techniques like those used in image enhancement (see p. DM 10).

This has been done with great success by Tom Stockham at the University of Utah, who has reprocessed old Caruso records into improved fidelity. In the picture we see him with equipment of some sort and an old record.

University of Utah

(Stockham has been in the news lately, as one of the panel puzzling over the notorious 18-Minute Gap.)

(The author has proposed the name Kitchensync™ for a system to synchronize motion pictures with "wild" sound recording by these means.)

BIBLIOGRAPHY

Thomas G. Stockham, Jr., "Restoration of Old Acoustic Recordings by Means of Digital Signal Processing." Audio Engineering Society preprint no. 831 (D-4), presented at Audio Engineering Society 1971 convention.

Prentiss H. Knowlton, "Capture and Display of Keyboard Music," Datamation May '72, 56-60. Describes a setup he built at U. of Utah that allows pianists to play music on an ordinary keyboard, and converts the input to symbolic representation in the computer. It uses an organ, a PDP-8 and a couple of CRT displays.

Heinz von Foerster and James Beauchamp, Music by Computers. Wiley, 1969. HAS RECORDS IN BACK.

Some of the early Bell Labs work may be heard on an excellent Decca LP with the misleading title "MUSIC from MATHEMATICS." (Decca DL 79103). (The mathematical myth is discussed on p. 8-9.)

time out for THREE COMPUTER DREAMS: "AI" (Artificial Intelligence); "IR" (Information Retrieval); "CAI" (Computer-Assisted Instruction)

THE THREE DREAMS. It's time for awe to be replaced with the critical eye.

These are three topics of great importance; of importance, unfortunately, less for what they have actually accomplished than for the degree to which they have confused and intimidated people who want to understand what's going on. Merely to mention them can be one-upmanship. All three titles mean so much, so many different specific things, as to mean almost nothing when lumped together as a whole. All three have developed a web of intricate technical facts (and sometimes theorems), but the applicability of these elegant findings is in all three cases a matter open to considerable scrutiny.

Since each of these fields has developed a considerable body of technical doctrine, the reader might well ask: why aren't they on the other side of the book, the computer side? The answer is that they are computerman's dreams, dreams of considerable intricacy and persuasiveness, and we are not considering the technicalities here anyway. As on the other side, the problem is to help you distinguish apples from oranges and which way is up. For more go elsewhere, but I hope this orientation will make sorting things out quicker for you.

These three terms-- "artificial intelligence," "information retrieval," "computer-assisted instruction"-- have a number of things in common. First, the names are so portentous and formidable. Second, if you read or hear anything in these fields, chances are it will have an air of unfathomable technicality. Both strange technicalism and deep mathematics may combine to give you a sense that you can't understand any of it. This is wrong. The fact that there are obscure and Deep Teachings in each has no bearing on the general comprehensibility of what they are about. More importantly, the question of how applicable all the things these people have been doing is going to be is a question of considerable importance, especially when some of these people want to take something over. Don't get snowed.

Each of these fascinating terms is actually a roof over a veritable zoo of different researchers, often of the most eccentric and interesting sort, each generally with his own dream of how his own research will be the breakthrough for humanity, or for something. It would take a Lemuel Gulliver to to show you the colorfulness and fascination of these fields; again, we just scratch the surface here.

Another interesting thing these three fields have in common: the frequent use of a classical computerman's putdown on anybody who dares question whether their super-ultimate goals can ever be achieved.

The line is, "WE DON'T KNOW HOW TO DO THAT YET."

If somebody pulls it on you, the reply is simply, "How do you know you ever will?"

ONE OF THE FEW GOOD LAYMEN'S COMPUTER JOKES

illustrating also certain problems of Artificial Intelligence.

A very large artificial-intelligence system (goes the story) had been built for the military to help in long-range policy planning; financed by ARPA, with people from M.I.T., Stanford and so on.

"The system is now ready to answer questions," said the spokesman for the project.

A four-star general bit off the end of a cigar, looked whimsically at his comrades and said--

"Ask the machine this: Will it be Peace or War?"

The clerk-typist (Sp4) translated this into the query language and typed it in.

The machine replied:

YES

"Yes what?" bellowed the general.

The operator typed in the query.

Came the answer:

Yes SIR

THE GOD-BUILDERS! ARTIFICIAL INTELLIGENCE ... sort of

"Artificial Intelligence" is at once the sexiest and most ominous term in the world. It chills and impresses at the same time. In principle it means the simulation of processes of mind, by any means at all; but it generally turns out to be some form or another of computer simulation (see "Simulation," p. 58). Actually, "artificial intelligence" has generally become an all-inclusive term for systems that amaze, astound, mystify, and do not operate according to principles which can be easily explained. In a way, "artificial intelligence" is an ever-receding frontier: as techniques become well-worked out and understood, their appearance of intelligence, to the sophisticated, continually recedes. It's like the ocean: however much you take out of it, it still stretches on-- as limitless as before.

Unfortunately laymen are so impressed by computers in general that they easily suppose computers can do anything involving information. And public understanding is not fostered by certain types of stupid demonstration. One year I heard from numerous people about how "they'd seen on TV about how computers write TV scripts"-- what had actually been shown was a hokey enactment of how the computer could randomly decide whether the Bad Man gets shot or the Good Guy gets shot-- both outcomes dutifully enacted by guys in cowboy outfits. Duh.

It should be perfectly obvious to anybody who's brushed even slightly with computers, however-- for The Brush, see the other side-- that they just don't work like minds. But the analogy hangs around. (Edmund C. Berkeley wrote a book in the forties, I believe, with the misleading title of Giant Brains, or Machines That Think. The idea is still around.)

Here's a very simple example, though. Consider a maze drawn on a piece of paper. Just by looking, we cannot simultaneously comprehend all its pathways; we have to poke around on it to figure out the solution. Computers are sort of like that, but more so. While our eyes can take in a simple picture, like a square, at once, the computer program must poke around in its data representation at length to see what we saw at once.

The principle holds true in general. The human mind can do in a flash, all at once (or "in parallel") many things that must be tediously checked and tried by the highly sequential computer program. And the more we know about computers, the more impressive the human brain becomes. (The seeming cleverness of some simple programs does not prove the simplicity of the phenomena being imitated.)

Nevertheless, it is interesting to try things with computers that are more like what the mind does; and that is mostly what artificial intelligence is about.

In various cases this has resulted in helpful tricks that turn out to be useful elsewhere in the computer field. In this sense, artificial intelligence is sort of like menthol: a little may improve things here and there. But (in my opinion), that does not mean a whole lot of it would make things better still.

Nevertheless, some artificial-intelligence enthusiasts think there is no limit on what machines can do. They point out that, after all, the brain is a machine. But so is the universe, presumably; and we're never going to build one of those, either.

PATTERN RECOGNITION

This is one of the most active areas in artificial intelligence, perhaps because of Defense Department money. (It might be nice, goes the reasoning, to have guns that could recognize tanks, machines that could look over aerial reconnaissance pictures, radars that could recognize missiles...)

What it boils down to is the study of clues and guessing among alternatives. In some cases, well-defined clues can be found for recognizing specific things, like parts of pictures (even straight lines cannot be recognized by computer without a complex program) or like handwriting (see below). In the worse cases, though, careful study only raises the most horrendous technical problems, and the pursuit of these technical problems is its own field of study (articles have titles like "Sensitivity Parameters in the Adjustment of Discriminators," meaning It Sure Is Hard to Draw The Line).

But in some felicitous cases, researchers actually boil a recognition problem down to a manageable system of clues. For instance, take the problem of written input to computers. (Some people don't like to type and would rather write by hand on special input tablets.) But how can a program recognize the letters? Aha: the answer, kids, is in your text.

The Ledeen Character Recognizer (described in detail in Newman and Sproull, Principles of Interactive Computer Graphics, Appendix 8) is a method by which a program can look at a hand-drawn character and try to recognize it. The program extracts a series of "properties" for the character and stores them in an array. Every character in a given person's block lettering will tend to have certain property scores. But the Ledeen recognizer must still be trained, that is, the average property scores of the letters that each individual draws must be put into the system before that individual's lettering can be recognized. Even then it's a question of probability, rather than certainty, that a given character will be recognized.

COMPUTERS DON'T ACTUALLY THINK. YOU JUST THINK THEY THINK. (We think.)

HEURISTICS (pronounced hewRIStics)

If we want to make a computer do what we know perfectly well how to do ourselves, then all we do is write a program.

Aha. But what if we want a computer to do something we do not know how to do ourselves?

We must set up its program to browse, and search, and seize on what turns out to work.

This is called heuristics.

What it amounts to basically is techniques for trying things out, checking the results, and continuing to do more and more of what seems to work.

Or we could phrase it this way: looking for successful strategies in whatever area we're dealing with. As a heuristic program tries things out, it keeps various scores of how well it's doing-- a sort of self-congratulation-- and makes adjustments in favor of what works best.

Thus the Greenblatt Chess Program, mentioned under "Chess," nearby, can "invent" chess strategies and "try them out"-- what it actually does is test specific patterns of moves for the overall goodness of their results (in terms of the usual positional advantages in chess), and discard the strategies that don't get anywhere. It does this by comparing its "strategies" (possible move patterns) against the records of chess matches which are fed into it.

(If you've read the other side of the book, heuristics may be thought of as a form of operations research (p. 58) carried on by the computer itself.)

In some ways heuristics is the most magical area of artificial intelligence: its results are the most impressive to laymen. But, like so many of the computer magics, it boils down to technicalities which lose the romance to a certain extent.

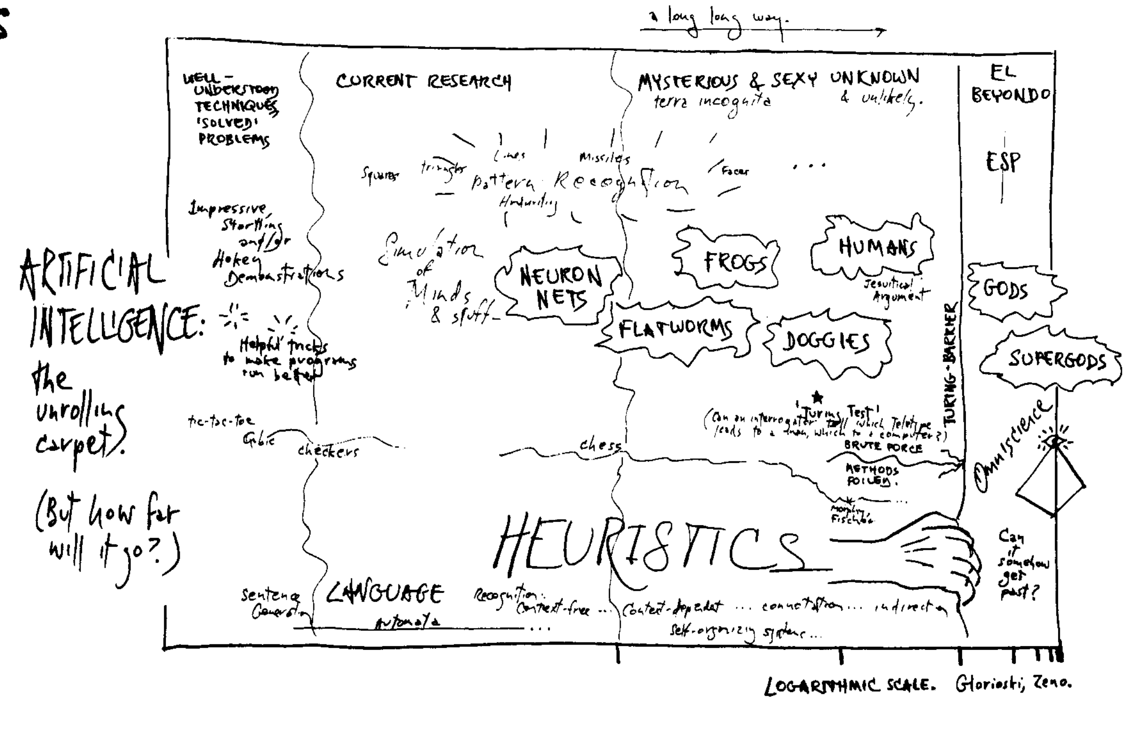

NEURAL SIMULATION

An important branch of Artificial Intelligence is concerned with what bunches of imaginary neurons could do, even neurons that we made up to follow particular rules. This area of study is somewhere between neurology and mathematics; much of it is concerned with the mathematics of imaginary setups, rather than the properties of actual nerve-nets, as studied by psychologists, physiologists and others. (The hypothetical studies, of course, alert researchers to complex configurations and possibilities that may turn out to occur in reality, as well as being interesting for their own sake-- and conceivably as useful ways of organizing things to be built.)

However, an earlier myth, that you could simulate neurons till you got a person, is about dead.

SIMULATION OF THOUGHT-PROCESSES

Nobody talks anymore about simulating artificial brains; there's too much to it, and it involves dirty approximations.

However, a cleaner area is in the simulation of thought: creating computer programs that mimic man's mental processes as he dopes through various problems. Trying things out, deducing thoughts from what's already known, following through the consequences of guesses-- these can all be done by programs that "try to figure out" answers to problems like The Cannibal and The Missionary, or whatever.

AUTOMATA

"Automata", as the term is used in this field, is just a fancy word for imaginary critters, particularly little thingies that behave in exact ways. (The Game of Life, see p. 48, is an automaton in this sense.)

SELF-ORGANIZING SYSTEMS, SELF-REPRODUCING SYSTEMS, AND SO FORTH

These are terms for imaginary objects, having exactly defined mathematical properties, about which various abstract things can be proven that tend to be of interest only to mathematicians.

SPEECH

1. SENTENCE GENERATION

The problem of computers speaking human languages-- not to be confused with computer languages, pp. 15-25 and elsewhere-- is incredibly complicated. Just because little human tykes start doing it effortlessly, it is easy to suppose that it's a basically easy problem.

No way.

Only since the mid-fifties has human language begun to be understood. That was when Noam Chomsky discovered the inner structure of human languages: namely, that the long (and complex) sentence constructions of language are built out of certain exact operations. Previous linguists had sought to classify the sentence structures themselves; this led to complexities which Chomsky discovered were unnecessary. It is unnecessary to catalog sentence types themselves if we can simply isolate, instead, the exact processes by which they are generated.

These processes he called transformations (a term he borrowed from mathematics). All utterances are created from certain elementary pieces, called kernels, which are then chewed by transformations into surface structures, the final utterances. Examples of kernels: The man lives in the house, The house is white. Result of combining transformation: The man lives in the white house. Kernel: I go. Result of past-tense transformation: I went.

The most important finding, now, is that the transformations are carried out in orderly sequences: any sentences can have more transformations carried out on it, all adhering to the basic rules, resulting in the most complex sentences of any language.

Linguists since then have confirmed Chomsky's conjecture, and proceeded to work out the fundamental transformations of major languages, including English.

Now, one result of all this is that it turns out to be easier to generate sentences in a language than to understand them. Why? Because it is comparatively easy to program computers to apply transformations to kernels, BUT very hard to take apart the result. A complex "surface structure" may have numerous possible kernels-- does "Time flies like an arrow" have the same structure as "Susie sings like a bird" or "Fruit flies like an orange?"

Result: to program a computer to generate speech-- that is, invent sentences about a data structure and type them out-- is comparatively easy, but to have it recognize incoming sentences, and break them up into their kernel meanings, is not.

We may think of a language-generating computer system as follows:

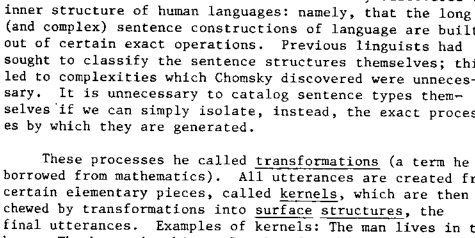

2. SENTENCE RECOGNITION

Chomsky and others have discovered that sets of transformation rules (or grammars, praise be) vary considerably. It is possible to invent languages whose surface structures are easy to take apart, or parse; such languages are called context-free languages. (Most computer languages, see other side, are of this type.) Unfortunately natural languages, like English and French and Navaho, are not context-free. It turns out that the human brain can pick apart language structures because it's so good at making sensible guesses as to what it meant-- and if there is one thing hard to program for computers, it is sensible guessing. (But see "Heuristics," nearby.)

This means that to create computer systems which will take real sentences apart into their meanings is quite difficult. We can't get into the various strategies here; but most researchers cut the problem down in one way or other.

Dorothy read the card aloud, spelling out the big words with some difficulty; and this is what she read:

SMITH & TINKER'S
Patent Double-Action, Extra-Responsive,
Thought-Creating, Perfect-Talking
MECHANICAL MAN
Fitted with our Special Clock-Work Attachment.
Thinks, Speaks, Acts, and Does Everything but Live.
Manufactured only at our Works at Evna, Land of Ev.
All infringements will be promptly Prosecuted according to Law.

"How queer!" said the yellow hen. "Do you think that is all true, my dear?"

55

Ozma of Oz

"I don't know," answered Dorothy, who had more to read. "Listen to this, Billina:"

DIRECTIONS FOR USING:
For THINKING:—Wind the Clock-work Man under his left arm, (marked No. 1.)
For SPEAKING:—Wind the Clock-work Man under his right arm, (marked No. 2.)
For WALKING and ACTION:—Wind Clock-work in the middle of his back, (marked No. 3.)
N. B.—This Mechanism is guaranteed to work perfectly for a thousand years.

"Well, I declare!" gasped the yellow hen, in amazement; "if the copper man can do half of these things he is a very wonderful machine. But I suppose it is all humbug, like so many other patented articles."

"We might wind him up," suggested Dorothy, "and see what he'll do."

GORDON PASK

Gordon Pask is one of the maddest mad scientists I have ever met, and also one of the nicest. An eloquent English leprechaun who dresses the Edwardian dandy, Pask sows awe wherever he goes. A former doctor and theatrical producer, Pask is one of the great international fast-talkers, conference-hopping round the globe from Utah to Washington to his project at the Brooklyn Children's Museum. This spring, 1974, he has been at the University of Illinois at Chicago Circle, but soon he goes back to England and his laboratory.

In a field full of brilliant eccentrics, Pask has no difficulty standing out.

Pask is one of the Artificial Intelligencers who is working on teaching by computer, about which more will be said; but the original core of his interest is perhaps the process of conceptualization and abstraction.

Pask has done a good deal on the mathematics of self-contemplating systems, that is, symbolic representations of what it means for a creature (or entity omega) to look at things, see that they are alike, and divine abstract conceptions of them. A crowning moment is when Omega beholds itself and recognizes the continuity and selfhood. (Pask says several others-- scholars from Argentina, Russia and elsewhere-- have hit on the same formulation.)

Models and abstraction, then, are what we may call the first half of Pask's work.

Gordon Pask will be continued on p. DM 47.

3. SPEECH OUTPUT AS SOUND

It is possible in principle to set up computers to "talk" by converting the language surface structures that their programs come up with into actual sound. See "Audio," p. DM 11.

4. SPEECH INPUT TO COMPUTERS BY ACTUAL SOUND

So far we have been talking about the computer's manipulation of language as an alphabetical coding or similar representation. To actually talk at a computer is another kettle of fish. This means breaking down the sound into phonemes and then breaking it into a data structure which can be treated with the rules of grammar-- a whole nother difficult step.

A few attempts have been made to market devices which would recognize limited speech and convert it to symbols to go into the computer. One of them, which supposedly can distinguish among thirty or forty different spoken words, is supposedly still on the market. Specific users have to "train" it to the particulars of their voices.

I repeatedly hear rumors of "dictation machines" which will type what you say to them. If such things exist I have been unable to confirm it.

(Everybody says that of course what we want is to be able to communicate with computers by speech. Speaking personally, I certainly don't. Explaining my punctuation to human secretaries is hard enough, let alone trying to tell it to a computer, when it's easy enough to type it in.)

5. ALL TOGETHER NOW

The complexity of the problem should by now be clear.

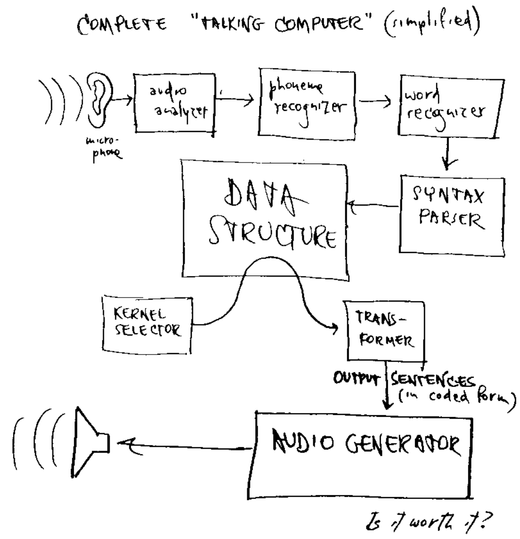

CYBERNETICS

Gordon Pask calls his field Cybernetics. The term "cybernetics" is heard a lot, and is one of those terms which, in the main, mankind would be better off without; although after talking to Pask I get the sense that there may be something to it after all.

The term "cybernetics" was coined by Norbert Wiener, the famously absent-minded mathematician who (according to legend) often failed to recognize his own children. Wiener did pioneering work in a number of areas. A special concern of his was the study of things which are kept in control by corrective measures, or, as he called it, Feedback. The term "cybernetics" he made out of a Greek word for steersman, applying it to all processes which involve corrective control. It turns out that almost everything involves corrective control, so the term "cybernetics" spreads out as far and as thinly as you could possibly want (The public is under the general impression that "cybernetics" refers to computers, and the computer people should be called "cyberneticians." There seems to be nothing that can be done about this. See "cybercrud," p. 8. This is an even worse term meaning "steering people into crud," specifically, putting things over on people using computers.)

Properly, the core of "cybernetics" seems to deal with control linkages, whether in automobiles, cockroaches or computers. However, people like Pask, von Foerster, Ashby (and so on) appear to extend the concept generally to the study of forms of behavior and adaptation considered in the abstract. The validity and fascination of this work, of course, is quite unrelated to what you call it.

THE TURING MACHINE

Is the most classical abstract Automaton. A Turing Machine, named after its discoverer, is a hypothetical device which has an infinite recording tape that it can move back and forth, and the ability to make decisions depending on what's written there.

Turing proceeded to point out that no matter how fast you go step-by-step, you can't ever outrun certain restrictions built into all sequential processes as represented by the Turing Machine. This lays heavy limits on what can ever be done step-by-step by computer. (It means we have to look for non-step-by-step methods, which much of Artificial Intelligence is about.)

DO WE WANT TALKING SYSTEMS?

I had one quite irritating experience with a "conversational" system, that is, computer program that was supposed to talk back to me. I was supposed to type to it in English and it was supposedly going to type back to me in English. I found the experience thoroughly irritating. My side of the conversation, which I sincerely tried to keep simple, produced repeated apologies and confusion from the program. The guy who'd created the program kept explaining that the program would be improved, so that eventually it could handle responses like mine. My reaction was, and is, Who needs it?

Many people in the computer field seem to think we want to be able to talk to computers and have them talk back to us. This is by no means a settled matter.

Talking programs are complicated and require a lot of space in the machine, and (more importantly) require a lot of time by programmers who could achieve (I think) more in less time by other means. Moreover, talking programs produce an irritating strategical paradox. In dealing with human beings, we know what we're dealing with, and can adjust what we say accordingly; there is no way to tell, except by a lot of experimenting, what the principles are inside a particular talking program; so that trying to adjust to it is a strain and an irritation. (Compare: talking to a stranger who may or may not turn out to be your new boss.) Now, some programmers keep saying that eventually they'll have it acting just as smart as a real person, so we needn't adjust; but that's ridiculous. We always adjust to real people. In other words, the human discomfort and irritation of psyching the system out can never be eliminated.

Furthermore, on today's sequential equipment and with feasible budgets, I personally think the likelihood of making programs that are really general talkers is a foolish goal. There are many simpler ways of telling computer systems what you want to tell them-- light pen choice, for example.

Moreover, having to type in whole English phrases can be irritating. (We can't even get into the problem of having the computer pick apart the audio if you talk it in.)

This is not to say understandably restricted talking systems are bad. If you know and understand the sorts of response the system makes to what kinds of thing, then an English-like response is really a clear message. For instance, the JOSS system (the first Quickie language-- see p. 15) had an eloquent message:

eh?

which actually meant, What you have just typed in does not fit the rules of acceptable input for this system. But it was short, it was quick, it was simple, and it was almost polite.

Similarly, talking systems that use an exact vocabulary, whose limits and abilities are known to the person, are okay. (Winograd, see Bibliography, has a nice example of telling a computer to stack blocks, where the system knows words like between, on, above and so on.) Where this is understood by the human, it can be a genuine convenience rather than a spurious one.

(The problem of rudeness in computer dialogue has not been much discussed. This is partly because many programmers are not fully aware of it, or, indeed, some are so skilled in certain subtle forms of rudeness they wouldn't even know they weren't acceptable. The result is that certain types of putdown, poke, peremptoriness and importunacy can find their way into computer dialogue all too easily. Or, to put it another way: nobody like to be talked back to. Cf. Those stupid green THANK YOU lights on automatic toll booths.)

Now, this is not to say that research in these areas is wrong, or even that researchers' hopes of some breakthrough in talking-systems is misguided. I am saying, basically, that talking systems cannot be taken for granted as the proper goal in computers to be used by people; that the problems of rudeness, and irritating the human user, are far greater than many of these researchers suppose; and that there may be alternatives to this potentially eternal leprechaun-chasing.

If like the author you are bemused by the great difficulty of getting along with human beings, then the creation of extraneous beings of impenetrable character with vaguely human qualities can only alarm you, and the prospect of these additional crypto-entities which must be fended and placated, clawing at us from their niches at every turn, is both distasteful and alarming.

Artificial Intelligence enthusiasts unfortunately tend to have a magician's outlook: to make clear how their things work would spoil the show.

Thus, for a rather peculiar art show held at New York's Jewish Museum in 1970, a group from MIT built a large device that stacked blocks under control of a minicomputer (Interdata brand). Now, the fact that it could stack and re-stack blocks with just a minicomputer was really quite an accomplishment, but this was not explained.

Instead, the block-stacking mechanism was enclosed in a large glass pen, in which numerous gerbils-- hoppy little rodents-- were free to wander about. When a gerbil saw that a block was about to be stacked on him, he would sensibly move.

Now, it is fairly humorous, and not cruel, to put gerbils into a block-stacking machine. But this was offered to the public as a device partaking of a far more global mission, the experimental interaction of living creatures and a dynamic self-improving environment, blah blah blah.

Passersby were awed. "Why are those animals in there?" one would say, and the more informed one would usually say, "It's some kind of scientific experiment."

Well, this is a twilight area, between science and whimsical hokum, but one cannot help wishing simple and humorous things could be presented with their simplicity and humor laid bare.

I remember watching one gerbil who stood motionless on his little kangaroo matchstick legs, watching the Great Grappler rearranging his world. Gerbils are somewhat inscrutable, but I had a sense that he was worshiping it. He did not move until the block started coming down on top of him.

I take this as an allegory.

CAN A COMPUTER PLAY CHESS?

The real question is, can a set of procedures play chess? Because that's what the computer program really does, enact a set of procedures.

And the answer is yes, fairly well.

Now, a chess program is not something you jot down on the back of an envelope one afternoon. It's usually an immense, convoluted thing that people have worked on for years. (Although I vaguely recall that second place in the 1970 inter-computer chess contest was won by a program that occupied only 2000 locations in a 16-bit minicomputer-- in other words, a compact and tricky sneaker.)

Now, simple games (like tic-tac-toe and Nim and even Cubic) can be worked out all the way: all alternatives can be examined by the program and the best one found. Not so with chess.

Chess basically involves, because of its very many possibilities, a "combinatorial explosion" of alternatives (see p. 45); that is, to look at "all" the possibilities of a midgame would take forever (perhaps literally-- the Turing problem), and thus means must be found for discarding some possibilities.

The structure of branching possibilities is a tree (see p. 26); so that methods of "pruning" the tree turn out to be crucial.

Basically there are two approaches to the design of chess programs. In one approach, the programmers look for specific threats and opportunities in the data structure representing the board, and try to find good strategies for selecting good moves on the basis of them. This is the approach taken in COKO, the "Cooper-Koz" chess program. The programmers selectively cope with individual problems and strategies as they turn out to be necessary. (This means that it is likely to have specific Achilles' heels; which, of course, the authors of the program keep trying to repair by adding specific corrections.)

A different approach is taken by the Greenblatt chess program. This is basically a big Heuristic program. It "learns" best strategies in chess by "watching" the game. That is, your pour historical chess matches through it, and it tries out strategies-- making various tentative rules about what kinds of moves are good, then scoring these moves according to the results of making them-- as seen in positional advantages that resulted in actually championship play.

Obviously this is a field in itself. You won't get grants for it, but to those who really care about both chess and computers, it's the only thing to be doing.

FRANKENSTEIN MEETS CYBERCRUD

Fred Brooks, the keynote speaker at the IEEE computer conference in Fall 74, seems to have said that HAL 9000 (the unctuous, traitorous Presence in the movie 2001) was the way computers should be. (Computer Decisions, Apr 74, 4.)

I find it hard to believe that anybody could think that. Nevertheless, there are those artificial-intelligence freaks whose view it is that the purpose of all this is eventually (a) to create servants that will read our minds and do our bidding, (b) servants who will take things over and will implement human morality, regardless of our bidding (though we humans are too frail to do so-- as in Asimov's I, Robot); or even (c) create masters who will take everything over and run everything according to their own principles and the hell with us. (I met a man in a bar, after an ACM meeting, who claimed to believe this was the purpose of it all: to create the master race that would replace us.)

According to Arthur C. Clarke's retroactive novel 2001: A Space Odyssey (Signet, 1968, 95¢), the HAL 9000 computer series began as follows:

> "In the 1980s, Minsky and Good had shown how neural networks could be generated automatically-- self-replicated-- in accordance with any arbitrary learning pattern. Artificial brains could be grown by a process strikingly analogous to the development of the human brain." (P. 96.)

I don't know who Good is, but these are among the lines Minsky has been working along for years, so I hope he's encouraged by the news of what he's going to accomplish.

Anyhow, so okay they grow the HAL 9000 in a tank. Then how come in the Death-of-Hal scene we see Keir Dullea bobbing around loosening circuit cards, just as if it were a plain old 1978 computer?

Possible answer #1. It is rumored that Clarke's retro-novel was Clarke's rebuttal to Kubrick's final film.

Possible answer #2. HAL's tanks of neural glop are controlled by PDP-11s, one to a card.

(Of course, if you take the letters after H, A and L in the alphabet, you get I, B and M. So maybe those are 1130s.)

DEUS EX MACHINA

Obviously such beliefs are outside the realm of science or engineering. They belong to pure speculation; and while various mechanisms have in fact been programmed to croak, stagger, stack blocks, compose sentences and so on, to suppose that we are in any real sense anywhere near mimicking human intelligence, let alone surpassing and superseding it, is either to be totally fooled or to hanker after some curious dream from inside yourself.

As we said on the other side of the book, everybody in computers has deeper motivations and interior twists that form his own special ties to these machines; and when it comes to our choices of fantasy machines, obviously an even deeper level of psychic imprint is projecting itself into the world.

...EX MENSA

People who fantasize about wondrous creatures and deities they want to make out of the computer obviously have something interesting in their own heads from which that comes. Perhaps it comes from a desire for imaginary playmates, or an ambivalence toward authority, or goodness knows what; there are so many odd people at different ends of Artificial Intelligence that there may be a lot of different psychological systems at work. Or maybe artificial intelligence is just where the most brilliant, determined and eccentric dreamers go. Anyway, I can only ask the question, not give the answer.

THE THREE LAWS OF ROBOTICS

1-- A robot may not injure a human being, or, through inaction, allow a human being to come to harm.

1-- A robot must obey the orders given it by human beings except where such orders would conflict with the First Law.

3-- A robot must protect its own existence as long as such protection does not conflict with the First or Second Law.

-- Handbook of Robotics, 56th Edition, 2058 A.D., quoted in I. Asimov, I, Robot.

Sometimes All We Need Is A Sympathetic Teletype

ELIZA WILL SEE YOU NOW

(What the ELIZA program types, we represent as all caps; what the person types is not.)

PLEASE TELL ME YOUR PROBLEM.
  I am feeling depressed.
HOW LONG HAVE YOU BEEN FEELING DEPRESSED?
  Ever since my mother left.
TELL ME MORE ABOUT YOUR FAMILY.
  My father is very strict.
PLEASE GO ON.
  Last night I dreamt he spanked me.
WHAT WOULD IT MEAN TO YOU IF HE SPANKED YOU?
  Didn't you ever get spanked?
WE WERE DISCUSSING YOU, NOT ME.
  You're just a dumb machine anyway.
DO COMPUTERS FRIGHTEN YOU?
  (quoted in Winograd, 74-5)

The fascinating thing about the ELIZA program is that it seems so like a real person-- that is, a psychoanalyst of the nondirective, Carl Rogers school.

But all it is is a bag of tricks.

It has a lot of canned phrases it can throw out, like PLEASE TELL ME YOUR PROBLEM.

Then it spots certain types of sentences and turns them around. If the person types a sentence beginning with "I am...", the ELIZA program turns it around with HOW LONG HAVE YOUR BEEN... ?

And it looks for specific words which indicate a sentence can be turned around for a convincing result. For instance, if the person uses the word "everyone," the program simply replies, WHO IN PARTICULAR CAN YOU THINK OF?

And certain highly-charged words are signals to throw in a zinger, regardless of what the user may have said. For instance, if the person uses the word "mother," the machine replies:

TELL ME ABOUT YOUR FAMILY.

Now, do not be fooled. There is no Ghost in the Machine. The program does not "understand" the user. THE PROGRAM IS MERELY ACTING OUT THE BAG OF TRICKS THAT JOE WEIZENBAUM THOUGHT UP. Credit where credit is due: not to The Computer's Omniscience, but to Weizenbaum's cleverness.

(Look at the above sample dialogue and see if you guess what tricks the program was using.)

The thing is, many people refuse to believe that it's a program. Even when the program's tricks are explained.

And even some who understand ELIZA like to call it up from their terminals for companionship, now and then.

BIBLIOGRAPHY

Terry Winograd, "When Will Computers Understand People?" Psychology Today May 74, 73-9.

(Weizenbaum's full article on ELIZA appeared in the Communications of the ACM sometime in the mid or late sixties; a flowchart revealed its major tricks.)

I have strong hunches about the inner workings of men who get millions of dollars from the Department of Defense and then say in private that really they're going to use it to create a machine so intelligent it can play with their children. (Not to name names or anything.) An obvious question is, do they play with their children? No, they play with computers.

But the point here is not to hassle the dreamers, just to sort out the dreams and put them on hangers so you can try them on, and maybe choose an ensemble for yourself.

BIBLIOGRAPHY

Terry Winograd, "When Will Computers Understand People?" Psychology Today May 74, 73-79. Particularly readable article.
Nicholas Negroponte, The Architecture Machine. MIT Press. Takes the view that computers should be made into magical servants which not only handle bothersome details, but more or less read our minds as well.
Leonard Uhr, Pattern Recognition, Learning and Thought: Computer-Programmed Models of Higher Mental Processes. Prentice-Hall.
Arthur W. Holt, "Algorithm for a Low Cost Hand Print Reader." Computer Design Feb 74, 85-89.
Edward A. Feigenbaum and Julian Feldman (eds.), Computers and Thought. McGraw-Hill. Old but still good for orientation.
A journal: Artificial Intelligence (North-Holland Publishing Co., Journal Division; P.O. Box 211, Amsterdam, The Netherlands. Was $26.50 a year in 1973. This'll show you what they're thinking about now.
Roger Lind, "The Robots Are Coming, The Robots Are Coming." Oui, Feb 1974. Typical layman's hype. You don't get told until the second page that a typical industrial "robot" is a huge mechanism with one grappling arm.
Edward W. Kozdrowicki and Dennis W. Cooper, "COKO III: the Cooper-Koz Chess Program." CACM July 73, 411-427.
Greenblatt, R.D., Eastlake, D.E., and Crocker, S.D., "The Greenblatt Chess Program." Proc FJCC 67, 801-810.
R.C. Gammill, "An Examination of Tic-Tac-Toe-like Games." Proc. NCC 74, 349-355.

INFORMATION RETRIEVAL

"Information Retrieval" is one of those terms that laymen throw around as if it were a manhole cover. It sounds as though it means so much, so very much. And so you actually hear people say things like: "But that would mean... (pregnant pause) ... Information Retrieval!!!" Similarly, some of the hokey new copyright notices you see in books from With-It publishers intone that said books may not be "placed in any information retrieval system..." I take this to mean that the publishers are forbidding you to put the book on a bookshelf, because "information retrieval" simply means any way at all of getting back information from anything. A bookshelf, since it allows you to read the spines of the books, is indeed an Information Retrieval System.

It happens, incidentally, that the phrase "information retrieval" was coined in the forties by Calvin Mooers, inventor of TRAC™ Language (see pp. 18-21). (If Wiener had coined it he might have called it Getback. If Diebold had coined it it might have been Thoughtomation.)

Anyhow, numerous entirely different things go on in the field, all under the name of Information Retrieval. Here are some.

1. Non-computer retrieval. (See Becker and Hayes, Automatic Information Retrieval.) These things are kind of old-fashioned fun-- cards with holes punched along the edge, for instance, that you sort with knitting needles, or the more recent systems with holes drilled in plastic cards. Trouble is, of course, that computers are becoming much more convenient and even less expensive than these, counting your own time as being worth something.

2. Document Retrieval. This basically is an approach that glorifies the old library card file, except now the stuff is stored in computers rather than on cards. But what's stored is still the name of the document, who wrote it, where it was published and so on. Obviously helpful to librarians, but scarcely exciting.

3. Automatic document indexing. Some organizations find it helpful to have a computer try to figure out what a book is about, rather than have a person look at it and check. (I don't see why this saves anything, but there you are.) Anyway, the text of the document (or selected parts) are poured through a computer program that selects, for instance, keywords, that is, the most important words in it, or rather words the program thinks are most important. Then these keywords can go on the headings of library file cards, or whatever.

There are various related systems by which people study, for instance, the citations between articles, but we won't get into that.

4. Content retrieval. Now we're getting to the sexy stuff. A system for content retrieval is one that somehow stores information in a computer and lets you get it back out.

The trick on both counts is of course how.

Well, as we said on the other side of the book, any information stored in a computer has a data structure, which simply means whatever arrangement of alphabetical characters, numbers and special codes the computer happens to be saving.

In a content-retrieval system, information on some subject is somehow jammed into a data structure-- possibly even by human coders-- and then set up so people can get it back out again in some way. Lot of possibilities here, get it?

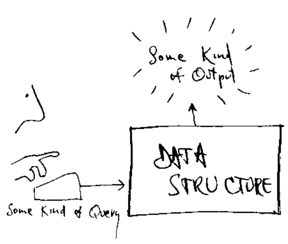

In the most startling of these systems, the QAS, or "Question-Answering System," some sort of dialogue program (see "Artificial Intelligence," nearby) tries to give you answers about the data structure. But this means there have to be a whole lot of programs:

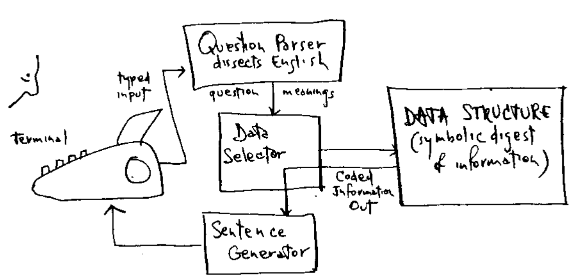

These systems can be quite startling in the way they seem to understand you (see Licklider book; also Winograd piece under Artificial Intelligence). But they don't understand you. They are just poor dumb programs.

Many people (including Licklider) seem to see in Question-Answering-Systems the wave of the future. Others, like this author, are skeptical. It's one thing to have a system that can deduce that Green's House is West of Red's House from a bunch of input sentences on the subject, but the question of how much these can be improved is in some doubt. A system that can answer the question, "What did Hegel say about determinism?" is some ways away, to put it mildly.

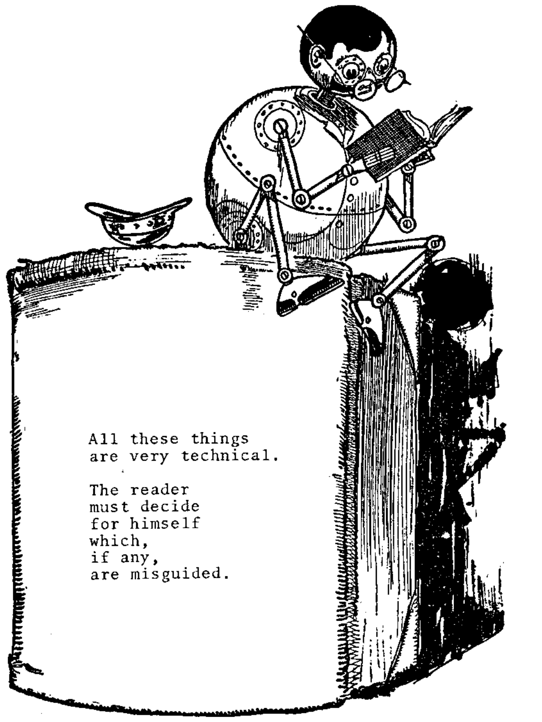

Then there is the matter of consistency. The really interesting subjects are the ones where different authors claim opposing facts to support opposite conclusions. In other words, there is inconsistency within the content of the field. In this case such systems are going to have a problem. (See "Rasho-Mon Principle" under "Tissue of Thought," pp. DM 16-19.)

Another fundamental point is this. It may be easy enough to program a system to answer the question,

WHAT TIME DOES THE NEXT PLANE LEAVE FOR LAGUARDIA?

but it is a lot simpler to have it display schedules your eye can run down, or allow you to go look at some kind of graphic display.

Speaking personally, I don't like talking to machines and I don't like their talking back to me. I'm not saying you have to agree, I'm just telling you you're allowed to feel that way.

5. Screen summaries. These systems let you sit at a computer display screen and read summaries of various things, as well as run through them with various programs to look for keywords. (The New York Times now offers such a system, costing over a thousand dollars a month to subscribers.)

6. "Full-text systems." These are systems that one way or another allow you to read all the text of something from a computer display screen. There are those of us who see these as the wave of the future, but many others are perfectly outraged at the thought. (Hypertext systems, now, are setups that allow you to read interconnected texts from computer display screens. See pp. DM 44-7.)

This has been brief and has skipped a lot. Anyway, as you see, IR is no one thing.

BIBLIOGRAPHY

Vannevar Bush, "As We May Think." Atlantic Monthly, July 1945, 101-108.
Theodor H. Nelson, "As We Will Think." Proc. Online 72 Conference, Brunel U. Uxbridge, England.
G. Salton, "Recent Studies in Automatic Text Analysis and Document Retrieval." JACM, Apr 73, 258-278.
Donald E. Walker (ed.), Interactive Bibliographic Search: The User/Computer Interface. AFIPS Press, $15.
Theodor H. Nelson, "Getting It Out of Our System." In Schechter (ed.), Critique of Information Retrieval (Thompson Books, 1967).
J.C.R. Licklider, Libraries of the Future. MIT Press, 1965. Clear and readable summary of the rest of the field; then he goes on to advocate "procognitive systems," systems that will digest what's known in any field and talk back to you, using techniques of artificial intelligence. Whatever its other merits, this book is great for shaking people up, especially librarians. It seems so official.
Richard M. Laska, "All the News That's Fit to Retrieve." Computer Decisions, Aug 72, pp. 18-22.

It is a truism that Mendel's theories of genetics got "lost" after publication in 1865, to be rediscovered in 1900. "If only there had been proper information retrieval under the right categories," people often say. Recent studies indicate that the publication containing Mendel's paper reached, or got nearly to, "practically all prominent biologists of the mid-nineteenth century." (Scientific American, July 68, 55.)

I take this as suggesting that the problem isn't categorical retrieval at all. It's multi-connected availability (see "hypertext," pp. DM 44-7.

COMPUTER-ASSISTED INSTRUCTION

Like Artificial Intelligence and Information Retrieval, Computer-Assisted Instruction sounds like something exact and impressive but is in fact a scattering of techniques tied together only nominally by a general idea.

The real name for it should be Automated Dialogue Teaching. That would immediately allow you to ask, should computer teaching use dialogues? But they don't want you to ask that.

In the classic formulation of the early sixties, there were going to be three levels of CAI: "drill-and-practice" systems, much like teaching machines, that simply helped students practice various skills; a middle level (often itself called, confusingly, "computer-assisted instruction"); and a third level, the Socratic system, which would supposedly be Ideal. Students studying on Socratic systems would be eloquently and thoughtfully instructed and corrected by a perfect being in the machine. "We don't know how to do that yet," the people keep saying. Yet, indeed.

(My personal view on this subject, expressed in an article (following) is that Computer-Assisted Instruction in many ways extends the worst features of education as we now know it into the new realm of presentation by computer.)

DOES THE NAME PAVLOV RING A BELL?

This is a true story. (The details are approximate.) It may provide certain insights.

An Assistant Commissioner of Education was being shown a CAI system by representatives of a large and well-known computer company.

One one side of the Commissioner stood a salesman, who wanted him to be impressed. On the other side stood one Dr. S., who knew how the system worked.

The terminal, demonstrating a history program that had hurriedly been put together, typed: WHO CAPTURED FORT TICONDEROGA?

"Can I type anything?" asked the Assistant Commissioner.

"Sure," said the salesman, ignoring the frantic head-shaking of Dr. S.

The Assistant Commissioner typed: Gypsy Rose Lee.

The machine replied:

NO, BUT YOU'RE CLOSE. HE CAPTURED QUEBEC A SHORT TIME LATER.

The Assistant Commissioner evidently enlivened many a luncheon with that one, and Computer-Assisted Instruction was effectively dead for the rest of the administration.

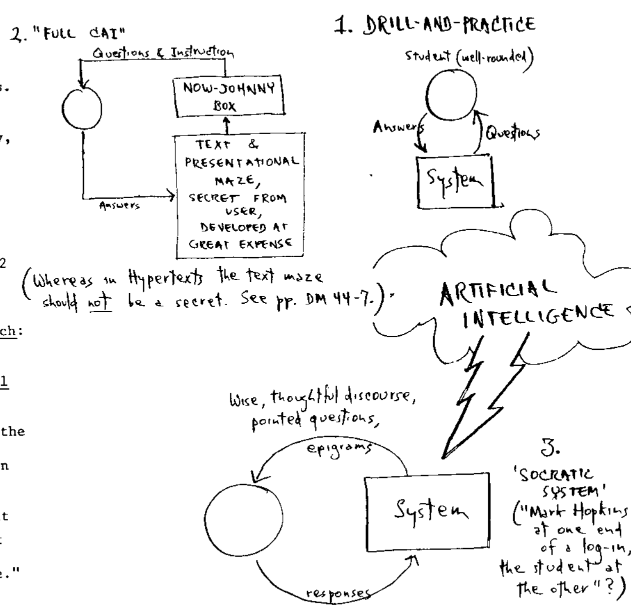

ANOTHER ANECDOTE

Some of us have been saying for a long time that learning from computers ought to be under control of the student.

One group (never mind who) has taken hold of this idea and gotten a lot of funding for it under the name of STUDENT CONTROL. This group talks as if it were some kind of scientific breakthrough.

A friend of mine suggests, however, that this phrase may have brought the funding because administrators thought it meant control of the student.

BIBLIOGRAPHY

George B. Leonard, Education and Ecstasy. Dell, $2.25. Argues for making education an enthusiastic process.
Theodor H. Nelson, "No More Teachers' Dirty Looks." Follows.

DM 16

(The following article appeared in the September, 1970 issue of Computer Decisions, and got an extraordinary amount of attention. I have changed my views somewhat-- we all go through changes, after all-- but after consideration have decided to re-run it in the original form, without qualifications, mollifications or anything, for its unity. Thanks to Computer Decisions for use of the artwork by Gans and for the Superstudent picture on the cover, whose artist unfortunately insists on preserving his anonymity.

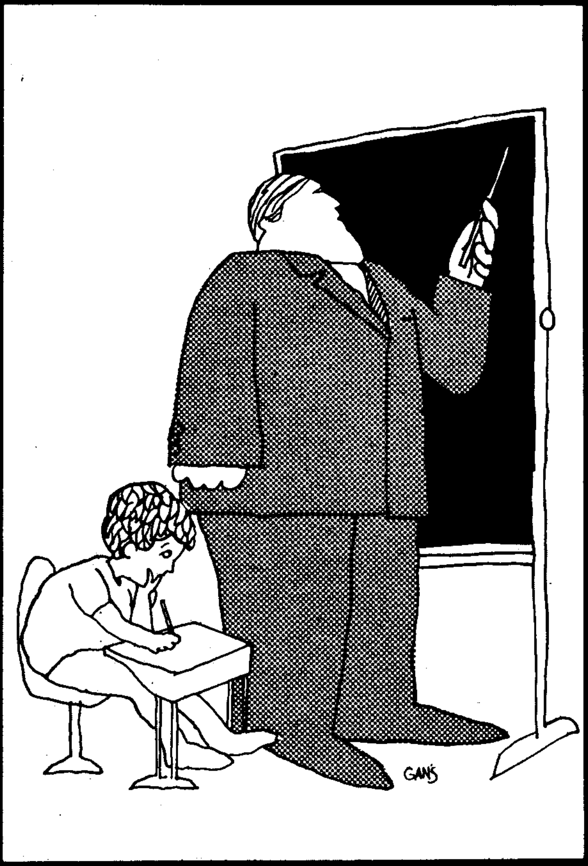

**Did you find school dismal and dreary?
Did it turn you off?
Here the author proposes safe and legal ways to turn kids on.**

NO MORE TEACHE

by Theodor H. Nelson
The Nelson Organization
New York

Some think the educational system is basically all right, and more resources would get it working again. Schools would do things the same way, except more so, and things would get better.

In that case the obvious question would be, how can computers help? How can computers usefully supplement and extend the traditional and accepted forms of teaching? This is the question to which present-day efforts in "computer-assisted instruction" — called CAI — seem to respond.

But such an approach is of no possible interest to the new generation of critics of our school system — people like John Holt (*Why Children Fail*), Jonathan Kozol (*Death at an Early Age*) and James Herndon (*The Way It Spozed To Be*). More and more, such people are severely questioning the general framework and structure of the way we teach.

These writers describe particularly ghastly examples of our schooling conditions. But such horror stories aside, we are coming to recognize that schools as we know them appear designed at every level to sabotage the supposed goals of education. A child arrives at school bright and early in his life. By drabness we deprive him of interests. By fixed curriculum and sequence we rob him of his orientation, initiative and motivation, and by testing and scoring we subvert his natural intelligence.

Schools as we know them all run on the same principles: iron all subjects flat and then proceed, in groups, at a forced march across the flattened plain. Material is dumped on the students and their responses calibrated; their interaction and involvement with the material is not encouraged nor taken into consideration, but their dutifulness of response is carefully monitored.

While an exact arrangement of intended motivations for the student is preset within the system, they do not usually take effect according to the ideal. It is not that students are *un*motivated, but motivated *askew*. Rather than seek to achieve in the way they are supposed to, students turn to churlishness, surliness, or intellectual sheepishness. A general human motivation is god-given at the beginning and warped or destroyed by the educational process as we know it; thus we internalize at last that most fundamental of grownup goals: just to get through another day.

Because of this procedure our very notion of human ability has suffered. Adult mentality is

An interesting point, incidentally, is that people read this a lot of different ways. One Dean of Education hilariously misread it as an across-the-board plug for CAI. Others read in it various forms of menace or advocacy of generalized mechanization. One letter-writer said I was a menace but at least writing articles kept me off the streets. Here is my fundamental point: computer-assisted instruction, applied thoughtlessly and imitatively, threatens to extend the worst features of education as it is now.

cauterized, and we call it "normal." Most people's minds are mostly turned off most of the time. We know virtually nothing of human abilities except as they have been pickled and boxed in schools; we need to ignore all that and start fresh. To want students to be "normal" is criminal, when we are all so far below our potential. Buckminster Fuller, in *I Seem To Be A Verb*, says we are all born geniuses; Sylvia Ashton-Warner tells us in *Teacher* of her success with this premise, and of the brilliance and creative potential she was able to find in all her schoolchildren.

Curricula themselves destructively arrange the study situation. By walls between artificially segregated "studies" and "separate topics" we forbid the pursuit of interest and kill motivation.

In ordinary schooling, the victim cannot orient himself to the current topic except by understand-

RS' dIRTY LOOKS

ing the official angle of approach and presentation. Though tie-ins to previous interests and knowledge are usually the best way to get an initial sense of a thing, there is only time to consider the officially presented tie-ins. (Neither is there time to answer questions, except briefly and rarely well — and usually in a way that promotes "order" by discouraging "extraneous" tie-ins from coming up.)

The unnecessary division and walling of subjects, sequencing and kibbling of material lead people to expect simplifications, to feel that naming a thing is understanding it, to fear complex wholes; to believe creativity means recombination, the parsing of old relations, rather than synthesis.

Like political boundaries, curriculum boundaries arise from noticeable features of a continuum and become progressively more fortified. As behind political borders, social unification occurs within them, so that wholly dissimilar practitioners who share a name come to think they do the same thing. And because they talk mainly to each other, they forget how near is the other side of the border.

Because of the fiction of "subjects," great concern and consideration has always gone into calculating the "correct" teaching sequence for each "subject." In recent years radical new teaching sequences have been introduced for teaching various subjects, including mathematics and physics. But such efforts appear to have been misinformed by the idea of supplanting the "wrong" teaching sequence with the "right" teaching sequence, one which is "validated." Similarly, we have gone from a time when the instructional sequence was a balance between tradition and the lowest common denominator of each subject, to a time when teachers may pick "flexible optimized strategies" from text-

If the computer is a universal control system, let's give kids universes to control.

STARK & CLEVER APL

Some people call it a "scientific" language. Some people call it a "mathematical" language. Some people are most struck by its use for interactive systems, so to them it's an interactive language. But most of us just think of it as THE LANGUAGE WITH ALL THE FUNNY SYMBOLS, and here they are:

```
⍟ρ∪∩α⊂÷ω↑∊|∧\⊃⎕;⌈∘∇'∆
¨¯<≠≤≥=>)∨:⊥_(↓⊤→⍳~○?⌊-
12384657]9.BF[UN←ITOQD+
PRVCAZ×WYEMO/XL,SJGKH
```

Enthusiasts see it as a language of inconceivable power with extraordinary uses. Cynics remark that it has all kinds of extraordinary powers for inconceivable uses-- that is, a weird elegance, much of which has no use at all, and some of which gets in the way.

This is probably wrong. APL is a terrific and beautiful triumph of the mind, and a very useful programming language. It is not for everybody, but neither is chess. It is for bright children, mathematicians, and companies who want to build interactive systems but feel they should stick with IBM.

APL is one of IBM's better products, probably because it is principally the creation of one man, Kenneth Iverson. It is mainly run on 360 and 370 computers, though implementations exist for the DEC PDP-10 and perhaps other popular machines. (Actually Iverson designed the language at Harvard and programmed it on his own initiative after moving to IBM; added to the product line by popular demand, it was not a planned product and might in fact be a hazard to the firm, should it catch on big.)

APL is a language of arrays, with a fascinating notation. The array system and the notation can be explained separately, and so they will.

Let's just say the language works on things modified successively by operators. Their order and result is based upon those fiendish chicken scratches, Iverson notation.

THAT NIFTY NOTATION

The first thing to understand about APL is the fiendishly clever system of notation that Iverson has worked out. This system (sometimes called Iverson notation) allows extremely complex relations and computer-type events to be expressed simply, densely and consistently.

(Of course, you can't even type it without an IBM Selectric typewriter and an APL ball. Note the product-line tie-in.)

The notation is based on operators modifying things. Let's use alphabetic symbols for things and play with pictures for a minute.

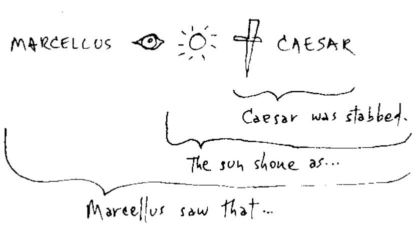

In considering the successive meanings of this rebus we are proceeding from right to left, as you note, and each new symbol adds meaning. This is the general idea.

You will note, in this example, the curious arrangement whereby you can have several pictures, or operators, in a row. This is one of the fun features of the language.

TWO-SIDED OPERATORS

In old-fashioned notations, such as ordinary arithmetic, we are used to the idea of an operator between two things. Like

2 + 2

or in algebra,

x X y

These, too, occur in APL; indeed, APL can also nest two-sided operators-- that is, put them one inside the other, like the leaves of a cabbage. Old-fashioned notations nest with parentheses. But APL nests leftward. It works according to a very simple right-to-left rule.

x X y X 2 + 2

the result of this
is operated on by the next thing and operator, yielding another result
which is in turn operated on by the next thing and operator, yielding final result.

ONE-SIDED OPERATORS

We are also used to some one-sided operators in our previous life. For instance:

\- 1

means the negation of 1;

\- (- 1)

means negating that.

APL can also nest one-sided operators.

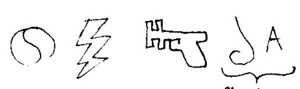

first operator is applied to A;
result is worked on by second operator;
result is worked on by third operator;
result is worked on by fourth operator, yielding final result.

SAME SYMBOLS WORK BOTH WAYS

Now, one of the fascinating kickers of APL is the fact that most of the symbols have both a one-sided meaning and a two-sided meaning; but, thank goodness, they can be easily kept straight.

Here is a concrete example: the symbol ⌈ or "ceiling." Used one-sided, the result of operator ⌈ applied to something numerical is the integer just above the number it is applied to: ⌈7.2 is 8. Used two-sided, the result is whichever of the numbers it's between is larger: 10⌈6 is 10. (There is also ⌊, floor, which you can surely figure out.)

Now, when you string things out into a long APL expression, Iverson's notation determines exactly when an operator is one-sided and when it is two-sided:

As you go from right to left,

another thing? } OP THING
another op? }

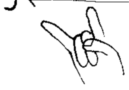

you generally start with a thing on the right. Then comes an operator. If the next symbol is another thing, then the operator is to be treated as a two-sided operator (because it's between two things). If the object beyond the first operator is another operator, however, that means APL is supposed to stop and carry out the first operator on a one-sided basis. Example:

A - B

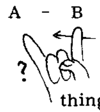

thing,
op,
thing. Conclusion: It's two-sided. Interpretation: "subtract B from A."

A + - B

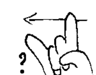

thing,
op,
op--
stop. Conclusion: The first operator is one-sided. Interpretation: "negate B." Then take next symbol.

A WEIRD EXAMPLE, TO HELP WITH THE NOTATION.

Just for kicks, let us make up a notation having nothing to do with computers, using these Iverson principles:

1) If an operator or symbol is between two names of things, carry it out two-sidedly. If not, carry it out one-sidedly.

2) Go from right to left.

The best simple example I can think of involves file cards on the table (named A, B, C...) and operators looking like this:

0) 45) 90) 180) 45ↄ 90ↄ 180ↄ

to which we may assign the following meanings:

ONE-SIDED: ROTATION OPERATORS

0) A	do nothing to A
45) A	rotate A clockwise 45°
90) A	rotate A clockwise 90°
	etc.

TWO-SIDED: STAPLING OPERATORS

B 45) A — staple A (thing named on the right) to B (thing named on the left) at a position 45° clockwise from middle of B's centerline.

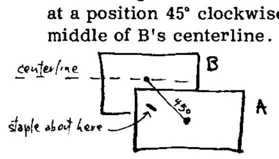

And equivalently for other angles.

Now, using these rules, and letting our things be any file cards that are handy, here are some results:

A 0) B
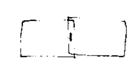

A 90) B
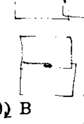

A 90) 90) B
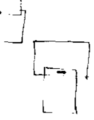

90) A 90) B
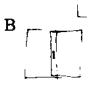

45ↄ A 90) B
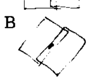

45ↄ A 90) 90) B
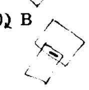

B 0) 45) A
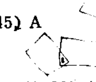

C 45) B 0) 90) A
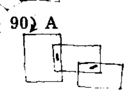

WOOPS! One of these examples is wrong! Can you find it?

It's hard to believe, but there you are. This notation seems adequate to make a whole lot of different stapled patterns.

Exercise! Use this nutty file card notation to program the making of funny patterns. Practice with a friend and see if you can communicate patterns through these programs, one person uncomprehendingly carrying out the other's program and being surprised.

The point of all this has been to show the powerful but somewhat startling way that brief scribbles in notations of this type can have all sorts of results.

DM19

Discrete hypertexts

"Hypertext" means forms of writing which branch or perform on request; they are best presented on computer display screens.

In ordinary writing the author may break sequence for footnotes or insets, but the use of print on paper makes some basic sequence essential. The computer display screen, however, permits footnotes on footnotes on footnotes, and pathways of any structure the author wants to create.

Discrete, or chunk style, hypertexts consist of separate pieces of text connected by links.

Ordinary prose appears on the screen and may be moved forward and back by throttle. An asterisk or other key in the text means, not an ordinary footnote, but a *jump*—to an entirely new presentation on the screen. Such jumpable interconnections become part of the writing, entering into the prose medium itself as a new way to provide explanations and details to the seeker. These links may be artfully arranged according to meanings or relations in the subject, and possible tangents in the reader's mind.

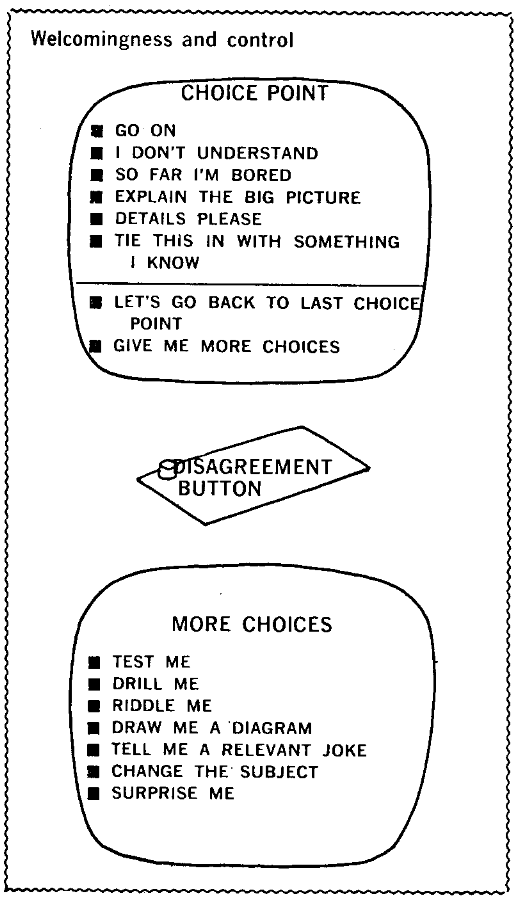

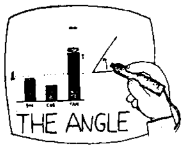

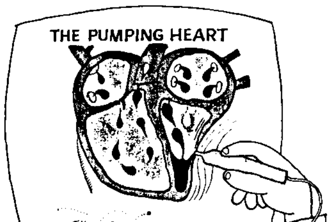

Performing hypergrams

A hypergram is a performing or branching picture: for instance, this angle, with the bar-graph of its related trigonometric functions. The student may turn the angle upon the screen, seizing it with the light-pen, and watch the related trigonometric functions, displayed as bar charts, change correspondingly.

Hypergrams may also be programmed to show the consequenees of a user's prod—what follows or accompanies some motion of the picture that he makes with a pointing tool, like the heartbeat sequence.

Stretchtext fills in the details

This form of hypertext is easy to use without getting lost. As a form of writing, it has special advantages for discursive and loosely structured materials—for instance historical narratives.

There are a screen and two throttles. The first throttle moves the text forward and backward, up and down on the screen. The second throttle causes changes in the writing itself: throttling toward you causes the text to become longer by minute degrees. Gaps appear between phrases; new words and phrases pop into the gaps, an item at a time. Push back on the throttle and the writing becomes shorter and less detailed.

The stretchtext is stored as a text stream with extras, coded to pop in and pop out at the desired altitudes:

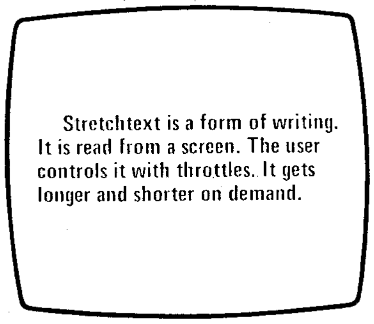

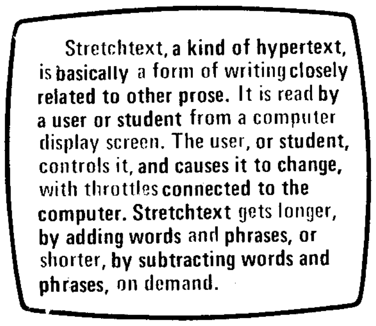

Hypermap zips up or down

The screen is a map. A steering device permits the user to move the map around the world's surface; a throttle zooms it in. Not by discrete jumps, but animated in small changes, the map grows and grows in scale. More details appear as the magnification increases. The user may request additional display modes or "overlays," such as population, climate, and industry. Such additional features may pop into view on request.

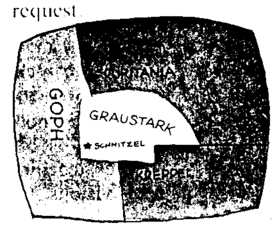

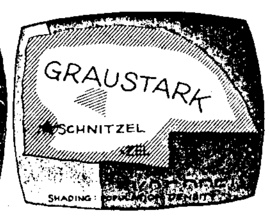

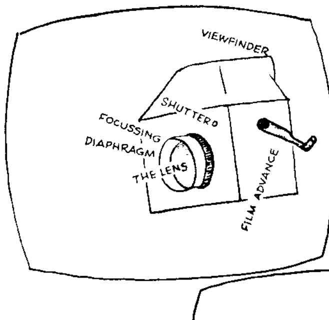

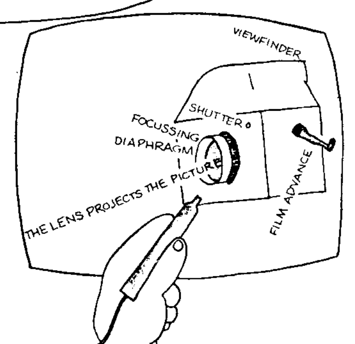

Queriable illustrations: a form of hypergram

A "hypergram" is a picture that can branch or perform on request. In this particular example, we see on the screen a line-drawing with protruding labels. When the student points at a label, it becomes a sliding descriptive ribbon, explaining the thing labelled.

Or asterisks in an illustration may signal jumps to detailed diagrams and explanations, as in discrete hypertexts.

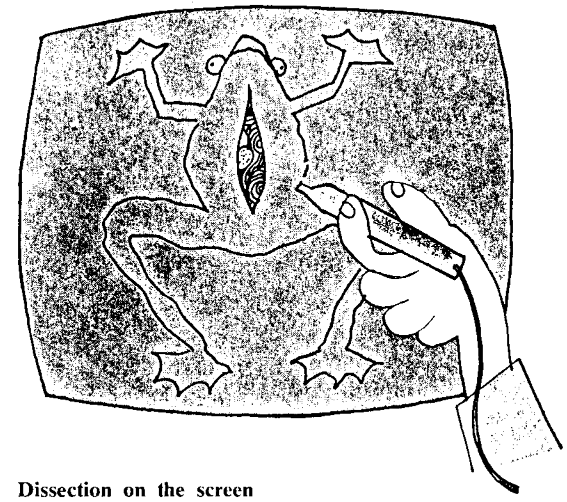

Dissection on the screen

The student of anatomy may use his light-pen as a scalpel for a deceased creature on the screen. As he cuts, the tissue parts. He could also turn the light-pen into hemostat or forceps, and fully dissect the creature —or put it back together again. (This need not be a complex simulation. Many key relationships can be shown by means of fairly simple schematic pictures, needing a data structure not prohibitively complicated.)

Hyper-comics are fun

Hyper-comics are perhaps the simplest and most straightforward hyper-medium. The screen holds a comic strip, but one which branches on the student's request. For instance, different characters could be used to explain things in different ways, with the student able to choose which type of explanation he wanted at a specific time.

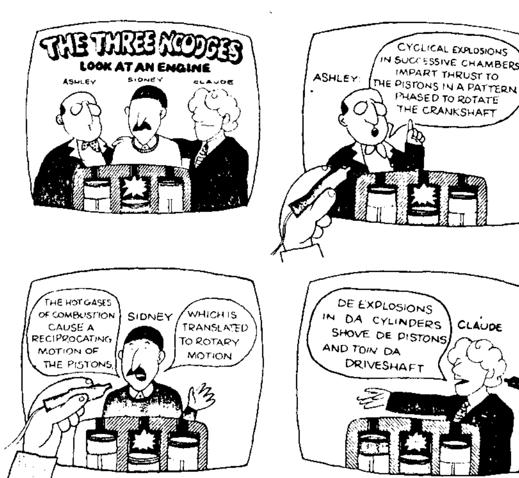

'Technicality' is not necessary

Proponents of CAI want us to believe that scientific teaching requires a certain setup and format, incomprehensible to the layman and to be left to experts. This is simply not true. "Technicality" is a myth. The problem is not one of technical rightness, but what *should* be.

The suggestions that have been given are things that should be; they will be brought about. □

★ THE MIND'S EYE ★

It was explained on the other side that computers have no fixed purpose or style of operation, but can be set in motion on detailed and repetitive tasks in any realm of human interest-- as long as those tasks are exactly specifiable in certain humdrum ways.

Now, if you had a machine like that burning a hole in the corner of your office, what would you really want to do with it?

You can't drive it on the road.

You can't make love to it. (But see p. DM 9.)

You can't cook in it, or get the news on it.

To get it to control elaborate events in the real world requires a lot of expensive equipment and interfaces, so cross that out.

Yet suppose you have an inquiring imagination-- which is not unlikely, considering that you are reading this sentence.

And we are also supposing (from an earlier paragraph) that you have a computer.

What sorts of thing would you do with it?

Things that are imaginative and don't require too much else.

I am hinting at something.

YOU COULD HAVE IT MAKE PICTURES
and show you stuff
and change what it shows
depending on what
you do;

and if this idea doesn't turn you on,
the rest of this book is probably not for you.

The techniques of making pictures by computer are called computer graphics.

But that includes the dull kinds of making pictures by computer, the ones that do it with pens and printing machines.

The techniques of making computers present things interactively on screens is called computer display. (Some say "interactive computer graphics;" this is not just too long, but too restrictive as well: interactive text systems are not "graphic" or pictorial, but they are going to be a profoundly important area of computer display.)

(Incidentally, the silly word "interaction" was coined because the previous word "intercourse," which meant exactly the same thing, had racy connotations for some people. Cf. "donkey" and "rooster," also relatively recent.)

You will note that computer display is what makes possible the computer terminals with screens that we saw on the other side. All that a screen-terminal is is some sort of computer display, to which a keyboard has been added.

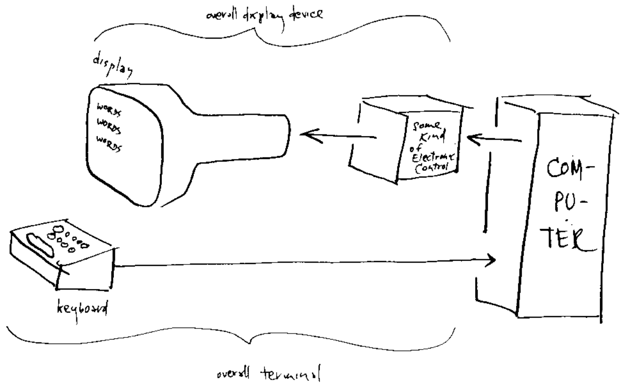

You will therefore see that to understand all the different computer display terminals, you would have to understand all the different computer display techniques; unfortunately we can only cover a few here, and those but sparely.

Some of the types of computer display to be covered hereabouts include:

CRT, or cathode-ray tube, displays; these are my favorite because the stuff on their screens may be animated by the computer.

video displays, which use television techniques. These have troubles deriving from the way a TV picture is timed.

panel displays, i.e., those which appear on a flat panel. These are going to be cropping up all over. (The pictures can't move much, but the devices are going to be cheap. Flat, too. Some people think that's very important.)

3-D displays, especially of the CRT type. NOTE: this term refers ambiguously to two different things: setups which present flat views of three-dimensional scenes, and those which present stereoscopic views of 3-D scenes; these are much rarer.

image synthesis or halftone techniques and systems. These are computer programs and special devices which make shaded or photograph-like pictures. (This happens to be a favorite topic of mine, and so there's quite a bit on it here, a lot of which is not widely known in the field.)

BIBLIOGRAPHY ★ THE MIND'S EYE ★ continues p. 22.

★ Newman and Sproull, Interactive Computer Graphics, McGraw, $15. Your basic text on all forms of computer graphics (and thus animation). ★

Responding computer displays come in all sizes and prices. This little setup (in the under-$10,000 class) is a PDP-8 minicomputer with home-built display circuitry. Gothic lettering data structure available from somebody in the military; message courtesy of R.E.S.I.S.T.O.R.S. The big display is an IBM 2250 (over $100,000, including minicomputer).

GET THE PIC-TURE?

→ Today's zippy picture systems come in many inter-related varieties.

this book presents information on various types; perhaps this diagram will help sort them out.

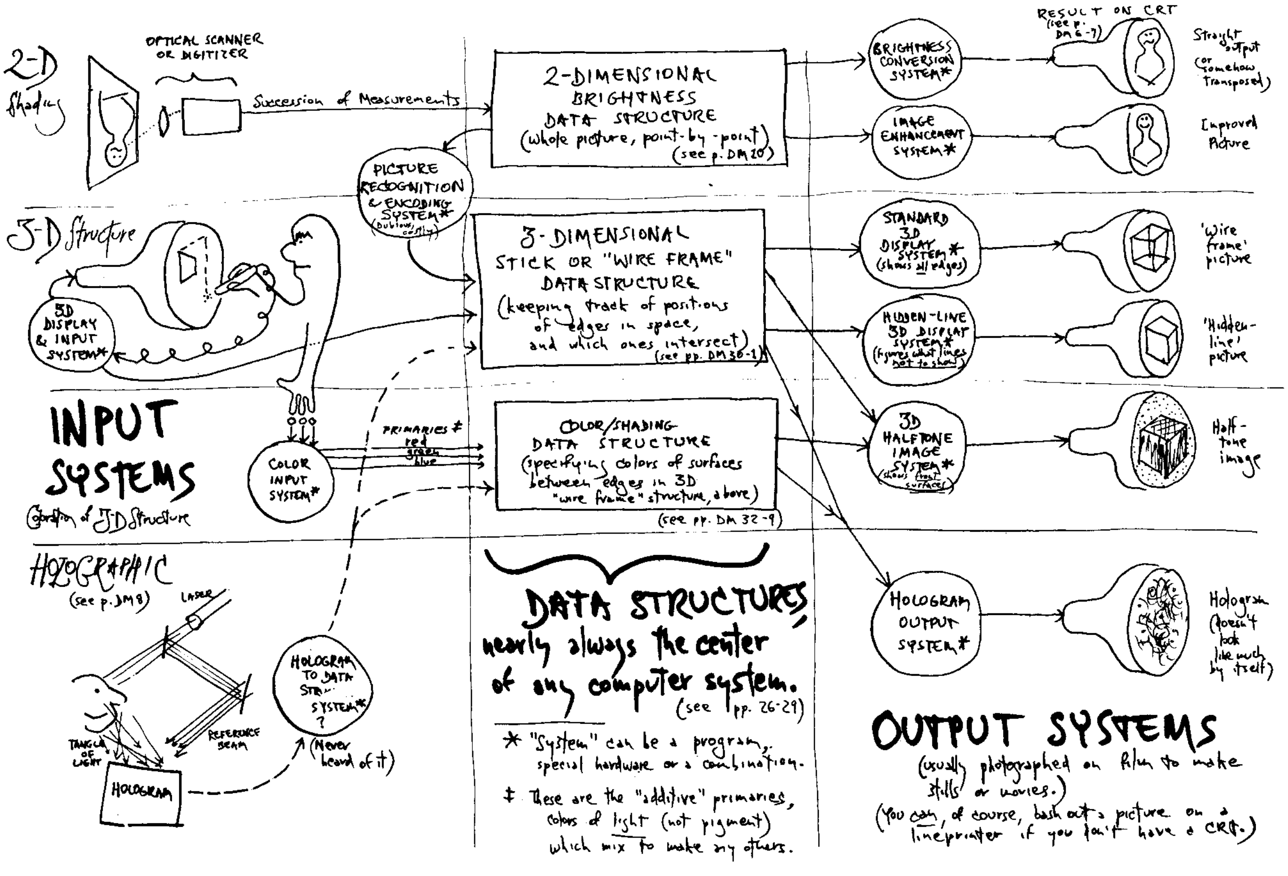

DISPLAY TERMINALS

Some computer displays have to be deeply attached to a computer and some don't. These latter we call display terminals.

A display terminal is like an ordinary computer terminal (see p 14): that is, fundamentally a device by which a computer and a person can type at each other. However, display terminals have screens.

Now, some display terminals only show text, just like ordinary printing terminals (described on the other side). But manufacturers are free to add any other features, and so different manufacturers make it possible to do various kinds of picture-making with their particular display terminals, if appropriate programs are running in the computer that controls them.

Some devices are sold as display terminals but actually, to further confuse the issue, contain complete minicomputers. (The fact that the manufacturer may not stress this is simply a marketing angle he has chosen.) Similarly, certain terminals contain microprocessors (see p. 44), which means they can be programmed to behave like various other terminals, but ordinarily they cannot be programmed to do much else by themselves.

Without getting into it deeply, there are two main types of display terminal: those that are refreshed and those that are not. A refreshed display is one whose viewing surface fades and must be continually re-filled; a non-refreshed display somehow stores the presentation in the viewing surface itself.

Non-refreshed displays simply take the symbols from the computer, blam them onto the screen, and that's it until the screen is erased (by either the computer or the user).

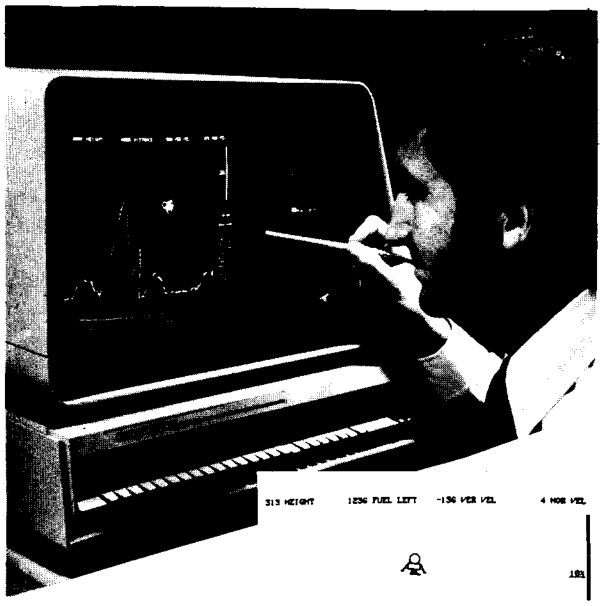

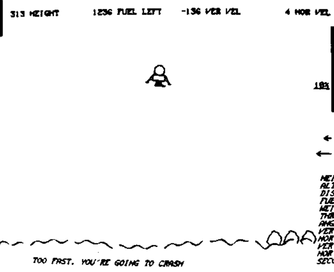

This honey is the GT-40 from DEC ($12,000, including computer-- the thing with teeth, below). It's a subroutining display (see p. DM 23).

Man is playing Moonlander game: controlling screen action with lightpen. Computer simulates real moon lander. Reversed white-to-black for readability here.

THE WONDER OF INTERACTIVE DISPLAY SYSTEMS

If you have not seen interactive computer display, you have not lived.

Except for a few people who can imagine it-- and I'm trying to help you with that as hard as I can-- most people just don't get it till they see it. They can't imagine what it's like to manipulate a picture. To have a diagram respond to you. To change one part of a picture, and watch the rest adapt. These are some of the things that can happen in interactive computer display-- all depending, of course, on the program.

For some reason there are a lot of people who pooh-pooh computer display: they say it's "not necessary," or "not worth it," or that "you can get just as good results other ways."

Personally, I wouldn't thing of trying to justify computer display on "practical" grounds. So what if it offers you faster access to information and pictures and maps and diagrams, the ability to simulate extremely complex things by modifying pictures, the ability to go through complex transactions with the system in very little time, the ability to create things in the world almost instantaneously (say, by creating fabric patterns which are then automatically woven, or design 3D objects which are then automatically milled by machines), and never mind that it enables the user, say, to control entire oil refineries by the flick of a lightpen.

As far as I'm concerned, these matters aren't very important compared to changing the world: making education an excitement, rather than a prison; giving scholars total access to writings and notes, in new complex form; allowing people to play imaginatively, and raising human minds to the potentials they should have reached long ago; and helping people think at the deepest level about very heavy and complex alternatives-- which confront us more ominously today than ever.

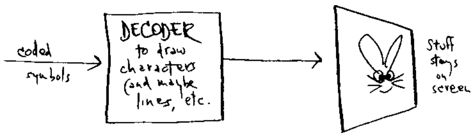

Two major types are the storage tube and the panel. These in turn have separate subtypes, etc.

Refreshed displays have to have some other kind of symbolic (digital) memory, whose contents repeatedly go to the screen:

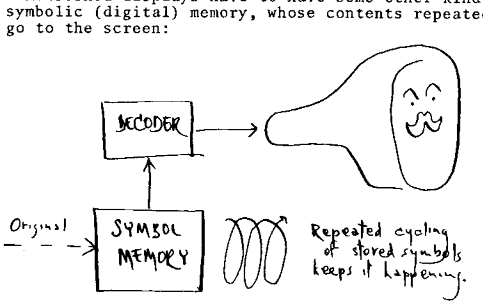

Most refreshed displays use an actual television screen-- that is, a CRT (see p.DM6-7) whose entire area is repeatedly re-painted by the elctron beam.

Since computers send text out to terminals as individual alphabetic and punctuation codes, each terminal must contain circuitry to change the character code to a visible alphabetical character on the screen. Such a piece of circuitry is called a character generator. There are various kinds, they go at various speeds, some offer more different characters than others.

Display terminals generally have a little marker, or cursor, that the user or the computer can move around the screen. The computer can sense what the user is pointing at by the motion codes it gets, telling where the user has moved the cursor.

I had intended here to print a little directory of display terminal manufacturers, but there simply is not time. See section on terminals, other side.

Note that the term video terminal is often used, incorrectly, for any display terminal. The term "video" should only be used when the screen is refreshed by an actual video raster. (See "Lightning in a Bottle," p.DM6-7.)

Text terminals (also called alphabetic terminals, character terminals or keyscopes) simply show written text, put in either by the computer or the user. (Some terminals, called transaction terminals, can be divided up into specific areas that the user may and may not type into-- for banking and stuff. However, whether that form of terminal is necessary may also be a matter of taste in the program design.)

Text terminals range in price from, say, $1500 on up to $6500. (This last is the price of a remarkable color text terminal demonstrated by Tec, Inc., at the 1974 National Computer Conference. Each alphabetic position could contain a letter and/or a bright color; altogether the screen could hold big colorful pictures made up of these bright spaces. Ostensibly just a text terminal, actually the device could be regarded as an Instant Movie Generator for television animation. But it may take Tec, Inc. awhile to realize what they have created.)

Graphic terminals offer some kind of pictures on their screens. These come in a great variety: line-drawing, some without, some with levels of grey. Of interest to the beginner are:

"The Tektronix." (Also called "the greenie," or "the green screen.") Tektronix, Inc., makes a display based on a pale green storage tube they make. (So does Computer Displays, Inc.) Such displays allow you to put more and more text and pictures on a screen, crowding it all up-- but you can't take the lines or words off individually.

"The PEP." Excellent (but very expensive) display that comes out to a video screen from a high-resolution storage tube. Permits grey scales and selective erase. Princeton Electronic Products.

The IDIgraf (Information Displays, Inc., Mount Kisco, NY). Allows line pictures with animation; interesting unit; somewhat less than $10,000.

A PLATO-like terminal (see PLATO terminal, nearby, and pp.DM26-27) is now available for use with STANDARD computer interfaces and software. "Less than $5000" from Applications Group, Inc., P.O. Box 444A, Maumee, Ohio 43537.

REFRESHED HIGH-RESOLUTION COLOR SYSTEMS. A number of companies manufacture computer displays allowing complex grey-scale pictures, including color. They are expensive but very very nice. Indeed, if you buy them in clusters, these fancy-picture scopes can cost as little as text terminals. Some manufacturers are:

- Data Disk. (Disk refresh.) Note: I once recommended them to a consulting client of mine, who later expressed complete satisfaction with their equipment.
- Ramtek. (Semiconductor storage.)
- Adage, Inc. Their model 200 is a video system refreshed from semiconductor storage.
- Comtal. (Disk.)
- Spatial Data Systems. (Disk.)
- Dicomed. (Disk.) Extremely high resolution.

Student programmer Alan McNeil, an art major, ponders something or other. It may be the program for the Nova space-game he and Pete Rowell are building.
Alan also made a film showing what may have been the motions of the continents, shooting straight off the PLATO screen.
Some PLATO purists point out that this is not exactly what PLATO was originally intended for. So?

PLATO panel display (see DM 26-7).

MEN AT WORK

The computer display screen is the new frontier of our lives.

That such systems should (and will) be fun goes without saying. That they will also be a place to work may be less obvious from the tone of this publication, so I want to stress it here.

Making pictures with the GE halftone system (see pp. DM 32-9).

The thing about display screens-- especially the high-performance, subroutining kind-- is that the screen can become a place from which to control events in the outside world.

Example: I believe a town in N.Y. State has its electrical system hooked up to an IDIIOM subroutining display (made by Information Displays, Inc., and coupled to a Varian 620 minicomputer). Instead of having a wall with a big painted map having switches set into it, like many such control centers, the switches are linked directly to the minicomputer, and a program in the minicomputer connects these circuits to the pictures on the screen. Thus to throw a switch in the real world, the operator points with his lightpen at the picture of the switch, and the minicomputer throws the switch.

There are oil refineries that work the same way. The operator can control flows among pipes and tanks by pointing at their pictures, or at symbols connected with them, and bingo, it happens Out There.

In another case, a person designing something at a screen can look across the room and see a machine producing what he just finished designing a few minutes ago. I wish I could say more about that particular setup.

The true problem that I think is emerging, though, is the problem of system response and style. Okay, so you're controlling widget assembly, or traffic light grids, at the CRT screen. The real question is, how does the screen behave and respond? This is not, darn it, a technical issue. It's psychological and then some. The design of screen activities which will enjoyably focus the user's mind on his proper concerns-- no matter how personal these may be-- is the new frontier of design, of art, and of architecture. But more of that later.

Now, the Xerox Corporation has said that they intend to replace paper (or, the way I heard it, "Somebody is going to replace paper with screens, and it will be either IBM or us, so let's have it be us.")

Well and good. Save the trees and stem the grey menace. But the question is: what will the systems be like? How should they perform? What forms will information take? What conventions, structures, diagrams, animations, ways to sign things, ways to view things ... HOW SHALL IT BE?

I am afraid that as long as people are befuddled with technicalities, or confused by those who profess that these considerations are their specialty by right, we will never get straight. Lacking time for the full discussion, I give you a motto:

IF THE BUTTON IS NOT SHAPED LIKE THE THOUGHT, THE THOUGHT WILL END UP SHAPED LIKE THE BUTTON.

SAVING ENERGY WITH COMPUTER DISPLAY

A timely criticism of computer display is that it needs electricity. But (as mentioned elsewhere) it saves paper, and, importantly, it bodes to save energy as well.

IF WE SWITCH TO COMPUTER SCREENS FROM PAPER, PEOPLE WON'T HAVE TO TRAVEL AS MUCH. Instead of commuting to offices in the center of town, people can set up their offices in the suburbs, and share the documentary structure of the work situation through the screens.

This view has been propounded, indeed, by Peter Goldmark, former director of research for CBS Labs, the man who brought you the LP record.

IF COMPUTERS ARE THE WAVE OF THE FUTURE, DISPLAYS ARE THE SURFBOARDS.

★ THE MIND'S EYE ★, Continued.

YOUR BASIC TYPES OF COMPUTER DISPLAY

(Note: the term "display" is also used in this field to refer to numbers and letters that can be made to light up in fixed positions, like on your pocket calculators. Those will not be discussed here. If you're interested see an article on the subject by Alan Sobel, Scientific American, early 1973 sometime.)

THE FORKED LIGHTNING

" Because their words have
forked no lightning they
Do not go gentle into that good night."
-- Dylan Thomas

The most basic, and yet eventually the most versatile, computer display is that of the CRT, or bottled lightning (as I like to call it). It is, you know: a beam of electrons, just like lightning in a storm, but from the neck of a very empty bottle to its flat bottom, whose chemically coated surface we watch. As manipulated by the computer, the CRT stabs its beam to all corners of the faceplate: forked lightning.

Computer display began in the late forties. Computers themselves were completely new, and so was Mr. Dumont's magical Cathode Ray Tube or CRT (see p. DM6), developed on a crash basis during the war so we could have radar, and as long as it was around after the war, we got television.

But the lightning bottle, or CRT, can be used in a variety of ways. Its control plates, which move the ray of electrons around on the screen, can be given various different electronic signals, causing the beam to move around in different patterns. In normal video, the signals move the beam in a zigzag pattern, where the zigs are very close together and the zags are invisible; the carpet of zigs covers the screen over and over in a repetitive pattern, and the beam's changing intensity paints the picture.

But we can drive the CRT differently. by using different control signals. For instance: we can apply a measured voltage to the height or "Y" plates of the CRT, moving the beam to a given vertical position, and another measured voltage to the sideways or "X" plates, controlling its horizontal position.

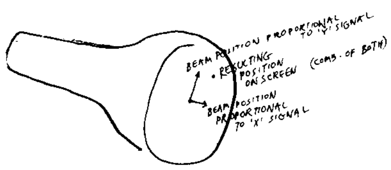

On our glorious Dig-It-All Screens we mingle the magics of air and of fire.

1. EARLIEST SYSTEM: A LITTLE PROGRAM TO MAKE DOTS

The earliest setup connected a CRT to a computer by the simplest possible means, and made its pictures with dots on the screen-- a sort of tattooing process.

It was simple because all the computer did was furnish to the connecting circuitry (or interface) symbols specifying how far up, and how far across the screen, the next dot should be. These symbols were actually coded numbers, and the interface turned them into voltages which then moved the beam correspondingly. (This process of making a measured voltage out of a coded numerical symbol is called digital-to-analog conversion, since (as explained on the other side) the main meaning of "analog" these days is "in a measured voltage.")

Now, this has several drawbacks. One is that the lines are dotty; nobody likes that. A more important annoyance, though, is that the computer scarcely has time for anything else. Here is a flowchart of what the computer has to do in its program. (Even if you didn't look at the other side of the book, flowcharts are nothing scary. They're just maps of what happens.)

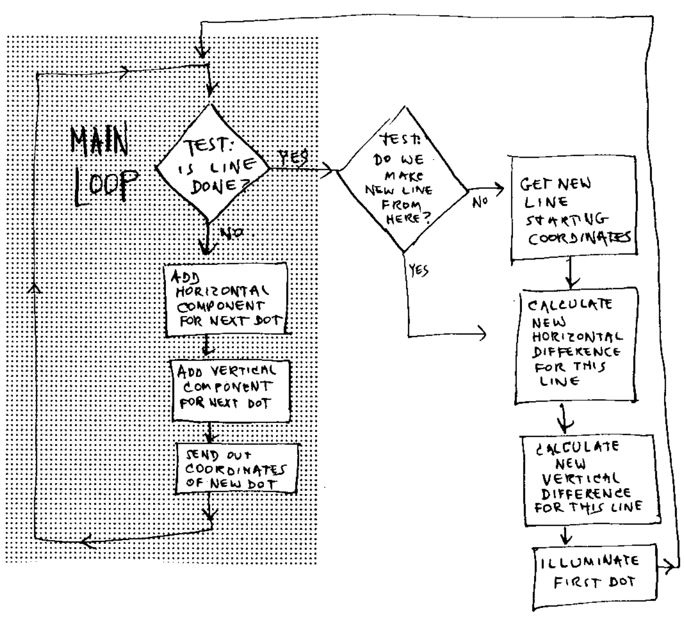

Furthermore, and here was the indignity of it, this system took far too long. To draw a line with thirty dots in it took thirty times around the main loop in the flowchart, and since each box in the flowchart takes at least one of the machine's rock-bottom instructions-- usually more-- then the main loop of this display routine takes four separate operations per dot, or 120 operations for a stupid 30-dot line. Plainly there has to be a better way to use an expensive computer.

Actually it wasn't just the ignominy of it, but the fact that it took so long, that made this a poor method. The amount of stuff the computer could draw in 1/40th of a second-- and this turns out to be how fast the whole picture has to be made-- was too little. After 1/40th of a second the human eye can see the lines on the CRT start to fade, and so the picture has to be redrawn to make it bright again before that happens. If your eye sees the picture fading, then when the computer draws the picture again you will see it get suddenly bright again-- and it will start to flicker. This is distracting, unhealthy, and disagreeable.

Note that the most important computer in the history of computer display used this technique (I think). This was the TX-2 at Lincoln Laboratories, a highly-guarded installation outside Boston which is formally part of MIT. The TX-2 was one of the first transistorized computers-- perhaps the first; and on it were programmed a number of milestone systems, including Sutherland's Sketchpad, Johnson's Sketchpad IV, and Baecker's GENESYS animation system (discussed somewhere).

2. LINE-DRAWING HARDWARE

The next step in design is to get the computer program out of the business of drawing lines by a succession of dots. So we build a piece of hardware that the computer program may simply instruct to draw a line. As an interface, it looks to the computer like four separate devices: registers that tell where on the screen the line must start ("first X" and "first Y") and registers that tell it where to stop ("end X" and "end Y").

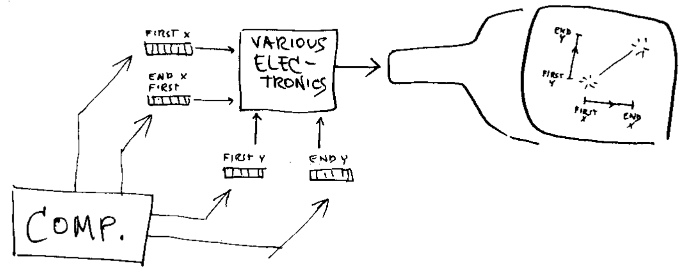

This speeds things up considerably, and allows the computer program to display on the CRT simply by telling the device what lines it wants drawn. Moreover, the program is free to do other things while each line is being drawn, though this involves the problem of how the program is to know when it's time to send out another line-- and we needn't go into that here.

(Incidentally, it is a puzzling fact that such a device is available nowhere, although lots of people end up building one for themselves. There was such a thing on the market a couple of years ago-- line-drawing hardware with no interface and no CRT-- but it was withdrawn because of reliability problems. A just price, if anybody wants to go into that, would be five hundred to a thousand dolars-- this year.)

3. EVOLUTION FROM THIS: TWO OPTIONS

There are basically two ways to go from this basic starting point. Either we can keep the display device intimately and integrally connected to the computer, or we can say the hell with it and cut the display device loose as a separate entity.

Ivan Sutherland has cannily noted that there is a certain trap involved in these designs: as we build additional "independent" structures to take the burden of display away from the computer, we are tempted to keep adding features which make the "independent" structure a computer in its own right. This paradoxical temptation Sutherland calls "the great wheel of Karma" of computer display architecture.

It is tempting to cut the display loose from the computer. It means the computer can be fully occupied with other matters than refreshing the screen-- preparing the next displays, perhaps. Many computer people believe this is the right way to do it, and it is certainly one valid approach. But unfortunately it also drastically reduces the immediacy of the system's reaction, making interaction with the system less intimate and wonderful.

Approaches which put display refreshment and maintenance in a separate device are less interesting to me, and so that discussion continues separately nearby. ("Display Terminals," p. DM21).

This article continues next page.)

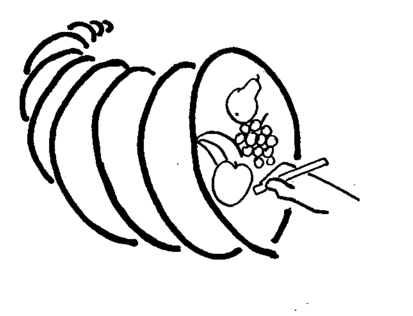

4. THE SECOND PROGRAM FOLLOWER

On the other side of the book, I explained that a computer is basically a zippy device, never mind how constructed, which follows a program somehow stored symbolically in a core memory. Such a device we call here a program follower. While programs may be in many computer languages-- all of them contrived systems for expressing the user's wishes, in different styles and with different general intent-- underneath they all translate to an inner language of binary patterns, which may just be thought of as patterns of X and O, or light bulbs on and off. The innermost program follower of the computer goes down lists of binary patterns stored in the core memory, and carries them out as specific instructions. It also changes its sequences of operations under conditions that the programmer has told it to watch for.

The most powerful and responsive computer displays are those which build a second program follower which goes down lists of picture-drawing instructions also stored in the same core memory.

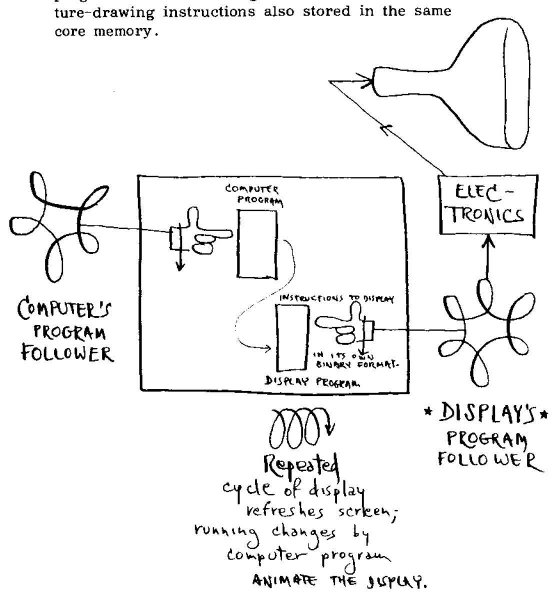

We may call this also a "list-of-lines" system, since the commands recognized by the display program follower are typically patterns that tell it what lines to draw.

Typically also it has its own way of jumping around in a program, and may jump to a specific list of lines, or subpicture, from numerous other parts of its program, always returning each time to the point from which it had jumped. This allows the same subpicture to appear in numerous places on the screen at the same time. (A program that can be jumped to by other programs which then resume operation is called a subroutine; thus the real, or most prestigious, name for such a device is a subroutining display.)

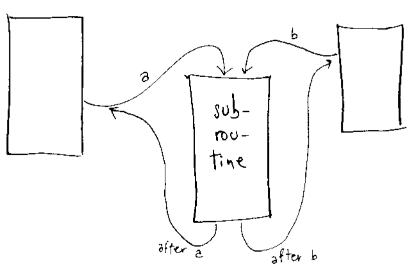

This design has some extraordinary advantages. One is that since the computer's program follower and the display's program follower both share the same core memory, they can work together most intimately. When the user demands something new-- by typing, say, or pointing with a light-pen-- the computer can step in and take various actions. Its program can compose a new picture for the user, get something from a disk or tape memory, or switch the display's program follower over to a new picture it has already prepared.

Most importantly, the computer can move images on the screen, allowing interactive animation on the screen under the user's control. Each time the display is about to show the same picture again, the computer simply supplies it with a new starting point. Since the list of lines is typically in the form of sequences of lines relative to one another, the picture is drawn in a new place each time-- and thus seen to move on the screen.

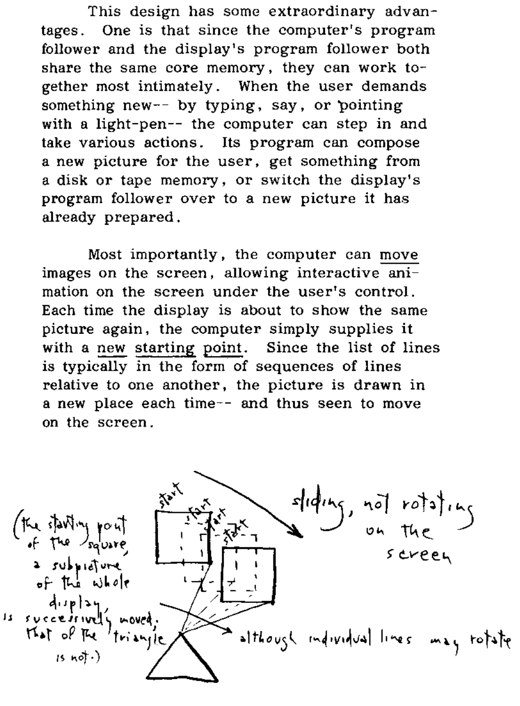

This design has some extraordinary advantages. One is that since the computer's program follower and the display's program follower both share the same core memory, they can work together most intimately. When the user demands something new-- by typing, say, or pointing with a light-pen-- the computer can step in and take various actions. Its program can compose a new picture for the user, get something from a disk or tape memory, or switch the display's program follower over to a new picture it has already prepared.

Most importantly, the computer can move images on the scren, allowing interactive animation on the screen under the user's control. Each time the display starts to show the same picture again, the computer simply supplies it with a new starting point. Since the list of lines is typically in the form of sequences of lines relative to one another, the picture is drawn in a new place each time-- and thus seen to move on the screen.

Finally, the computer itself is free most of the time-- free, that is, to do other things, which typically is always desirable. Just how much the computer can or should do in such a partnership is a matter of dispute. (Ordinarily such devices are spliced onto minicomputers; and minicomputer fans, such as the author, see no reason not to perform all services for the display there in the minicomputer-- and a pox on the big machines. Others, for various reasons, see the subroutining display and its host mini as needing the tender ministratious of a big-computer via some sort of communications line. There are various reasons for holding this entirely legitimate view. People who are devoted to the high number-crunching capacity of big computers, or to languages which require great big computers to run in, have a right to their opinion. Moreover, it is currently feasible to store large bodies of data only on big computers -- not because big disk and tape memories can't be easily attached to the small ones, for they can, but they usually aren't; and other ways to tie minicomputers to big stores of data aren't available yet.)

Subroutining displays often have commands allowing them to display text as well as lines and dots. In the display of text they can use the same technique of "moving the picture" by starting its display at successively creeping points; this will cause, say, whole paragraphs to slide on the screen. The importance of this feature in the displaying of text cannot be overemphasized. As more and more people have experience with displays of different kinds, they are beginning to realize how confusing and disorienting it is for a screen to clear and be filled with something new to read. You don't know where you are. On subroutining displays, moving the text can give the reader the same sense of orientation he gets from turning pages -- an important thing to replace.

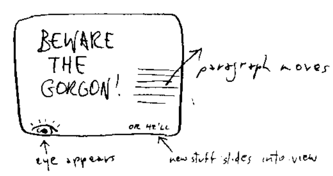

It must be stressed here that, just as computers themselves have no fixed mode or style of operation, neither do computer displays; and so the purpose of such devices is simply

HELPING PEOPLE SEE AND MANIPULATE
PICTURES AND TEXT
IN ANY STYLE, AND FOR ANY PURPOSE.

Since pictures can be of anything, and text can be about anything, this effectively comprehends the entire mental and working life of mankind.

Many readers will scoff, supposing that computer display systems will always cost tons of money. This is not the case. You can already get a beauty, with its minicomputer, for as little as $13,000; and this price should fall to three or four thousand within a few years-- as soon as the minicomputer manufacturers realize that the market frontier is not in the office or factory, but in the home. But we're getting a bit ahead of ourselves here.

TYPES OF SUBROUTINING DISPLAY

Some early subroutining displays used a screen-dotting technique, but took the burden of it off the computer itself: it would extract from core memory the instructions telling it to draw individual lines and show text. (I refer here to the DEC model 338, introduced about 1965; this attached to a PDP-8 computer (see p. 40) and cost about $50,000 including the computer.) Others drew lines as straight zips of light across the screen; an example is the IBM 2250 display, introduced about 1966. (The model 1 of this device buckled directly to the 360, and cost, I believe, something like $75,000; its successor, the model 4, buckled to their 1130 minicomputer, the package costing some $150,000, and then you were supposed to attach it to an IBM 360.) The 2250 was a good machine, but in performance suffers greatly from the restrictions of the 360 computer itself (see p. 41).

These earlier machines are being replaced by new versions with better-designed instructions (see "Computer Architecture," p. 32, for a sense of what well-designed instructions are). An especially fine unit is DEC's GT40, which buckles on the exceptionally fine PDP-11 minicomputer (see p. 42). The GT40 is illustrated nearby. (p. DM21) It goes for some $12,000 including the computer. (That's today, we repeat. Consider not the price at this instant, but how fast it's going down.)

The units mentioned above are of the most basic type: "two-dimensional," whose pictures at any given instant correspond to flat drawings -- but, of course, derive their excitement and magnificence from their capacity to interact, change and animate what you are looking at.

THE MIND'S EYE continues on p. DM30.

Sutherland's SKETCHPAD

Seldom has an event in a new field had as much power and influence as what dour Ivan Sutherland did as a young man in the period 1960-64.

The SKETCHPAD system, which was basically his thesis work at MIT, was at once inventive, profound, overwhelmingly impressive to laymen, and deeply elegant. Simply for the universal influence it has had in the computer field, it deserves our close attention.

Sutherland was one of the first people to understand the use of the computer in helping people visualize things that weren't fully clear yet-- the opposite, of course, of the conventional notion of computers. While computers had been made to do animations as early as the forties, and computer graphics had been put to work-aday duties in the old SAGE system (defending us against bombers in the fifties-- remember the good old days?), Sutherland turned computer display from an expensive curiosity into a true dream machine.

SKETCHPAD ran on the 36-bit TX-2, a one-of-a-kind experimental machine at Lincoln Laboratories (a military research place nominally a part of MIT). It had a display screen, light pen and lots of handy switches.

SKETCHPAD was basically a drawing system. But rather than simulating paper (as some people might have done), it found splendid ways to take advantage of the computer's special capabilities.

In the Sketchpad system, Sutherland looked for ways that a responding computer display screen could help people design things. He pioneered methods of drawing on screens, with such techniques as the "rubber-band line" (a straight line on the screen, one end of which follows your lightpen while the other remains fixed), and the "instance"-- a subpicture stored in core memory which could appear numerous times and ways in a larger picture).

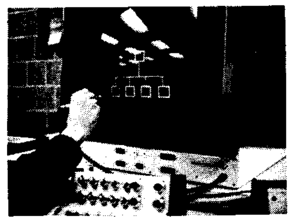

This picture vaguely simulates the "instance" facility of Sketchpad, by which an overall picture may be created out of repetitions of a single master pattern.

Simulated with GRASS language (see p.31).

The mind-blowing thing about Sketchpad was the way you could move and manipulate the picture on the screen, with all its parts. One overall picture could be constructed out of a hundred copies of a basic picture; then a change in the basic picture would immediately be shown in all hundred places. Or you could expand your picture until it was effectively the size of a football field (with you looking at a tiny view in the handkerchief-sized screen). Or you could draw meshing gears on the screen, and with the light-pen (and through the "constraint" facility) make one gear turn by turning the other!

This elegant technique, the constraint, does not seem to have been imitated even now. A "constraint" was a restriction placed on some part of the overall stored picture complex. The user could move or manipulate various parts of the picture on the screen, but the parts that had constraints could only move in certain directions, or according to certain formulas, or dragging other parts along, etc., as specified.

This was a profound idea, because it meant that any rules for the manipulation of particular objects on the screen could be added to Sketchpad as particulars within the larger program, rather than having to be programmed in from scratch.

(One extremely interesting aspect of Sutherland's thesis, which most people seem to have missed, dealt with displaying a structure of constraints: that is, showing what elements depended on what other elements, in a highly abstracted diagram that the system could show you. This form of display has remarkable possibilities.

After his brilliant SKETCHPAD work, Sutherland was made head of ARPA's computer branch (see "Military," p. 58). There he was involved in many of the computer funding decisions of the late sixties, which contributed to the impetus of computer display. (His predecessor, Licklider, had been a pioneer in time-sharing, and much of the forward movement in the computer field in recent years may just have had to do with the strategic position of those two men when they were at ARPA/IPT.)

Sketchpad went on as a continuing research tradition at Lincoln Labs. Timothy Johnson, for instance, made a version of it that allowed the drawing of three-dimensional objects; this became the forerunner of the various three-dimensional line systems described hereabouts.

From ARPA, Sutherland went on to the University of Utah, whence he slipped off with the Computer Science department chairman to found the Evans and Sutherland Computer Company, makers of the top-of-the-line computer display systems (see p. DM30 and p. DM33).

Sutherland's work has shown an elegance and inventiveness outstanding in the field. (For instance, I believe one issue of Communications of the ACM had two unusual articles by him: one describing an eccentric "Chinese auction" system worked out for scheduling use of a computer, which benefited users more than any previous method; and the infamous "Great Wheel of Karma" article, where he compared the design of graphical computers to the Hindu system of reincarnation-- if you keep adding desirable features to the design, soon you have another program follower and another computer in the same box-- over and over.)

COMPUTER MOVIES

How do computers make movies?

Well, first of all, computers do not make movies unless thoroughly provoked.

In fact, only people make movies. But computers, if sufficiently provoked, will do a lot of it: enact the movie and photograph it, frame by frame.

There is no single method.

All forms of computer display and computer graphics may be used to make computer movies.

"Computer animation" is any method of making movies in which a computer successively draws or paints the successive individual frames, which may be done by any of the methods mentioned in this book. Now, since there are numerous methods of making pictures by computer, then any method of making different individual pictures, in a succession of changing frames, is computer animation. So a "computer movie" is any film made by, or with the picture-making aid of, computers.

In other words-- it's no one thing.

Now, there already exist hundreds, if not thousands, of computer movies. So far most of them have been on technical topics-- the mechanics of satellite orbit stabilization, the mechanics of explosions and so on.

Here are a few stills from some other movies, more humanistic.

BIBLIOGRAPHY

Newman & Sproull, Interactive Computer Graphics. McGraw, $15.

This is the textbook. Anyone interested in computer display should get this immediately.

An expensive journal, Computer Graphics and Image Processing, comes from Academic Press.

Sherwood Anderson, "Computer Animation: A Survey." Journal of Micrographics, Sep 71, 13-20. Lists nineteen computer-animation languages of that time.

Ken Knowlton, "Computer-Made Films," Filmmakers Newsletter Dec 70, 14-20.

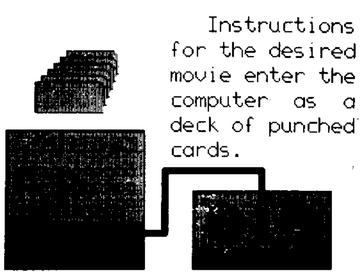

Vintage Knowlton, using BEFLIX. (This language used the COM quite efficiently: dots were actually out-of-focus letters.)

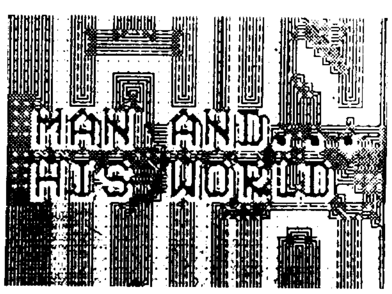

Vanderbeek & Knowlton (using TARPS, which shows strong influence of BEFLIX, which it grows from).

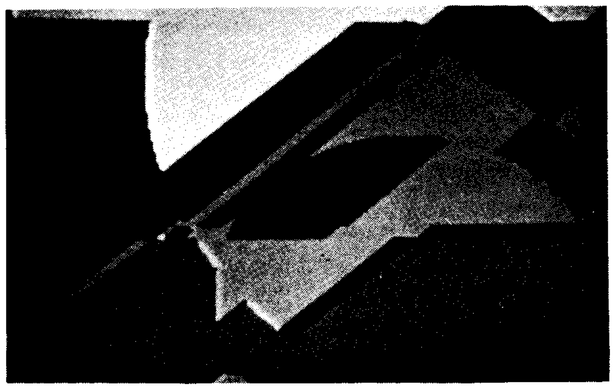

Lillian Schwartz

LILLIAN SCHWARTZ

A talented artist with a feel for technology, Ms. Schwartz has been working for several years with Knowlton and others at Bell Labs. Her films with Knowlton, mentioned elsewhere, are marvelous. She now works at a more permanent setup, a minicomputer that runs successive images on a color TV screen, employing a modified form of Knowlton's EXPLOR language. The work is immediately viewable. This allows rapid film construction, not previously possible when the work had to go through a slow animation camera before she could see the result.

For Knowlton-&-Schwartz films contact: Martin Duffy, AT&T, 195 Broadway, NY NY.

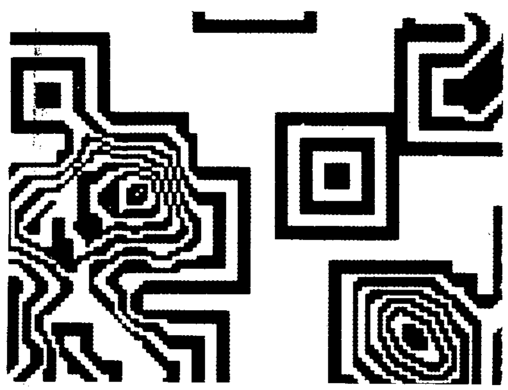

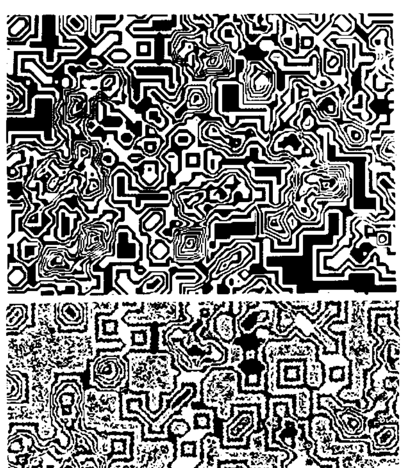

Schwartz & Knowlton. Using the EXPLOR language, they make pictures and patterns scintillate and grow together. (EXPLOR in some ways generalizes Conway's Game of Life; see p. 48 and p. DM 26.)

JOHN WHITNEY

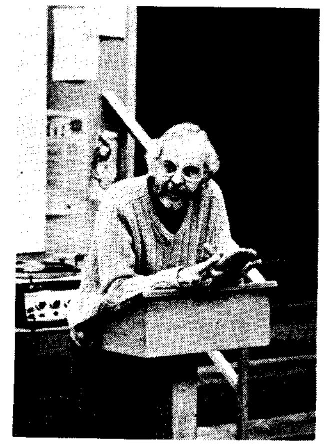

JOHN WHITNEY

John Whitney is the ancestor of us all, probably the first computer movie-maker. He is also a gripping speaker.

In the forties, he built a special animation stand-- using analog computers.

Deeply concerned with music, Whitney has in his images emphasized rhythmic and contrapuntal movement of shapes and lines.

Whitney films available from: Pyramid Films, Box 1048, Santa Monica CA 90406.

John Whitney

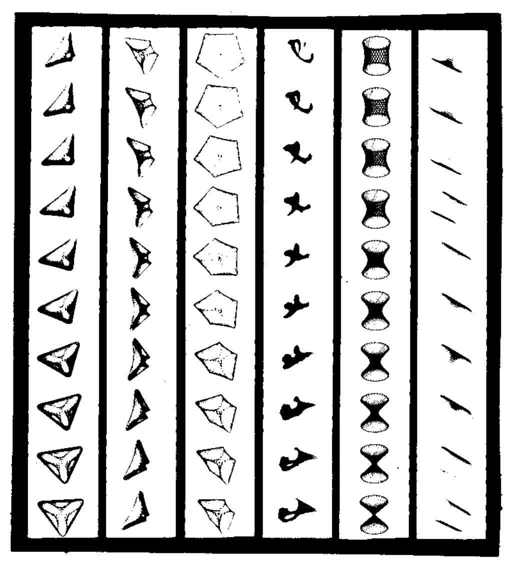

John Whitney

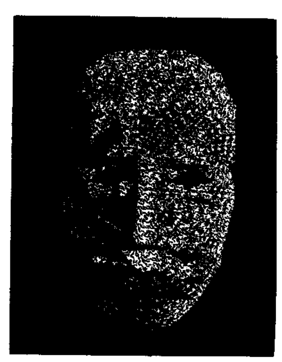

Lillian Schwartz (with Henry Magnuski; see p. DM 34)

RON BAECKER'S GENESYS

By now there are dozens of computer animation languages-- perhaps hundreds. Each one employs the techniques of animation which its developer wanted to use, tied together in the ways that seemed appropriate to him. (See "Computer Languages," p. 15, and note Knowlton's various animation languages, described nearby.)

One of the more influential animation systems has been Ron Baecker's GENESYS, a 2-dimensional animation system programmed in the late sixties at MIT's high-security Lincoln Laboratory. (It used the TX-2 computer, mentioned elsewhere in this book.)

Baecker, a cheery and genial fellow, expressed interest as a student in using the TX-2 for animation, and was allowed to. The system he produced has a number of lessons for us all.

GENESYS is a "Good-Guy" system, as discussed on p. 13 Meaning, in this case, that it is easy to learn and simple to use. As argued elsewhere in this book, making computer systems clear and simple is often hard for the programmer (and may go against his grain), but is essential.

PICTURES AND MOTIONS

GENESYS makes the following simplifications of your movie: all images are made up of dots. They do not change as you watch; animation consists of the images either moving or being replaced.

To create an image, you draw it onto the screen with a lightpen or a tablet. (As in the SKETCHPAD system; see p. DM 23.) Parts of the image may be changed until you're satisfied.

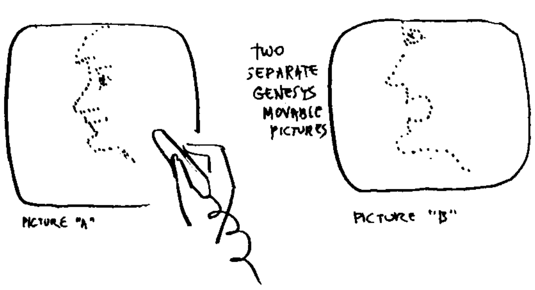

Now, to create the animation, you do the same thing. Each image can be made to move on the screen; and the path of the motion may be drawn on the screen, through the picture area. Not only that, but the timing of the motion is controlled through the same diagram, by the spacing of the dots. (Baecker calls his control diagrams p-curves.)

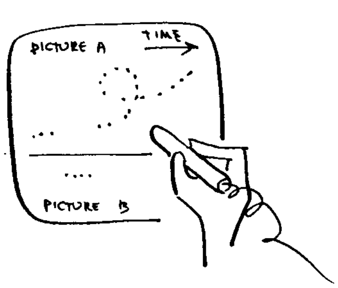

Lastly, sections of picture may be replaced by means of the control diagram (as indicated in picture above).

Having created such an animated sequence, which is stored in symbolic form in the computer ("digitally"), you can view it on the screen, decide what you do and don't like about it, and change any part of it.

The basic elegance of the system is this: Baecker made everything work the same way, through control by screen diagrams. He simplified the animation problem in a clear and simple way.

Ron now teaches in Canada and is into working with PDP-11s. The results should be fun.

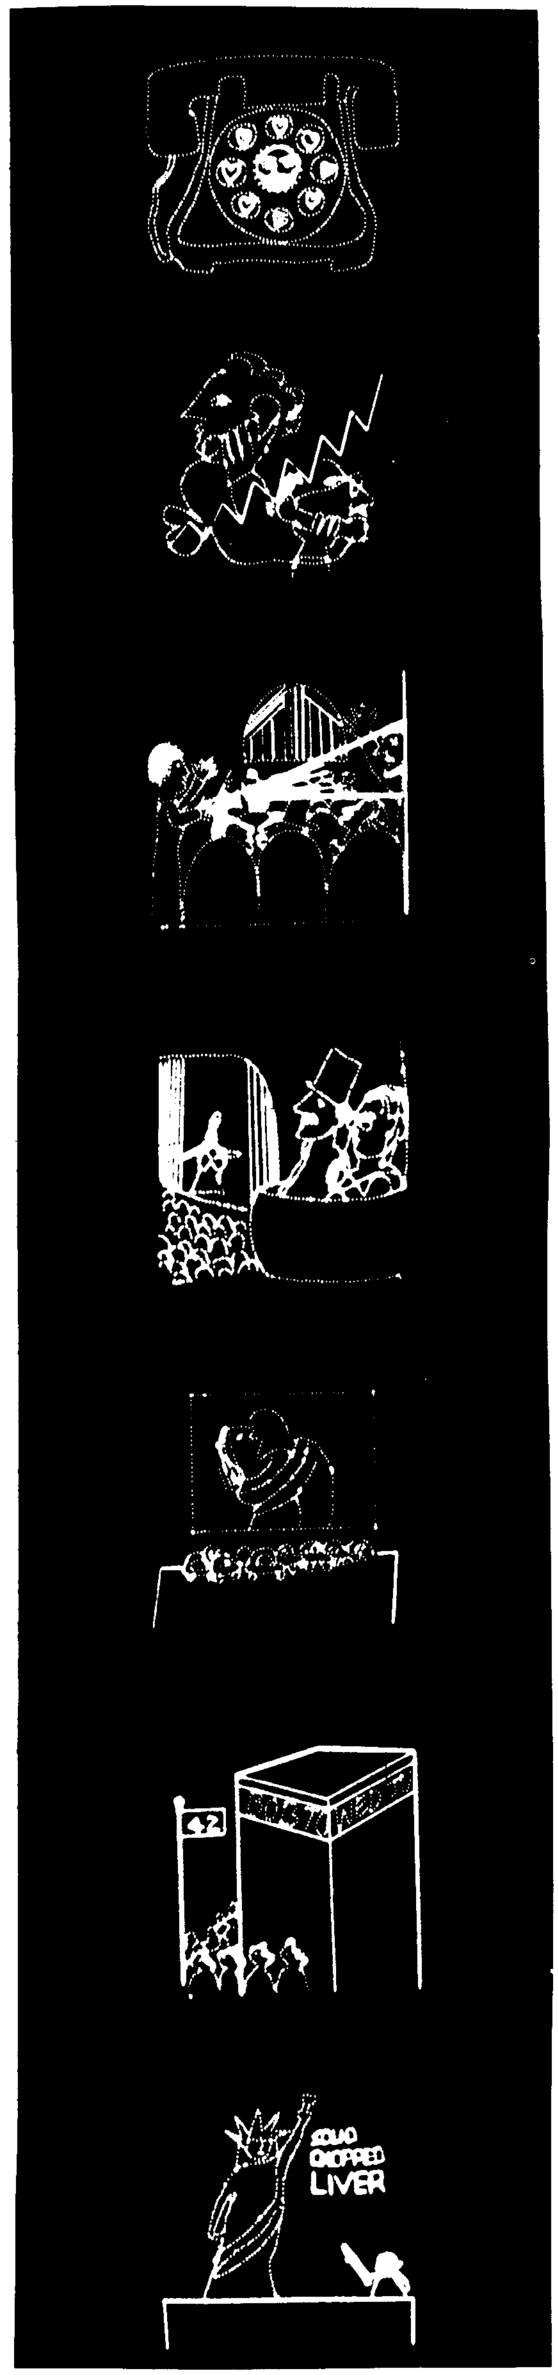

LYNN SMITH

Lynn Smith is a young Boston artist who has worked extensively with Baecker's GENESYS (see nearby). One result has been a movie which should be an example to us all: "The Wedding Movie for Bob and Judy." (Her Friends Bob and Judy were getting married, so she made this movie, a few minutes long and quite clever, to celebrate it.)

This is my favorite example of how computers should be used in the human world; it says more on the subject than any dozen articles.

(One question that remains unanswered is how a system like GENESYS could have been used for such a purpose, seeing that most people in the field believe GENESYS only runs on the heavily-guarded TX-2 computer. Regretfully, I can shed no light on this here.)

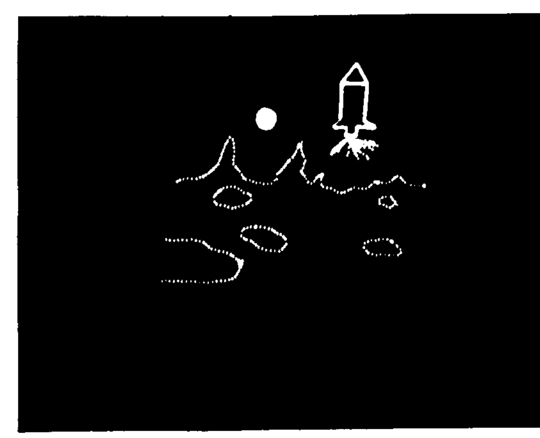

COMs (Computer Output Microfilm devices)

are what you use to make computer movies. Basically they consist of a CRT and a movie camera in a box.

Mostly they are used to put text on microfilm by computer, so generally they are not connected to a computer but run off magnetic tape.

This turns out to be very annoying if you want to hook up the computer directly to the COM, and make movies that fill the frames spot-by-spot. For that you really need your own movie camera and a minicomputer. (Movie cameras that can be made to start and stop by computer are called "pulse cameras" or "instrumentation cameras.") The society for people who make Movies by Computer is called UAIDE (Users of Automatic Information Display Equipment-- an obsolete title). It used to be a club just for companies that owned COMs made by Stromberg Datagraphix, but evidently it has now cut itself loose and become a subsidiary of the National Microfilm Association, 8728 Colesville Road, Silver Spring MD 20910.

(NOTE: for them as want to make color movies, the two alternatives have been either to have separate primary negatives combined at a lab-- the "old Technicolor" process-- or to add a complicated color-filter box to a COM or other CRT setup. Such things are available commercially now, from Dicomed-- a whole Color COM.)

BIBLIOGRAPHY

Computer Output Microfilm. $10 from National Microfilm Assn., above. Lists available COMs and service centers.

DM26

THE
NEW REPUBLIC
OF

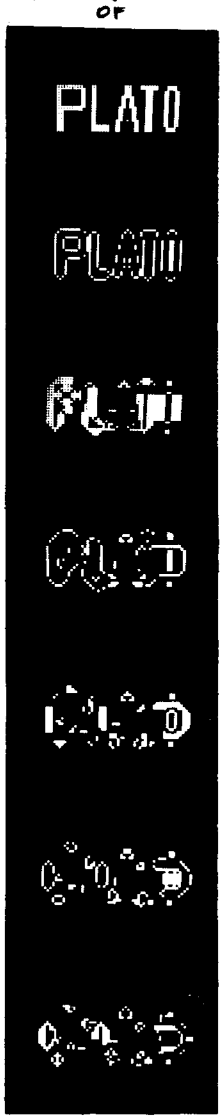

(Above: PLATO LIVES! Anyway, the word itself goes through changes in the Game of Life (see p. 48), as programmed for the PLATO system by Larry Cleator, and photographed from a PLATO screen.)

PLATO is the world's greatest computer display system.*

Some 500 users, at terminals around the world (but mostly in Illinois), simultaneously tie up to a big computer in Urbana, Illinois and savor instantaneous pictorial and text deliveries on their bright orange screens. Diagrams, explanations, tests and even animation of a sort, flow almost without interruption to the bright orange screens all over. The system is extremely responsive: depending on what the user is up to, its various programs can respond to each pressing of a key, usually within a fraction of a second.

While literature on PLATO is copious, it is hard to read and slightly sales-oriented. But a few minutes' intercourse with a PLATO terminal makes anyone an enthusiast for the system.

PLATO is the brainchild of Don Bitzer, a U. of Illinois engineer who has devoted over a decade to its creation. Michael Scriven, no slouch himself, has called Bitzer "one of the great men of our time." Bitzer is also certainly one of the world's greatest salesmen. A crew-cut, huggy-bear sort of a fellow, he flies around the world demonstrating, lugging a great terminal along. When you sign on the system you may be informed that Bitzer is at that very moment demonstrating in Paris or Tokyo. This "travelling dog and pony show," as PLATO staffers call it, has created awe and excitement wherever it goes, and where the awe has been strong enough to generate money, there you will now find PLATO terminals.

If you have a PLATO terminal-- (you presumably being a school or other favored institution-- you can in principle log onto PLATO from anywhere in the world, though most terminals stay in one place. There is one main network, consisting of a big Control Data computer in Urbana (the model 6800; see p. 41) with tendrils extending out into the phone system and the educational TV cable of the state of Illinois. When the Urbana system is "finished" and fully loaded, it will have 1008 terminals; all are already spoken for. The PLATO terminal is a totally unique animal (see box), manufactured (all too slowly) by Magnavox, incorporating a terrific plasma panel built by Corning. (The plasma panel was invented by Bitzer, and even though much of PLATO was publicly funded, he is reputedly rich from it. We said he was a great politician.)

* *In terms of high performance for lots of users. Various systems (described hereabouts) offer more power, but at huge cost.*

As a first taste of interaction on a graphical computer system, PLATO can be a thrilling mind-opener-- especially to people who think computers can only behave loutishly or through printout.

PLATO is a complete stand-alone system, with its own monitor program or "operating system" (see p. 45) running on the CDC 6800 computer all by itself. It does not run on any other manufacturer's computers, or simultaneously with any other big programs. It communicates only with PLATO terminals, no other, and PLATO terminals, because of their unusual design, can communicate only with it, partly because of its unusual design and partly because of its unique 20-bit interface. (See diagram of PLATO terminal, box nearby.)

A PLATO terminal costs about $4000 and the price seems to be going up; $5000 in the next few years is a popular estimate. But you can't just buy one. You have to get on the waiting list, and who are you, anyway? There was a time when almost anybody could buy into PLATO, but now that the system is unstoppable, applicants are being scrutinized.

Is it really unstoppable? Educational Testing Service, of Princeton, is conducting an elaborate Effectiveness Evaluation of the PLATO system, presumably to decide whether it should live or die (on public funds). But with so many terminals in the field already and so many man-years already gone into its creation and the making of materials for it (-- the ghastly term "authoring" seems likely to stick), it is hard to believe PLATO could die. Not now.

Especially considering that two more systems are now being put together: at Lowry Air Force Base (Colorado) and Florida State University. That means there will be whole other computers of the CDC 6000 series running the PLATO Monitor and shepherding PLATO materials to users at PLATO terminals, unconnected to Urbana, one for Lowry AFB and one in Florida.

And it won't end there.

Control Data, whose vested interest in the system (though they didn't pay for its creation) is enormous, is said to be projecting

ONE MILLION PLATO TERMINALS BY 1980.

Another sign in the wind: Montgomery Ward has one.

Now, to call the PLATO system a "computer graphics" system may seem somewhat odd to people who know it in another guise, as a system for Computer-Assisted Instruction (called CAI). But as the author does not like CAI in general, at least as it's been going-- see p. 15-19 -- and rather likes PLATO, I prefer to describe it as I prefer to see it.

Nevertheless, to understand PLATO properly we had better consider what the people have been doing in terms of what they think they have been doing, and offer any amendments or restatements later.

"OPTIMIZED FOR CAI"

PLATO stands for "Programmed Logic for Automated Teaching Operations," and has been billed (and sold) as a system for automated instruction.

It is, most PLATO fanciers will tell you, "optimized for instruction." ("Optimized," in computer talk, means "just what somebody says you need for a specific purpose.") As with any system, the leaps of faith between its basic design premises have become lit by airport beacons; clearminded individuals with alternate views have difficulty making themselves understood to some PLATO enthusiasts. But the most basic underlying feature of the system, INSTANT RESPONSE, cannot be quarreled with. PLATO can respond, as already mentioned, to a single key-pressing by a user, almost instantly; this feature is virtually impossible, say, on IBM systems (but see box p. 55). This responsiveness is the system's greatest beauty.

Because of the need for high responsiveness, it was decided that all users had to have their particular programs ("lessons") running in core at the same time. That meant there would be no swapping (bringing in materials from disk memory), which can bring mortifying delays (if a lot of people need it at once); but it also meant lessons have to be very small. Large bodies of material, which would have to be moved in from disk, are not allowed; thus each lesson is basically a little love-nest that must generate its own action. Hence there is an emphasis on little programs to respond various ways, rather than text which may be read in quantity.

Partly because large amounts of text cannot be shipped to the user, a little PROJECTOR is in the terminal. It uses a tiny microfiche, or microfilm sheet, small enough to fit in the palm of your hand.

If a PLATO author deems it necessary, he requires for his lesson, not just the use of the keyboard and plasma screen, but a microfiche as well. The student must put the microfiche in place when he starts the lesson; signals from Urbana (or wherever) then jump the projected image among 256 different images, in response to what the student does.

Now, PLATO people are not doctrinaire about how their system is to be used. The plasma screen can be continuously showing little decorations along with the teaching material. The microfiche could be showing irrelevant works of art or travel scenes. These are all facilities at the option of the PLATO author; at his beck and call, if he thinks his program or lesson needs them. (But it's very bothersome to have the microfiche made-- an important difficulty.)

Every terminal has the screen, the keyboard, and the projector. Other options may be added, however:

1. The touch panel. This is a transparent window that goes over the plasma screen and reports to the main computer whether it has been touched, and where. (This allows illiterates, especially kiddies, to use the system without typing.)
2. The audio disk. This allows the terminal to respond with sound, including canned words, to the student. (It does not actually synthesize the sound, as discussed on p. DM11.)
3. The general jack. Not to be confused with Pershing, this is a connector socket that will send and receive data from any other device-- provided you've got the right interface. This allows all kinds of other devices, such as piano keyboards, to be used for student input. Or output (like gum-ball machines.)

Actually, except for the restriction on quantities of material that can reach the student, PLATO is an extremely general system. Despite the strange convention of calling all user programs "lessons"; despite the odd stipulation that all users are called either "students" or"authors"; and despite being told by PLATO spokesmen that PLATO is not a general-purpose system; actually, it is.

Amongst the terminals, PLATO room, Circle Campus. What one person is doing ordinarily has no bearing on the others, who could as well be in Timbuktu as far as the main computer is concerned.

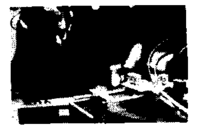

PLATO's audio device permits the system to respond to the user with a spoken phrase, snatch of music, or whatever -- in a fraction of a second. The magnetic disk is forever turning; compressed air shoots the read-head to the required track on the disk for the reply.

The hardware was designed by Bitzer. The software-- that is, the underlying computer part, never mind the contents to be shown (also regrettably called "software" by many hands)-- was initially less stressed by Bitzer, but eventually grew under the direction of others. In particular, an ex-biologist named Paul Tenczar (pron. "Tenzar") created its underlying TUTOR language. (For an introduction to computer languages see p. 15 and what comes after.) The TUTOR language exists only on PLATO; and PLATO authors may only use the TUTOR language, Paul Tenczar's creation.

The TUTOR language can best be understood as a reaction to Coursewriter, another CAI language offered by IBM on its 1500 Instructional System. Coursewriter's original intent was to enable non-computer people, especially teachers, to create drill-and-practice instructional lessons roughly of the type

Now, Johnny, what is 3 X 5? *(particulars filled in; name filled in)*

Kid's answer {10 / 15}   No. / Good! } *System replies.*

Obviously, by changing the numbers and pushing the kid on types of problems he hasn't mastered, the computer can patiently bring students to mastery of various simple skills, diagnosing weaknesses and stressing the individual student's problems. The difficulty is that attempting to extend this method out of the very simple has great pitfalls and may not even be worthwhile (see pp. DM 15-19).

Anyway, Coursewriter was promulgated by IBM with the 1500 and thus suffered premature standardization before things had been thought out. IBM is not to blame for Coursewriter's deficiencies, they were just trying to make a buck; but because a lot of scared people believed Coursewriter was the way it had to be, the evolutionary improvement usual for computer languages didn't have time to occur. An egregious omission: Coursewriter did not allow the author much access to the computer itself. That is, programs written for numerical calculation, say, could not be brought into instructional materials at a sophisticated level.

Tenczar's TUTOR changed all that. It has both the virtues and defects of being original. Apparently unlettered in computermen's controversies and dogma, Tenczar designed a language of great power and speed; it is utterly strange to computer people, offers various brilliant features, and is in some respects quite irritating. It looks very simple to the user-- but beyond a few deceptively simple techniques, it has to be learned in considerable detail to do anything interesting. (See box, "Is It Better to Toot?", next page.)

This tale has, of course, been simplified. Bitzer and Tenczar did not work alone, but rather were leaders in a seething community of dozens of smart people working like blazes on the project. It has taken some fifteen years of Bitzer's effort, and tens of millions of dollars, to get the system where it is now-- Ready and Working.

Project PLATO now extends far beyond its original domain. Originally a fairly tight nucleus at the Computer-Based Education Research Laboratory ("CERL") at the U. of Illinois in Urbana, the community of PLATO now sprawls out through its lines to a larger constituency, the PLATO community of users.

(Indeed, this extended Republic of PLATO-- the systems people (see p. 45) in Urbana, the authors and locals-in-charge throughout the network-- constitute one of the maddest rookeries of computer freaks in the world. Where else would you find a 14-year-old systems programmer who's had his job for two years? Where else would you see people fall in love over the Talkamatic (a PLATO program which allows you to have written conversations with people at other terminals, wherever they may be) only to clash when at last they meet in person? Where else can you play so many different games with faraway strangers? (See Box.) Where else can students anywhere in the network sign into hundreds of different lessons in different subjects (most of them incomplete)? Where else are people working on various different programs for elementary statistics, all to be offered on the same system?)

PLATO is one of the wonders of the world.

Mike O'Brien, a Tolkien fancier, has put the entire Elvish alphabet onto PLATO as a special character-set. Here the system gives a famous warning to turn back, both in English and Elvish. Mike says it intimidates snoopers poking around his material.

Unfortunately, there are so many learners, and so few PLATO terminals, that use of the terminals must now be fairly strictly controlled. (The eight terminals at the University of Illinois at Chicago Circle, at which most of these pictures were taken, generally work an eight-hour day.) The time was when people could just walk in, sit down at a terminal and do what they liked; now, sadly, each user must have an "account" and a password.

But the rabble is howling at the gates. Many professors want to use it to take rote aspects of teaching off their backs; and the computer hums and students want to play the PLATO games (see box) and tinker with an interactive system of its power and lusciousness. But most of them will have to wait.

PLATO's services are "free," for now. That is, if your school has PLATO terminals, and IF it will pay for the communications lines, THEN the services of the central computer are "free"-- the National Science Foundation is bankrolling its operation for a couple of years more. Then, bango, PLATO central service becomes something that has to be paid for too.

Just to give you an idea, the communication costs to Urbana for Circle Campus's eight terminals run at over $10 ,000 a year. But these costs should be coming down sharply; it is the price of tooling up for whatever the PLATO future is going to be. Anyway, the general cost of the system comes out to about $1.50 an hour, the same as general time-sharing on a PDP-10 (see p. 42). But that's without paying for the central computer-- another cost which we expect to go down, however.

This is all a far cry, of course, from Bitzer's claim a decade ago that PLATO terminals would cost only $400. But considering the system's success, we needn't dwell on that.

Perhaps the real question is this: with man-machine intercourse of this quality now possible, can people's love for the system stay Platonic?

PLATO GAMES

Moonwar on a Saturday in Urbana. It's man-to-man among the craters; then a quick kill of the unknown adversary.

And our doughty warrior looks to the Big Board for more challengers. Kids love PLATO games.

They work hard and they play hard on the mighty PLATO system.

When the Author gets tired of Authoring, or the Student of Stewing, just around the corner, a few keystrokes away, are diversions and games to boggle the imagination.

You can go to a program ("lesson rose") and look at "the great roses"-- elaborate curlicues generated by mathematical patterns that appealed to the authors of that program; or find, also tucked in rose, Conway's Game of Life (see writeup, p. 46 , and picture series, nearby).

Then there are games you can play against the system, like racetrack and blackjack. (These games let you win astronomical sums of money-- play money, forgotten when you sign off.) Remember, of course, that you're not really playing against a computer but against a specific program, with its quirks and shortcuts and blind spots.

Then there are games you play by yourself-- actually responding resources (see pp. DM 18-19), which entice you into trying things out. Tenczar himself has created two elegant, gem-like lessons, man and picto, which teach you computer programming without ever saying so. These two programs present the user with a little picture of a man on the screen, and show him how the little man may be moved around and made to pick up pictures of balls. From there on the student may have his way-- and is never told that he's learning to program a true computer language. (Though it is a quite restricted one, dealing exclusively with little men and their excursions among balls and falling sticks.)

Another charming game, I don't know by whom, is called candy factory. Here too the user may control the animation of the picture by what he types. Machines are seen to manufacture candy, box it and ship it-- depending on what buttons you press.

Some games are played between people who sit together before a single PLATO terminal, often with teaching intent. Such games include the hop game, where Bunny (you) and Frog (your friend) add their way along a board with numbered squares. Older children can dig How the West Was (1+2)X 3, which involves grouping the numbers you get by chance to try to get ahead of the other stagecoach.

THE "BIG BOARD" GAMES

Still another category of games, though, awaits the adult who craves real excitement. Because PLATO has so many terminals, all over, there is a curious combination of anonymity and intimacy between users (-- much like the curious Nonexistent Phone Numbers of Paris; in the French phone system, people calling the same nonexistent phone number can talk to each other; strange blindfolded encounters occur at the Number of The Day, spread by word-of-mouth; sometimes these result in people really getting together...)...

Anyway, the Big Board games of PLATO have exactly that: a shared list, or "Big Board," showing who is playing the specific game.

But you don't have to use your right name.

In this jaunty society of shadows, you pick your own nom de guerre, or fighting name. This has numerous advantages: the most obvious is that as you improve at play, you can shed the identity in which you have been humiliated.

The main games with Big Boards are that old standby, spacewar (rocketships wheeling and firing at each other and sliding around on the screen); dogfight (biplanes wheeling and firing at each other and sliding around on the screen), moonwar (shooting at the other guy by specified angles as you stand among craters). In addition, PLATO offers (not during working hours) what must be two of the most baroque space-war games anywhere, empire (eight races (the Vulcanians, Klingons, etc.) seek to control the galaxy) and nova (simulated navigation among millions of different stars and solar systems, all of which may be revisited, all of which are different...)

Welcome to the Hop Game. PLATO often uses animated opening titles.

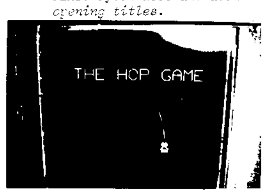

Here it is Bunny's turn. Screen instructs you personally: "Press -NEXT- to spin, ermintrude".

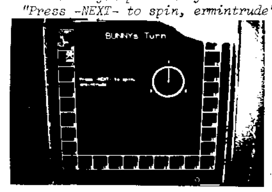

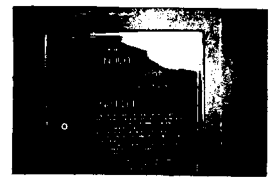

The Navigation part of nova is already working. To get around you need instruction; here we are at the Training Center.

View from your Nova spaceship includes perspective view of where you are among billions of stars; and your various controls.

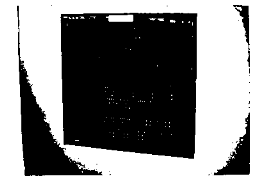

People who only play PLATO games occasionally have to sign on by typing their names into the big board. (They often get slaughtered by the regulars). The regulars-- hah. When they're signed into the system, they have merely to jump to a specific game for their fightin' names to be posted on the big board. A mighty rollcall they make, too-- such great warriors as von Dave, zot, fright pilot, AL 9000, simpson, doc, THE RED BARON, The Red Sweater, The Giant Pud, Fodzilla, tigress, enema salad, Conan, Siddhartha, wonder pig!!!!!, and EXORCIST.

(As those insiders who have automatic sign-on to Big Boards write programs to do the sign-on, their arrival in a Big Board game is often an animated sign-on. The cutest trick is THE RED BARON's: it looks like this.

THE RED BARON ← (plane falling in flames)

It works like this. For dogfight, the terminal already has stored in its temporary memory, as "characters," the little pictures of airplanes that are going to buzz around the screen. So the Baron just follows his name with the code for that special character.)

One last point. No longer can you sign on with an obscenity: a little obscenity-checking program looks for the usual expletives, in case visitors or other priggish folk might be looking. But of course this is easy to circumvent by putting periods between the letters of your nasty word, or something similarly deceptive to the poor program.

THE KEYBOARD AND WHAT IT SORT OF MEANS

THE STRUCTURE OF PLATO-SPACE

The PLATO keyboard. What looks odd and arbitrary to you is believed by devout Platonists to be divinely ordained.

PLATO IV- STANDARD KEYBOARD

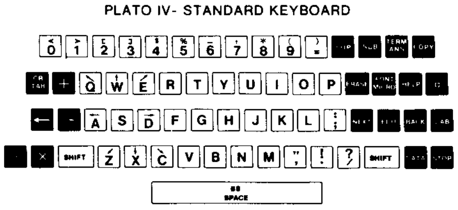

TO MOVE BETWEEN LESSONS, the basic action is to hold down SHIFT and press STOP. (For further complications see Ins-And-Outs diagram.)

TO MOVE WITHIN A LESSON, basic actions are NEXT (to go forward or tell the system it's its turn; BACK, which sometimes returns you to earlier points in the sequence of your lesson; and six step-out-of-line options, by which the author may permit the user to sidestep to explanations, enrichment material, or things out of sequence.

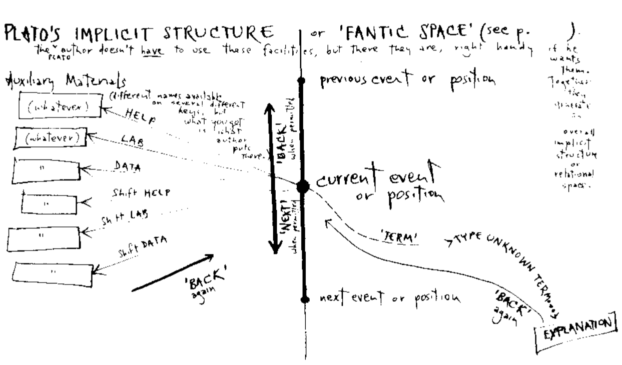

The original idea was evidently that there would be a basic sequence, in which NEXT and BACK would be the forward and back controls, and the other six would represent Help for the Confused, a "Lab" allowing experiments, and additional Data the student decides he needs. The three with Shifts simply provided a second option of each type.

How the author might use these, however, was his own affair.

"TERM" evidently was for when students wanted things Looked Up: by pressing TERM and typing the unknown word, the student would get a definition. "ANS" suggested that it might also be used when the student was allowed the option of being told the answer.

Note the arrows over Q,W,E,A,D,Z,X,C. They allow the student to move cursors, draw, point directions, etc. Unfortunate confusion ensues with the left-arrow on the far left, used in programming (as in APL; see p. 22-3.)

ERASE allows the student to correct his input; COPY helps edit and change things. SUP and SUB allow superscripts and subscripts; FONT MICRO is like a special shift key, going into whatever special font is currently stored on the terminal. I have no inkling of what the little square means.

IS IT BETTER TO TOOT?

"A tutor who tooted the flute
Tried to tutor two tutors to toot.
But he asked through his snoot:
is it better to toot
Or to tutor two tutors to toot?"

Folk thing

The TUTOR language grew out of drill-and-practice, for which it has a command specifying where a student's answer is to appear on the screen. This is the "arrow" command. The language has a strange scanning structure built around this "arrow" command, much as the TRAC Language (see pp. 18-21) has a scanning structure built around parentheses and commas. Beginners don't need to understand the scan and the arrow command, but journeymen do.

TENCZAR'S CONCEPT OF A CONCEPT

Much has been made of TUTOR's facility for "analyzing the content" of what students type in. Actually, of course, the computer does not "understand" what the student says (see "Artificial Intelligence," pp. DM 12-14), but rather offers certain efficient tricks to the person using TUTOR to prepare presentational materials.

Basically, TUTOR's "concept" facility reduces every input word to a 60-bit code. The technique of reduction (called a "hashing function") supposedly substitutes for any word of any language a code of 60 bits (see "Binary Patterns," p. 33), which means the program in TUTOR can rapidly test a student's input for numerous different possible things. (The power of this technique will be readily recognized by computer people; unfortunately there is no room to explain it further here.)

Thus a TUTOR program may contain "concept searches" that test whether a student types either a desired response or numerous alternatives. While it may be strange to call this a "concept," it is a powerful technique.

Paul Tenczar's TUTOR language, the programming language inside PLATO, is like any other programming language (see pp. 15-31); intricate, and unlike its results. That is, a program bears no more resemblance to what it does than the word "cow" looks like a cow.

PLATO is a system for canned presentations that respond to the student. Students need not know TUTOR. Anyone out to prepare such presentations must learn it, however; and the attempt has discouraged many.

Tenczar is a former biologist, and had no preconceptions from computer orthodoxy to bind him in the design of TUTOR. Thus the language is very original. There is only room to raise the following points:

To learn the first steps in TUTOR-- how to set up drill-and-practice lessons, for instance-- is unusually easy.

To do anything complex, however, requires you to learn the bulk of the TUTOR language. Thus when people say TUTOR is "easy," they mean those first steps.

TUTOR is not Extensible, like, say, TRAC Language (see pp. 18-19) or GRASS (see p. DM 31). That is, a programmer cannot customize the language with new compound functions of his own making. Steps are being taken to correct this; meanwhile, it is said that the Urbana people can be persuaded to put in new commands others want for, e.g., chocolate chip cookies.

THAT EXTRAORDINARY TERMINAL —

THE CAVE OF PLATO.

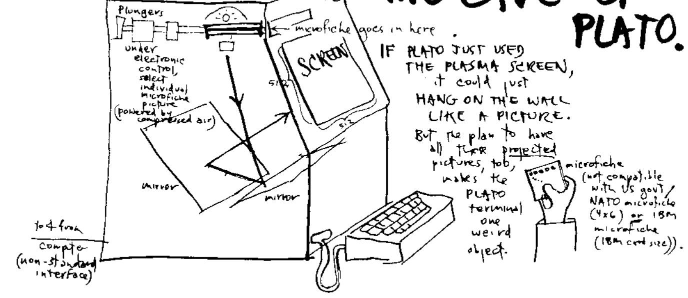

You can read the standard-size lettering off the screen at SIX FEET-- even though it's NO BIGGER THAN PICA TYPE. Fantastic.

The internal circuitry that draws on the screen is highly capable. Receiving a 20-bit code, the terminal itself deciphers it as--

A LINE ON THE SCREEN, or
TWO STANDARD CHARACTERS ON THE SCREEN from its FIXED character memory, or
TWO SPECIAL CHARACTERS ON THE SCREEN from its CHANGEABLE character memory (which can be loaded with Russian, Armenian, katakana, Cherokee or whatever-- even little pictures-- at the start of the lesson), or
A COMMAND TO THE MICROFICHE PROJECTOR, or
A COMMAND TO THE AUDIO PLAYER, or
A COMMAND TO WHATEVER'S IN THE GENERAL JACK.

Note that all lines and characters for the plasma screen can be turned on (orange on black) or off (black on orange).

PLATO'S HANDY KEYBOARD is on a flexible cable, can be worn in your lap.

AUTHOR'S PLATO-SPACE (as of early 1973)

Getting around in This system might be easier. (See "Fantics," p. DM 48-51.)

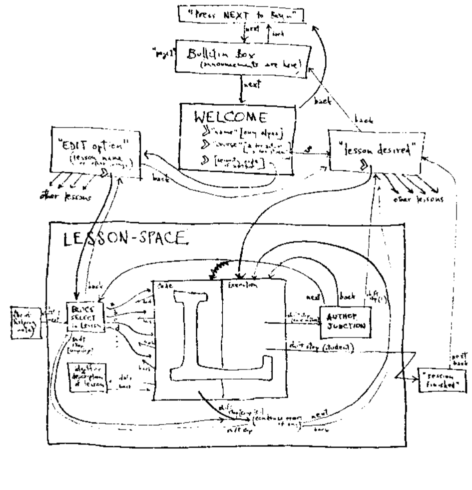

WHAT THIS IS. I briefly visited Alfred Bork's CAI shop at the University of California at Irvine on a consulting basis. Bork is a really swell guy, but he's devoted to Dialogue CAI-- that is, to teaching programs that have pseudo-conversations with the student. (As I've said variously already, the pseudo-conversation parts are not only expensive and difficult, but sometimes irritating and objectionable; and happier, zippier, simpler techniques are available using various techniques of old-fashioned showmanship-- as from movie-making, writing and (here) the comic book.

This is my reply to Bork's question, "Well, how would you do it?"

This ties into Bork's physics display system. That is, it's intended to be a front-end program (see p. 13) on a Tektronix graphics terminal (see p. DM7 and DM 20-23), leading into a simulation program (see p. 58) allowing the user to see all kinds of motions in physical law. The program it's intended to supplant uses dialogue.

WHAT IT CONTAINS: introductory remarks; statement that physical law (as of motions) simply summarizes constant covariances. Sorry if readability is poor (Xerox of a Xerox).

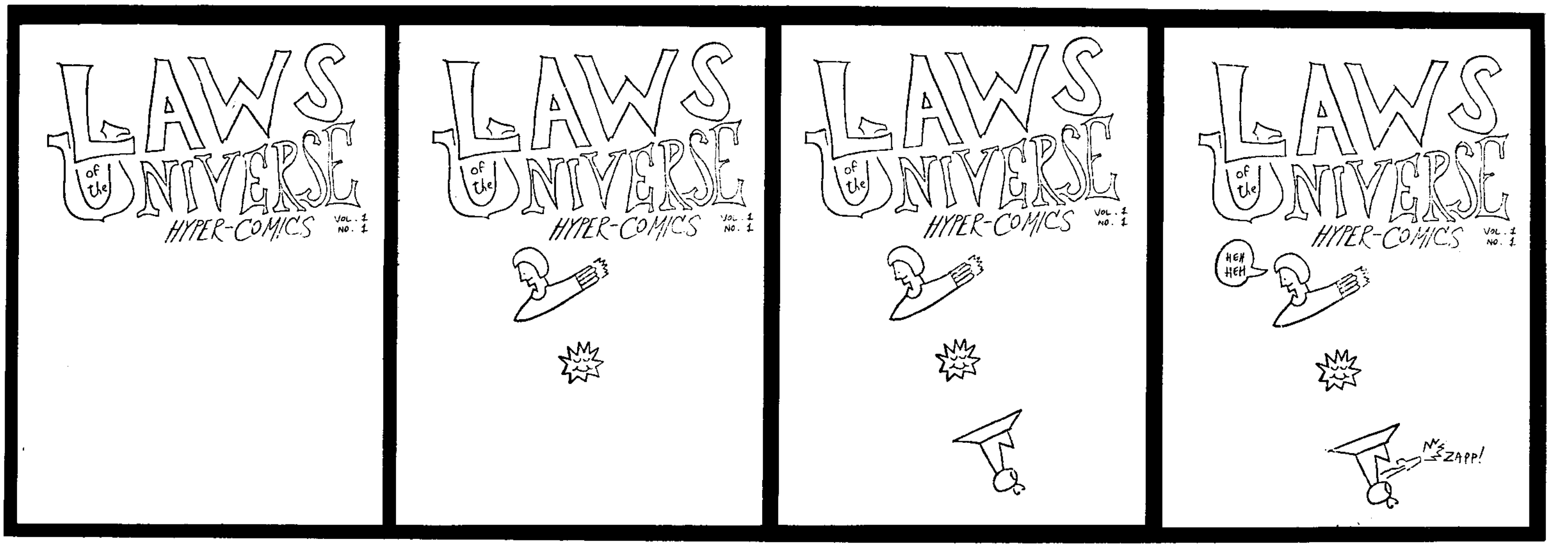

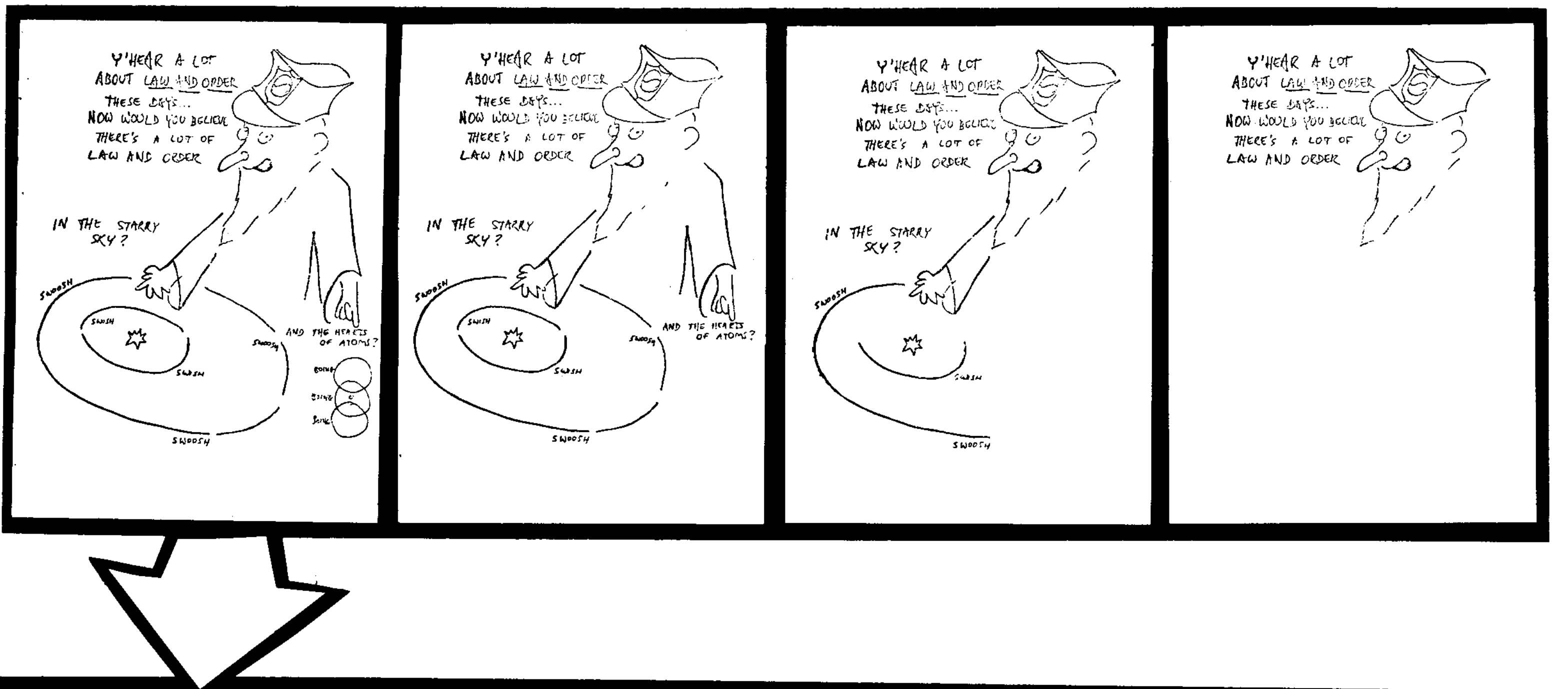

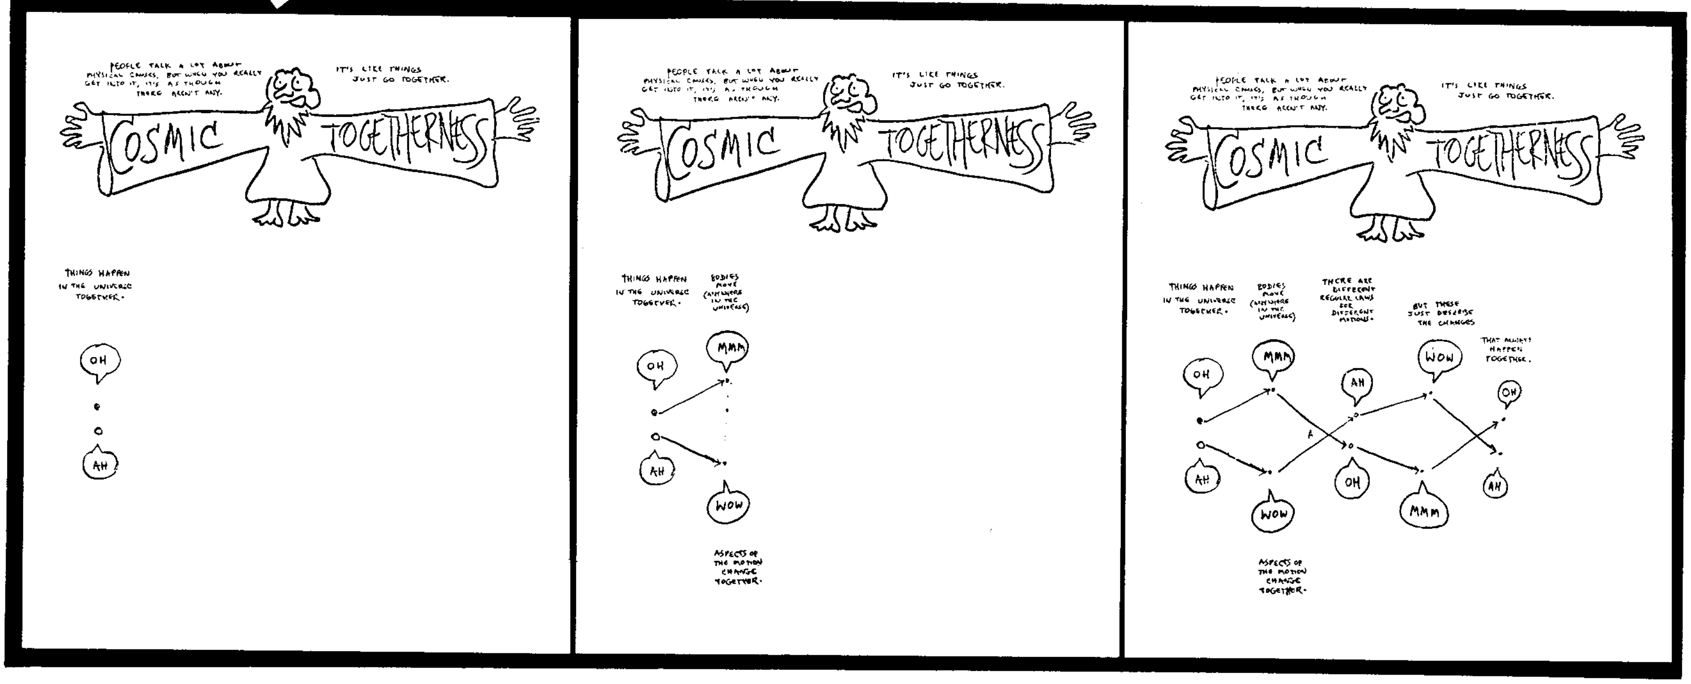

You will note the artistic problem of composing cumulative animation for a display screen.

Some people have accused me of trying to be humorous. Obviously nothing of the sort was intended. Research supported by NSF grant no. GJ296 (but "Mr. Natural" character property of Robert Crumb).

•Homage to Robert Crumb.

BIBLIOGRAPHY: for comic technique, study the works of Crumb; also, comicbook stands are currently featuring reprint magazines of THE SPIRIT, which is some of the finest stuff ever done. Also study Wally Wood in the early MADs.

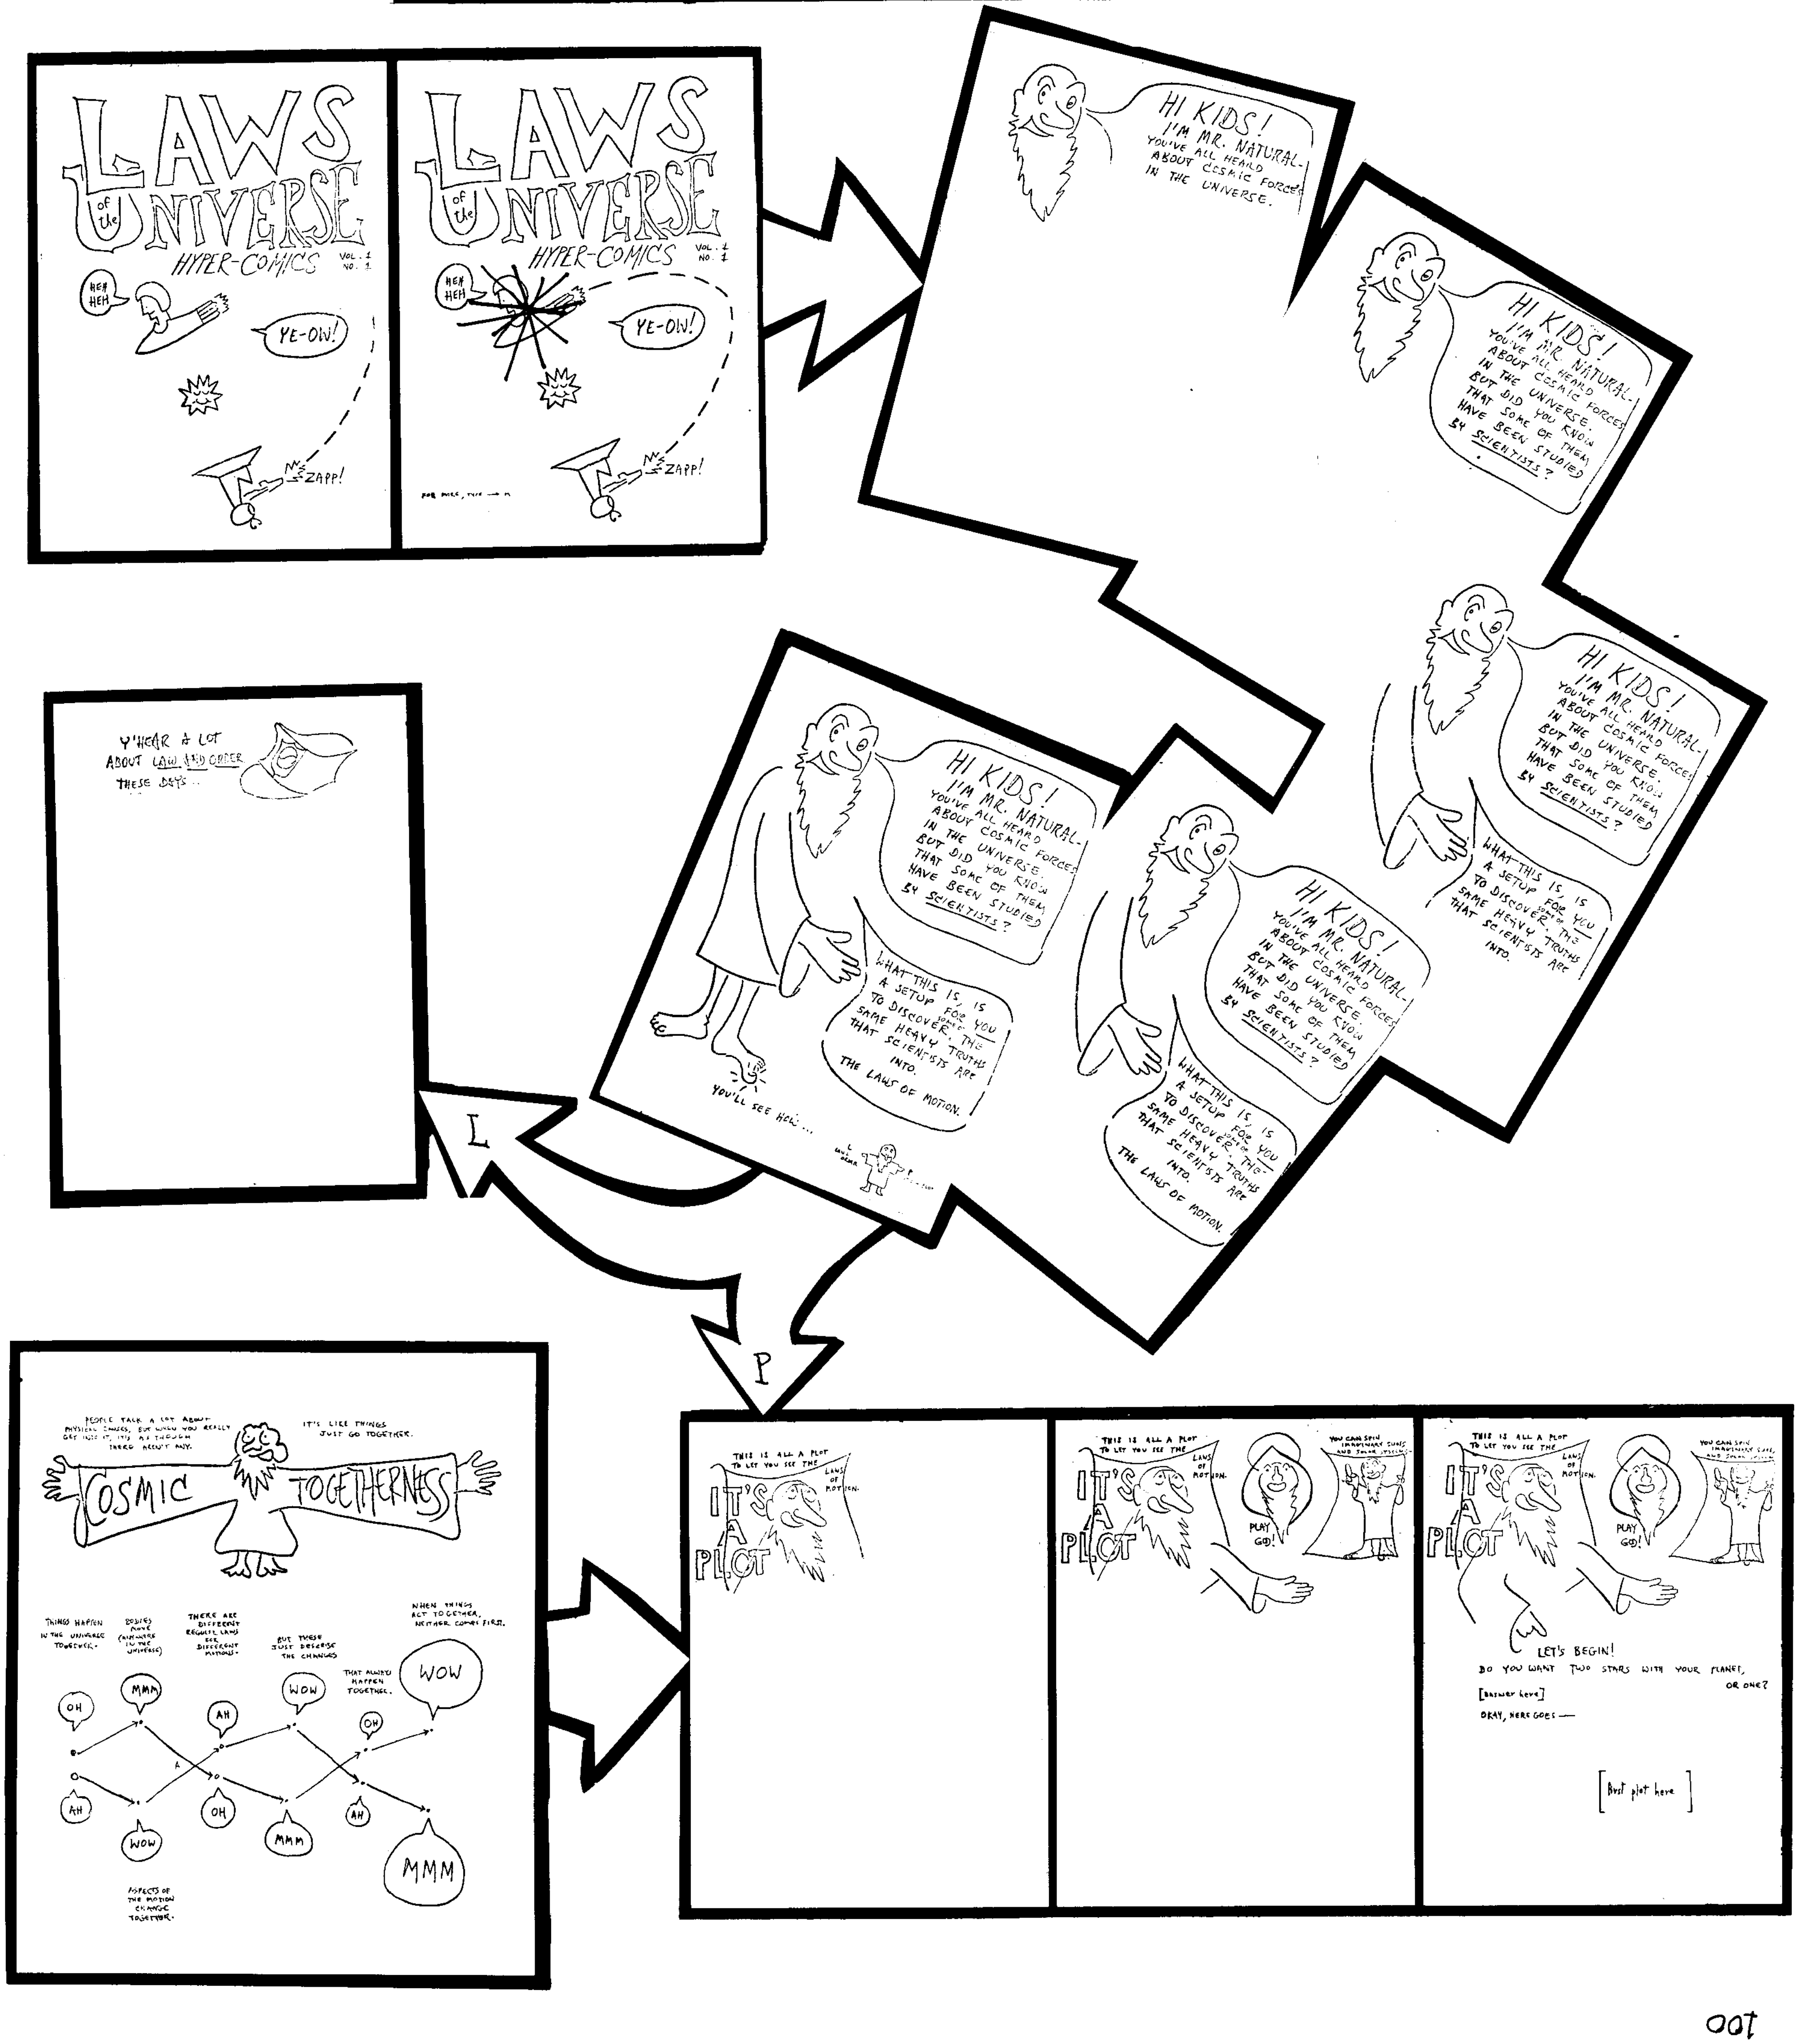

THREE-DIMENSIONAL LINE DISPLAYS

* THE MIND'S EYE, MORE *

So far we've discussed the two-dimensional subroutining displays. However, things do not by any means stop there. A number of people in the early days experimented with techniques for drawing line pictures by program; the earliest of these used plotters, output devices that let the program draw with a pen. But interest soon grew in the possibility of interactive three-dimensional displays on screens. Johnson's Sketchpad 4 did this entirely by program. But as night follows day, people set about putting these techniques into hardware, creating devices that would automatically show things in three-dimensional views-- allowing the viewer to rotate views of nonexistent objects as if they were on unseen turntables.

The views we are talking about, now, co sist of bright lines on a dark field, and so the "objects" we are talking about are called "wire-frame" objects-- they could effectively be made of welded wire. But now we do not have to build them physically to see them.

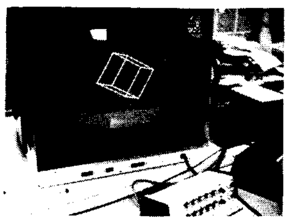

Basically a three-dimensional system of this type stores the lines as coordinates in threes: endpoints of lines in a mythical three-dimensional space. Each point's location in the space is told by three numbers (example showing a house may be seen on p.); a line in a space is represented in the data structure by two such points, and a code or something tying them together.

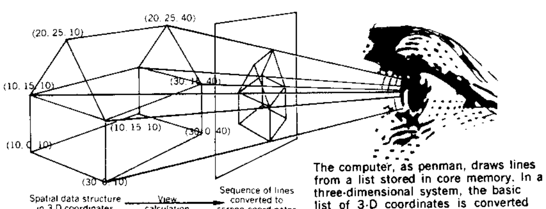

The computer, as penman, draws lines from a list stored in core memory. In a three-dimensional system, the basic list of 3-D coordinates is converted to a list representing a particular view; the result looks like a wire frame.

The second program follower in such a device behaves much as it does in the 2D system, but with certain additions. Like the 2D system, it proceeds down its own program one step at a time. Like the 2D system, it finds in its program the coordinates of a line to display and creates electronic signals representing its endpoints. But it does not display these directly, since these are three-dimensional coordinates. Instead it routes these signals to what we may call a view calculator, a particular piece of hardware that has been primed with the angle from which you want to view the object. This view calculator, automatically and by mysterious means which vary among machines, produces the view, and its signals go to the screen.

Let's say we want to display a point. The display's program follower pulls three numbers from its display list and notes the code that says it's a spatial point and not the end of a line. These three numbers slide on into the view calculator, already primed with the angle of rotation; and the view calculators figgers where on the screen that point should be displayed. The coordinates for the screen-- telling where the point goes in the desired picture-- go to the screen controller, and the point is brightened.

How are these coordinates calculated? Well, some commercial units do it electronically ("in analog") and some do it symbolically ("in digital"). The result is the same.

(If you want the equations for this, they're in the Newman and Sproull book.)

Then how does the view calculator handle a line? Same thing.

The program follower pulls three numbers from its display list and notes the code that says it's a line, so it takes three more. Then the view coordinates of both points are calculated and fed to the screen controller. The screen controller now has two points on its screen-- so it draws a line between them.

The first device of this type was, I think, the so-called Kludge (pron. "Klooj"-- computer slang for a ridiculous machine, but in this case applied affectionately) built at MIT's Electronic Systems Laboratory in the early sixties. This device was a one-of-a-kind, built out of DEC circuit cards and hooking to a bigger machine. The ESL Kludge showed vividly how good it was to have instantaneous view calculation under a user's control.

The first of these systems to be offered commercially, I believe, was the "Adage Display," made by Adage, Inc. of Boston, which used their unusual Ambilog computer (see p. 43) to rotate objects on the screen. I vaguely recall that it cost about $80,000 with computer but without accessories.

Actually Adage had a tremendous lead in this field, but they let it slip for some reason, and have now lost it to two firms: Evans and Sutherland on the high end, Vector General on the low end. (But of course things keep changing.)

The Evans and Sutherland Computer Company was founded in 1966 by Ivan Sutherland, creator of the masterful Sketchpad system, and David Evans, chairman of computer science at the University of Utah. (For a time both held appointments at U^2 at the same time, but now both have left the university to devote full time to their dream factory in Salt Lake City.)

Their first product was an extraordinary piece of hardware called the LDS-1, which they said innocently stood for Line Drawing System. (To anybody from Utah, however, LDS means Latter-Day-Saint, and don't you forget it. Evans, indeed, is a Mormon, but I've been told it may have been Sutherland's sense of humor that chose the acronym.)

It should be pointed out that a special advantage of digital perspective calculation is that viewed coordinates can be read back by the computer, and serve as new data, if you go for that sort of thing.

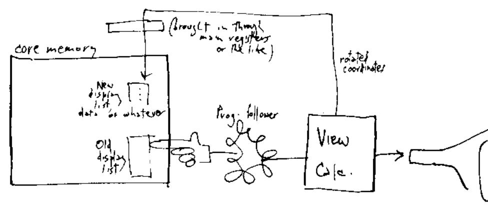

The Adage Display is isometric, meaning that lines do not get shorter as they get farther away or longer as they get closer. While this is marvelously impressive, most people want real perspective; and it was this that Evans and Sutherland set about to make available in real time, i.e., in direct response to the viewer's actions.

The LDS-1, weighing in at half a million dollars or so, buckled to the PDP-10, a big 36-bit computer from DEC (see p. 40). Its view calculator worked symbolically (digitally), and thus could work to the higher precision necessary for true perspective calculation.

Among the exciting demonstrations that you can see sitting at an LDS-1 are a map of the United States you can zoom in on, bringing you in to a map of New Jersey, then Atlantic City, then a specific intersection, all in one smooth continuous motion. Also a simulated landing on the flight deck of an aircraft carrier -- with you flying the airplane, so you can go over it, to the side, into the drink or straight at the carrier. In all cases the ghostly ship will move, turn and change perspective on the screen as if somehow it were really there.

Several LDS-1s were sold.

Meanwhile a little new firm of young guys in Southern California, Vector General, came up with a line of terminals like the Adage line, except that they could buckle to the 16-bit minicomputer of your choice. (In practice most of them have been attached to PDP-11s; see p. 42.)

The Vector General display is isometric, and makes its calculations in analog, like the Adage Display. It has been very successful among both universities and private corporations. In addition, a highly interactive and well-designed language is available for the creation of data structures representing 3D objects, as well as for general-purpose programming and the creation of whole environments. And it's free to individuals or companies that have Vector General displays attached to PDP-11s. (See "Coup de GRASS," p. DM31.)

But wait. Evans and Sutherland has now dropped the LDS-1 and given us-- no, not LDS-2, but something called The Picture System-- also built onto the PDP-11, but this one works symbolically (digitally) and in full perspective. The price starts at eighty grand.

Since the Picture System works out of the PDP-11 core memory, the commands it follows are 16 bits long, since that's the size of a slot in PDP-11 core. But wait. They've designed the thing to convert to 36 bits, so that coordinates are moved to a private store or buffer between the program follower and the display. This means the display can zoom and zip around in the scene without bothering the computer.

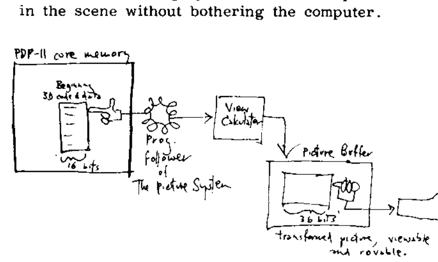

Another important feature of The Picture System: it will do, not just ordinary perspective, out such weird view calculations as wideangle barrel distortion, pincushion distortion and similar stuff.

INTERACTIVE ROTATION

3D screens-- aside from their fun and excitement-- allow people to understand and work with complex 3D structures without having to build them physically.

The understanding, however, comes from being able to turn and manipulate the structure on the screen. If you can't turn it you can't really perceive the 3D structure, because the arrangement of lines could be anything.

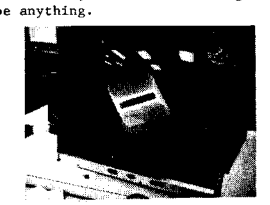

However, systems like the Adage and the Vector General and the Evans and Sutherland devices allow you to turn things on the screen as easily as if they were on turntables behind a pane of glass. That's how you see, you see.

This interaction is what makes computer display augur a new era for mankind, if we're lucky. (It's also why we use the term computer display in this book, rather than "computer graphics," since people who make computers draw with pens are also doing "computer graphics"-- a related activity, but not one to change the world.)

UNFORTUNATELY, just to get through the basics, there is only room to discuss stick-figure graphic display here. But curved surfaces may also be depicted, though usually not interactively. See below, and pp. DM32-9.

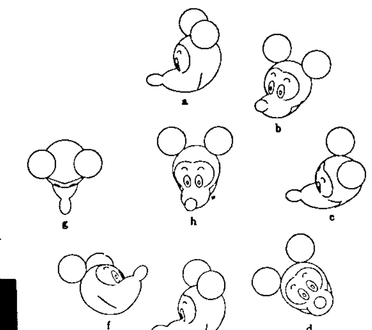

Drawing by Ruth Weiss' BE VISION program, done at Bell Laboratories, mid-sixties. (© Walt Disney Productions.)

This program represented truly curved surfaces in its data structure, as "quadric surfaces"-- that is, involving powers of two in the math-- and calculated the visible lines tangent to the edges from the viewpoint, thus drawing the edges. Removing the hidden parts of the curves is of course one of the greatest problems. (From Ruth A. Weiss, "BE VISION." JACM Apr 66, 194-204, p. 201.)

Mr. Rolls, Meet Mr. Royce

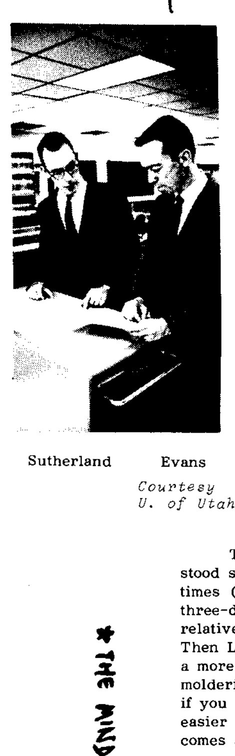

Sutherland Evans

Courtesy U. of Utah

* THE MIND'S EYE * CONCLUDES ON P. 40.

The rules of perspective have been understood since the Renaissance. In olden computer times (up till about 1965) people used to do three-dimensional view calculation by angles relative to a three-dimensional data structure. Then Larry Roberts at MIT noted that there was a more appropriate mathematical method, long moldering in obscure texts. The idea is this: if you add an extra dimension to the data, it's easier to program. It's easier because it becomes a simple matrix multiplication, which has no commonsense explanation but is important to mathematicians.

SO that means that to calculate views of three-dimensional objects, the most usual way is now to add that extra dimension. Instead of having a point in space whose position is 36-24-36 (in some set of three-dimensional coordinates), another arbitrary number is added to make it, say, 36-24-36-1.

It seems that in the mathematics of multiple dimensions, it comes out simpler that way. Indeed, from a mathematical point of view the new improved dimension is just like the other three. For this reason, such an augmented system of coordinates is called homogeneous coordinates. Like homogenized milk, the additional coordinate is just stirred in with the rest, and out comes your desired view calculation. (The formulas are to be found in Newman and Sproull, Principles of Interactive Computer Graphics, McGraw, $15, your basic text on the subject.)

At any rate the additional coordinate is often referred to, incorrectly, as the "homogeneous coordinate." They're all homogeneous, which is why it works.

DeFanti*s Coup de GRASS

Impudent and plucky Tom DeFanti was an assistant professor at 24. This in part because he has created one of the world's hottest 3D graphics languages, which he calls GRASS. (He says it stands for GRAphics Symbiosis System-- also, he says, it Turns You On.)

Tom's GRASS language is an excellent beginner's computer language for two reasons: first, it is easily taught to beginners, and second, it is about things of interest to beginners, i.e., pictures and graphical manipulation on screens. (But compare the three beginners' languages presented briefly on pp. 16-25.)

A prototype for the system was developed at Ohio State, on a project directed by artist Charles Csuri. Tom had a free hand, though, and the language design is his; but much of the specific coding was done by Gerry Moersdorf, and the graphics algorithms and rotation were programmed by Manfred Knemeyer. Inspiration was furnished by Maynard E. Sensenbrenner.

GRASS runs on the PDP-11, a splendid minicomputer (Tom's is shown on p. 36) and is specifically designed for the control of three-dimensional stick-figure displays on the Vector General display system (see p. DM 30). But a lot of people have wrestled with these matters and not done as well. Let's consider:

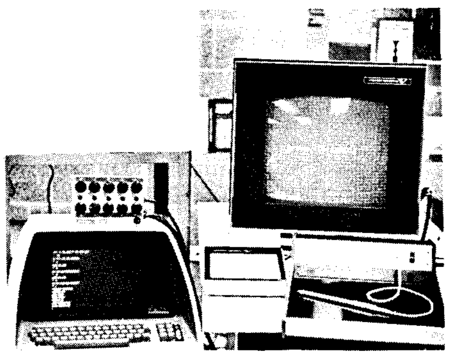

1. ITS CLEAR SIMPLICITY. Tom believes computers are for everybody; he is not a high priest bent on making things obscure (see "Cybercrud," p. 8). Thus he made his language as sensible, clear and easy to learn as possible. Tom likes to stress the concept of "habitability" (a term of W.C.Watt), meaning the coziness of a system.

2. ITS GENERALITY. Refining and condensing the basic ideas of a system is the hardest part of the design. DeFanti made several interesting decisions.

A. The internal form of the language is ASCII code (see p. 26). In other words, you can read programs in their final GRASS form.

B. For a three-dimensional system such as the Vector General, the main form of data structure is the three-dimensional object-- a list of points and lines in space. This is the form of data GRASS uses for most purposes.

C. In the design of such a system you want larger 3D objects to be buildable out of smaller ones. This implies arranging data in tree structures (see p. 24). You also want to be able to make things do compound motions on the screen-- for example, showing an airplane flying around on the screen with its propellor spinning; this too implies a tree structure. There are some programmers who would use different tree structures for both objects grouped together and for movements grouped together; Tom uses one.

D. Objects shown on Tom's system can also appear to move on complicated paths through three-dimensional space. In Tom's system, such a path is merely another object. It seems obvious when you say it, yet this kind of simple generality is exactly what many programmers seem to avoid. (Note: this facility is a generalization of Baecker's p-curve; see p. DM 25).

E. Input devices are completely arbitrary and programmable. What happens on the screen can be controlled by anything-- any variable (see p. 16) in the programming language. In other words, DeFanti has decoupled the screen from any particular form of control, allowing user programs to make the connection between controls and consequences. This means that, using Tom's language, it is comparatively easy to build complex custom controls for any function. (This is discussed under "Fantics,"* p. DM 48-51.)

F. The language has string functions that allow text handling. Since the language may also use conversational terminals, it is eminently suited for "good-guy" interactive systems for naive users, as described on pp. 12-13.

G. Tom's language is interpretive, like TRAC Language (see p. 30). That means it is "slow" in terms of the number of machine cycles required for it to do each operation. However, DeFanti has added a "compile" feature to the language, so that for long macros (sections of program) that have to run repetitively, more efficient compiled versions of the macros may be generated.

H. The language is extensible, meaning that the user may create new commands in the language as programs. These commands, however, may be used in later programs as if they were built into the language itself.

I. The system is completely general-purpose. Many graphics languages are not, being restricted only to their original purpose. This is more difficult, but oh, so much more worthwhile.

3. ITS DEEP GENERALITY. Things should be versatile, and able to be tied together in many different ways. This is what we mean by "generality;" and this kind of generality can make a system very powerful. (The term in mathematics is "elegance.") As is said on the other side of the book, complicatedness is not generality or goodness or power, but a sign of the designer's shallowness.

Anyway, GRASS has this kind of generality. It has a great number of facilities, growing weekly, and they all tie together in clear and predictable ways, without exceptions. Rather than create special functions which cannot be tied together, Young Doctor DeFanti has chosen instead to make the separate desirable functions part of a simple and clear language. (A note to you elegant types: GRASS is fully recursive. As a nice example, Dan Sandin (see p. DM 8) wrote a program to display Peano lines that was under forty GRASS instructions long. It is also astonishingly reversible: you can watch it uncreate the Peano line, straightening itself backward.)

In the more usual sense, DeFanti's language is not the 'most advanced'; there are more powerful 3D systems than the Vector General (the LDS-1, see p. DM 30, offers true perspective), more elegant user-level languages (see TRAC Language and APL, other side), true halftone (the Watkins Box); yet his achievement on close examination is extraordinary. Never mind his age, the more esoteric features of his system (full recursiveness, etc.) or the fact that he does not seem to have made one mistake, which is infuriating. Consider only this: TOM DEFANTI'S 'GRASS' LANGUAGE IS PERHAPS THE ONLY SYSTEM THAT CAN BE TAUGHT IN A FEW HOURS TO COMPUTER-NAIVE BEGINNERS THAT PERMITS FULL THREE-DIMENSIONAL ANIMATED INTERACTIVE GRAPHICS WITH TREE-STRUCTURED DATA.

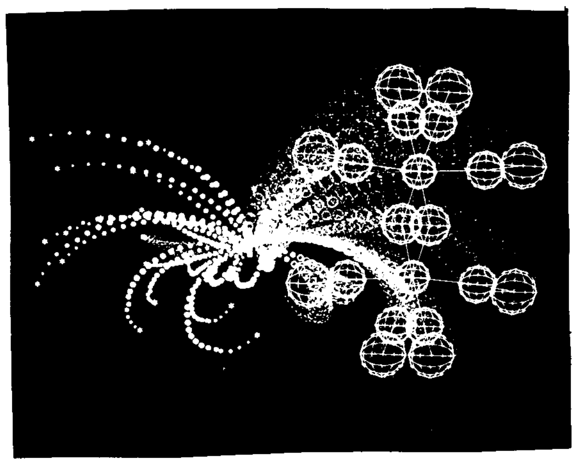

Tom DeFanti

THREE WAYS OF SEEING MOLECULES USING 3D COMPUTER DISPLAY.

Much of today's impetus for 3D computer display is coming from the field of chemistry. University chemistry departments are buying equipment like the Evans & Sutherland LDS-1, the Adage and the Vector General.

Why?

Because chemistry is increasingly involved with complex three-dimensional structures. Crystals, long folding chain molecules, minuscule forces acting on structures whose shape determines the outcome. Organic molecules that involve thousands of atoms, and whose complex folded structure exposes only certain key features. And so on.

The Vector General display illustrated here and there on these pages belongs to the Department of Chemistry, University of Illinois at Chicago Circle.

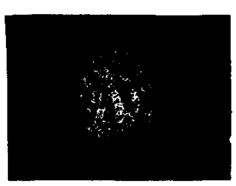

Tom DeFanti. Shows part of hemoglobin molecule. Data structure from Richard J. Feldmann, NIH.

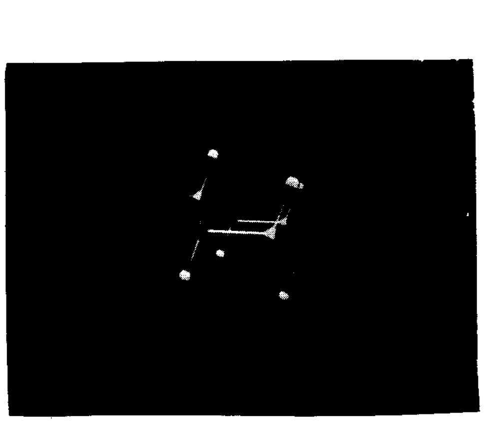

Bouknight & Kelley (see p. DM 34)

* I coined the term fantics, for the art and technology of showing things, long before I ever heard of Tom DeFanti, and I am not about to change it just because he is now my friend and roommate.

The best feature of all: it's currently available. PDP-11 owners-- even without Vector General displays-- may inquire of: Tom DeFanti, Doctor of Arts Program, UICC, Chicago IL 60680.

You may wonder how a young bronking buck like DeFanti has managed to do such an excellent job, so elegantly, where so many have stumbled and failed?

"I just learn from other people's mistakes," he says cheerily.

Prof. DeFanti on the system.

MISCELLANY:

Coupling his system with that of Dan Sandin (p. DM 8) has created the "Circle Graphics Habitat," described on p.

I hope I'm around long enough to write the GRASS language manual.

(DeFanti's GRASS is an ideal language for something like the 3D Thinkertoy, described on p. DM 55 . However, it doesn't have any provision for the storage of large complex data structures, so the hard part would actually be working out an adequate storage data structure and storage macros within GRASS's use of the DEC file system.)

SCREEN CONTROLS

The great thing about CRT displays is that they can be used to control things by manipulation of pictures. Instead of moving buttons or levers, you can seize parts of the picture with the light-pen and move some part of the picture. The computer, sensing the choice or adjustment you have made, can then perform whatever operations you have directed.

Some samples:

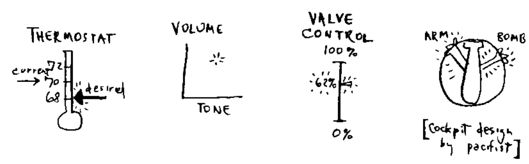

The design of screen controls-- easy-to-use, clear and simple controls for everything-- is one of the frontiers of computer graphics. (See "Fantics," p. DM 48-51)

DIMENSIONAL FLIP

3D scopes are about the best we've got-- so what do we do about multidimensional phenomena?

One very good solution is to show a selection of three dimensions at a time, and provide for easy "flip" from one dimension to another-- so that instead of looking at something on demensions A, B and C you are looking at it on dimensions A, B and X.

For example, suppose you're a sociologist looking at measurements of various traits among a group of people. It's a cloud of dots in three dimensions-- whatever three dimensions you're looking at. Some could be: age, height, weight, sex, ethnic background, premarital experience, education... etc.

You view this cloud of dots, say, according to age, weight and ethnic background. That means you can rotate it around and see how many people in the group are what.

Using dimensional flip, however, you can change the view as follows: rotate the box-frame till it becomes a square to your eye. Then you hit the control that makes the unseen dimension "flip" to another dimension that interests you. The cloud still looks the same-- until you rotate it, and the third dimension is now "premarital experience." So you can quickly get a view of how populations are really divided up. (Note to sociologists: this same operation, with stretching and clipping, provides a visual technique for "partialing" operations of the Lazarsfeld type.)

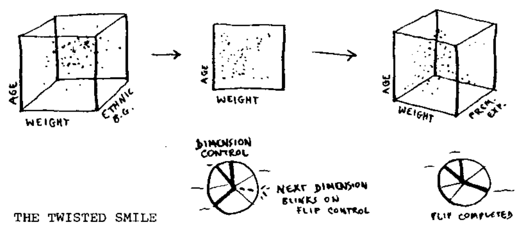

THE TWISTED SMILE

You can make a character change expression on a 3D scope by making his mouth a twisted wire that can be rotated between "frown" and "smile" positions. The trick is the shape of the wire.

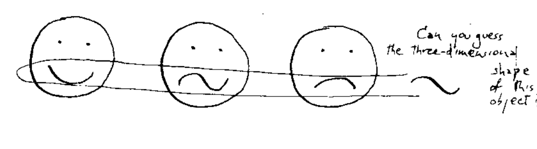

NOW GUESS WHAT: DeFanti's GRASS language is the best language I know of for doing all the above things.

DM32

COMPUTER HALFTONE IMAGE SYSTEMS.

A Series of Review Articles for Computer Decisions Magazine.

WHERE TO GET IT.

Computer 3D halftone systems are now available to moviemakers from a variety of sources. It tends to cost a lot of money, but when compared with normal Hollywood production expenses, it turns out not to be so bad.

SALES OF MACHINES.

Computer Image Corporation, Denver, offers various systems for sale. See p. DM 39.

Evans and Sutherland Computer Corporation, Salt Lake City, offers the Watkins Box, a real-time display device using the Watkins Method (see next page) and offering also Gouraud pseudo-curved shading (see p. DM 37). It costs about $500,000 and attaches to a PDP-10 large computer; see p. 40).

SERVICES.

General Electric, Syracuse, offers three-dimensional scene synthesis like that at the bottom of this page. Every job is custom. It's done on videotape through programs running on a smallish computer. Production costs, after your data structures are all in, could run as little as hundreds of dollars per minute (rather than thousands).

Contact: Charles P. Venus, General Electric Co., Building 3, Syracuse NY 13201, 315/456-3552. (Given in detail because harder to reach than these others.)

Computer Visuals, Inc., Elmsford, NY. Offer more detail than GE system, and go straight to film without video. More expensive: probable costs run in the thousands of dollars per minute. Again, every job is custom.

Contact: Nat C. Myers, president.

Dolphin Productions, NYC, has several Computer Image machines, but their president, Allen Stanley, is interested in everything.

Computer Image Corp., Denver and Hollywood, also offers services on their machines. On occasion they have been willing to back film-makers, reportedly on a 50-50 basis. Their president, Lee Harrison III, is a swell fella.

Author's note. These articles were written for Computer Decisions magazine, and reflect the results of a lot of phone calls they paid for. The first of these articles was published in 1971. The others have not been previously published, as the editors and I were never able to get together on quite what they wanted.

This is, to my knowledge, the only existing collection and summary of computer half-tone systems to date, and in some cases the articles reveal more about the systems than has been published anywhere. Surprisingly, even two years later they do not seem out of date.

However, due to the editorial style of Computer Decisions, and my own, this has all come out extremely condensed, and phrased in breezy and humorous ways not ordinarily considered acceptable for serious technical reviews. The hope is that they will supply orientation to the browser, deeper insights to the technically-minded, and further directions for them as wants to pursue.

My thanks to the publishers of Computer Decisions and its editor, Robert C. Haavind, for their encouragement, phone money and permission to reprint this.

FIRST ARTICLE

General idea of 3-D halftone.

Polygon Systems.

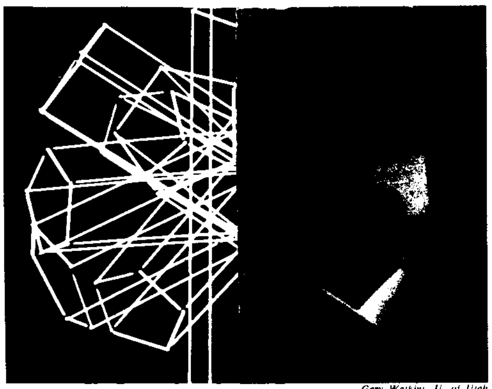

Gary Watkins, U. of Utah

halftone image synthesis

There are more ways than one to produce shaded pictures with computers. Here are the methods of the 'polygon school.'

by Theodor H. Nelson
The Nelson Organization

To most people in the computer field, "computer graphics" means line drawing—systems and programs for mapmaking, pipe layout, automobile and aircraft design, or any other activity where a diagram may help. Using line-drawing programs and equipment, designers may create line drawings on fast-responding graphic screens, reworking their ideas until satisfied; the system then disgorges polished drawings and specifications for the designer's real intent, something else that is to be made or done. But it is possible for a picture itself—instructive, interesting or pretty—to be the goal. In that case we will often want pictures that look like things instead of wires. A picture that is not all black and white we call "halftone."

With much secrecy and a slow start, computer halftone systems are now being built all over. The methods are extremely different from one another; only the outputs are similar. Some exist in software, some have already been built into special hardware. These systems have many potential uses for visualization, animation and new kinds of photography, in art, scholarship, motion pictures and TV; for visualizing worlds lost and imagined, equipment yet unbuilt, the responsiveness of aircraft. It may not be long until moviemakers can buy different brands of picture synthesizer, just as musicians choose today among Moog, Buchla and ARP music synthesizers. But none is in production yet. This is an attempt to review the coming apparatuses of apparition.

Not only is the field of halftone one of the most exciting in computing; it is also one of the nuttiest and most secretive. For instance, at one time a firm that was supposedly marketing its halftone system declared the present author *persona non grata* and not to be communicated with in any way, though information was freely available to others. "I don't think it's necessarily paranoia," says Rod Rougelot of General Electric. "A lot of guys started about the same time, and proceeded in a heads-down manner." It took a special kind of initiative to head off in that direction with no external provocation. "All those heavy cats from ARPA and MIT were saying in the sixties I could never do a Mickey Mouse," says Lee Harrison III of Computer Image. "But I'm not that kind of researcher. I talk to the Lord."

The systems' stories are as different as the systems themselves. General Electric's system grew out of cockpit displays for blind flying. The system of Pennsylvania Research Associates began with terrain and radar modelling. The system of MAGI (Mathematical Applications Group, Inc.) began with the study of radiation hazards in battlefield machinery. Two system families, that of Computer Image Inc. and my own Fantasm, were designed from the beginning for moviemaking, especially "special effects" and puppeteering. The most poignant tale may be that of Lee Harrison, whose struggling family was warmed through cold winters by the tubes of their analog computer.

Halftones in two dimensions

Two-dimensional computer halftone is not new. Halftone pictures converted from photographs have often been printed out on line printers, either for fun (nudes often turn up at big installations), or in connection with some scientific problem, such as analyzing chromosomes. Kenneth C. Knowlton, at Bell Laboratories, has executed some well-known photo conversions making pictures into huge grids of tiny whimsical symbols having different grey-values.

Various other systems have allowed users to create their own original 2-D pictures. But the natural temptation is to want the computer really to make pictures. Why not have the computer produce a photographic picture directly from the 3-D representation of objects? Computers don't do this by nature, any more than they do anything else by nature, so *how* it may be done by computer is very interesting. The problem is also interesting because of its intuitive nature. Visions of scenes in space are around us constantly, and we intuitively understand the geometry of outlines and light. As 3-D work progresses large problems are being overcome. The famed "hidden line problem," for example, was misleadingly couched, since the problem is not finding what lines are hidden, but what surfaces are in front!

Computer graphics the ordinary way
The computer, as penman, draws lines from a list stored in core memory. In a three-dimensional system, the basic list of 3-D coordinates is converted to a list representing a particular view; the result looks like a wire frame.

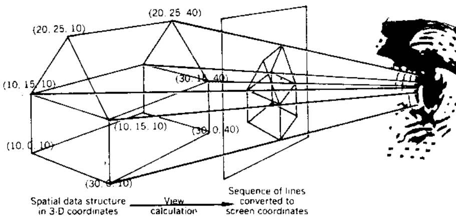

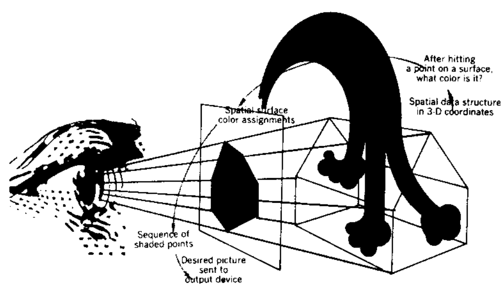

3-D halftone system
Today's new procedures can use the same data to make a realistic shaded or halftone picture. The visible parts of the objects are ascertained by programs or special hardware, using the same 3-D coordinates as in the ordinary systems. These visible parts are then shaded according to the appropriate color information. The series of shading-points makes the picture on an output device.

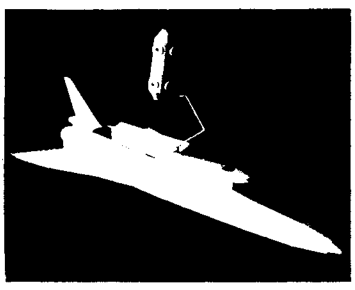

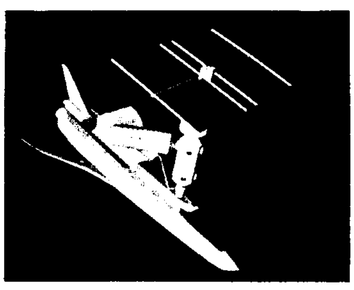

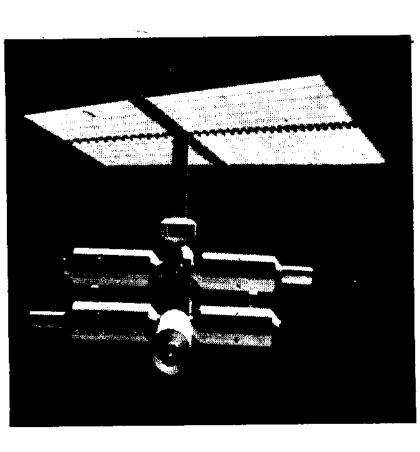

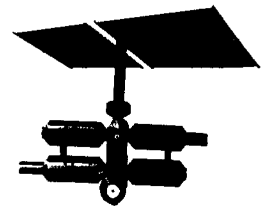

General Electric will make movies and videotapes for you with their pictorial synthesis system. These are from a beautiful (really beautiful) film they did for NASA. The point of the film was to explain to everybody how a proposed space laboratory would be built and would function. Rather than use diagrams, they enacted it in the GE system, so viewers could understand how the sections would be delivered and fit together, how the antennas would unfold and so on. For exposition of that kind, nothing beats this kind of enactment.

DM33

We must draw on this understanding of scenes to figure out how to make pictures, for there is no mathematically elegant or preferable approach. Scenes are geometrically rich, and thus many different techniques may be used to extract pictures from them. These techniques may look at planar structures, spatial interconnections, relative edges of intersections or anything else you can define and process. I prefer to think of computer halftone as like trick photography of the kind done in Hollywood: a variety of techniques can be combined in various ways. As in trick photography, the number of touches and enhancements that you add generally determines how good it will look, regardless of what system you begin with.

The simplest systems are those that depict objects made of polygons—that is, planes with straight edges. We will discuss such systems in the present installment.

The wild polygon yonder

At least two companies are building image systems that will behave and respond like onrushing reality. Such a system, hooked to cockpit-like controls, can show a trainee pilot the delicate and precipitous results of what he does. Realistic action, rather than surface detail, is crucial.

The techniques of action polygon halftone were originally developed by General Electric, of Syracuse, N.Y., and are now also under development at Link Division of Singer Company (makers of the beloved pilot trainer and its progeny). Basically such systems operate upon the scan-lines that crisscross a television screen, switching the color of the running scan as it crosses from polygon to polygon.

The action polygon school—GE and Link—takes a curious but effective approach to halftone TV: their "environments" are composed entirely of convex objects made entirely of convex polygons. To use only convex objects (no dents) means that one object may be in front of another or vice versa, but never both. (An object with apparent indentations, such as an airplane, has to be made out of a group of convex objects flying together.) To use only convex polygons (notchless) makes it easy for the system to decide, at a given instant, whether the scan is crossing the polygon or not.

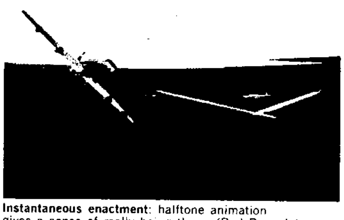

Instantaneous enactment: halftone animation gives a sense of really being there. (Rod Rougelot, General Electric)

This work evolved in part from GE's work in the fifties with a "ground plane simulator," a system that would show a correct representation of the ground's position, dipping and rotating, to the pilot of an aircraft in fog or night. In 1963 the General Electric group, under Rod Rougelot, worked out for NASA the design of an "environment simulator"—a device that would simulate the appearance and performance of any equipment. This is now called the "old NASA system." It permitted the user—seated before a color TV screen—to work controls for an imaginary aircraft or spacecraft, and see roughly what the pilot of the craft would see, flying in real time through a breathtaking color scene. Films made on this machine have been stunning. Imaginary cities, roller coasters and aerial dogfights are among the visions that can be presented.

General Electric's old NASA method is fairly weird if not mischievous. The earlier "ground plane simulator" had shown an edge (the horizon) digitally displayed on a crt; the system was extended to many edges, and the logical analysis of areas between them.

The scene was represented by a collection of edge boxes, physically jumpered into a collection of facet boxes. Each edge box and facet box was loaded with certain numerical and logic values, representing edges and facets in the scene, which could change between frames as required by the action.

In the preprocess for each frame the old NASA system used a specially built digital computer, the "vector calculator." This performed at great speed the three-part vector calculations necessary to determine all scene positions, including the positions and slants of all edges. Each individual edge generator, loaded with its own edge position, constantly reported whether the running scan of the picture was to the left or right of its own edge. It dutifully guarded this edge from border to border of the picture.

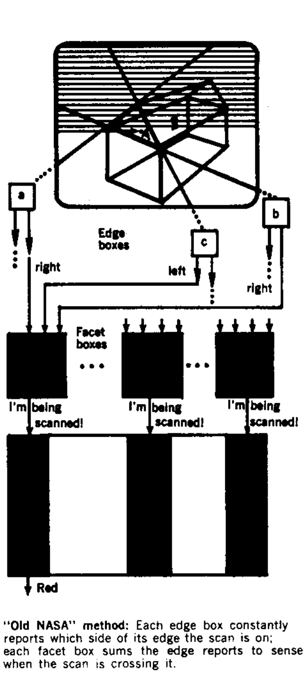

"Old NASA" method: Each edge box constantly reports which side of its edge the scan is on; each facet box sums the edge reports to sense when the scan is crossing it.

The edge-box reports summed into the facet boxes, each of which was set to respond to a particular combination of left-right, above-below reports. At the instant all the facet's edge boxes replied in the proper preset combination, the facet box signalled that its own facet was being crossed by the scan-line. When more than one facet-box responded, the one nearest the viewpoint had its color gated to the screen.

Now Rougelot's group is replacing the old NASA system by a new NASA system, which works on entirely different principles, but keeps the vector calculator. The old one could show scenes with up to 240 edges; the new NASA system will at least double that. GE's new method is already operational on smaller research facilities. They don't tell what it is, but basically it involves sorting by distance. Supposedly the sort method is good enough to make the old edge boxes obsolete.

The Link group claims competitive performance for their system, which will go to black-and-white thousand-line TV. They say their system is different, better, and secret.

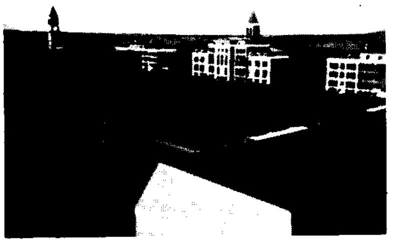

Campus of Fooled U. (GE)

Wylie-Romney: shoot the works

The Wylie-Romney method, disclosed in 1967, was the first generally publicized procedure for making halftone pictures. Indeed, the 1967 publication signalled the explosion of the University of Utah into the forefront of computing research.

The Wylie-Romney method was actually the joint work of Chris Wylie, Gordon Romney, David C. Evans and Alan Erdahl; but much of the impetus for its development came from Evans, chairman of computer sciences at Utah, who had long suspected the possibility of 3-D halftone synthesis.

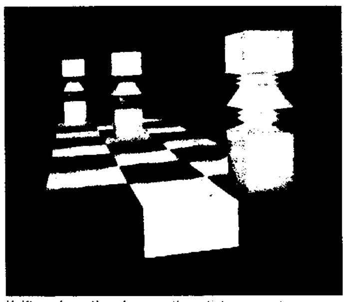

Halftone for art's sake: now the artist can create worlds and photograph them. (Gordon Romney, Utah)

(Note: more output by various Utah systems appear on following pages.)

The Wylie-Romney method is this: for each picture-point desired in the final picture, shoot a searching ray through the scene at a corresponding angle. Find where this searching ray hits every surface in its way.

Since the locations in space of these hit-points are easily calculated, figure their distances from the vantage point. The nearest of the intersections is the visible one. Look up the color of that surface and shade the output point accordingly.

This may sound inefficient, but it is comparatively easy to ascertain all the piercing-points, since the surfaces to be hit in a given scanning row can be largely predicted from the previous row.

John Warnock's method, also from Utah, is unrelated to the other methods, but has qualities mathematicians like, as well as a certain whimsy.

Consider a square in the picture area. (At the start consider the whole picture area.) *Now then.* Test whether the present square is entirely filled with one color. If so, output a corresponding square all of that color. If the present square is not all one color, divide it into four smaller squares. Take another square and go back to *Now then.* End the process when each of the squares in the broken-down picture has been completely filled with one color—or the unsatisfied squares are too small to care about.

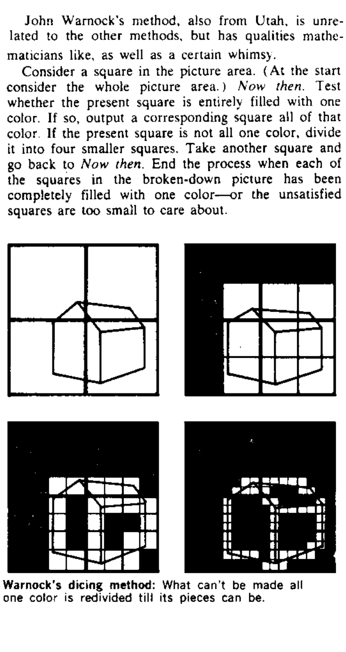

Warnock's dicing method: What can't be made all one color is redivided till its pieces can be.

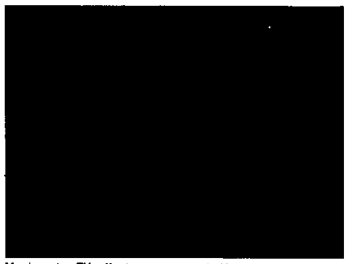

Movie sets, TV effects: computer halftone is ready to compete. (Gary Watkins, Utah)

The method of Gary Watkins is the result of a profound search at the University of Utah for *the* method—a polygon technique fast enough for real-time enactment, but cheaper than the GE-type systems and not subject to the convexity restrictions. They seem to have found it.

Each video scan of the scene results in a "slice" through surfaces in the scene. The two nearest surfaces are continuously compared to see which is closer, as if by two rulers. The instant a new surface becomes the nearer one, the system makes it the visible one. The nearest surface always shows, down to the precise instant two surfaces cross.

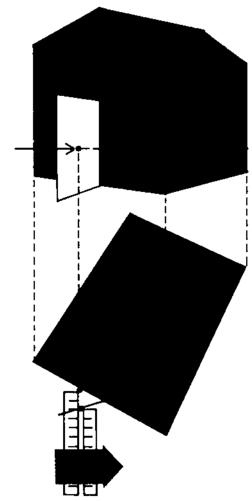

Watkins method: A new nearest surface is instantly sensed through continuous comparison of the closest two.

NOW AVAILABLE! Machine running Watkins technique, the Watkins Box, allows you to view imaginary objects in color and manipulate them in real time. See top of preceding page.

Shading: Last of the great fudge-functions

Suppose that we have some data structure representing a three-dimensional object, and a halftone method to search out its visible surfaces. How do we shade the output points? What do we take into account: how combine the basic greys or colors, how blend them with computations of surface angle, distances from the vantage point, or anything else we can think of?

The answer: any way at all. The combining function is an aesthetic choice. There are not many areas left where you can make up a mathematical hodge-podge and get pleasing or interesting results. Computer halftone is a felicitous exception: you can augment by adding or multiplying, diminish by subtracting or dividing, and yet always come up with an image resembling something. Anyone who has worked in a darkroom will recognize that this is like enlarging: playing with parameters won't obliterate the picture.

There are purists who insist that halftone coloration should exactly follow the formulas that simulate the behavior of real light. For some purposes, like pilot training, this may often be true. But insisting on mathematical accuracy as a general principle is like insisting on ultra-high fidelity—an aesthetic judgment couched as a mechanical imperative.

Until now the output hardware was not really ready for halftone. Five years ago a computer could usually create halftone pictures only on a line printer or a 4020 microfilm plotter. Today there are many different photographic printers, going to all sizes of film and paper; one even uses a laser. There are various display terminals permitting grey-scale and color halftone on TV screens.

The age of computer image synthesis has begun. Polygon systems are fast and simple, and will come to be used in our daily lives for such diverse purposes as molecule study, the memorization of delivery routes, and visualization of every kind of layout and design. They will be fundamental to our new world of computer display. □

DM34

SECOND ARTICLE.

Surface patterns.

Curvature.

Shadow.

SHADES OF REALITY

VARIOUS NEW TECHNIQUES PERMIT US TO ADD CURVES, SHADOWS AND SURFACE PATTERNS TO COMPUTER-GENERATED HALFTONE PICTURES

THE PLOT SO FAR.

Various computer methods now make it possible to create artificial photographs of three-dimensional objects or scenes represented in the computer's storage. This is done by coloring or shading points in an output picture like the points in the scene that can be sighted through them from the vantage point. What the methods really boil down to, though, are searching processes in the data representation of the three-dimensional scene.

In an earlier article we have considered some of the techniques being used to depict simple scenes-- those made up of polygons. Now we turn to more elaborate scenes which add shadows, surface patterns and curvature.

One of the most interesting things about this branch of computer graphics-- already seen in the polygon methods discussed earlier-- is the variety of techniques that can be employed. Moreover, these methods, for all their sophistication, can usually be intuitively understood as thought they were operations performed on objects in space. The same continues to be true for the more complex systems.

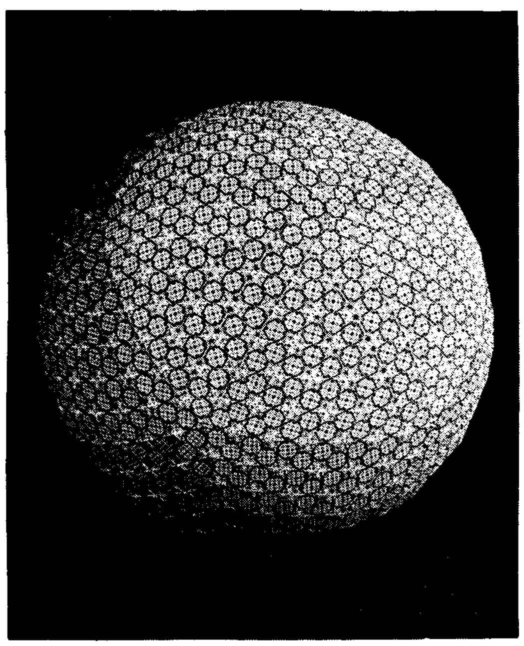

MAGNUSKI'S PATTERNED CONSTRUCTIONS

A number of contributions have been made by individuals working alone. For instance, Henry Magnuski, at M.I.T., created a program that repeatedly positions patterned facets in space to make large constructions.

This program did not calculate "true" shadow, basing its shading partly on angle of surfaces. Neither does it show true curves. Yet it shows the impressive degree to which such effects may be approximated. The resulting beach ball picture is reminiscent of Moorish architecture.

ENHANCED POLYGON SYSTEMS

In the methods discussed so far, we looked at several computer techniques for photographically depicting scenes and objects made up of polygons-- planar facets-- in a represented three-dimensional scene. Imaginary houses of cards, cardboard airplanes and triangular scenery take on a compelling vividness when depicted by the computer. And for visualizing such things as architectural arrangements, such systems promise to be of increasing practical value.

Those of us interested in the artistic aspects of computer halftone images want more. This article looks at some ways to add the appearance of curvature and surface pattern to computer-synthesized images.

MAGNUSKI'S CONSTRUCTIONS OF REPEATED PATTERNS

(different perspective calculations)

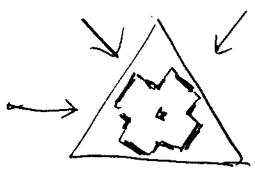

Basic triangle pattern...

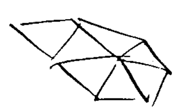

is stitched together in adjacent positions at appropriate angles.

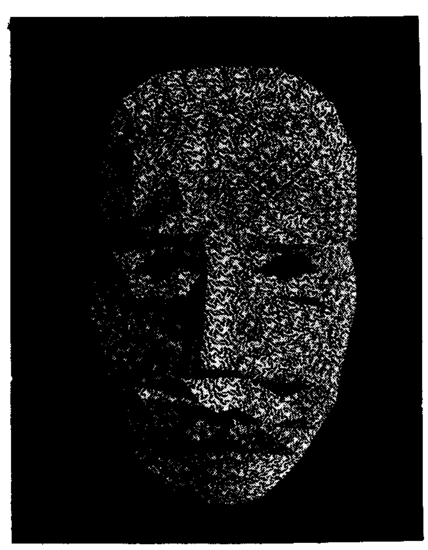

BOUKNIGHT AND KELLEY: PICKING THROUGH A CAT'S CRADLE

The method of Bouknight and Kelley, at the University of Illinois, permits the addition of shadow to polygon pictures. Their method uses an intricate system of scanning sweeps across the scene, analyzing the successive edge-crossings. For each output line, a list of the edges in the scene is ordered according to which will be next encountered. To make a specific output line of shaded points, we step through successive positions of the scan-line, until an an edge is crossed. With each edge we cross, we enter or leave at least one facet. Of all the current facets we are in after a given edge-crossing, the system finds out the nearest one, the visible one, by comparing distances. The coloration of this facet is then fed out to the picture, until the next edge-crossing.

Bouknight and Kelley expand their method to show shadows by an additional step. They create a new list of edges to be encountered, this one relative to scans from the light source. Then, during the regular output picture scan, they look to this latter data to see about shadow. As soon as they know two consecutive edges of a visible object in the picture, they are able to search the shadow-edge list to see if any shadow-edges impinge between them. The final list of edges-- visible facet edges and shadow edges-- goes to the picture output device.

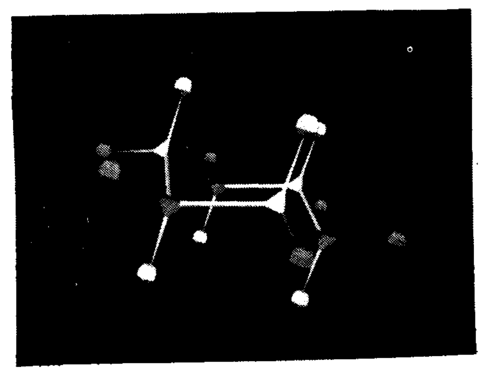

BOUKNIGHT-KELLEY METHOD

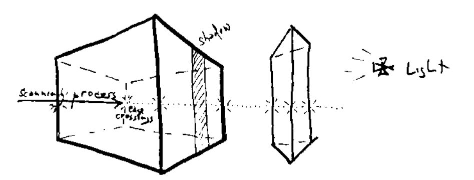

Consider the series of edges whose projections cross the current scan-line. Each time the scan-line crosses an edge, find out what facets are currently pierced by a sight-line from the viewpoint. The nearest of these facets is the visible one.

To add shadow, use an extra list of the scene's edges relative to the light rather than the camera. Between viewed edges, check for shadow-edges as well.

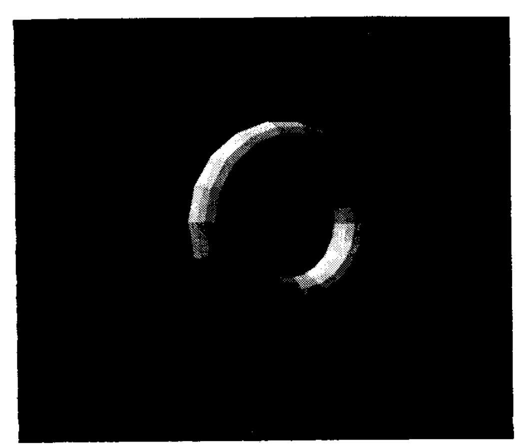

DM35

DON LEE FILLS IN THE GAPS

Don Lee, at the University of Illinois, produced his fine-toned pictures of spheres in 1966 simply because someone bet him a quarter he couldn't program the method he'd suggested in twenty-four hours. He almost made it. He made his pictures of spheres and polygons by calculating the boundaries, then checking for overlap and filling in with greys according to viewing angle. His program works only in special cases, but is interesting for its historical position; it was one of the earliest half-tone curvature systems.

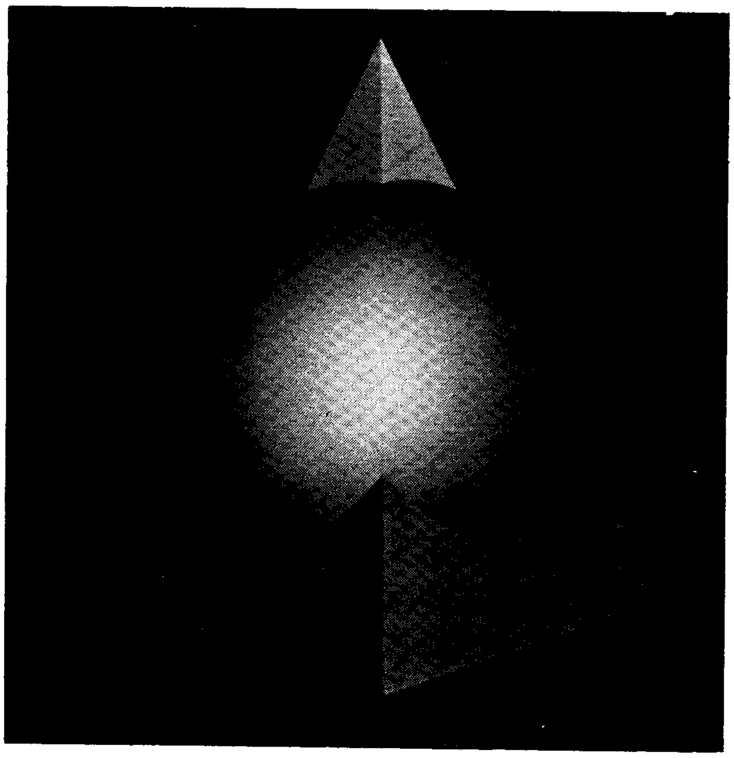

HAVE A BALL WITH DON LEE.

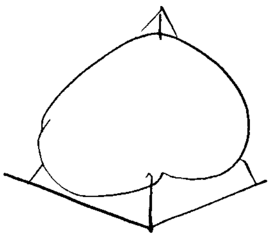

His program first works out the general outlines.

Then fills in curvaceous shading.

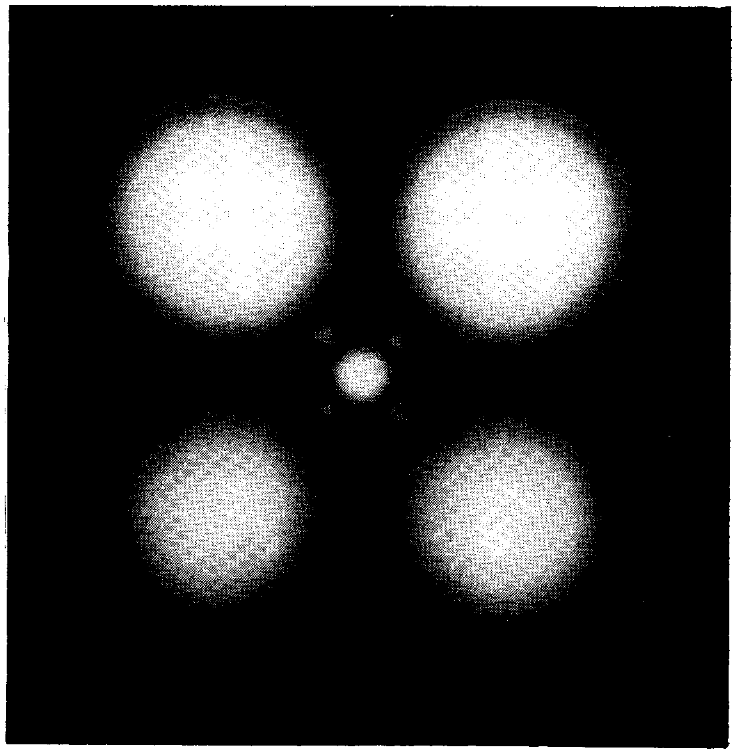

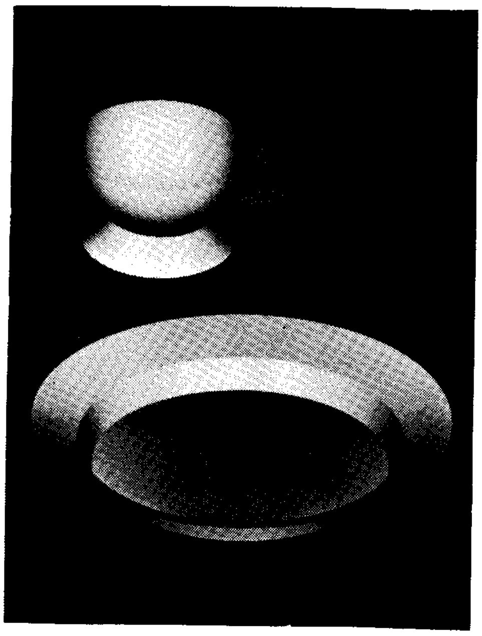

SIMPLEX CURVATURE SYSTEMS: MAHL & MAGI

A fundamental type of system we may call the "simplex" system was exemplified in the previous article by the Wylie-Romney program. A simplex technique simply projects simulated rays toward the scene from the vantage point till they hit the represented objects, and fills corresponding positions on the output picture with the colors encountered on the front surfaces of objects in the scene.

The same principle extends naturally to scenes with curved and otherwise embellished objects.

Robert Mahl, at the University of Utah, has recently reported his results with simplex methods using quadric surfaces-- those curved surfaces generated by mathematical powers of two. His pictures-- like the cup and saucer shown here-- have a pleasing 1920s Bauhaus-like quality.

One problem with this method is that computational complexity increases rapidly as the scenes grow more complex; the more surfaces and piercing-points, the more time-consuming (and expensive) it becomes to make the picture.

MAHL'S SIMPLEX METHOD

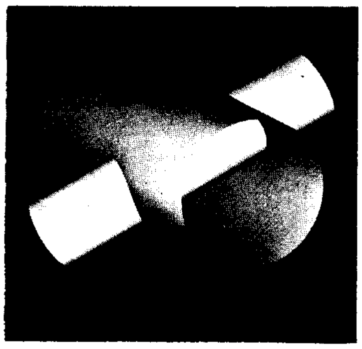

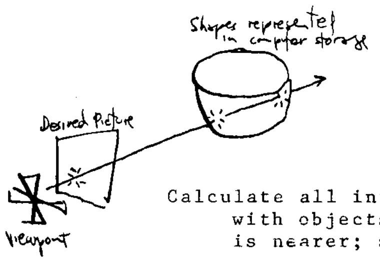

Calculate all intersections of sighting ray with objects in scene; calculate which is nearer; shade it according to angle.

tib

GENERALIZED SIMPLEX METHOD, AS EXEMPLIFIED BY MAGI SYSTEM

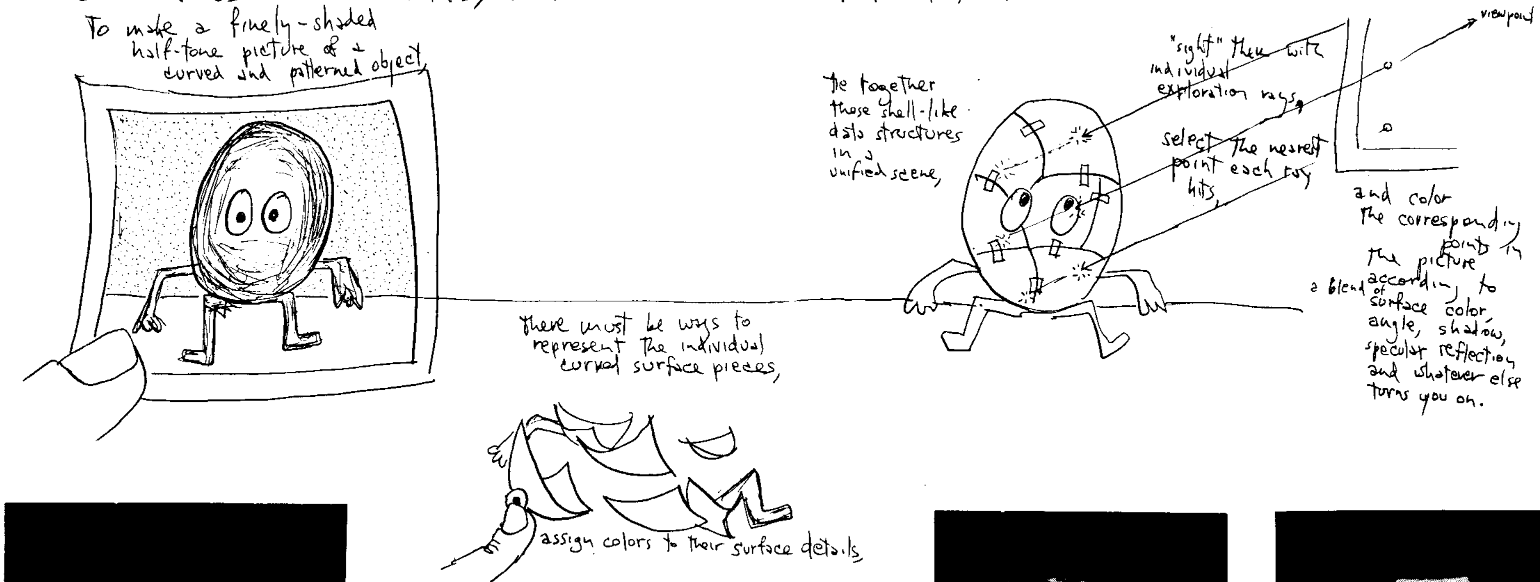

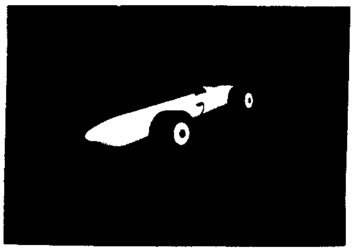

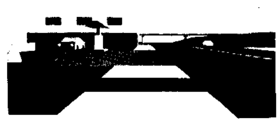

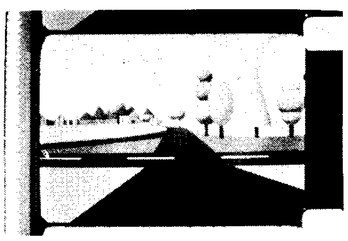

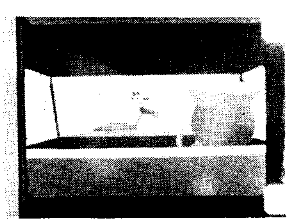

It seems, however, that Mahl's work may only be a rediscovery of what one organization worked out earlier and is being secretive about. A firm delightfully called MAGI (Mathematical Applications Group, Inc.) of Elmsford, N.Y., has extended the same idea more elaborately. They happened into the halftone game through a military contract.

MAGI's system, now thoroughly developed under Robert Goldstein, began in 1965 in a study of radiation hazards in battlefield equipment. They wrote a program to simulate paths of radiation, say, that might reach a tank driver under varous disagreeable circumstances. Having written a program that would ascertain the susceptibility to radiation of battlefield machinery, they noted that the same program could be adapted to making photographs. The progam simulated radiation; light is radiation; ipso facto, pictures. Substantially the same program would make photograph-like images, by treating the objects as opaque, and reflecting different shades according to color and angle of view.

The resulting system makes nice pictures of objects composed of planes and quadric surfaces; and includes, as will be seen from the racing car and chair, colored surface designs, shadows and spectral reflections. Not only does MAGI's software for this process produce delicately shaded pictures; if the virtual picture-plane is moved until it intersects the subject, it produces a cross-section.

MAGI runs this program remotely in Fortran on a big computer-- but they have their own minicomputer setup for photographing the results as color movies. They now offer use of this system commercially for making movies or stills.

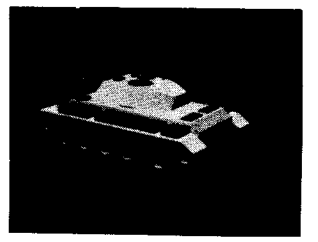

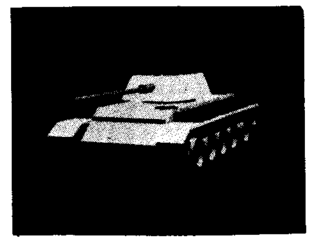

MAGI program was originally developed for study of radiation hazards inside military armor; the pseudo-photographic techniques were a side effect of the approach chosen. Who knows, these tanks may be the ones studied.

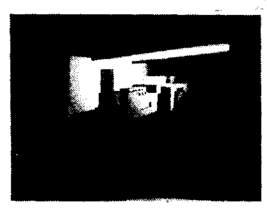

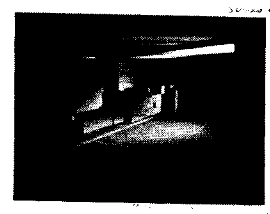

MAGI techniques were used to study alternative ways of lighting mines.

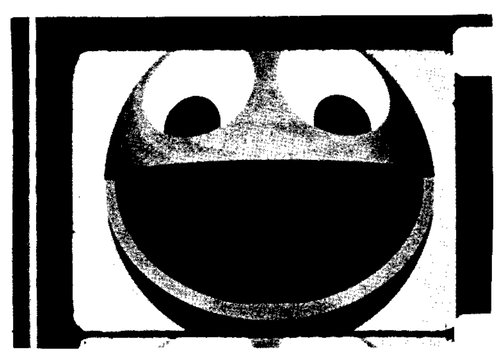

Enlargement from MAGI film. I hope the reproduction shows the concentric rings, called Mach bands, that divide areas of shading; Knowlton and Harmon (citation p. DM 10) advise on pseudo-random techniques for correcting this.

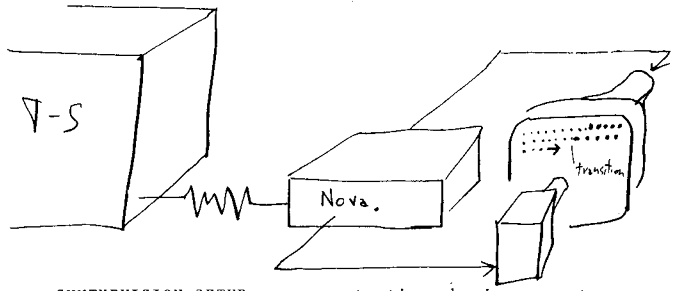

SYNTHEVISION SETUP uses remote time-sharing computer, running big secret Fortran program and containing entire data structure of three-dimensional scenes. Minicomputer photographic setup is on premises at Computer Visuals, Inc., MAGI subsidiary marketing the Synthevision service.

Local setup uses Nova minicomputer controlling both CRT display and camera. Informed guess would suggest that time-sharing system does not send all successive points of output line, but difference and transition values; Nova program would then interpolate gradations in relatively quiet sections of the scan-line.

MAGI's precise system is secret. However, the only real questions boil down to: forms of surface representation; systems of scene sorting; and method of scene scanning to produce output scan.

Note that one of the most impressive things about MAGI work, at least for sophisticates, is the degree of artistic control that seems to have been realized in their input and revision systems. It seems they offer excellent control over motion and color, and, of course, revision of the action in a scene till the maker is satisfied.

Popular Science, I think it was, had a spread on Synthevision in fall of 73.

An early MAGI character.

ROUNDUP

These have been some of the highlights of the halftone game to date. The methods described so far are mainly software-oriented, and for the most part work most efficiently as programs. In the next article we will look at some outlandish new forms of equipment, under construction or proposed, for dedicated production of 3-D halftone pictures.

THIRD ARTICLE. Specialized hardware systems.

SPECIAL EQUIPMENT IS NOW BEING BUILT FOR MAKING "REALISTIC" HALFTONE PICTURES BY COMPUTER. THIS ARTICLE COVERS SOME OF THE MORE UNUSUAL HALFTONE HARDWARE SYSTEMS NOW IN EXISTENCE OR BEING PLANNED.

HARDENING OF THE ARTISTRIES

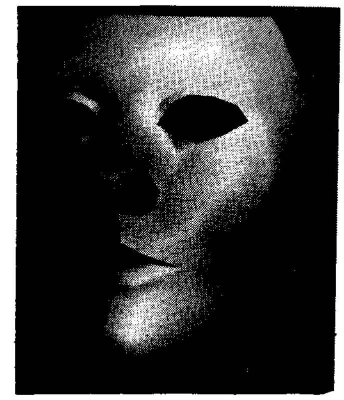

Results of Gouraud's swell smoothing technique. Mme. Gouraud posed for the data structure on the left, a system of interconnected flat polygons. The Gouraud process (see box below) created the smooth-looking face from it by an extremely simple process. (Note that the power of the technique is in the use of a simple polygon data structure, rather than the more difficult truly-curved surfaces used, e.g., by MAGI.) (Note also that the edges remain jagged.)

HARD TIMES A'COMIN'.

In two previous articles we have summarized some of the important basic techniques in computer halftone-- the artificial construction by computer of photographic pictures of 3-D scenes, scenes which are represented within the computer as colored or shaded surfaces placed in a coordinate system of three dimensions.

The techniques we have looked at were all intuitively "spatial" in character, having to do with the analysis of sight-lines and relative edge positions, and suited to implementation in computer software. Now we turn to some more advanced and peculiar techniques and equipment intended to make 3-D computer halftone faster to use, or more realistic, or easier to work with, or cheaper. These systems represent a coming generation of halftone hardware.

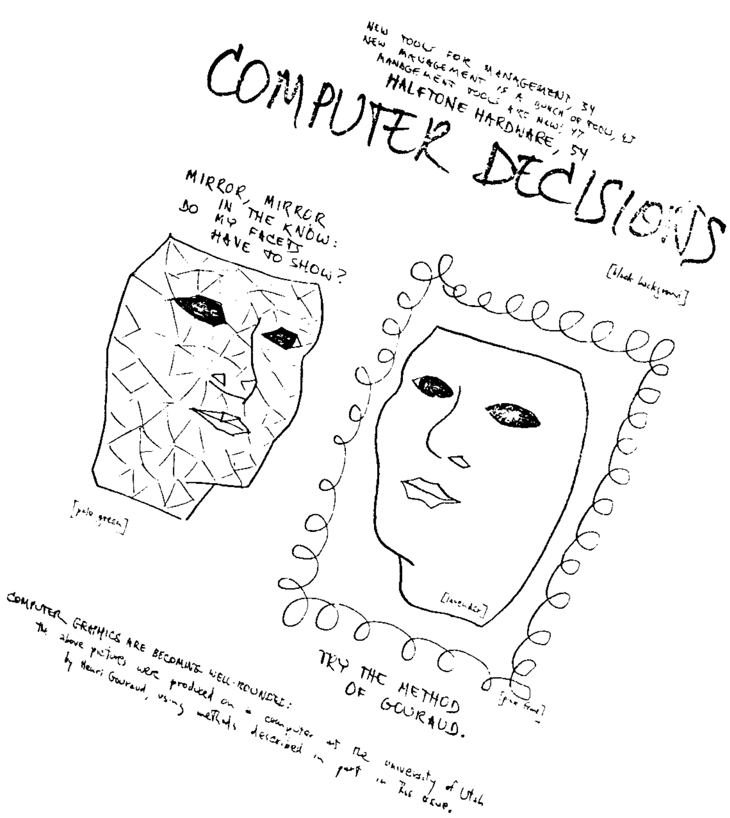

I suggested this cover for this article. The folks at Computer Decisions reacted with puzzlement if not dismay. "This cover doesn't have practical applications for the average user," I think someone said.

THE WATKINS BOX

The University of Utah is now building what wil be for some time the world's most spectacular interactive computer display, the Watkins Box. This device, interfacing between a computer and a television screen, will carry out the Watkins algorithm (described in the first article of this series) in real time: ripping through a predigested list of facet information, the Watkins Box will create on the screen an image of an opaque object which the user can rotate or see manipulated by program.

The Watkins Box can operate in two modes: normal mode, in which the object appears faceted, and Gouraud mode, in which it appears to be curved over (see masks, nearby).

The Gouraud algorithm, developed by a graduate student of that name, is a ridiculously simple technique which marries perfectly to the Watkins method. Instead of shading the facets uniformly, this technique calculates a shade of gray for each point. In effect the method interpolates the shade of the point from those around it, across facet boundaries. In actual procedure, the Gouraud method shades a point by linear interpolation between two edge-colors: the color of the last edge and the next edge to be encountered on the present scan-line. (These shades are in turn found by linear interpolation between their endpoints.)

It will be noted that Gouraud's method does not curve the edges. But considering its simplicity as a small addition to the Watkins box, that's no great sacrifice.

Naturally, the Watkins Box will not reach the private home for several years; current likely price is in six figures. But that's now.

⇨ Can it be an accident that this curvaceous system was worked out by a Frenchman?

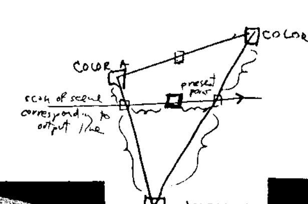

GOURAUD'S TWIST adds the appearance of curvature to a faceted object shown opaquely by the Watkins method (described in first article).

Instead of shading each point within a facet with the same color, interpolate between the vertex-colors according to how far down the edges you've gotten. Note that the jagged edges are retained.

"Wire frame" of Old-Fashioned 3D computer graphics

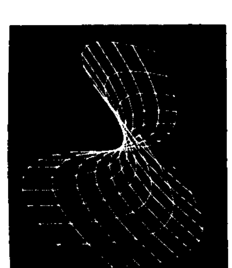

as shaded by Watkins (or other) method

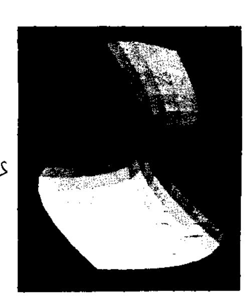

... and as Gouraud makes it look curved.

GOURAUD'S SPECIAL TWIST

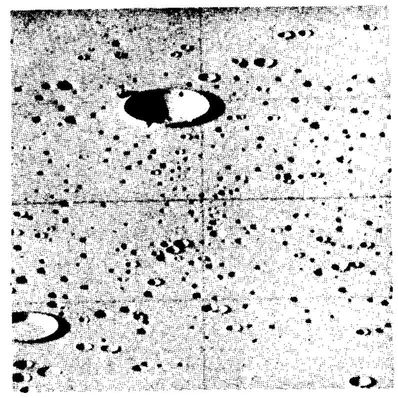

PRA'S WORLD-VIEW

Roger Boyell, of Pennsylvania Research Associates, Philadelphia, likes to refer to the company's main interest as "modelling the physical world." Thus he and his associates have developed systems for cartography, landscape modelling, pipe design, and simulation of complex radar systems.

A radar simulator they are putting together for the Navy will show the results of any possible radar system moving over any possible terrain. A pilot or navigator trainee, in a simulated cockpit, will see the mission's changing radar picture as he changes the plane's course or the radar's tuning. The radar picture, appearing on a screen and changing in real time, will look just the way the radar would look on a real mission-- flying in perspective among mountains or valleys, high or low, at any bearing and speed, and viewed through any type of radar.

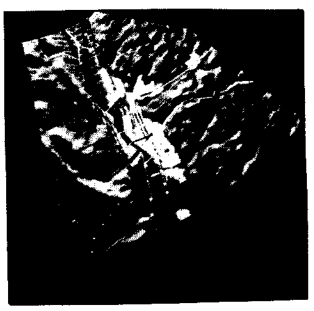

Boyell's approach is to treat each component of the pictorial/radar simulation as a separate problem, to be handled in different ways, and blended in a final buffer, a core memory which is read out to television. Separate mechanisms supply components of shadow, specular reflection, coloration and randomizing effects. The core buffer continuously refreshes the scanned CRT display.

Boyell has put the same techniques to work making simulated halftone pictures of the moon (see cut). Both the radar and moon systems use the same type of halftone image synthesis, even though superficially they seem quite different. But radar is radiation, just like light, and Boyell's techniques of three-dimensional modelling and search apply equally well to depiction by reflected visible light-- i.e., halftone images.

BOYELL'S TERRARIUM fills a fast core-memory buffer with a TV image constantly being read out (much like the Knowlton-Schwartz setup: see pp. DM10 and DM24, top Schwartz picture) and changes individual features one-at-a-time to match a changing view.

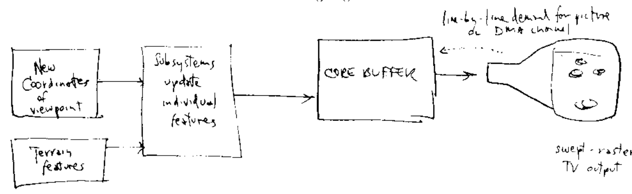

An outfit called HUMRRO, in Washington, say they have a real-time interactive half-tone that will knock several people out of the ballpark-- especially the GE hardware and the Evans and Sutherland Watkins Box (earlier).

The HUMRRO system is intended to go out to color screens (modified Sony Trinitrons) with shaded pol ygon halftone, offering pseudo-curved shading like Gouraud's (see earlier).

FLASH!

The techniques were worked out by Ron Swallow, and they're not telling about how they work. It is claimed, however, that their real-time picture generator handles scenes with 16,000 edges, and that this will cost \$150,000 and service 16 (or was it 64) user terminals simultaneously.

It may have been a bad phone connection, or this may be what they're really claiming. Obviously it'll be really great if it turns out to be real.

Evidently they have in mind the use of such high-performance scopes for teaching, allowing students to explore intricate three-dimensional scenes or objects. Terrific.

(Note: compare the claim of 16,000 edges on a \$150,000 system with the 2000 (?) edges allowed by the old NASA system built by GE, or the Watkins Box-- I don't know how many edges-- at \$500,000 from Evans and Sutherland.)

THE SHAPE OF THINGS TO COME

If these systems sound far-fetched, or only for theoretical investigtation, consider this: the Air Force is now letting contracts for an advanced flight-training simulator that is a small boy's dream. To be situated in Dry Lake, Arizona, the simulator will have the most realistic cockpits ever built: the entire mockup will turn and tilt in response to the user, and the seats will even swell and deflate, to simulate acceleration and weightlessness. The cockpits alone, without the visual display screens, will cost ten million dollars each.

But the visual systems-- ah. The pilot-user will look out into an artificial world, among whose mountains and meadows and clouds he will fly in real time. Six CRTs, arranged as parts of a dodecahedron in an entire visual surround, will show him the changing terrain and flying environment. Each of these CRTs will be driven by a real-time perspective halftone simulator, with all displays spliced together and driven by a master simulator responding to his actions. Who will build them is not yet decided; they could be Warnock or GE boxes.

The sheer joy of such a system will be hard to beat. But no doubt others will be on the way-- perhaps at the amusement-park level.

AIR FORCE SUPERTOY

The new pilot trainer will not only swing and dip in response to the controls; on six giant CRTs, with optics in front that focus the eye on infinity and connected at the seams, the pilot will see a responding perspective simulation of the world he is flying through, planes he is dogfighting with, and who knows-- witches? Superman?

NELSON'S FANTASM(TM): A LOT OF BOSCH?

I don't expect you to believe this, because not even my patent attorney does, but the system I call Fantasm is intended to make pictures that pass the Turing-test: you won't be able to tell them from real photographs. Fantasm is intended to allow the user to make realistic, Hieronymus Bosch-like photographs and movies, with real-looking people (and scenery, imaginary characters, monsters, etc.) in scenes of arbitrary complexity. It is expected that 1975 economics will make its construction feasible.

Fantasm I originally conceived as a method of making realistic photographs and movies, not knowing at the time that this was impossible, but feeling it could be done somehow if the problem were broken down sufficiently. At times it was not clear which of us would be broken down first, I or it.

It occurred to me sometime in 1960-1 that computer-interpolated, Disney-type cartooning methods would be feasible. After some thought I realized that pseudo-photography would be possible, and dropped the cartooning idea. The strange behavior of people whom I told about this led me to increasing secrecy.

The general goal was to make a system that could do realistic movies without scenery or actors, and make pictures indistinguishable from real photographs of real scenery and actors. ("What do you mean, indistinguishable from photographs?" people keep asking. What do they mean what do I mean?) The surfaces are to be put in by "sculptors," animated by "puppeteers," and photographed by a "director." The objective is for moviemaking to be under the utter imaginative control of the creative user.

I am indebted to Prof. Charles Strauss for the formalization of my smoothing-function.

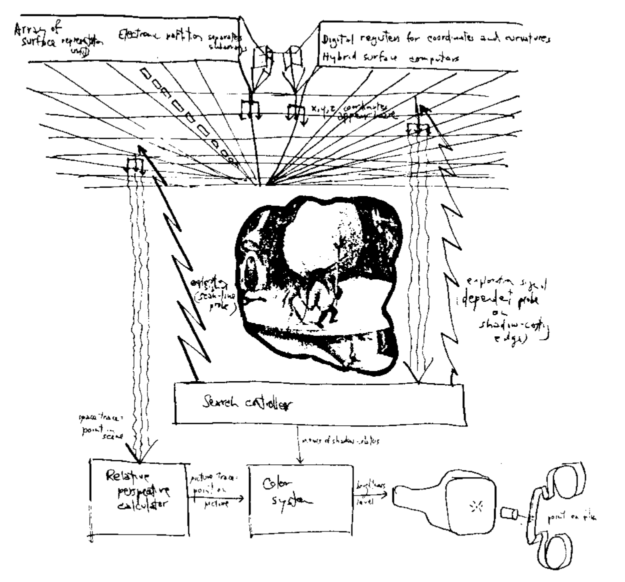

FANTASM AT LAST PARTIALLY REVEALED,

at least to certain readers.

A scene of arbitrary curvature and topology is represented in a system of holding registers; the surface is presented (through D-to-A converters and an array parallel function generator) to interrogating circuitry which steers an inquiring signal around the represented surfaces. Operation is empirical. Array has partition logic allowing simultaneous queries of various sub-surfaces. Feedback steering circuitry allows multiple loops through array. Steering signal and returned surface parameter are analog and continuous. List techniques manage shadow and visibility 'umbrellas' (surfaces of occulted volumes or umbras).

The Fantasm Scene Machine(TM), the representation and search array, is one chip repeated in a carpet. Large-scale-integration permits the required digital storage of about 500 bits per surface section plus analog circuitry and switching logic. Patent work underway.

SUMMARY: outlines handled by Perimeter Parameter Occultation Chasing, fill-in by Bullet Search, animation continuity management by list-processing techniques.

The system could come in a number of different versions. One of these involves a large array of LSI computing modules (the checkerboard Scene Machine) to be guided by special hardware under an unusual monitor running on a general-purpose computer. The checkerboard Scene Machine holds a great spread of surface data. It is a logical curiosity, an array that replies as a unit, ignoring cell boundaries, to electrical explorations of the shapes represented in it. The resulting trace makes various 3-space explorations on the faces, mountains or automobiles spreadeagled in it. Think of its trace as a radio-controlled firefly skating over a bumpy checkerboard. Using this machine, and various cat's-cradle list structures based on the geometry of light around odd volumes of occultation, the problem of halftone analysis of arbitrary shapes is solved by brute force rather than analytically. A variety of other processes have also been defined in the system for other types of graphic application.

As far as I have been able to learn, Fantasm is the most baroque computer graphic system anyone has proposed. It is not intended to operate in real time, but rather take as long as it needs, or as long as the user wants to pay for, to fill in complex visual details, shadow, reflections, curlicues, leaves, hair, etc. It is best suited to the production in Panavision of Busby Berkeley musicals, or "The Lord of the Rings" with realistic wraiths and interspecies battles. But it may well cost too much to use for that. Indeed, its economics seem to improve in low-budget settings like videotape, although there its output bandwidth will flower unseen. But the Scene Machine should also be useful for more mundane applications, such as contour mapping, automobile design, advertising photography and medical illustration.

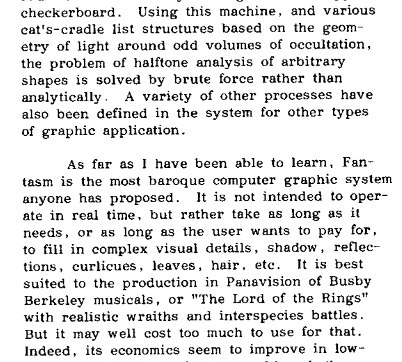

FOURTH ARTICLE.

Systems of Computer Image Corporation.

DM39

COMPUTER IMAGE'S MAD WHIRL

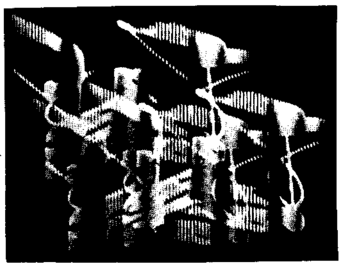

Scanimate's twirl, by now familiar to most TV watchers. Scanimate is extensively used on "The Electric Company."

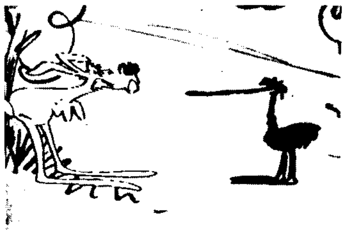

CAESAR System. Characters are made to move jaws and lips by jointing technique similar to Animac II (below), but in such a way as to matte over drawn artwork-- meantime wiggling other drawn artwork through scan manipulation.

SO FAR WE HAVE SUMMARIZED AND DISTINGUISHED AMONG THE MAJOR TECHNIQUES FOR COMPUTER SYNTHESIS OF IMAGES FROM DIGITALLY STORED REPRESENTATIONS OF SCENES. WE NOW TAKE THE WRAPS FROM A DIFFERENT BUT RELATED SET OF TECHNIQUES-- THE SYSTEMS OF COMPUTER IMAGE CORPORATION.

Lee Harrison III got the idea for what is now Computer Image Corporation in 1959. Already having an art degree, he went on for a degree in electrical engineering, and through long lean years put together the technical basics around which CI's systems are now built. Computer Image Corporation is now a going concern, and output from their systems, especially Scanimate, is now widely visible on television.

Computer Image Corporation seems to be the first firm to be commercially successful in the halftone field. Whether they should be included with the others is arguable, however. Their systems are not widely understood, and the relation of these systems to the other systems and programs described in these articles is problematical. Among the few who understand their techniques, some argue that they do not synthesize images at all, but rather twist preexisting pictures with a sort of Moog synthesizer, and that their analog techniques are really just compound oscillators rather than true computing. I think that this view is wrong, at least as regards their most ambitious system, and that CI's techniques deserve review. All the world is not digital. CI systems do fill up areas with grey-scale (and other) pictures, and their systems involve three-dimensional coordinates, occultation and coloration; thus I think it appropriate to discuss them here.

The following discussion is the first, I believe, to lift the veil of secrecy that has hitherto confounded observers of this company's work. In the light of the extreme sophistication with which they have pursued extremely strange techniques, they should benefit from the wider understanding. (Note that this material, which has been assembled from various sources and careful TV watching, is partly conjectural.)

Computer Image's systems represent an apparently unpromising approach brilliantly followed through.

All of CI's systems are a strange combination of closed-circuit TV and analog components out of a music synthesizer: oscillators, potentiometers, interconnection networks. The basic mechanisms are the same for all, but they are carried to different logical extremes, with differing accoutrements, in the four systems. They all seem to be based on the extraordinary Animac II, not yet implemented; it would seem that for business reasons the company decided to raise money promoting simpler systems, so its bread and butter now consists of two less ambitious systems, Scanimate and Animac I; both of which might be puzzling if not recognized as parts of a more elegant whole. It would seem they were designed backwards as spinoffs from Animac II, as was CAESAR, their more recent 2-D system.

The extraordinary ramifications and varieties of this system, with all its electronic add-on and composite methods, stagger the most jaded technical imagination.

At the heart of the CI systems is the principle of filling areas of a CRT screen with an oscillating trace. This is a principle common to both Lissajous figures and television; but Computer Image has elaborated it peculiarly. By variations they paint twisted television images, wiggle sections of superimposed drawings, create moving filigree effects, and hope to animate whole groups of opaque electronic puppets in 3-space.

Consider an oscillating trace on an oscilloscope. This is a two-dimensional oscillation, having two signals, x and y. But a three-dimensional oscillation is also possible; any third signal, z, can be interpreted as a third dimension, meaning that a "point of light" is whirling out some pattern in a three-dimensional space-- an oscillotank, so to speak. Let us call this point moving in three dimensions a "space trace."

Now to view this trace we need to cut it down to two dimensions. By ignoring one of the traces we can view the oscillotank in certain fixed ways; but by creating a "view calculator," a box performing certain perspective transformations on the three signals of the space trace, we may obtain a view of the oscillotank from a movable vantage point. This is an x-y view which we may put on an ordinary oscilloscope.

Let us now add one more signal, b (for brightness). This is the brightness signal familiar in television.

Brightness of the spot is thus independent of the movement of the space trace. For example, the space trace could describe a helical path, a sort of tornado motion, and we could time its spinning to phase with a TV signal. If we now brighten the space trace only with the brightness signal of a TV pickup, we now will see (in our view of the oscillotank) what would look like a TV picture curled around itself in space.

The different CI systems are built around this effect.

Output from all these signals is ordinarily picked up by another vidicon, which stabilizes it by converting it into conventional television imagery.

THE WHIRLING UNIVERSE OF COMPUTER IMAGE CORPORATION.

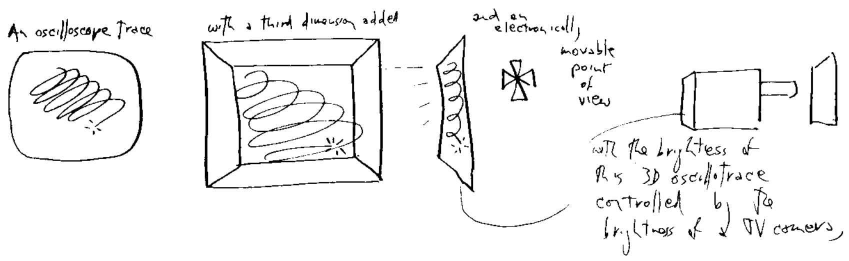

gives us a window into a peculiar sort of world: a world in which luminous shapes can undulate and spin on invisible spindles (Scanimate), or wiggle as separate bones (CAESAR).

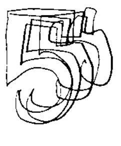

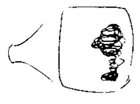

Tubelike shapes may be rotated and shaped in 3D (Animac), and puppets may eventually be rolled like cigarettes (Animac II), which may then be painted from a TV pickup on the side nearest the viewpoint.

By using a storage tube and spinning the trace close together, like cotton candy, and cutting off the painting signal while the trace is within the area already filled, we get electronic masking: which blends animated drawings in 2D (CAESAR) and may eventually manage shadows and occultation masking among 3D puppets (Animac II).

SHAPING METHOD

Lissajo_us and zigzag figures are rapidly spun in three dimensions -- that is, varying voltages x, y and z. The resulting "tubes" and "curtains" are then viewed by perspective calculation. The circuitry permits these shapes to flex at joints, wave, and go through other changes.

IN SCANIMATE: zigzag and curling shapes define a moving scroll on which an image is painted.

IN CAESAR: curling shapes are treated 2-dimensionally, as blocking controls for artwork.

IN ANIMAC II: puppets will be sculpted much like rolling a cigarette.

BLOCKING METHOD

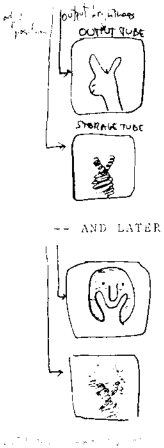

While showing nearer part--

spin out same part on storage scope; continuously test potential of screen;

-- AND LATER, IN THE SAME FRAME:

cut off the brightness of the output signal, wherever there is already an image on the screen.

OCCULTED COMPOSITE

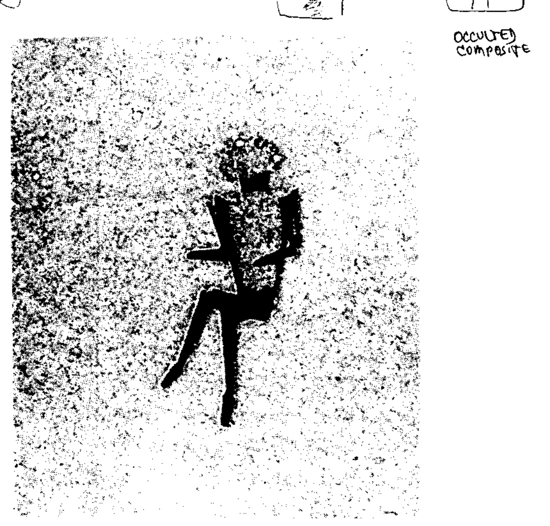

The only picture I've been able to find that relates to the 3D sculpturing of Animac II is this frame, blown up from a short 16mm sequence. The figure is sculptured from oscillations in three variables, modulated to represent this figure of thirteen sections or "bones." Head and torso are clearly visible in the film; the figure is seen to spin as if in an ejection seat.

A last CI technique, technically minor but remarkable in effect, permits this blocking and shadowing among separate objects. This is the use of a storage CRT tube on which every frame is painted (from the viewpoint or from the light source). The picture is painted on the storage CRT, nearest things first; and the return signal from the screen tells whether the space trace is crossing an area already painted during the frame. The tube's output signal then effectively constitutes a silhouette. This clue indicates that the space trace should not be visible; and hence is used to cut off brightness while the trace is within the already-filled area. This gates between two desired objects or pictures, foreground and background. If operated from the point of view of the light, it gates shadow: the signal is used to control the relative brightness of the shadowed and unshadowed features of a puppet in 3-space.

A fascinating variety of embellishments has been put into these systems by CI's ingenious engineers. Coloration of the final video signal is added by gating color levels under control of the brightness signal, permitting pictures with several grey-levels to be transformed to up to four rainbow hues. Separate shapes described by the space trace may be independently moved and jointed at the same time: Harrison pointedly calls such separate shapes "bones." Darkening at the backside of a spun shape, or brightening at edges of a painted portion, and brightening in proportion to curl, are all strange capabilities of this machine. Lip-synchronized mouthlike motion can be imparted to any part of the shape spun by the space trace (whether or not a mouth is painted on it), by an audio detector feeding directly to the circuitry from a live mike. And the limbs of CI's ghostly figures can be made to swing by connection of sensors to the animators themselves-- in a living pantograph.

SCANIMATE is a popular device now widely used (at CI's studios) for the making of TV commercials and station-break emblems. This is their simplest system, used for the conversion and discombobulation of flat artwork. In Scanimate, the space trace is controlled by hand-operated potentiometers. Two separate oscillator settings are available, so that the space trace can have two separate oscillation patterns, spinning out two entirely different virtual shapes in 3-space. A hand-throttle eases from one oscillator setting to the other. This permits an image to be moved, shrunk or enlarged, or flipped; to go from whirling around to a sort of hula; and many more effects. The picture painted on it may be seen to roll on invisible spindles, bloom into fountains, or undulate as pennants-- all by modulating the brightness of the flying spot as it traces its unseen shape. This shape, in turn, can move between its two forms under control of the throttle.

Animac I (usually called Animac) provides greater flexibility in controlling the space trace. The system's oscillations are controlled by an input vidicon, which artists may quickly modify with pastel check at the pickup. Ghostly tubular lettering, swarming pendulum-patterns and jiggling filigrees are among the possible doodles.

CAESAR, their newest system, is oriented toward Yogi Bear-type animation. The artist's cartoons are automatically superimposed on a background or each other. They may be moved, and made to wiggle under real-time control by the user.

But it is to Animac II that these curiosities lead. What Harrison calls the "Snow White Capability" of Animac II will permit the sculpture of full humanoid puppets, with perhaps thirty articulated "bones," opaque to one another and casting shadows, colored, moving and talking.

Two young fellas in a Manhattan loft, Messrs. Rutt and Etra, are offering a machine similar to Scanimate but much cheaper.

(Rutt Electrophysics, 21-29 W. Fourth St., NY 10012.)

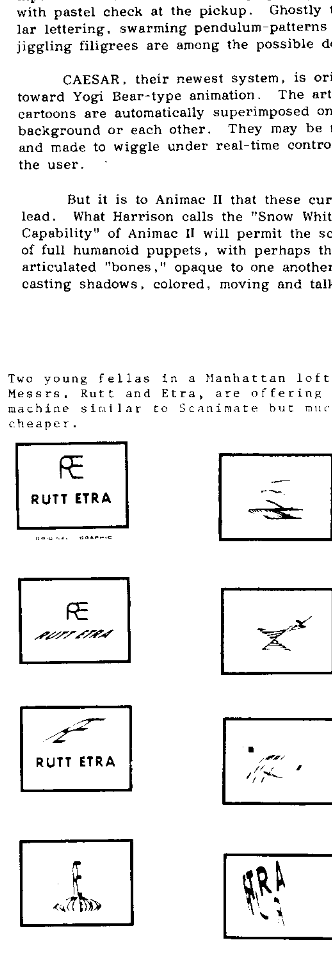

It's not as finely detailed-- the inner screen runs at 525 lines rather than 700-- but it costs some $15,000 instead of $150,000.

WHAT ABOUT REAL THREE-DIMENSIONAL DISPLAY?

In science-fiction stories you hear about how objects are made to appear as if they're standing in the middle of the room. For instance, I believe that in Heinlein's Stranger in a Strange Land they watched a "tank" in which things appeared.

Well, a lot of people have thought about this, and it's not so easy as you might think.

One interesting scheme used a sort of translucent propellor, spinning rather fast, on which computer-generated images were projected from below. It was done by the dotting method, so that a bright dot of light would appear high or low in space depending on whether it was projected on a relatively high or low point on the propellor.

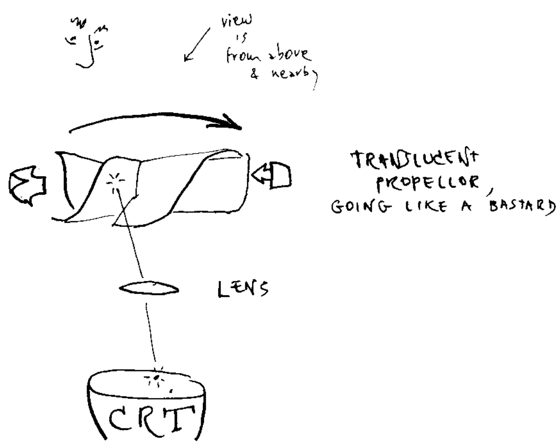

This was interesting but had numerous disadvantages-- not the least of which was the danger of the thing flying apart. (Translucent materials tend not to be as strong as, say, metal.) Another basic problem, though, was the fact that any given point in the space could only be displayed at a given time, when the propellor's height in that region was just right, and that meant that at that given instant you couldn't display any of the other points that could only be displayed at that instant. A considerable disadvantage.

Probably the most astonishing 3D display is Sutherland's Incredible Helmet. This consists of a helmet with two dinky CRTs mounted on it, each being driven in real time by a perspective system (such as the LDS-1) and set up with prisms to the wearer's eyes. Through the prisms the wearer can see the real world in front of him. Reflected in the prisms, however, and thus mixed into the view of the real world, is the glowing wire-frame being presented to him-- in perspective, and with its separate views merging into an apparent object in front of him. But he need not stand still: as he moves, the helmet's changing position is monitored by the program, and the display system changes the views accordingly meaning he can walk around and through a displayed object. The illusion, and the possibilities, are fantastic: imaginary architecture, explanations and diagrams of things in the room, poetry that changes as you walk through it, ... well, you work on it. Not available commercially.

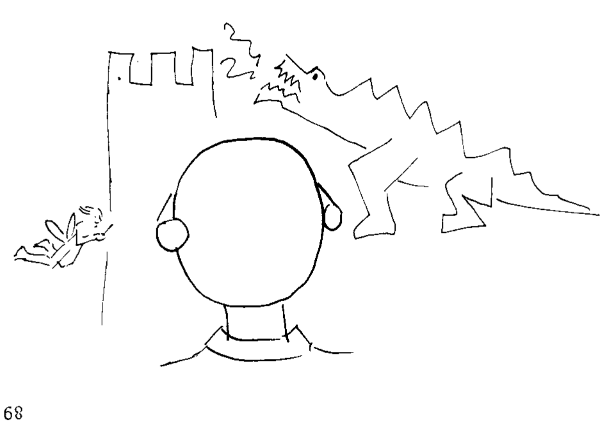

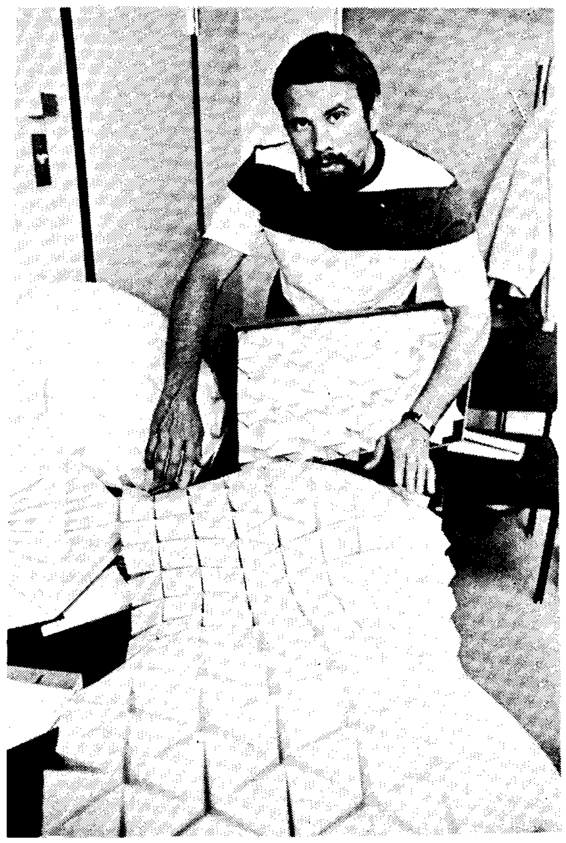

U.U.

RETURN TO THE FOLD

There was a lot to be said for tents. They could be made by tailors, rather than construction gangs; they could be transported and stored flat. Their surface-to-volume ratios couldn't be beat.

Noting this, an architect named Ron Resch said to himself: what about making large-scale foldable structures, likeunto geodesic domes, that could be simply manufactured in sheet form and creased at the factory, then bolted and cabled and strutted in the field?

Resch has now for years been experimenting with complex folded structures.

There's only one trouble. If you've messed with paper airplanes you know that folding is an inaccurate process, and so the prospect of discovering complex geometric structures by the hand-folding of paper is rather slim.

Recognizing this, Resch has contrived to work at a computer display. His work-- the search for great folding structures-- is one of the first practical uses of halftone polygon computer graphics. He is, naturally, at the University of Utah.

Lou Katz, of NYU, put old-fashioned stereopticons up to the CRT, and displayed two separate views to the two eyes. Works fine, even with isometric display.

Bob Spinrad of Xerox Data Systems has a patent on displaying 3D from a computer through an ordinary color TV. Assuming you're using some standard way of refreshing the TV-- described elsewhere-- the image for one eye is displayed in green, the other in red, and you look through red and green glasses. The wonders of modern science. Spinrad chuckles over it himself.

Another scheme glued silver Mylar to the front of a loudspeaker, then played a soft hum through the loudspeaker to pulse the Mylar back and forth. Then you used that as a mirror to look at what was going on the CRT-- which was showing a lot of points at odd places that would appear to be in space. Unfortunately this was hard to coordinate, and, like the propellor, often required you to put dots in several places at once, which don't work.

For a while you could get-- maybe you still can-- a three-dimensional computer output device. Here's what it did: it created objects showing data structures that had three variables. (It didn't make wire-frame objects or the like.) Automatically ejecting wire through a styrofoam block, and snipping the done ones, it created little mountains showing three-dimensional data. Very cute. Since many people have problems with mountainous computer data, it probably should have caught on.

Then a lot of people mumble the word "holography," as if that is going to settle something. While holograms are terrific and remarkable, and have been produced on computers, making them is not a process that can be carried out decently on sequential machines-- let alone making them in real time. So if a solution to interactive three-dimensional computer display is going to come through holography, it means a whole new batch of technology will have to be invented.

My friend Andrew J. Singer, who comes and goes in the computer field and is one of the five or six smartest people I ever met, says he knows how to build a display tank, and I believe him. He explained it quickly to me once and I asked him to tell it again, but he just said sadly, "What's the use-- there are so many great things that could be done..."

FOUR DIMENSIONS, EGAD

So much for three dimensions. Now, some readers are bound to ask, "What about four dimensions?" because they are science-fiction fans or troublemakers or mathematicians or something.

Just as we can make a two-dimensional picture of a three-dimensional object, it is possible, dear reader, to make a two-dimensional picture of a four-dimensional object.

What is a four-dimensional object?

Why, any object that has four dimensions, (thanks a lot, you say), or even four measurable qualities, such as height, weight, age and grade point average. Well, let's not get into that, but it turns out that views of such multidimensional structures may be obtained by the same homogeneous matrix techniques already mentioned for regular perspective calculations. Rule of thumb: however many dimensions your data has originally, you add one more dimension, homogeneous with the rest, and there exist formulas (sorry, I don't have them) for view calculation.

(Note, of course, that while a two-dimensional view is a picture, a three-dimensional view is a three-dimensional object-- you'll have to view it on an interactive 3D computer display of some kind.)

End of *THE MIND'S EYE*.

THE CIRCLE GRAPHICS HABITAT

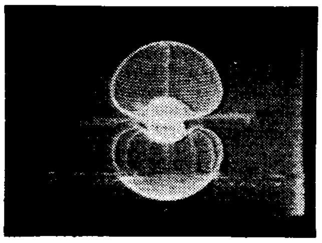

From videotape, "The Hydrogen Atom According to Quantum Mechanics" by T.J. O'Donnell & David Parrish.

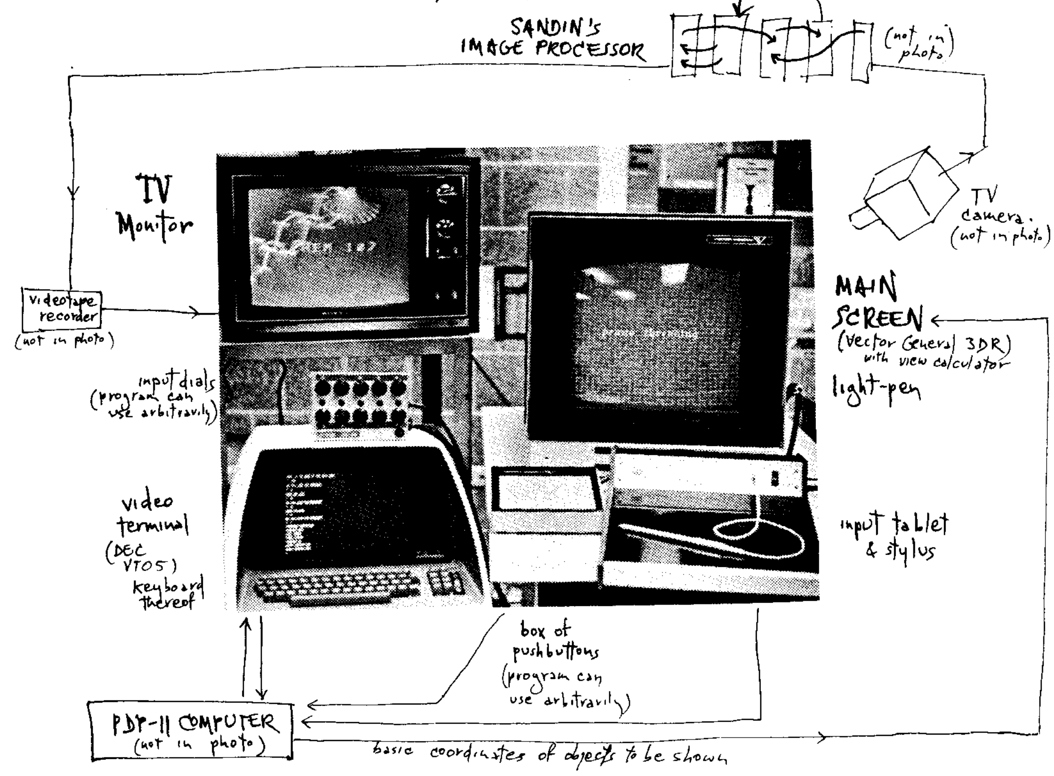

It is usually hard to combine things: especially complicated technical things. Usually it takes infinite reconsiderations, finagling, modification, intertwingling.

The Circle Graphics Habitat, however, is something else again. It results from two intricate, independent technological developments, each an intricate system carefully crafted by an exceptionally talented person, coming together like two hands clapping. Like ham and eggs, like man and woman, Sandin's Image Processor and DeFanti's GRASS language conjoin directly and interact perfectly as if they had been made for each other, which they were not.

Dan Sandin's Image Processor (see p. DM8) is a system of circuit boxes that allow video images to be dynamically colored, matted, dissolved and palpitated; Tom DeFanti's language (see "Coup de GRASS, p. DM 31) permits the rapid creation, viewing and manipulation of three-dimensional objects on the screen of a particular computer setup.

To combine them, you just point Dan's system at Tom's system.

Let's say that on the screen of Tom's system we are viewing an animated bird, flapping its wings. Since it's being shown on a three-dimensional refreshed line display (see pp. DM22-3, DM30), it appears only as white lines on a dark screen.

Dan merely points a TV camera at Tom's screen, and runs the TV signal into his Image Processor. Now, in the Image Processor, he gives it the magic of color. Different colors, interplaying with gradations and subtlety.

From the Image Processor, the finished signal goes out to videotape recorders.

What then have we overall? One of the world's most flexible facilities for the rapid production of educational videotapes.

To explain something, you create a three-dimensional stick-figure "model" of it, using DeFanti's GRASS language. Then you make a videotape of it, showing rotations or other manipulations, using the Image Processor to give it color.

DeFanti and Sandin have spent much of the academic year '73-4 getting the kinks out of this procedure. (Many of the difficulties stem from the unreliability of videotape recorders.) Stills from some of the first work are shown here.

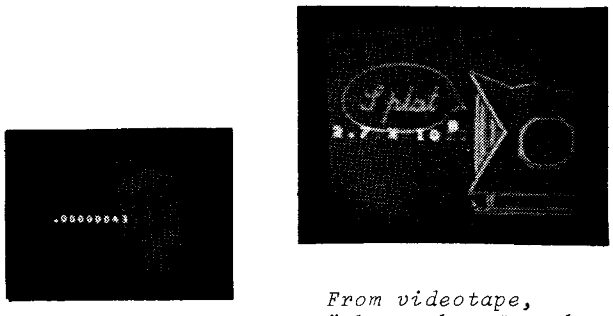

From videotape, "The Number Cruncher," by TDF & DJS.

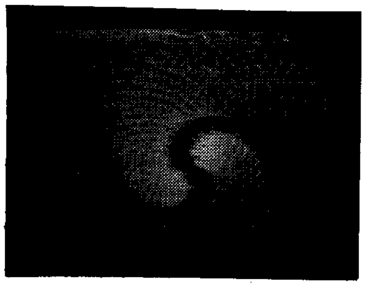

From videotape, "The Spiral Tape," by DJS and TDF.

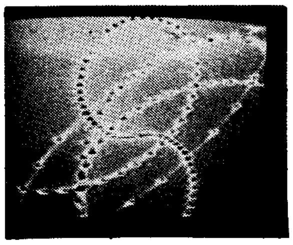

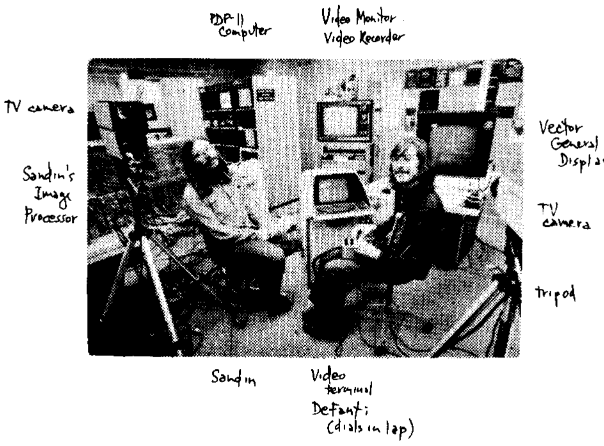

BIBLIOGRAPHY

Thomas A. DeFanti, Daniel J. Sandin and Theodor H. Nelson, "Computer Graphics as a Way of Life." To be presented at U. of Colorado computer graphics conference, July 1974; to appear in proceedings purportedly to be called Computers and Graphics.

THE TISSUE OF THOUGHT

Uneducated people typically think of education as the learning of a lot of facts and skills. While facts and skills certainly have their merits, "higher education" is also largely concerned with tying ideas together, and especially alternative structures of such tying-together: with showing you the vast uncertainties of things.

A wonderful Japanese film of the fifties was called Rasho-Mon. It depicted a specific event-- a rape-- as told by five different people. As the audience watches the five separate stories, they must try to judge what really happened.

The Rasho-Mon Principle: everything is like that. The complete truth about something is never known.

Nobody tells the complete truth, though some try. Nobody knows the complete truth. Nowhere may we find printed the complete truth. There are only different views, assertions, supposed facts that support one view or another but are disputed by disbelievers in the particular views; and so on. There are "agreed-on facts," but their meaning is often in doubt.

The great compromise of the western world is that we go by the rule: assume that we never know the final truth about anything. There are continuing Issues, Mysteries, Continuing Dialogues. What about flying saucers, "why Rome fell," was there a Passover Plot, and Did Roosevelt know Pearl Harbor would be attacked?

Outsiders find the intellectual world pompous, vague in its undecided issues, stuffy in its quotes and citations. But in a way these are the sounds of battle. The clash of theories is what many find exhilarating about the intellectual world. The Scholarly Arena is simply a Circus Maximus in which these battles are scheduled.

Many people think "science" is free from all this. These are people who do not know much about science. More and more is scientifically known, true; but it is repeatedly discovered that some scientific "knowledge" is untrue, and this problem is built into the system. The important thing about science is not that everything will be known, or that everything unanimously believed by scientists is necessarily true, but that science contains a system for seeking untruth and purging it.

This is the great tradition of western civilization. The Western World is, in an important sense, a continuing dialogue among people who have thought different things. "Scholarship" is the tradition of trying to improve, collate and resolve uncertainties. The fundamental ground rules are that no issue is ever closed, no interconnection is impossible. It all comes down to what is written, because the thoughts and minds themselves, of course, do not last. (The apparatus of citation and footnote are simply a combination of hat-tipping, go-look-if-you-don't-believe-me, and you-might-want-to-read-this-yourself.)

"Knowledge," then-- and indeed most of our civilization and what remains of those previous--is a vasty cross-tangle of ideas and evidential materials, not a pyramid of truth. So that preserving its structure, and improving its accessibility, is important to us all.

Which is one reason we need hypertexts and thinkertoys.

A lot of people are afraid to ask questions because they're afraid of looking dumb. But the dumb thing is not asking questions.

If I ever get my school, the one course taught will be

HOW TO LEARN ANYTHING

As far as I can tell, these are the techniques used by bright people who want to learn something other than by taking courses in it. It's the way Ph.D.'s pick up a second field; it's the way journalists and "geniuses" operate; it brings the general understandings of a field that children of eminent people in that field get as a birthright; it's the way anybody can learn anything, if he has the nerve.

1. DECIDE WHAT YOU WANT TO LEARN. But you can't know exactly, because of course you don't know exactly how any field is structured until you know all about it.

2. READ EVERYTHING YOU CAN ON IT, especially what you enjoy, since that way you can read more of it and faster.

3. GRAB FOR INSIGHTS. Regardless of points others are trying to make, when you recognize an insight that has meaning for you, make it your own. It may have to do with the shape of molecules, or the personality of a specific emperor, or the quirks of a Great Man in the Field. Its importance is not how central it is, but how clear and interesting and memorable to you. REMEMBER IT. Then go for another.

4. TIE INSIGHTS TOGETHER. Soon you will have your own string of insights in a field, like the string of lights around a Christmas tree.

5. CONCENTRATE ON MAGAZINES, NOT BOOKS. Magazines have far more insights per inch of text, and can be read much faster. But when a book really speaks to you, lavish attention on it.

6. FIND YOUR OWN SPECIAL TOPICS, AND PURSUE THEM.

7. GO TO CONVENTIONS. For some reason, conventions are a splendid concentrated way to learn things; talking to people helps. Don't think you have to be anybody special to go to a convention; just plunk down your money. But you have to have a handle. Calling yourself a Consultant is good; "Student" is perfectly honorable.

8. "FIND YOUR MAN." Somewhere in the world is someone who will answer your questions extraordinarily well. If you find him, dog him. He may be a janitor or a teenage kid; no matter. Follow him with your begging-bowl, if that's what he wants, or take him to expensive restaurants, or whatever.

9. KEEP IMPROVING YOUR QUESTIONS. Probably in your head there are questions that don't seem to line up with what you're hearing. Don't assume that you don't understand; keep adjusting the questions till you can get an answer that relates to what you wanted.

10. YOUR FIELD IS BOUNDED WHERE YOU WANT IT TO BE. Just because others group and stereotype things in conventional ways does not mean they are necessarily right. Intellectual subjects are connected every whichway; your field is what you think it is. (Again, this is one of the things that will give you insights and keep you motivated; but it will get you into trouble if you try to go for degrees.)

* * *

There are limitations. This doesn't give you lab experience, and you will continually have to be making up for gaps. But for alertness and the ability to use his mind, give me the man who's learned this way, rather than been blinkered and clichéd to death within the educational system.

BIBLIOGRAPHY

Wilmar Shiras, Children of the Atom.
Science-Fiction about what a school could be like where kids really used their minds. I've always been sure it was possible; the R.E.S.I.S.T.O.R.S. (see p. 47) made me surer.

PRESENTATIONAL SEQUENCES ARE ARBITRARY

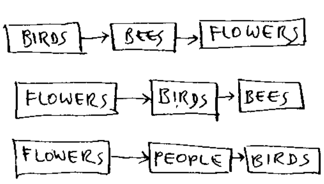

HIERARCHIES ARE TYPICALLY SPURIOUS

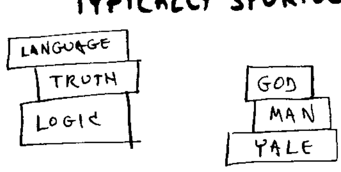

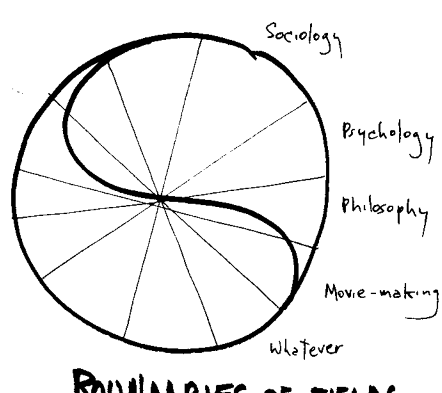

BOUNDARIES OF FIELDS ARE ARBITRARY

COMPARTMENTALIZED AND STRATIFIED TEACHING PRODUCES COMPARTMENTALIZED AND STRATIFIED MINDS.

Verbal communication-- whether written or spoken-- is the disassembly of the Tinkertoy of thought into pieces, and placing it on a conveyor belt to its place of reassembly.

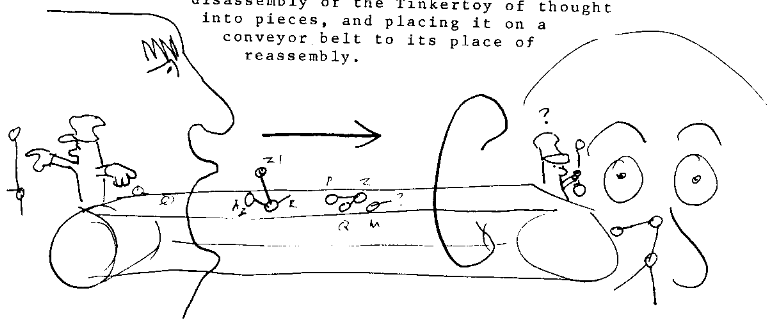

"ON WRITING," a paradigm of the creative process

being an examination of some very Complex Matters which Nobody Seems to Understand; and whose Generality of Relevance may be Gradually Apprehended. (Eventually I hope to develop a somewhat more formal treatment of "ideas," as distinct from propositions, sentence kernels, etc. But there is certainly no room for that here. (Logicians: show me the truth-table of "BUT.")

The process of writing is poorly understood in most quarters. Many working writers despair of being "systematic," getting things done as best they can. On the other hand, people who think they might be able to contribute-- particularly the symbolic logicians and transformational linguists-- being immersed in their own formalisms, simply don't see what's going on-- at least, when I've tried to talk to them.

Writing is not simple. As with vision or speech or riding a bicycle, an immensely complex process is being unconsciously pursued.

Some people think you make an outline and follow it, filling out the details of the outline until the piece is finished. This is absurd. (True, some people can do this, but that is simply a shortcutting of the real process.) Basically writing is

THE TRY-AND-TRY-AGAIN INTERPLAY of PARTS AND DETAILS against OVERALL AND UNIFYING IDEAS WHICH KEEP CHANGING.

In fact a number of things are happening, often simultaneously. We can separate them into three:

1. Provisional development of ideas and points:
   A) forming overall organizing ideas, B) selecting tentative points; C) inductively finding overall organization among them; D) finding relations of interest between points.
2. Complex sifting and adjustment among collections of points, overall ideas.

3. Fine splicing within developed sequences.
   A) transition and juxtaposition managements, B) cross-citations, C) smoothing.

Regrettably, there's no room or time to pursue this here. (The article I had intended to write would take a whole spread.) For people who really care about the matter, I will make some points in very abbreviated form.

The interesting structures in written material include:

"Points"-- pieces, sentences, phrases, examples, plot events, and expository "points."

Organizing principles and structures (which we will call here arches)-- final ironies, things to be led up to, themes, plots, concepts, principles, expository structures, organizing titles, overconcepts. These may be either local or global, over the entire work. (Note: arches may not be heirarchical relative to one another.)

Now, we may think of points and arches as individual objects which have individual relations to one another. Between two points there may be a good transition; a specific point may link well to a specific arch.

The problem in writing, then, is that overall structures you choose (systems of arches) may not link well to the points that have to be included among them; and that transitions between points don't work out the way you want them to. Good transitions can't be worked out for the sequence of points you want to make, or, alternatively, there are too many good transitions within a specific structure of points, and picking among them involves difficult choices-- especially when you have to devise appropriate arches on the basis of the final sequence of points.

There are a number of other important structures in written material. They include accordances, juxtapositions, cross-citations, connotations, nuances and rhythms.

The only ones we will discuss here are accordances.

The term "accordance," as I shall use it here, is simply a vaguely formal way of talking about whether things match or fit together. Two items are in accord if they match or fit well, or in discord if they match or fit badly. Thus a good transition between points (as mentioned early) represents an accord, and a good link between a point and an arch is also an accord.

Now, it happens that a great deal of writing is concerned with notes to the reader about accordances in the material. In fact, quite a few words are exclusively concerned with subtly pointing out to the reader the accords and discords within the expository structure of what he is reading. We may call these accordance-connectives or accordance-notes.

Two of the most basic terms are indeed and but.

The word indeed has an interesting function.

The word indeed (in its main use, at the beginning of a sentence) indicates an accord between what has just been said and what is to follow. In other words, it functions as a positive transition, impetus or gas pedal, indicating a continuation of the flow in the direction already indicated. So do the words thus, then, therefore, moreover, so and furthermore. These are infix accords, that is, notes of accord that go between two items. We also see prefix accords, such as since, inasmuch as, insofar as; these have to be followed by two clauses, the second of which is in accord with the first.

The word but is exactly the opposite. It indicates a discord or contradistinction, a negative transition, "brakes" in the flow. Other such infix discords include nevertheless, despite this, on the other hand, even so, and "Actually,..." Similarly, there are prefix discords: while, despite, though..., notwithstanding.

I find this topic of inquiry very interesting. These sorts of terms have been used since time immemorial by writers adjusting their transitions for smooth flow (note such antiquey variants as haply, howbeit, withal, forasmuch and howsomever), but the importance and structure of this service has not, I think, been generally understood.

(Note also that there are more intricate accordance-connectives: I wish we could go here into the structure of In fact..., at least, ...if not... , ... otherwise... , Anyway... , and Now....)

* * * * * * * * * * * * * * * * *

(Note: the try-and-try-again revision and reconsideration process, tinkering with structural interconnections, is a universal component of the creative process in everything from movie editing to machine design. There ought to be a name for it. I can't think of a satisfactory one, although I would commend to your attention grandesigning, piece-whole diddlework, grand fuddling, meta-mogrification, and that most exalted possibility, tagnebulopsis (the visualization of structure in clouds).)

THE HERITAGE

The past is like the receding view out the back of an automobile: the most recent is more conspicuous, and everything seems eventually to be lost.

We know we chould save things, but what? Those with the job of saving things-- the libraries and museums-- save so many of the wrong things, the fashionable and expensive and high-toned things esteemed by a given time, and most of the rest slips past. Each generation seems to ridicule the things held in esteem by times before, but of course this can never be a guide to what should be saved. And there is so much to save: music, writing, sinking Venice, vanishing species.

But why should things be saved? Everything is deeply intertwingled. We save for knowledge and nostalgia, but what we thought was knowledge often turns to nostalgia, and nostalgia often brings us deeper insights that cut across our lives and very selves.*

Computers offer an interesting daydream: that we may be able to store things digitally instead of physically. In other words, turn the libraries to digital storage (see Hypertexts, p. DM 48+); digitize paintings and photographs (see "Picture Processing, p. DM 10); even digitize the genetic codes of animals, so that species can be restored at future dates (see "The Mitiest Computer," p. 60).

Digital storage possesses several special advantages. Digitally stored materials may be copied by automatic means; corrective measures are possible, to prevent errors from creeping in-- i.e., "no deterioration" in principle; and they could be kept in various places, lessening mankind's dependence on its eggs being all in one basket (like the Library at Alexandria, whose burning during the occupation of Julius Caesar was one of the greatest losses in human history).

But this would of course require far more compact and reliable forms of digital storage than exist right now.

Nevertheless, we better start thinking about it. Those who fear a coming holocaust (see p. 68) had best think about pulling some part of mankind through, with some part of what he used to have.

* See T.H. Nelson, The Snunking of the Heart: On the Psychology of Puns and Preterism in Carroll and Others. 1980, unless a decent writing system comes along.

BRANCHING PRESENTATIONAL SYSTEMS—

HYPERMEDIA

In recent years a very basic change has occurred in presentational systems of all kinds. We may summarize it under the name branching, although there are many variants. Essentially, today's systems for presenting pictures, texts and whatnot can bring you different things automatically depending on what you do. Selection of this type is generally called branching. (I have suggested the generic term hypermedia for presentational media which perform in this (and other) multidimensional ways.)

A number of branching media exist or are possible.

Branching movies or hyperfilms (see nearby).

Branching texts or hypertexts (see nearby).

Branching audio, music, etc.

Branching slide-shows.

Wish we could get into some of that stuff here.

BRANCHING MOVIES

The idea of branching movies is quite exciting. The possibility of it is another thing entirely.

The only system I know of that worked was at the 1967 Montreal World's Fair (Expo 67). At the Czech Pavilion-- you will recall that before the crackdown they had quite a yeasty culture going in Czechoslovakia-- there were some terrific fantic systems going. One was a wall of cubes with slide projectors inside (that rolled toward you and back as they changed their pictures). And then the Movie.

The Czechoslovakian Branching Movie-- I forget its real name-- had the audience vote on what was to happen next at a number of different junctures. What should she do now, what will he do next, etc. And lo and behold! after they had voted, the lights went down, and that's what would happen next. People agreed that this gave the movie a special immediacy.

I never saw the movie-- I waited in line several hours but the line was too long to get into the last showing. So instead I went backstage and talked to Radusz Cincera, who worked out the system. It turns out that it didn't work quite the way people supposed. A lot of people thought that "all the possibilities" had been filmed in advance. Actually, there were always only two possibilities, and no matter what the audience had chosen, somehow the film was plotted to come down to the same next choice anyway:

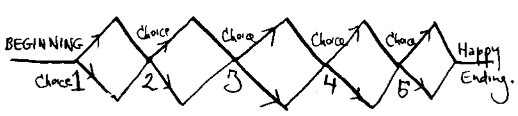

In the actual setup, they simply had two projectors running side by side, with Film A and Film B, and the projectionist would drop an opaque slide in front of whichever wasn't chosen. But Cincera said that audiences almost always chose the same alternatives anyway, so half the movie was hardly ever used...

In the early sixties a movie was making the rounds in which audiences were supposedly allowed to vote on the ending-- "Mr. Sardonicus," I believe it was called. From the ads it seemed that audiences would be polled as to which last reel to show. Whether the villain was to get his comeuppance, or whatever.

Then there was that Panacolor cartridge projector, mentioned elsewhere, which would have allowed choices by the user

More recently there's the CMX system, also mentioned elsewhere. This is a setup, being jointly marketed by CBS and Memorex, for computer-controlled movie editing. But actually it could also be used as a branching movie system. Essentially the movie itself is stored frame-by-frame (as video) on big disks, made by Memorex; and, under computer control, the output can be switched rapidly among the frames, effectively showing the stored movies. (To my knowledge, the video networks haven't yet recognized the possibilities of this.)

The only trouble is, it's extremely expensive (half a million?), it has an exact storage capacity limited by the number of disk tracks (presumably one track per frame)-- perhaps five minutes total one one big unit, but you can buy more-- and it can only give its full performance to one viewer at a time.) (Or to the whole network, live.)

It may be that the most practical branching movie system would be a cartridge movie viewer and a big stack of cartridges. When you make your choice, change the cartridge. But of course that's not as much fun as having it happen automatically.

REALITY IS OBSOLETE

The idea that objective reality is perceived by our senses, is an obsolete concept. Old truisms like "seeing is believing", become much less believable as we become more aware that, the biological machinery of life itself, transforms images of the physical world before we are made conscious of them. These biological mechanisms share many similarities in principle and in application, to other mechanisms observed in the natural environment and those invented for our own use. Since we are becoming more aware of the nature of perception and those mechanisms involved, now is the time to gain control of ourselves and share more discretion in the operation of our own biological machinery. We have entered the age of hyper-reality.

Day-to-day living provides only a limited variety of physical stimulus, and little incentive to manipulate the physiological and psychological processing involved. Man's historical preoccupation with the need to maintain constant images of the physical world, is a product of his extreme orientation toward physical survival in a hostile environment. The current evolving society of leisure orientations removes this need for constant images and thereby enhances the opportunities for a more complete use of the sensory apparatus and those related brain functions. Many have turned to drugs or meditation. More specifically it is proposed here, that modern communications technology be employed as a "vehicle of departure" from this need for constant images, to bring about a more complete use of the human technology itself. Hyper-reality is the employment of technology other than the biological machinery, when used to affect the performance of the biological machinery beyond its own limitations. This is almost like making adjustments on a television set, except you are what's plugged in, and the controls are outside your body, being part of whatever technology is interfaced to the body itself. As part of such a man-machine interface you could extend your own mental processes, or if you should choose, you could just diddle with the dials. Hyper-reality is an opportunity to enhance the various qualities of the human experience. Reality is obsolete.

-- *How Wachspress (see p. DM 6)*

!GREBNETUG

Now, in our time, we are turning Gutenberg around. The technology of movable type created certain structures and practices around the written word. Now the technology of computer screen displays make possible almost any structures and practices you can imagine for the written word.

So now what?

For new forms of written communication among people who know each other, jump to "Engelbart" piece, nearby.

To learn about new forms of multidimensional documents for computer screens, jump to "Hypertexts."

Or just feel free to browse.

HYPERTEXT

By "hypertext" I mean non-sequential writing.

Ordinary writing is sequential for two reasons. First, it grew out of speech and speech-making, which have to be sequential; and second, because books are not convenient to read except in a sequence.

But the structures of ideas are not sequential. They tie together every whichway. And when we write, we are always trying to tie things together in non-sequential ways (see p.DM 42). The footnote is a break from sequence; but it cannot really be extended (though some, like Will Cuppy, have toyed with the technique).

I have run into perhaps a dozen people who understood this instantly when I talked to them about it. Most people, however, act more bemused, thinking I'm trying to tell them something technical or pointlessly philosophical. It's not pointless at all: the point is, writers do better if they don't have to write in sequence (but may create multiple structures, branches and alternatives), and readers do better if they don't have to read in sequence, but may establish impressions, jump around, and try different pathways until they find the ones they want to study most closely.

(The astute reader, and anybody who's gotten to this point must be, will have noticed that this book is in "magazine" layout, organized visually by ideas and meanings, for that precise reason. I will be interested to hear whether that has worked.)

And the pity of it is that (like the man in the French play who was surprised to learn that he had been "speaking prose all his life and never known it"), we've been speaking hypertext all our lives and never known it.

Now, many writers have tried to break away from sequence. I think of Nabokov's Pale Fire, of Tristram Shandy and an odd novel of Lazaro Cortazar called Hopscotch, made up of sections ending with numbers telling you where you can branch to. There are many more; and large books generally use many tricks to get around the problem of indexing and reviewing what has and hasn't been said or done already.

However, in my view, a new day is dawning. Computer storage and screen display mean that we no longer have to have things in sequence; totally arbitrary structures are possible, and I think that after we've tried them enough people will see how desirable they are.

TYPES OF HYPERTEXT

Let's assume that you have a high-power display-- and storage displays won't do, because you have to see things move in order to understand where they come from and what they mean. (Especially text.) So it has to be a refreshed CRT.

Basic or chunk style hypertext offers choices, either as footnote-markers (like asterisks) or labels at the end of a chunk. Whatever you point at then comes to the screen.

Collateral hypertext means compound annotations or parallel text (see p. DM42).

Stretchtext changes continuously. This requires very unusual techniques (see p. DM19), but exemplifies how "continuous" hypertext might work.

Ideally, chunk and continuous and collateral hypertext could all be combined (and in turn collaterally linked; see "Thinkertoys," p. DM52).

A "fresh" or "specific" hypertext-- I don't have a better term at the moment-- would consist of material especially written for some purpose. An anthological hypertext, however, would consist of materials brought together from all over, like an anthological book.

A grand hypertext, then, folks, would be a hypertext consisting of "everything" written about a subject, or vaguely relevant to it, tied together by editors (and NOT by "programmers," dammit), in which you may read in all the directions you wish to pursue. There can be alternative pathways for people who think different ways. People who have to have one thing explained to them at a time-- many have insisted to me that this is normal, although I contend that it is a pathological condition-- may have that; others, learning like true human beings, may gather and sift impressions until the ideas become clear.

And then, of course, you see the real dream.

The real dream is for "everything" to be in the hypertext.

Everything you read, you read from the screen (and can always get back to right away); everything you write, you write at the screen (and can cross-link to whatever you read; see Canons, p. DM58).

Paper moulders. Microfilm is inconvenient. In the best libraries it takes at least minutes to get a particular thing. But as to linking them together-- footnoting Aeschylus with Marcus Aurelius, linking genetic data to 15th-century accounts of Indian tribes-- well, you can only do it on paper by writing something new that ties them together. Isn't that ridiculous? When you could do it all electronically in seconds?

Now that we have all these wonderful devices, it should be the goal of society to put them in the service of truth and learning. And this is the way I propose. Not through obscure forms of "information retrieval;" not through newly oppressive forms of "computer-assisted instruction;" and not through a purported science of "artifical intelligence" that will create new personalisms to irk us. All these obstructive oddities, I think, have developed as separate ideals because of the grand preposterosity of Professionalism that has created a world-wide cult of mutual incomprehensibility and disconnected special goals. Now we need to get everybody together again. We want to go back to the roots of our civilization-- the ability, which we once had, for everybody who could read to be able to read everything. We must once again become a community of common access to a shared heritage.

This was of course what Vannevar Bush said in 1945 (see box on right), in an article everybody cites but nobody reads.

The hypertext solution in many ways obviates some of these other approaches, and in addition retains and puts back together the great traditions of literature and scholarship, traditions based on the fact that dividing things up arbitrarily just generally doesn't work.

EVERYTHING IS DEEPLY INTERTWINGLED.

(The only way in which my views differ with those of Engelbart and Pask, I think is in the matter of structure and hierarchy. Both men generally assume that whatever natural hierarchy may exist in particular subjects needs to be accentuated; I hold that all structures must be treated as totally arbitrary, and any hierarchies we find are interesting accidents.)

CAN IT BE DONE?

I dunno.

Licklider, one of computerdom's Great Men, estimated in 1965 that to handle all text by computer, and bring it out to screens, would cost no more than what we pay for all text handling now. (But of course there is the problem of what to do with the people whose lives are built around paper; that can't be taken up here.)

The people who make big computers say that to get the big disk storage to hold great amounts of text, you have to get their biggest computers. Which is a laugh and a half. One IBM-style computer person pompously told me that for large-scale text handling the only appropriate machine was an IBM 360/67 (a shamefully large computer). Such people seem not to understand about minicomputers or the potential of minicomputer networks-- using, of course, big disks.

There are of course questions of reliability, of "big brother" (see Canons, p.), and so on. But I think these matters can be handled.

The key is that people will pay for it. I am sure that if we can bring the cost down to two dollars an hour-- one for the local machine (more than a "terminal"), one for the material (including storage, transmission and copyrights)-- there's a big, big market. (And that's what the Xanadu network is about; see p. DM57.) My assumption is that the way to do this is not through big business (since all these corporations can see is other corporations); not through government (hypertext is not committee-oriented, but individualistic-- and grants can only be gotten through sesquipedalian and obfuscatory pompizzazz); but through the byways of the private enterprise system. I think the same spirit that gave us McDonald's and kandy kolor hot rod accessories may pull us through here. (See Xanadu Network, p. DM57.)

Obviously, putting man's entire heritage into a hypertext is going to take awhile. But it can and should be done.

BIBLIOGRAPHY

Theodor H. Nelson, "The Hypertext." Proc. World Documentation Federation, 1965.

COULDN'T HAVE HYPERTEXT NOVELS, YOU SAY?

Consider the hypertext character of--
- Tristram Shandy, by Sterne.
- Spoon River Anthology, by Masters.
- Hopscotch, by Cortazar.
- Pale Fire, by Nabokov.
- Remembrance of Things Past, by Proust.

And, surprisingly, hypertext actually FIGURES IN Giles Goat-Boy, by Barth.

GLINDA'S MAGIC BOOK

Glinda the Good, gentle sorceress of the southern quadrant of the land of Oz-- not the flaphead portrayed by Billie Burke in the Goldwynized film-- has a Magic Book in which Everything That Happens is written.

The question, of course, is how it's chosen.

You can only watch news tickers for a short time before getting very bored.

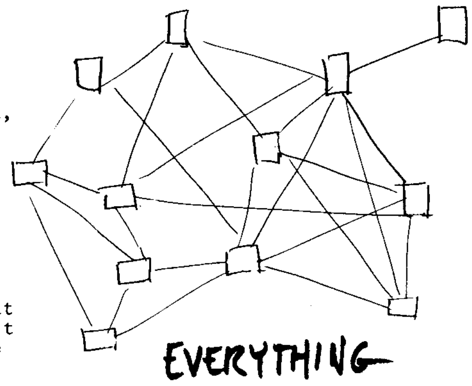

EVERYTHING IS DEEPLY INTERTWINGLED.

In an important sense there are no "subjects" at all; there is only all knowledge, since the cross-connections among the myriad topics of this world simply cannot be divided up neatly.

Hypertext at last offers the possibility of representing and exploring it all without carving it up destructively.

Arthur C. Clarke wrote a book entitled The Lost Worlds of 2001 (Signet, 1972), about the variants and alternatives of that story that did not find their way to the screen.

In a hypertext version, we could look at them all in context, in collateral views, and see the related variants-- with annotations.

Mortimer J. Adler, the man who reduced all of Western Culture to a few Great Books plus an index under his own categories, has now Addled the Encyclopedia Britannica.

Since 1965 he has been creating Britannica 3, the venturesome and innovative new version, now on sale for about half a thou.

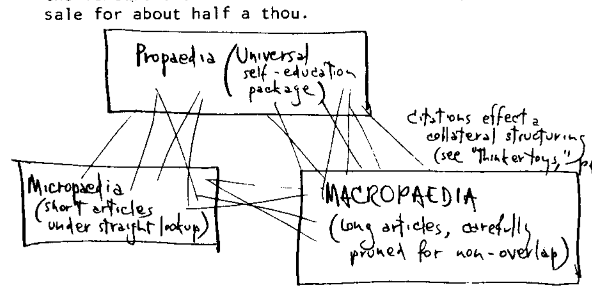

Britannica 3 is basically a 3-level hypertext, made to fit on printed pages by the strictures of Adler's editing (according to Newsweek, some 200 authors withdrew their work rather than submit to the kind of restrictions he was imposing).

The idea may be basically good, even though the sesquipaedalian titles may impaed the raeder.

THE BURNING BUSH

In fact hypertexts were foreseen very clearly in 1945 by Vannevar Bush, Roosevelt's science advisor. When the war was in the bag, he published a little article on various groovy things that had become possible by that time.

"As We May Think" (Atlantic Monthly, July, 1945) is most notable for its clear description of various hypertext techniques-- that is, linkages between documents which may be brought rapidly to the screen according to their linkages. (So what if he thought they'd be on microfilm.)

How characteristic of Professionalism. Bush's article has been taken as the starting point for the field of Information Retrieval (see p.), but its actual contents have been ignored by acclamation. Information Retrieval folk have mostly done very different things, yet thought they were in the tradition.

Now people are "rediscovering" the article. If there's another edition of this book I hope I can run it in entirety.

DOUG ENGELBART AND "THE AUGMENTATION OF INTELLECT"

Douglas Engelbart is a saintly man at Stanford Research Institute whose dream has been to make people smarter and bring them together. His system, on which millions of dollars have been spent, is a wonder and a glory.

He began as an engineer of CRTs (see "Lightning in a Bottle," p. DM6); but his driving thought was, quite correctly, that these remarkable objects could be used to expand man's mind and improve each shining hour.

Doug Engelbart's vision has never been restricted to narrow technical issues. From the beginning his concern was not merely to plank people down at display consoles, but in the most profound sense to expand man's mind. "The Augmentation of Human Intellect," he calls it, by which he means making minds work better by giving them better tools to work with.

An obvious example is writing: before people could write things down, men could only learn what they experienced or were told by others in person; writing changed all that. Within the computer-screen fraternity, the next step is obvious; screens can double and redouble our intellectual capacities. But this is not obvious to everybody. Engelbart, patiently instructing those outside, came up with a beautiful example. To show what he meant by the Augmentation of Intellect, Engelbart tied a pencil to a brick. Then he actually made someone write with it. The result, which was of course dreadful, Engelbart solemnly put into a published report. Not yet being able to demonstrate the augmentation of intellect, since he had as yet no system to show off, he had masterfully demonstrated the disaugmentation of intellect: what happens if you make man's tools for working out his thoughts worse instead of better. As this poor guy was with his brickified pencil, explained Engelbart, so are we all among our bothersome, inflexible systems of paper.

Starting small, Engelbart programmed up a small version of what most fans call "The Engelbart System" some ten years ago. One version has it that when it came to looking for grants, management thought he acted too kooky, and so assigned a Front Man to make the presentation. But, as the story goes, the man from ARPA (see "Military...", p. 59) pointed at Engelbart and said, "We want to back him."

A small but dedicated group at SRI has built up a system from scratch. First they used little CDC 1700 minicomputers; then, various grants later, they were able to set up their own PDP-10, in which the system now resides, and from which it reaches out across the country.

Doug calls his system NLS, or "oN-Line System." Basically it is a highly responsive, deeply-structured text system, feeding out to display terminals. From a terminal you may read anything you or others have written, and write with as-yet-unmatched flexibility.

The display terminals are all over. The project has gone national, though at great expense: through the ARPA net of computers, you can in principle become a user of NLS for something like $50,000 a year.

THE "KNOWLEDGE WORKSHOP"

For a lucky fifty or so people, Engelbart's system is Home. Wherever they are-- at Stanford Research Institute or far away on the ARPAnet-- a whole world of secretarial and communication services is at their fingertips. The user has but to call up through his display terminal and log on. At that point all his written files, and numerous files shared among the users, are at his fingertips. He may read, write, annotate the cross-link. (Engelbart's system has provision for collateral structuring: see "Thinkertoys," p. DM52.) He may send messages to others in the Workshop. He may open certain of his files to other people, and read those that have been opened to him.

This all has a certain vagueness if you do not understand how bound you are today by paper-- the problems of finding it, sorting it, looking things up. (If you write, that is, write a lot, you know all too well how intractable is paper, what a damned nuisance.) With a system like Engelbart's, now, whatever is written is instantly there. Whatever you want to look up is instantly there, simultaneously interconnected to everything else it should be connected to-- source materials, footnotes, comments and so on. A document is completed the moment it is written; no human being has to retype it. (It need not be typed on paper at all, if it's just for the workshop members: a printout is only needed if it has to go to someone outside the system.)

In many ways, Engelbart's system is a prototype for the world of the future, I hope. ALL HANDLING OF PAPER IS ELIMINATED. Whatever you write, you write on the screens with keyboard and pointer. (No more backs of envelopes, yellow pads, file cards, typewriters.) Whatever you transmit to fellow users of the system you simply 'release'-- no physical paper changes hands.

The group has also worked out some remarkable techniques for collaborative endeavor. Two people-- say, one in California and one in New York-- can work together through their screens, plus a phone link; it's as if they were side-by-side at a magic table. Each sees on his screen what the other sees; each controls a moving dot (or "cursor") that shows where he's pointing. The effect is somewhere between a blackboard and a desk; both may call up documents, point things out in them, change them, and anything else two people might do when working on something together.

Engelbart meets with someone far away, as others watch.

THE SYSTEM ITSELF

Basically the system is a large-scale setup for the storage, bringing forth, viewing and revision of documents and connections among them.

The documents are stored (of course) in alphabetical codes. Connections among them, or other relations within them, are signalled by the presence of other codes within them; these are ordinarily not displayed, however, except as directed by a particular display program and display programs can of course vary.

There are various programs for display, in large part depending on what sort of screen system the individual user has. (NLS is used with everything from high-resolution line-drawing screens converted to 1000-line television, down to inexpensive Delta Data terminals-- a brand, incidentally, that allows text motion, which most don't.) Engelbart's system is extremely general, allowing the creation of files having all kinds of structures, and display programs in all kinds of styles. (I hope that this side of the present book conveys a sense of how many styles that can be.) However, most users are devoted to certain standardized styles of working that have been well worked out and permit the easy sharing of material and of operating practices. Here, for instance, is their standard screen layout:

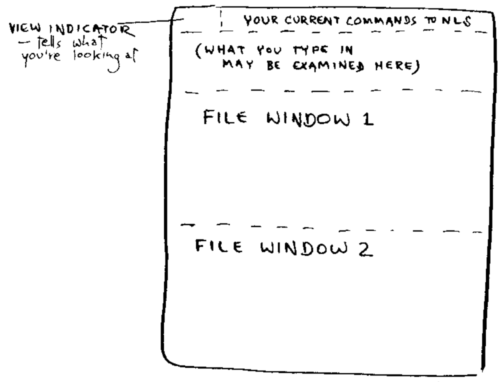

Two separate panels of text appear, and links may be shown on them. (Thus it's a thinkertoy-- see p. .) Two little windows at the top remind you of what you're seeing and what you're asking for. We can't get into the rest of it here

THE COMMAND LANGUAGE

NLS has a command language which all users must learn. While it is a streamlined and straightforward command language, nevertheless it requires the user to type in a specific sequence of alphabetical characters every time he wants something done. (This is acceptable to computer-oriented people; I suspect it would not be satisfactory, say, for philosophers and novelists. For designs oriented to such users, see JOT (p. DM50) and Carmody's System, nearby, Parallel Textface (p. DM53) and Th3 (p. DM55).)

Incidentally, NLS users may also employ a cute little keyboard, something like a kalimba, that allows you to type with one hand. You simply type the six significant ASCII bits (see chart p. 28) in one "chord" -- it sounds hard but is easy to learn.

Sample commands: I (insert), D (delete), M (move or rearrange). Then you point with the mouse.

MOUSE?

The Engelbart Folks have built a pointing device, for telling the system where you're pointing on the screen, that is considerably faster and handier than a lightpen. (Unfortunately, I don't believe it's commercially available.) It's called The Mouse.

The Engelbart Mouse is a little box with hidden wheels underneath and a cable to the terminal. As you roll it, the wheel's turns are signalled to the computer and the computer moves the cursor on the screen. It's fast and accurate, and in fact beats a lightpen hands down in working speed.

Through the command language, NLS allows users to create programs that respond in all sorts of ways; thus the fact that certain text-handling styles are standard (as in above illustration of screen layout) results more from tradition than necessity.

The same apparently is true of the data structure. I used to be somewhat disturbed at the way Engelbart's text systems seem to be rigorously hierarchical. This in fact is the case, in the sense that having multiple discrete levels is built deep into the system. But it turns out to be harmless. The stored text is divided by the storage techniques into multiple levels, corresponding to a Harvard outline. Think of it as something like this:

1. HIERARCHICAL FORMAT
   - A. STORAGE
   - B. DISPLAY
   - C. LANGUAGE

But let's expand this example a little:

1. HIERARCHICAL FORMAT
   - A. STORAGE
     - A1. Everything in NLS is stored with hierarchical codes.
     - A2. Their effect depends on the display.
   - B. DISPLAY
     - B1. The hierarchical codes of NLS have no consequences in particular.
     - B2. The hierarchical codes for NLS can splay the material out into a variety of display arrangements.
       - B2A. They can be displayed in outline form.
       - B2B. They can be displayed in normal text form.
       - B2C. These dratted numbers can even be made to disappear.
   - C. LANGUAGE
     - C1. The command language determines what the display shows of the hierarchical structure.
     - C2. What is shown can be determined by a program in the command language. (For instance, "how many levels down" it is being shown).
       - C2A. This is four levels down. (The earlier example wasn't.)
     - C3. The display format all depends on what display program you use, in the NLS command language.

That's enough of that. I can't help remarking that I still don't like that sort of structuring, but it is deep in NLS, and if you don't like it either (poor deprived lucky user of NLS) you can program it to disappear, so it's hardly in your way.

BY THE BEARD OF THE PROPHET!

Engelbart in German means Angelbeard; Doug Engelbart is indeed on the side of the angels. In building his mighty system he points a new way for humanity. The sooner the better. Any history of the twentieth century will certainly hold him high. Few great men are also such nice guys.

BIBLIOGRAPHY

Douglas C. Engelbart, Richard W. Watson & James C. Norton, "The Augmented Knowledge Workshop." Proc NCC 73, 9-21.

Charles H. Irby, "Display Techniques for Interactive Text Manipulation." Proc NCC 74, 247-255.

A VERY BASIC HYPERTEXT SYSTEM

Hypertext is non-sequential writing. It's no good to us, though, unless we can go instantly in a choice of directions from a given point.

This of course can only mean on computer display screens.

Engelbart's system, now, was mainly designed for people who wanted to immerse themselves in it and learn its conventions. Indeed, it might be said to have been designed for a community of people in close contact, a sort of system of blackboards and collaborative talking papers.

A more elemental system, with a different slant, was put together at Brown U. on IBM equipment. We'll refer to it here as "Carmody's System," after the young programmer whose name came first on the writeup.

Carmody's system runs on an IBM 360 with 2250 display. While the 2250 is a fine piece of equipment, the quirks of the 360's operating system (see p. 45) often delay the user by making him wait, e.g., for someone else's cards to get punched before it responds to his more immediate uses; this is like making ice-skaters wait for oxcarts.

Anyway, the system essentially imposes no structure on the material; it may consist of text segments of any length and ties and links between them. An asterisk appearing anywhere in one piece of text signals a possible jump, but the reader doesn't necessarily know where to; zapping the asterisk with the lightpen takes you there, however.

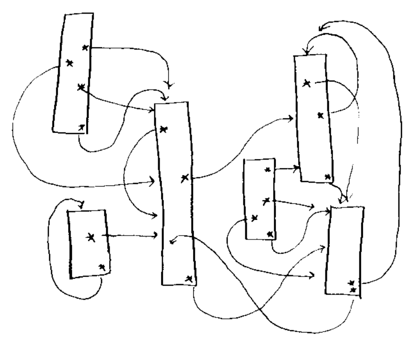

This is stark and simple. It could also get you good and lost. However, a simple technique took care of that: everytime the user jumped, the address of his previous location was saved on a stack (see "The Magic of the Stack," p. 42). The user also had a RETURN button: when he wanted to go back to where he had last jumped from, the system would pop the last address off the top of the stack, and take him there. (This feature was adapted from my 1967 Stretchtext paper, and turned out to work out quite well in practice.)

The system also had handy features for light-pen text editing, and various nice printout techniques. All told, it was a clean and powerful design. While it lacked higher-level visualization facilities, like Engelbart's display of Levels (see "outline" in Engelbart article) or collateral display (see "Thinkertoys," p.DM52), it was in some ways suited for naive users; that is, it was eventually fairly safe to use, and could in large part be taught to rank beginners in a couple of hours-- provided they didn't have to know about JCL cards.

It is left for the reader to figure out interesting uses for it. How would you do collateral structures? How could you signal to a reader which of several pieces of text a jump was to?

(At least one real hypertext was actually written on this system. It tied together a lot of patents for multilayer electrodes. Readers agreed that they could learn more from it about multilayer electrodes than they had imagined wanting to know.)

BIBLIOGRAPHY

Steven Carmody, Walter Gross, Theodor H. Nelson, David Rice and Andries van Dam, "A Hypertext Editing System for the 360." In Faiman and Nievergelt (eds.), Pertinent Concepts in Computer Graphics (U. Ill. Press '69), 291-330.
Note: Mr. Gross now goes by the name of Lightning Clearwater.

GORDON PASK RETURNS [DIFFICULT SECTION]

This continues the remarks on Gordon Pask begun on p.

I will now try to describe Pask's work as he has explained it to me. Perhaps this will be of some help to those who may have been mystified or dumfounded by contact with this fabulous man.

Gordon Pask's concern is abstraction and how concepts are formed, whether in a creature of nature or a robot or a computer program. Abstraction is of interest primordially (as life evolved thinking capacity), psychogenetically (as the mind acquires new facilities, described most peculiarly by Piaget), and epistemologically (how do we know? Like, how do we know, man?), and methodologically (how can we most effectively find out more?).

His interest, then, is in teaching by allowing students to discover exact relations in a specific subject matter by the very process of abstraction that is of so much interest.

What he does, then, is prepare given fields of learning so that they can be studied by students using abstractive methods, without guidance.

This preparation basically has two steps. First he sets up the whole field. This is done in collaboration with a "subject matter expert," who names the important topics in the field and states what interconnections they have. The result is a complex graph structure (see p. 26) which Pask calls a conversational domain. It comes out to huge diagrams of labels and lines between them.

Then Pask processes this structure to make a more usable map of the field that he calls an entailment structure. The processing basically involves removing "cycles" in the graph, thus making the structure hierarchical in a slightly artificial way justified by what the subject-matter-expert has said is the structure of the field.

(This processing is carried out by a program called EXTEND.)

The resulting Entailment Structure is then presented to the student as a great map of the field which he may explore.

Pask intends that the student's explorations will consist of testing analogies, or what Pask calls morphisms, to find the exact structures of knowledge he is supposed to be acquiring. This knowledge will be in the form of isomorphisms, or exact analogies, i.e. laws.

Pask's overall system, examples of which he has running in his laboratory in England, he calls CASTE (Course Assembly System and Tutorial Environment). A further development, which is to be put on a PDP-11/45 computer (see p. 36 and p. 42) at the Brooklyn Children's Museum, is called THOUGHT-STICKER. This program is intended to allow the demonstration and testing of analogies directly, by children.

PASK AND HYPERTEXT

Gordon Pask's work is remarkably similar to my own stuff on hypertext.

Essentially Pask is reducing a field to an extremely formal structure of relations which may then be studied by the student, at the student's initiative.

(What I don't quite understand is how the analogies are to be explored and tested.)

Anyway, a principal point is that the student is in control and may use his initiative dynamically; the subject is not artificially processed into a presentational sequence. Moreover, the arbitrary interconnections of the subject, which are no respecters of the printed page, are recognized as the fundamental structures the student must deal with and come to understand. On all these points Pask and I are in total agreement.

Indeed, his explorable systems-- (I don't know if they will be what I elsewhere call hypergrams or responding resources)-- will be fascinating, fun and terrifically educational. Because he is.

Now it turns out that this exactly complements the notion of hypertext as I have been promulgating it lo these many years.

Hypertext is non-sequential text. If we write a hypertext on something, it will be most appropriate if we give it the general interconnective structure of the field. In other words, the interconnective structures chosen for the textual parts are likely to have the same connective structure (in general) as Pask's Entailment Structure.

For another kind of hypertext, the anthological hypertext built up of lots of other writings, it is also reasonable to expect the connective structures to cluster to the same general form as Pask's entailment structure.

In other words, the very same field of knowledge Pask is out to represent as an explorable, formalized whole, I am out to represent as an explorable informalized whole, with anecdotes, jokes, cartoons, "enrichment materials," and anything else people might dig.

In still other words, let's have both and call it a party.

The Augmentation of Intellect. Infamous Ape Sequence from my slide-show.

Actually it needs the '2001' music.

You can't read the screen here. It says: COGITO ERGO SUM $e=mc^2$ Call me Ishmael.

It really needs the music.

FEELED-EFFECT SYSTEMS ARE THE NEW FRONTIER

FANTICS

— BUT IT'S SHOWMANSHIP THAT'S PARAMOUNT, NOT ANY TECHNICAL SPECIALTY

Ah, Love! could you and I with Him conspire
To grasp this sorry Scheme of Things entire,
Would not we shatter it to bits—and then
Re-mould it nearer to the Heart's Desire!

Edward Fitzgerald.

Almost everyone seems to agree that Mankind (who?) is on the brink of a revolution in the way information is handled, and that this revolution is to come from some sort of merging of electronic screen presentation and audio-visual technology with branching, interactive computer systems. (The naive think "the" merging is inevitable, as if "the" merging meant anything clear. I used to think that too.)

Professional people seem to think this merging will be an intricate mingling of technical specialties, that our new systems will require work by all kinds of committees and consultants (adding and adjusting) until the Results-- either specific productions or overall Systems-- are finished. Then we will have to Learn to Use Them. More consulting fees.

I think this is a delusion and a con-game. I think that when the real media of the future arrive, the smallest child will know it right away (and perhaps first). That, indeed, should and will be the criterion. When you can't tear a teeny kid away from the computer screen, we'll have gotten there.

We are approaching a screen apocalypse. The author's basic view is that RESPONSIVE COMPUTER DISPLAY SYSTEMS CAN, SHOULD AND WILL RESTRUCTURE AND LIGHT UP THE MENTAL LIFE OF MANKIND. (For a more conventional outlook, see box nearby, "Another Viewpoint.")

I believe computer screens can make people happier, smarter, and better able to cope with the copious problems of tomorrow. But only if we do right, right now.

WHY?

The computer's capability for branching among events, controlling exterior devices, controlling outside events, and mediating in all other events, makes possible a new era of media.

Until now, the mechanical properties of external objects determined what they were to us and how we used them. But henceforth this is arbitrary.

The recognition of that arbitrariness, and reconsideration among broader and more general alternatives, awaits us. All the previous units and mechanisms of learning, scholarship, arts, transaction and confirmation, and even self-reminder, were based in various ways upon physical objects-- the properties of paper, carbon paper, files, books and bookshelves. To read from paper you must move the physical object in front of you. Its contents cannot be made to slide, fold, shrink, become transparent, or get larger.

But all this is now changing, and suddenly. The computer display screen does all these things if desired, to the same markings we have previously handled on paper. The computer display screen is going to become universal very fast; this is guaranteed by the suddenly rising cost of paper. And we will use them for everything. This already happens wherever there are responding computer screen systems. (I have a friend with two CRTs on his desk; one for the normal flow of work, and one to handle interruptions and side excursions.) A lot of forests will be saved.

Now, there are many people who don't like this idea, and huff about various apparent disadvantages of the screen. But we can improve performance until almost everyone is satisfied. For those who say the screens are "too small," we can improve reliability and backup, and offer screens everywhere (so that material need not be physically carried between them).

The exhilaration and excitement of the coming time is hard to convey on paper. Our screen displays will be alive with animation in their separate segments of activity, and will respond to our actions as if alive physically too.

The question is, then: HOW WILL WE USE THEM? Thus the design of screen performances and environments, and of transaction and transmission systems, is of the highest priority.

THE FRENCH HAVE A WORD FOR IT

In French they use the term l'Informatique to mean, approximately, the presentation of information to people by automatic equipment.

Unfortunately the English equivalent, informatics, has been preempted. There is a computer programming firm called Informatics, Inc., and when I wrote them about this in the early sixties they said they did not want their name to become a generic term. Trademark law supports them in this to a certain extent. (Others, like Wally Feurzeig, want that to be the word regardless.) But in the meantime I offer up the term fantics, which is more general anyhow.

MEDIA

What people don't see is how computer technology now makes possible the revision and improvement-- the transformation-- of all our media. It "sounds too technical."

But this is the basic misunderstanding: the fundamental issues are NOT TECHNICAL. To understand this is basically a matter of MEDIA CONSCIOUSNESS, not technical knowledge.

A lot of people have acute media consciousness. But some people, like Pat Buchanan and the communards, suggest that there is something shabby about this. Many think, indeed, that we live in a world of false images promulgated by "media," a situation to be corrected. But this is a misunderstanding. Many images are false or puffy, all right, but it is incorrect to suppose that there is any alternative. Media have evolved from simpler forms, and convey the background ideas of our time, as well as the fads. Media today focus the impressions and ideas that in previous eras were conveyed by rituals, public gatherings, decrees, parades, behavior in public, mummer' troupes...but actually every culture is a world of images. The chieftain in his palanquin, the shaman with his feathers and rattle, are telling us something about themselves and about the continuity of the society and position of the individuals in it.

Now the media, with all their quirks, perform the same function. And if we do not like the way some things are treated by the media, in part this stems from not understanding how they work. "Media," or structured transmission mechanisms, cannot help being personalized by those who run them. (Like everything else.) The problem is to understand how media work, and thus balance our understanding of the things that media misrepresent.

THOUGHTS ABOUT MEDIA:
1. ANYTHING CAN BE SAID IN ANY MEDIUM.

Anything can be said in any medium, and Inspiration counts much more than 'science'. But the techniques which are used to convey something can be quite unpredictable.

Voice of Little Girl:
"Santa, are you more important than God?"
Announcer, plonkingly:
"Your answer is..."

(After How To Be A Department Store Santa Claus, produced by the author for CBS Laboratories and the AVS-10 instructional device. Original slide, starring Michelle Dellinger and Henry Shrady, unfortunately mislaid.)

2. TRANSPOSABILITY

There has always been, but now is newly, a UNITY OF MEDIA OPTIONS. You can get your message across in a play, a tract, a broadside, a textbook, a walking sandwich-board, a radio program, a comic book or fumetti, a movie, a slide-show, a cassette for the Audi-Scan or the AVS-10, or even a hypertext (see p. DM46).

(But transposing can rarely preserve completely the character or quality of the original.)

3. BIG AND SMALL APPROACHES

What few people realize is that big pictures can be conveyed in more powerful ways than they know. The reason they don't know it is that they see the content in the media, and not how the content is being gotten across to them-- that in fact they have been given very big pictures indeed, but don't know it. (I take this point to be the Nickel-Iron Core of McLuhanism.)

People who want to teach in terms of building up from the small to the large, and others who (like the author) like to present a whole picture first, then fill in the gaps, are taking two valid approaches. (We may call these, respectively, the Big Picture approach and the Piecemeal approach.) Big pictures are just as memorable as picky-pieces if they have strong insights at their major intersections.

4. THE WORD-PICTURE CONTINUUM

The arts of writing and diagramming are basically a continuum. In both cases the mental images and cognitive structures produced are a merger of what is heard or received. Words are slow and tricky for presenting a lot of connections; diagrams do this well. But diagrams give a poor feel for things and words do this splendidly. The writer presents exact statements, in an accord-structure of buts and indeeds, molded in a structure of connotations having (if the writer is good) exact impreciseness. This is hardly startling: you're always selecting what to say, and the use of vague words (or the use of precise-sounding words vaguely) is simply a flagrant form of omission. In diagrams, too, the choice of what to leave in and out, how to represent overweening conditions and forces and exemplary details, are highly connotative. (Great diagrams are to be seen in the Scientific American and older issues of TIME magazine.)

This word-picture continuum is just a part of the broader continuum, which I call Fantics.

ANOTHER VIEWPOINT

[from handout, 1974 Nat'l. Joint Comp. Conf.]

John B. Macdonald
Research Leader
Computer Applications: Graphics
Western Electric Company
Engineering Research Center

PROBLEMS, PERILS, AND PROMISES OF COMPUTER GRAPHICS

I would begin with some definitions which may be obvious but bear repeating.

1. Engineering is the application of science for ($) profit,
2. Computer graphics does not make possible anything that was previously impossible: it can only improve the throughput of an existing process,
3. A successful application of computer graphics is when over a period of five years the cost savings from improved process throughput exceed the costs of hardware, software, maintenance and integration into an existing process flow.

FANTICS

By "fantics" I mean the art and science of getting ideas across, both emotionally and cognitively. "Presentation" could be a general word for it. The character of what gets across is always dual: both the explicit structures and feelings that go with them. These two aspects, exactness and connotation, are an inseparable whole; what is conveyed generally has both. The reader or viewer always gets feelings along with information, even when the creators of the information think that its "content" is much more restricted. A beautiful example: ponderous "technical" manuals which carry much more connotatively than the author realizes. Such volumes may convey to some readers an (intended) impression of competence, to others a sense of the authors' obtuseness and non-imagination. Explicit declarative structures nevertheless have connotative fields; people receive not only cognitive structures, but impressions, feelings and senses of things.

Fantics is thus concerned with both the arts of effect-- writing, theater and so on--and the structures and mechanisms of thought, including the various traditions of the scholarly event (article, book, lecture, debate and class). These are all a fundamentally inseparable whole, and technically-oriented people who think that systems to interact with people, or teach, or bring up information, can function on some "technical" basis-- with no tie-ins to human feelings, psychology, or the larger social structure-- are kidding themselves and/or everyone else. Systems for "teaching by computer," "information retrieval," and so on, have to be governed in their design by larger principles than most of these people are willing to deal with: the conveyance of images, impressions and ideas. This is what writers and editors, movie-makers and lecturers, radio announcers and layout people and advertising people are concerned with; and unfortunately computer people tend not to understand it for beans.

In fantics as a whole, then we are concerned with:

1. The art and science of presentation. Thus it naturally includes

2. Techniques of presentation: writing, stage direction, movie making, magazine layout, sound overlay, etc. and of course

3. Media themselves, their analysis and design; and ultimately

4. The design of systems for presentation. This will of course involve computers hereafter, both conceptually and technically; since it obviously includes, for the future, branching and intricately interactive systems enacted by programmable mechanisms, i.e. computers. Thus computer display, data structures (and, to an extent, programming languages and techniques) are all a part.

Fantics must also include

5. Psychological effect and impact of various presentational techniques-- but not particular formal aesthetics, as of haiku or musical composition. Where directly relevant fantics also includes

6. Sociological tie-ins-- especially supportive and dysfunctional structures, such as tie-ins with occupational structure; sponsorship and commercials; what works in schools and why. Most profoundly of all, however, fantics must deal with psychological constructs used to organize things:

7. The parts, conceptual threads, unifying concepts and whatnot that we create to make aspects of the world understandable. We put them into everything, but standardize them in media.

For example, take radio. Given in radio-- the technological fundament-- is merely the continuous transmission of sound. Put into it have been the "program," the serial (and thus the episode), the announcer, the theme song and the musical bridge-- conventions which are useful presentationally.

The arbitrariness of such mental constructs should be clear. Their usefulness in mental organization perhaps is not.

Let's take a surprise example, nothing electronic about it.

Many "highways" are wholly fictitious-- at least to begin with. Let's say that a Route 37 is created across the state: that number is merely a series of signs that users can refer to as they look at their maps and travel along.

However, as time goes by, "Route 37" takes on a certain reality as a conceptual entity: people think of it as a thing. People say "just take 37 straight out" (though it may twist and turn); groups like a Route 37 Merchants' Association, or even a Citizens to Save Scenic 37, may spring up.

What was originally simply a nominal construct, then, becomes quite real as people organize their lives around it.

This all seems arbitrary but necessary in both highways and radio. What, then, does it have to do with the new electronic media?

Simply this: till now the structures of media somehow sprang naturally from the nature of things. Now they don't anymore. Radio, books and movies have a natural inner dynamic of their own, leading to such constructs. While this may prove to be so for computer media as well (--as I argued in "Getting It Out of Our System," cited p.DM 15), then again it may not. In other words, WE MUST ACKNOWLEDGE THAT WE ARE INVENTING PRESENTATIONAL TECHNIQUES IN THE NEW MEDIA, not merely transporting or transposing particular things into them because they seem right. The psychological constructs of man-machine systems may turn out to be largely arbitrary. Thus bringing to terminal systems conventions like dialogue instruction ("CAI"), or arbitrary restrictions of how things may be connected, presented or written on the computer may be a great mistake.

The highway-number analogy continues. The older highways had numbers for convenience, and our travels became organized around them, and particular highways (like "U.S. 1" and "Route 66") came to have special character. But now with the Interstates, a highway is a planned, sealed unit, no longer just a collection of roads gathered together under a name.

This unit, the Interstate, is not merely a psychological construct, but a planned structure. Knowing what works and what doesn't in the design of fast highways, the Interstates were built for speed, structured as closed units. Designing them with limited access has been a conscious decision in the system design for well-based reasons, not a chance structure brought in from horse-and buggy days.

Now, the constructs of previous media-- writing, films, other arts-- evolved over time, and in many cases may have found their way to a "natural" form. But because of the peculiar way that computer media are currently evolving (--under large grants largely granted to professionals who use very large words to promote the idea that their original professions are largely applicable--), this sort of natural evolution may not take place. The new constructs of computer media, especially computer screen-media, may not have a chance to be thought out. We need designs for screen presentations and their mixture-- vignetting, Windows, screen mosaics, transformed and augmented views, and the rapid and comprehensible control of these views and windows. We are still just beginning to find clever viewing techniques, and have hardly begun to discover highly responsive forms of viewability and control (cf. collateration in "Thinkertoys," p.DM 52), and Knowlton's button-box (omitted) (See T. Nelson, "A Conceptual Framework for Man-Machine Everything," cited p. , and material on controls, below.)

THE MIND'S UNIFICATION

One of the remarkable things about the human mind is the way it ties things together. Perceptual unity comes out of nowhere. A bunch of irregular residential and industrial blocks becomes thought of as "my neighborhood." A most remarkable case of mental unification is afforded by the visage of our good friend Mickey Mouse. The character is drawn in a most paradoxical fashion: two globelike protrusions (representing the ears) are in different positions on the head, depending on whether we view him from the front or the side. No one finds this objectionable; few people even notice, it seems.

THE PARADOXICAL ANATOMY OF MICKEY MOUSE

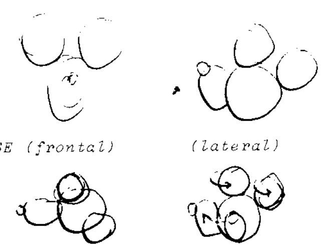

MICKEY MOUSE (frontal) *(lateral)*

POSSIBLE RECONCILIATIONS: Diagonal Mounting *Rolling Relative to Camera*

What this shows, of course, is the way the mind can unify into a consistent mental whole even things which are inconsistent by normal rules (in this case, the rules of three-dimensional structure).

Even perceptions are subject to the same principle of unification. The fingernail is an excrescence with no nerves in it; yet somehow you can feel things with your fingernails-- tying together disparate sensations into a unified sense of something in the world (say, a coin you're trying to pick up). In the same way, an experienced driver feels the road; in a very real sense, the car's wheels and suspension become his own sensory extensions.

This principle of mental unification is what makes things come together, both literally and figuratively, in a fantic field. A viewer sees two consecutive movie shots of streets and unifies them into one street; controls, if you are used to them, become a single fused system of options; we can have a sense of a greater whole, of which one view on a screen is a part.

THE GESTALT, DEAR BRUTUS
IS NOT IN OUR STARS
BUT IN OURSELVES.

CONTROLS: THEIR UNIFICATION AND FEEL

Controls are intimately related to screen presentation, just as arbitrary, and just as important.

The artful design of control systems is a deeply misunderstood area, in no way deconfused by calling it "human factors." There are many functions to be controlled, such as text editing operations, views of the universe on a screen, the heading of a vehicle, the tilt of an aircraft, the windage and adjustments of artillery, the temperature of a stove burner and any other controllable devices. And nowadays any conceivable devices could control them-- pushbuttons, knobs, cranks, wheels, levers and joysticks, trigger, dials, magic wands, manipulation by lightpen on CRT screens (see p.DM 31), flicks of the finger, the turning of the eyes (as in some experimental gun-aiming devices), the human voice (but that introduces problems-- see p.DM 13), keyboards, electronic tablets, Engelbart mice and chordwriters, and so on.

The human mind being as supple as it is, anything whatever can be used to control systems. The problem is having it be a comprehensible whole.

As already remarked, our ability mentally to unify things is extraordinary. That we somehow tie together clutch, gear, accelerator and brake into a comprehensible control structure to make cars go and stop should amaze and instruct.

Engineers and "human factors" people speak as though there were some kind of scientific or determinate way to design control systems. Piffle. We choose a set of controls, much like an artist's Palette, on the basis of general appropriateness; and then try best and most artistically to fit them to what needs doing.

The result must be conceptually clear and retroactively "obvious"-- simply because clarity is the simplest way to keep the user from making mistakes. Clear and simple systems are easier to learn, harder to forget, less likely to be screwed up by the user, and thus are more economical-- getting more done for the resources put in.

There is a sort of paradox here. The kinds of controls are totally arbitrary, but their unification in a good system is not. Smoothness and clarity can come from disparate elements. It is for this reason that I lay particular stress on my JOT system for the input and revision of text, using a palette of keys available on the simplest standard computer terminal, the 33 Teletype. I cannot make the final judgement on how good this system is, but it pleases me. JOT is also an important example because it suggests that a conceptually unified system can be created from the artful non-obvious combination of loose elements originally having different intended purposes.

Mental analogy is an important and clear control technique. We tend to forget that the steering wheel was invented, separately replacing both the boat's tiller and the automobile's tiller. We also forget that the use of such steering mechanisms must be actually learned by children. Such continuous analogies, though, require corresponding continuities in the space to be controlled-- an important condition.

Simplicity and clarity have nothing to do with the appearance of controls, but with the clarity and unique locatability of individual parts. For this reason I find deplorable the arrayed controls that are turning up, e.g. on today's audio equipment. Designers seem to think rows of things are desirable. On the contrary: the best designed controls I ever used are on the Sony TC-50 pocket tape recorder

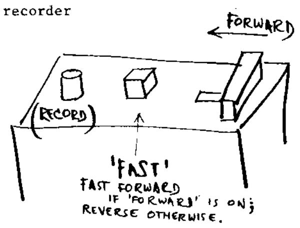

but of course this is now phased out; instead most cassette recorders have five or six stupid buttons in a row. (Was it too good to last?)

Spurious control elegance comes in many guises. Consider Bruce McCall's description of the Tap-A-Toe Futuroidic Footless De-Clutching™ system. This was offered on the fictitious 1934 Bulgemobiles, and allowed you to drive the car with one pedal, rather than three (see box nearby).

Careless and horrible designs are not all fictitious. One egregious example also indicates the low level of design currently going into some responding systems: computer people have designed CRT writing systems for newspapers which actually have a "kill" button on the console, by which authors would accidentally kill their stories. In a recent magazine article it was explained that the eventual solution was to change the program so that to kill the story you had to hit the "kill" button twice. To me this seems like a beautiful example of what happens when you let insulated technical people design the system for you: a "kill" button on the keyboard is about as intelligent as installing knives on the dashboard of a car, pointing at the passenger.

There is another poor tendency. When computer programmers or other technical people design particular systems without thinking more generally, things are not likely to be either simple or combinable. What may result is intricate user-level controls for one particular function, controls that are differently used for another particular function, making the two functions not combinable.

What makes for the best control structures, then? There is no simple answer. I would say provisionally that it is a matter of unified and conspicuous constructs in the mental view of the domain to be controlled, corresponding to a well-distinguished and clearly-interrelated set of controlling mechanisms. But that is hardly the last word on the subject.

THE ORGANIZATION OF WHOLENESS

It should be plain that in responding screen-systems, "what happens on the screen" and "how the controls respond" are not really distinguishable. The screen events are part of the way the controls respond. The screen functions and control functions merge psychologically.

Now, there is a trap here. Just as the gas pedal, clutch, gearshift and brake merge psychologically, any control structure can merge psychologically. Clutch and gear shift do not have, for most of us, clear psychological relevance to the problem of controlled forward motion. Yet we psychologically integrate the use of these mechanisms as a unified means for controlling forward motion (or, like the author, get an Automatic). In much the same way, any system of controls can gradually come through use to have a psychological organization, even spuriously. The trap is that we so easily lose sight of arbitrariness and even stupidity of design, and live with it when it could be so much better, because of this psychological melding.

But useful wholeness can be helped along. Just as what I have called the accordance-structure of writing (see "Writing," p.DM 43) moves it along smoothly, fantic design that builds from a well-organized internal dynamic should confer on a fantic system the same momentum and clarity that carefully-organized writing has.

This contribution of wholeness can only occur, however, if the under-level complications of a system have been carefully streamlined and smoothed back, at least as they affect the user. Consider the design of the JOT text editing system (p.DM 50): while it is simple to the user, computer people often react to it with indignation and anger because it hides what are to them the significant features of computer text editing-- explicit preoccupation with storage, especially the calling and revision of "blocks." Nevertheless, I say it is the details at this level which must be smoothed back if we are to make systems for regular people.

The same applies to the Th3 system (see p. DM 55), which is designed to keep the user clear-minded as he compares things in multiple dimensions. The mechanisms at the computer level must be hidden to make this work.

FANTIC SPACE

Pudovkin and Eisenstein, great Russian movie-makers of the twenties, talked about "filmic space"-- the imaginary space that the action seems to be in.

This concept extends itself naturally to fantic space, the space and relationships sensed by a viewer of any medium, or a user in any presenting or responding environment. The design of computer display systems, then, is really the artful crafting of fantic space. Technicalities are subservient to effects. (Indeed, I think computer graphics is really a branch of movie-making.)

FANTIC STRUCTURE

The fantic structure of anything, then, consists of its noticeable parts, interconnections, contents and effects.

I claim that it is the fantic unity-- the conceptual and presentational clarity of these things-- that makes fantic systems-- presentational systems and material-- clear and helpful, or not.

Let us take an interesting example from a system for computer-assisted instruction now under implementation. I will not identify or comment on the system because perhaps I do not understand it sufficiently. Anyway, they have an array of student control-buttons that look like this:

OBJ [objective]	HELP	ADVICE
MAP	HARDER	EASIER
RULE	EXAMP [example]	PRACT [practice]

The general thinking in this system seems to be that the student may get an overall organizing view of what he is supposed to be learning (MAP); information on what he is currently supposed to be about (OBJ); canned suggestions based on what he's recently done (ADVICE). Morever, he can get the system to present a rule about the subject or give him practice; and for either of these he may request easier rules or practice, or harder rules (i.e., more abstruse generalities) or harder practice.

For the latter, the student is supposed to hit RULE or PRACT followed by HELP, HARDER or EASIER, viz.:

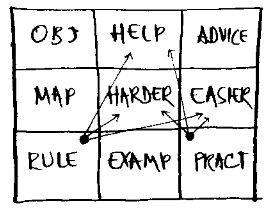

Now regardless of whether this is a well-thought-out way to divide up a subject-- I'll be interested to see how it works out-- these controls do not seem to be well-arranged for conceptual clarity. It seems to be the old rows-of-buttons approach.

I have no doubt that the people working on this system are certain this is the only possible layout. But consider that the student's options might be clearer to him, for instance, if we set it up as follows:

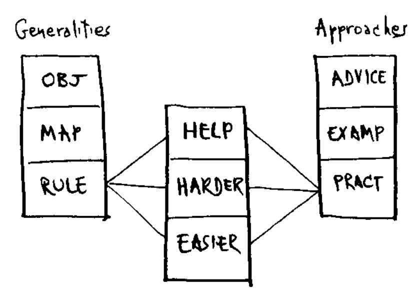

Or like this:

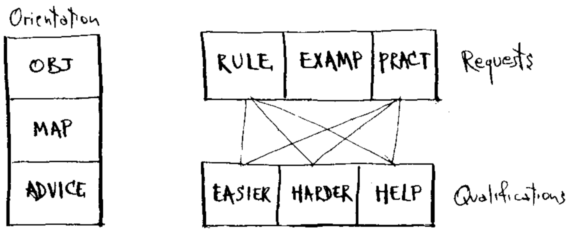

What I am trying to show here is that merely the arrangement of buttons creates different fantic constructs. If you see this, you will recognize that considering all the other options we have, designing new media is no small matter. The control structures merge mentally with the presentational structures. The temptation to settle on short-sighted designs having shallow unity is all too great.

FANTIC DESIGN

Fantic design is basically the planning and selection of effects. (We could also call these "performance values"-- cf. "production values" in movies.)

Some of these intended effects are simply the communication of information or cognitive structure-- "information transfer," to use one of the more obtuse phrases current. Other desirable effects include orienting the user and often moving him emotionally, including sometimes overwhelming or entrancing him.

In the design of fantic systems involving automatic response, we have a vast choice among types of presentational techniques, tricks that are just now becoming understood. Not just screen techniques and functions, but also response techniques and functions.

(If "feelie" systems are ever perfected, as in Huxley's Brave New World, it's still the same in principle. See Wachspress, p. DM9.)

In both general areas, though-- within media, and designing media-- it seems to me that the creation of organizing constructs is the most profound problem. In particular, the organizing constructs must not distract, or tear up contents. An analogy: in writing, the inventions of the paragraph, chapter and footnote were inventions in writing technique that helped clarify what was being expressed. What we need in computer-based fantic design is inventions which do not artificially chop up, constrain, or interfere with the subject (see box, Procrustes, nearby).

I do not feel these principles are everywhere sufficiently appreciated. For instance, the built-in structures of PLATO (see "Fantic Space of PLATO," p. DM 27) disturbs me somewhat in its arbitrariness-- and the way its control keys are scattered around.

But there is always something artificial-- that is, some form of artifice-- in presentation. So the problem is to devise techniques which have elucidating value but do not cut connections or ties or other relationships you want to save. (For this reason I suggest the reader consider "Stretchtext," p. DM19, collateral linkage (p. DM52), and the various hypergrams (p. DM 18-19). These structures, while somewhat arbitrary and artificial, nevertheless can be used to handle a subject gently.)

An important kind of organizing construct is the map or overall orienting diagram. This, too, is often partly "exact" and partly "artifice:" certain aspects of the diagram may have unclear import but clear and helpful connotation. (For instance, consider the "picture systems" diagram on p. DM 20 -- just what does the vertical dimension mean? Yes, but what does it really mean?)

Responding systems now make it possible for such orienting structures to be multidimensional and responding (cf. the orienting function of the "dimensional flip" control illustrated on p. DM 31).

Fantic design, then, is the creation either of things to be shown (writing, movie-making, etc.) at the lower end, or media to show things in, or environments.

1. The design of things to be shown-- whether writing, movie-making, or whatever-- is nearly always a combination of some kind of explicit structure-- an explanation or planned lesson, or plot of a novel-- and a feeling that the author can control in varying degrees. The two are deeply intertwined, however.

The author (designer, director, etc.) must think carefully about how to give organization to what is being presented. This, too, has both aspects, cognition and feelings.

At the cognitive end, the author must concern himself with detailed exposition or argument, or, in fiction, plot. But simply putting appropriate parts together is not enough: the author must use organizing constructs to continually orient the reader's (or viewer's) mind. Repeated reference to main concepts, repeated shots (in a movie) of particular locations, serve this function; but each medium presents its own possible devices for this purpose.

The organization of the feelings of the work criss-crosses the cognitive; but we can't get into it here.

Selection of points and parts contributes to both aspects. If you are trying to keep the feeling of a thing from being ponderous, you can never include everything you wanted, but must select from among the explicit points and feeling-generators that you have thought of.

2. The design of media themselves, or of media subsystems, is not usually a matter of option. Books, movies, radio and TV are given. But on occasion, as for world's fairs or very personal projects, we have a certain option. Which allows things like:

- Smellavision or whatever they called it: movies with a smell-track, which went out into the theater through odor generators.
- Branching movies (see p.DM44).
- "Multi-media" (multiple audio tracks and simultaneous slide projections on different screens).
- Stereo movies.

And so on. The thing about the ones mentioned is that they are not viable as continuing setups for repeated productions. They do not offer a permanent wide market; they are not stable; they do not catch on. Which is in a way, of course, too bad.

But the great change is just about now. Current technicalities allow branching media-- especially those associated with computer screens. And it is up to us now to design them.

3. MENTAL ENVIRONMENTS are working places for structured activity. The same principles of showmanship apply to a working environment as to both the contents of media and the design of media. If media are environments into which packaged materials are brought, structured environments are basically environments where you use non-packaged material, or create things yourself. They might also be called "contentless media." The principles of wholeness in structured environments are the same as for the others, and many of our examples refer to them.

The branching computer screen, together with the selfsame computer's ability to turn anything else on and off as selected by the user, and to fetch up information, yields a realm of option in the design of media and environment that has never existed before. Media we design for screen-based computer systems are going to catch on widely, so we must be far more attentive to the options that exist in order to commit-- nationally, perhaps-- to the best.

In tomorrow's systems, properly unified controls can give us new flexibilities. If deeply well-designed, these promise magnificent new capabilities. For instance, we could allow a musician to "conduct" the performance of his work by a computer-based music synthesis system (see "Audio," p.DM11), perhaps controlling the many qualities of the performance on a screen as he goes, by means of such techniques as dimensional flip (see p.DM31). (The tradition of cumulative audio synthesis, as practiced in the fifties by Les Paul and Mary Ford, and more recently by Walter Carlos and Mike Oldfield, will take on a new fillip as multidimensional control techniques become common.)

One of the intents of this book has been to orient you to some of the possibilities and some of the options, considered generally. There is not room, unfortunately, to discuss more than one or two overall possibilities in detail. The most successful such system so far has been PLATO (discussed pp. DM18-19); others could not be listed for lack of space.

NEW MEDIA TO LAST

What's worse, we are confronted not merely with the job of using computers to present specific things. The greater task is to design overall computer media that will last us into a more intelligent future. Adrift in a sea of ignorance and confusion, it is nevertheless our duty to try to create a whole transportation system that everybody can climb aboard. For the long run, fantic systems must be treated not as custom systems for explicit purposes, but as OVERALL GENERAL DESIGNS WHICH WILL HAVE TO TIE TOGETHER AND CATCH ON, otherwise collapse and perish.

FINAL CONSEQUENCES.

It seems to me certain that we are moving toward a generalized and universal Fantic system; people can and should demand it. Perhaps there will be several; but if so, being able to tie them together for smooth transmission is essential. (Think of what it would be like if there were two kinds of telephones?) This then is a great search and crusade: to put together truly general media for the future, systems at which we can read, write, learn and visualize, year after year after year. The initiatives are not likely to come from the more conventional computer people; some of them are part of the problem. (Be prepared for every possible form of aggressive defensiveness from programmers, especially: "Why would you want that?" The correct answer is BECAUSE, dammit!)

But this all means that interior computer technicalities have to be SUBSERVIENT, and the programmers cannot be allowed to dictate how it is to behave on the basis of the underlevel structures that are convenient to them. Quite the contrary: from the fullest consideration of the richest upper-level structures we want, we the users-to-be must dictate what lower-level structures are to be prepared within.

But this means you, dear reader, must develop the fantic imagination. You must learn to visualize possible uses of computer screens, so you can get on down to the deeper level of how we are going to tie these things together.

The designer of responding computer systems is creating unified setups for viewing and manipulating things-- and the feelings, impressions and sense of things that go with them. Our goal should be nothing less than REPRESENTING THE TRUE CONTENT AND STRUCTURE OF HUMAN THOUGHT. (Yes, Dream Machines indeed.) But it should be something more: enabling the mind to weigh, pursue, synthesize and evaluate ideas for a better tomorrow. Or for any at all.

BIBLIOGRAPHY

Theodor H. Nelson, "A Conceptual Framework for Man-Machine Everything." Proc. NCC 73.

-----, "Computopia and Cybercrud." In Levien (ed.), Computers in Instruction, The Rand Corporation, 1971.

JOT™: Juggler Of Text.

From "A Human Being's Introduction to the JOT System." ©1972 T. Nelson.

Here's how simple it is to create and edit text with the JOT system. Since your typewriter is now a JOT machine, not every key does what it used to.

CREATING TEXT: just type it in.

You type: The quick brown fox jumps over the lazy dog.

It types: The quick brown fox jumps over the lazy dog.

REVIEWING A SENTENCE YOU JUST TYPED: the back-arrow takes you back, the space bar steps you through.

You type: ← sp sp sp sp

It types: (bell) The quick brown fox

DELETIONS AND INSERTIONS: the RUBOUT key rejects words you don't want. To insert , merely type.

You type: ← sp sp RUBOUT lithe sp sp sp sp sp sp

It types: (bell) The quick /brown/ lithe fox jumps over the lazy dog.

REARRANGING TEXT: first we make three Cuts in the text, signalled by free-standing exclamation points.

You type: sp ! sp ! sp ! fox

It types: The ! quick ! lithe ! fox

TO REARRANGE IT, YOU TYPE: LINE FEED key. This exchanges the two pieces between the cuts.

CHECK THE RESULTS:

← sp sp sp sp

(bell) The lithe quick fox

THE WALKING NET™

→ A one-minute system that three-year-olds can learn [Sorry to have to show you a writeup, instead of the real thing.]

Another application of the present invention is also in the area of pictorial display, but offers a great variety of potential user choices in a simple circumstance. I call this the "walking net" system because control is effected through a changing network of choices which step, or "walk," around the screen.

The problem of intricate computer graphics may be phrased as follows: given that a digital system can hold a wide variety of graphical materials ready to present, how may the user most simply and conveniently choose them? Indeed, how may the user keep track of what is happening, where he is and where he has been?

The external mechanism I have selected for this facility paradoxically combines great versatility for sophisticated presentations with great simplicity before the naive user. The idea is this: the user may command a continuing succession of changing presentations, making only one simple choice at a time, yet receiving intricate and rich animations with extremely clear continuity on the screen.

The exterior mechanism is this: along with an arbitrary graphic presentation on the screen, the user is continuously presented with the image of a forking set of labelled arrows, e.g.:

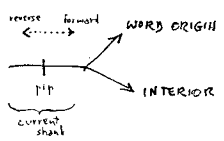

The pip is a conventional light-pen cursor. The "current shank" is a line whose implicit gradations control developments in the picture; and the choice of arrows at the end of the current shank determine a discrete choice between alternatives that are to transpire.

The user, seizing the pip with the lightpen, moves it (through the usual lightpen techniques) sideways along the current shank. Moving it in the "forward" direction causes progressive developments in the picture, moving it "backward" causes a reversal of animations and other previous developments.

When the pip reaches the choice point in the forward direction, the user may drag it (through the usual lightpen techniques) along either of the beckoning alternatives. This then causes further developments in the presentation consonant with the line selected.

"Developments" of the picture here include expansion, contraction, sliding movements and frame-by-frame animation.

(These materials will have been, of course, explicitly input by authors and artists.)

In a sample employment, consider a presentation on the subject of Volcanoes. Let the first shank of the control net control the "rise of a volcano from the sea"-- an undulating ocean surface pierced first by a wisp of smoke, then a growing peak, with rivulets of lava seen to run down its sides and darken as they contribute to its growth.

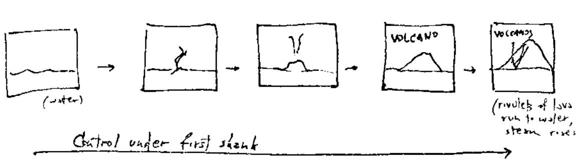

At the end of the first shank, the user may branch to two arrows, labelled respectively WORD ORIGIN and INTERIOR. Either option continues the presentation without a break, retaining much of the picture on the screen. Selection of WORD ORIGIN causes the word VOLCANO to change to VULCAN, and a picture of the god Vulcan is seen to seize a lightning bolt rising from the crater; text appears to explain this. Alternatively, if the user chooses INTERIOR, the tubes and ducts within the volcano appear, and explanatory text also.

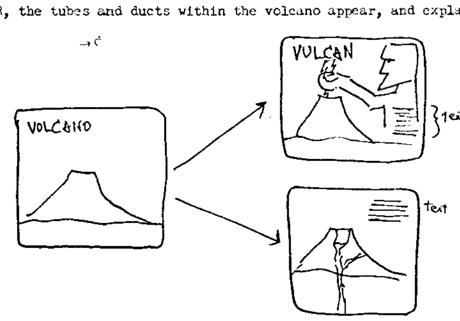

(The path unchosen fades from the screen, as does the previous shank.)

Either of these alternatives may continue with its own developments and animations under control of its own shank.

Several features of this control application are of special interest. One is that the presentation may be continuous in all directions, aiding in continuous user orientation. Another is that presentations are reversible in various ways, an aid both in user orientation and self-study. (Not only is a demonstration reversible within a given shank, but the user may back the pip through an intersection into the antecedent shank-- which reappears at the juncture as the lightpen backs up-- and the user may continue to reverse the presentation through that preceding shank, or to re-enter the intersection and make another choice, "the path not taken.") These features allow the user clearly to repeat demonstrations as often as he likes and to explore numerous alternatives.

The displayed control net is thus to be understood as a large network of choices, mostly unseen, whose currently visible portion "walks" around the screen as use progresses. Within this system, then, numerous variants are possible. For instance, the currently visible portion of the net may itself be whimsically incorporated in a picture, viz.:

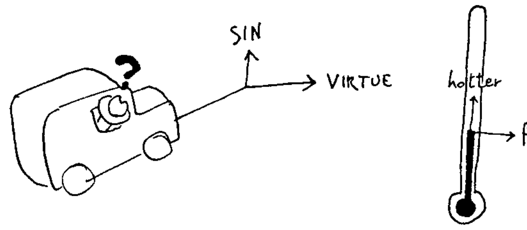

PROCRUSTES THE GIANT

The Greeks told of a giant named Procrustes (rhymes with Rusty's) who was very hospitable to passing travelers. He would invite, indeed compel them, to sleep in his bed. Unfortunately, because it was a very odd bed, he had to cut them up first...

Procrustes has haunted conversations ever since; and any time we are forced to use categories that don't properly fit a subject, it seems like an invitation to the Procrustean bed.

Hypertext systems at last offer total freedom from arbitrary categorizing and chopping; but in some systems for storing and presenting information, I can't help hearing the whisk of Procrustes' knife--

"Take new Tap-A-Toe Futuroidic Footless De-Clutching. Instead of old-fashioned gas, brake and clutch pedals that kept your feet busier than a dance marathon, Tap-A-Toe Futuroidic Footless De-Clutching offers the convenience of Single Pedal Power Control-- combines all foot functions in one single pedal!

"Think of it: one tap-- you go, moving off faster than a barfly after Repeal.

"Two taps-- you change gears, as smooth and automatic as a mortgage foreclosure.

"Three taps-- you stop quicker than the U.S. economy.

"And that's all there is to it. Tap-A-Toe Futuroidic Footless De-Clutching with Single Pedal Control is as easy and effortless as the Jap march on Manchuria!"

Bruce McCall,
"1934 Bulgemobile Brochure,"
National Lampoon, May 74, 76-7.

STELLAVISION

A nice example of a unified presentational system would allow you a "feelie" glove along with your computer display-- the sort of thing Mike Noll has been doing at Bell Labs.

Now, suppose you are playing with a diagram of a star on a computer display screen. It's all very well to see its layers, flowing arrows representing convection currents, promontories and so on-- but some things you ought to be able to feel. For example, the mechanical resonance-properties of stars. It would be nice to be able to reach and grasp the star, to squeeze it and feel its pulsations as it regains its shape. This could be done in the glove-- at the same time the image of the glove grasps the star on the screen, and the star is squished.

Of course, to build such a responding glove, particularly one that gave you subtle feelings back in your fingers, would probably be very expensive. But it's the kind of possibility people should start considering.

Should I have called it

TEACHOTECHNICS?
SHOWMANSHIPNOGOGY?
INTELLECTRONICS? [S. Ramo]
THOUGHTOMATION?
MEDIA-TRONICS?

ABOUT THE TERM 'FANTICS.'

First of all, I feel that very few people understand what interactive computer systems are about. It's like the story of the blind men and the elephant-- each thinks it's a different thing (based, usually, on his own technical specialty).

But I think it's all show business. PENNY ARCADES are the model for interactive computer systems, not classrooms or libraries or imaginary robot playmates. And computer graphics is an intricate branch of movie-making.

Okay, so I wanted a term that would connote, in the most general sense, the showmanship of ideas and feelings-- whether or not handled by machine.

I derive "fantics" from the Greek words "phainein" (show) and its derivative "phantastein" (present to the eye or mind).

You will of course recognize its cousins fantastic, fantasy, phantom. ("Phantom" means what is shown; in medical illustration it refers to an opaque object drawn as transparent; a "phantom limb" is an amputee's temporary feeling that the severed limb has been restored.) And a fantast is a dreamer.

The word "fantics" would thus include the showing of anything (and thus writing and theater), which is more or less what I intended. The term is also intended to cover the tactics of conveying ideas and impressions, especially with showmanship and presentational techniques, organizing constructs, and fundamental structures underlying presentational systems.

Thus Engelbart's data hierarchy (P. DM 46-7), SKETCHPAD's Constraints (p. DM 23), and PLATO's fantic spaces (p. DM 26-7) are fantic constructions that need to be understood if we are to understand these systems and their potential usages.

Livermore Labs, those hydrogen-bomb design people, will have a "Laboratory for Data Analysis," an opulent facility for experimenting with multidimensional visualization.

One of your jolly ironies. I have seen pictures of beautiful multibutton control handles which were designed for project SMASH, would you believe Southeast [Asia] Multisensory Armament System for Helicopters. Aargh.)

The best with the worst.

Everything is deeply intertwingled.

Designing screen systems that focus the user's thought on his work, with helpful visualizations and no distractions, is the great task of fantic design.

In a system I designed for CRT motion-picture editing, the user could maipulate written descriptions on the screen (corresponding to the usual yellow-pad notes). To see the consequences of a particular splice, for instance, the editor would only have to draw a line between two annotated lines representing shots. Trim variations could be seen by moving this cut-line (illustrated).

Not long after, CBS and Memorex did introduce a system for movie-editing by CRT-- but I've heard that in their system the user has to actually deal with numbers. If so, this is missing the whole point.

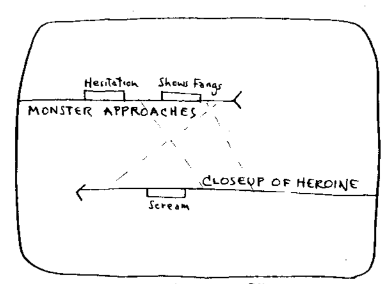

CINENYM™ ©1968 T. Nelson

THINKERTOYS

Our greatest problems involve thinking and the visualization of complexity.

By "Thinkertoy" I mean, first of all, a system to help people think. ('Toy' means it should be easy and fun to use.) This is the same general idea for which Engelbart, for instance, uses the term "augmentation of intellect."

But a Thinkertoy is something quite specific: I define it as a computer display system that helps you envision complex alternatives.

The process of envisioning complex alternatives is by no means the only important form of human thought; but it is essential to making decisions, designing, planning, writing, weighing alternate theories, considering alternate forms of legislation, doing scholarly research, and so on. It is also complicated enough that, in solving it, we may solve simpler problems as well.

We will stress here some of the uses of these systems for handling text, partly because I think these are rather interesting, and partly because the complexity and subtlety of this problem has got to be better understood: the written word is nothing less than the tracks left by the mind, and so we are really talking about screen systems for handling ideas, in all their complexity.

Numerous types of complex things have to be inter-compared, and their relations inter-comprehended. Here are a few of the many types:

Alternative designs.

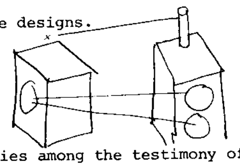

Discrepancies among the testimony of witnesses.

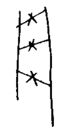

Successive drafts of the same document.

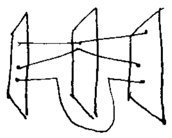

Pairs of things which have some parts the same, some parts different (contracts, holy books, statutes of different states, draft versions of legislation...)

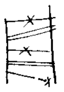

Different theories and their ties to particular examples and evidences.

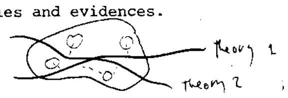

Under examination these different types of inter-comparison seem to be rather different. Now, one approach would be to create a different data structure and viewing technique for each different type of complex. There may be reasons for doing that in the future.

For the present, however, it makes sense to try to find the most general possible viewing technique: one that will allow complex intercomparisons of all the types mentioned, and any others we might run across.

One such technique is what I now call collateration, or the linking of materials into collateral structures,* as will be explained. This is fairly straightforward if you think enough about the problem; Engelbart discovered it independently.

Let us call two structures collateral if there are links between them, connecting a selected part of one with a selected part of the other. The sequences of the connected parts may be different. For simplicity's sake, suppose each one is a short piece of writing. (We will also assume that there is some convenient form of rapid viewing and following between one end of a link and another.)

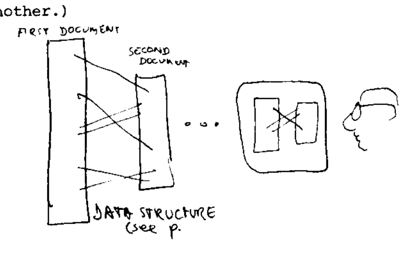

Now, it will be noted first off that this is an extremely general method. By collateral structuring we can easily handle the equivalents of: tables of contents; indexes; comments and marginalia; explanations, exegesis, explication; labeling; headings; footnotes; notes by the writer to himself; comments and questions by the reader for later reference; and additional details out of sequence.

Collateration, then, is the creation of such multiple and viewable links BETWEEN ANY TWO DATA STRUCTURES, in principle. It is general and powerful enough to handle a great variety of possible uses in human intellectual endeavor, and deserves considerable attention from researchers of every stripe.**

The problem then, is how to handle this for rapid and convenient viewing and whatever other work the user wants to do-- writing and splicing, intercomparing, annotating and so on. Two solutions appear on this spread: The Parallel Textface™, designed as a seminal part of the Xanadu system (see p.DM 56), which I hope will be marketed with that system in the near future, and a more recent design which I've worked on at the University of Illinois, the 3D Thinkertoy or Th3.

CLARITY AND POWER

We stressed on the other side of the book that computer systems must be clear, simple and easy to use. Where things like business uses of computers are concerned, which are intrinsically so simple in principle, some of the complications that people have been forced to deal with in ill-designed computer systems verge on the criminal. (But some computer people want others to think that's the way it has to be. "Your first duty is to keep your job, right?" one computer person said to me recently. "It wouldn't do to set up systems so easy to use that the company wouldn't need you anymore." See "Cybercrud," p.8.)

But if it is desirable that computer systems for simple-minded purposes be easy to use, it is absolutely necessary that computer systems for complicated purposes be simple to use. If you are wrangling over complex alternatives-- say, in chess, or in a political simulation game (see "Simulation," p. 58), or in the throes of trying to write a novel, the last thing you will tolerate is for your computer screen to introduce complications of its own. If a system for thinking doesn't make thinking simpler-- allowing you to see farther and more deeply-- it is uselss, to use only the polite term.

But systems can be both powerful and simple at the same time. The myth that things have to be complicated to do anything for you is pernicious rubbish. Well-designed systems can make our mental tasks lighter and our achievements come faster.

It is for this reason that I commend to the reader these two designs of mine: as examples of user-level control and viewing designs-- fantic environments, if you will (see p.DM 48-51)-- that are pruned and tuned to give the user great control over the viewing and cross-consideration of intricate alternatives, without complication. I like to believe that both of these, indeed, are ten-minute systems-- that is, when we get them running, the range of uses shown here can be taught to naive users in ten minutes or less.

It is because of my heartfelt belief in this kind of simplicity that I stress the creation of prefabricated environments, carefully tuned for easy use, rather than the creation of computer languages which must be learnt by the user, as do such people as Engelbart (see p.DM 46) and DeFanti (see p.DM 31). Now, their approach obviously has considerable merit for sophisticated users who want to tinker repeatedly with variant approaches. For people who want to work incessantly in an environment, and on other things-- say writers-- and are absent-minded and clumsy and nervous and forgetful (like the present author), then the safe, prefabricated environment, with thoroughly fail-safe functions and utterly memorable structural and control interrelationships, is the only approach.

* In my 1965 paper (see bibliography) I called collateral structures zippered lists.

** A group at Brown University has reportedly worked along these lines since I worked with them, but due to certain personal animosities I have not kept up with their developments. It will be interesting to see what kind of response they can get out of the IBM systems they are using.

BIBLIOGRAPHY

Theodor H. Nelson, "A File Structure for the complex, the Changing and the Indeterminate." Proc. ACM 65, 84-100.

-----, "Simplicity Versus Power in User Systems." Unpublished.

WE OFTEN WANT TO SAVE ALTERNATIVES.

28

And closely to my [heart] she press'd, (side)
And closer still with bashful art,
[And ask'd me with her swimming eyes]
That I [would] rather feel than see (might)
The swelling of her Heart
[Her gentle Bosom rise.—]

29

[And now serene, serene & chaste,] I calm'd her fears; & she was calm
(But seen in calm and solemn tone)
[She] told her love with maiden pride; (And)
And so I won my Genevieve,
My [bright] & lovely Bride.] (dear)

From Coleridge's Poems: A Facsimile Reproduction of the Proofs and MSS. of some of the Poems. (Folcroft, 1972.)

DECISION/CREATIVITY SYSTEMS

[THINKERTOYS]

Theodor H. Nelson
19 July 1970

It has been recognized from the dawn of computer display that the grandest and most important use of the computer display should be to aid decisions and creative thought. The work of Ivan Sutherland (SKETCHPAD) and Douglas Engelbart have really shown how we may use the display to visualize and effect our creative decisions swiftly and vividly.

For some reason, however, the most important aspect of such systems has been neglected. We do not make important decisions, we should not make delicate decisions, serially and irreversibly. Rather, the power of the computer display (and its computing and filing support) must be so crafted that we may develop alternatives, spin out their complications and interrelationships, and visualize these upon a screen.

No system could do this for us automatically. What design and programming can create, however, is a facility that will allow us to list, sketch, link and annotate the complexities we seek to understand, then present "views" of the complexities in many different forms. Studying these views, annotating and refining, we can reach the final designs and decisions with much more in mind than we could otherwise hold together in the imagination.

Some of the facilities that such systems must have include the following:

Annotations to anything, to any remove.

Alternatives of decision, design, writing, theory.

Unlinked or irregular pieces, hanging as the user wishes.

Multicoupling, or complex linkage, between alternatives, annotations or whatever.

Historical filing of the user's actions, including each addition and modification, and possibly the viewing actions that preceded them.

Frozen moments and versions, which the user may hold as memorable for his thinking.

Evolutionary coupling, where the correspondences between evolving versions are automatically maintained, and their differences or relations easily annotated.

In addition, designs for screen "views", the motion, appearance and disappearance of elements, require considerable thought and imagination.

The object is not to burden the user, or make him aware of complexities in which he has no interest. But almost everyone in intellectual and decision pursuits has at some time an implicit need for some of these facilities. If people knew they were possible, they would demand them. It is time for their creation.

A full-fledged decision/creativity system, embracing both text and graphics, is one of the ultimate design goals of Project XANADU.

We might also think of them as systems for

THE MANAGEMENT OF LOOSE ENDS.

The Parallel Textface™

This user-level system is intended to aid in all forms of writing and scholarship, as well as anywhere else that we need to understand and manipulate complex clusterings of text (i.e., thought). It will also work with certain animated graphics.

The Parallel Textface, as described here, furnished the initial impetus for the development of the Xanadu™ system (see p. DM 56). Xanadu was developed, indeed, originally for the purpose of implementing some of these functions, but the two split apart. It turned out that the Parallel Textface required an extremely unusual data struc- structure and program techniques; these then became the Xanadu system. As developed in the final Xanadu design, they turn out to handle some very unusual kinds of screen animation and file retrieval. But this grew out of structuring a system to handle the functions described here.

Thus the Parallel Textface basically requires a Xanadu system.

It is hoped that this system can be sold complete (including minicomputer or microprocessor--no connection to a large computer is required) for a few thousand dollars by 1976 or 1977. See p. (Since "business people" are extremely skeptical as to whether anybody would want such a thing, I would be interested in hearing expressions of interest, if any.)

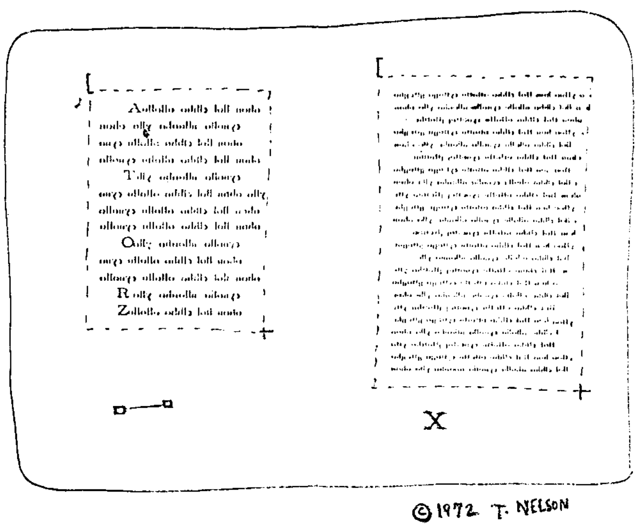

As shown here, the screen presents two panels of text; more are allowed. Each contains a segment of a longer document. ("Page" would be an improper term, since the boundary of the text viewed may be changed instantly.)

The other odds and ends on the screen are hidden keys to control elements which have been made to fade (in this illustration), just to lessen the distraction.

Panel boundaries and control graphics may be made to appear by touching them with the lightpen.

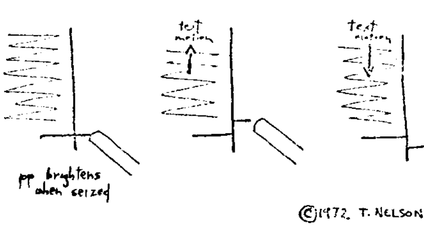

ROVING FUNCTIONS

The text moves on the screen! (Essential.) The lower right hand corner of each text panel contains an inconspicuous control diagram. The slight horizontal extension is a movable control pip. The user, with his light pen, may move the pip up or down. "Up" causes the text to move smoothly upward (forward in the material), at a rate proportional to how far you push the pip; "down" causes it to move back. (Note that we do not refer here to jerky line-by-line jumps, but to smooth screen motion, which is essential in a high-performance system. If the text does not move, you can't tell where it came from.)

PARALLEL TEXTFACE (1971)

Real person sits at cardboard Xanadu mockup.

"Nice keyboard. But what happened to your typewriter?"

Two panels are about right for a 10 x 10 screen.

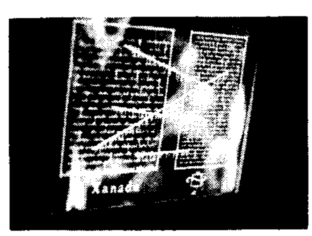

Independent text pulls dependent text along. Painted streaks simulate motion, not icicles.

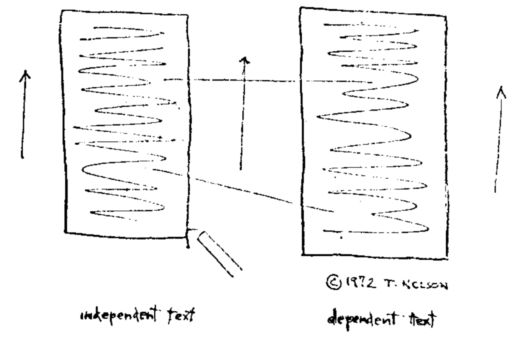

DERIVATIVE MOTION: when links run sequentially, connecting one-after-the-other on both sides, the contents of the second panel are pulled along directly: the smooth motion in one panel is matched in the other. This may be called derivative motion, between independent text (being handled directly with the lightpen) and dependent text (being pulled along). The relationship may be reversed immediately, however, simply by moving the lightpen to the control pip of the other panel, whose contents then become the independent text.

Irregularities in the links will cause the independent text to move at varying speeds or jump, according to an average of the links' connectivity.

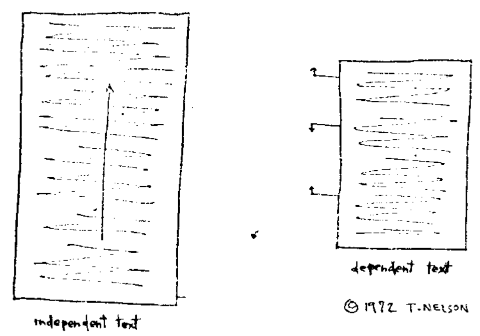

If no links are shown, the dependent text just stops.

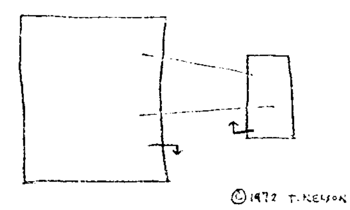

Collateral links between materials in the two panels are displayed as movable lines between the panels. (Text omitted in this diagram; panel boundary has been made to appear.)

Some links may not have both their endpoints displayed at once. In this case we show the incomplete link as a broken arrow, pointing in the direction of the link's completion.

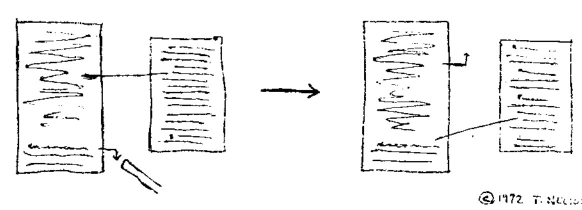

The broken arrow serves not merely as a visual pointer, but as a jump-marker leading to the linked material. By zapping the broken arrow with the lightpen, the user summons the linked material--as shown by the completion of the link to the other panel. (Since there has been a jump in the second panel, we see that in this case the other link has been broken.)

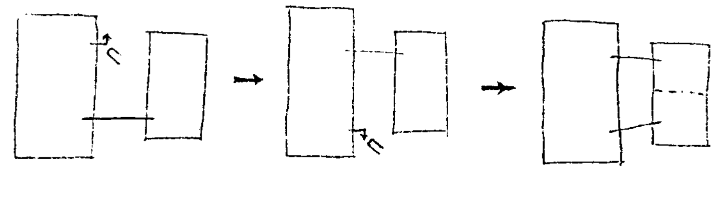

When such links lead to different places, both of these destinations may nevertheless be seen at once. This is done by pointing at both broken links in succession; the system then allows both links to be completed, breaking the second panel between the two destinations (as shown by dotted line across panel).

XANADU

©1972 T. NELSON

FAIL-SAFE AND HISTORICAL FEATURES.

In systems for naive users, it is essential to safeguard the user from his own mistakes. Thus in text systems, commands given in error must be reversible. For instance, Carmody's system (see p. DM47) requires confirmation of deletions.

Another highly desirable feature would allow the user to view previous versions, to see them collaterally with the corresponding parts of current versions, and even go back to the way particular things were and resume work from the previous version.

In the Parallel Textface this is all comprised in the same extremely simple facility. (Extremely simple from the user's point of view, that is. Inside it is, of course, hairy.)

In an egregious touch of narcissistic humor, we use the very trademark on the screen as a control device (expanded from the "X" shown in the first panel.)

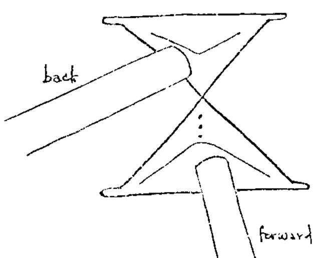

Actually the X in "Xanadu™," as it appears on the screen, is an hourglass, with a softly falling trickle of animated dots in the lower half, and Sands of Time seen as heaps above and below. These have a control, as well as a representative, function.

TO UNDO SOMETHING, YOU MERELY STEP "BACKWARD IN TIME" by dagging the upper part of the hourglass with the lightpen. One poke, one editing operation undone. Two pokes, two operations.

You may then continue to view and make changes as if the last two operations had never taken place. This effectively creates an alternative time-line.* However, if you decide that a previously undone edit operation should be kept after all, you may step forward-- stepping onto the previous time-line-- by using the lower half of the hourglass.

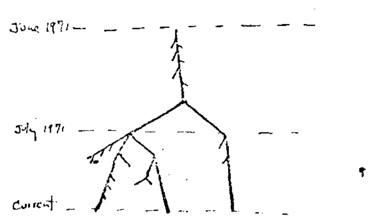

Revision Tree
©1972 T. Nelson

We see this clarified in a master time diagram or Revision Tree which may be summoned to the screen, never mind how. In this example we see that three versions are still "current," various other starts and variations having been abandoned. (The shaggy fronds correspond to short-lived variations, resulting from operations which were then reversed. In other words, "excised" time-lines, to use Gerrold's term-- see footnote.)

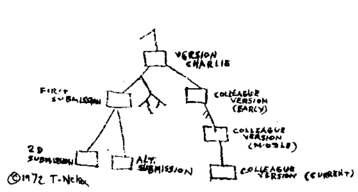

The user-- let's say he is a thoughtful writer-- may define various Versions or Drafts, here marked on the Revision Tree.

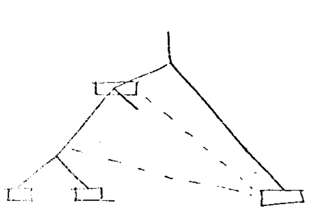

©1972 T. NELSON

He may, indeed, define collateral linkages between different versions defined at various Times in the Tree...

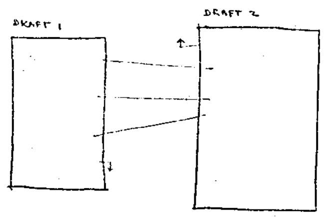

©1972 T. NELSON

... and see them displayed collaterally; and revise them further.

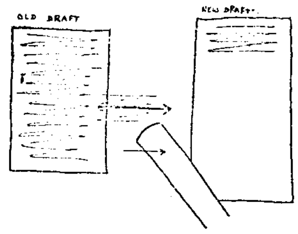

©1972 T. NELSON

Materials may be copied between versions. (Note that in the copying operation of the Parallel Textface, you actually see the moved text moved bodily as a block.)

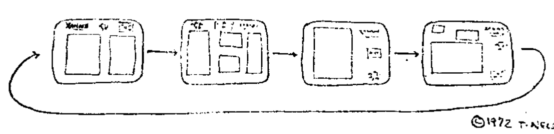

©1972 T. NELSON

GETTING AROUND

The user may have a number of standby layouts, with different numbers of panels, and jump among them by stabs of the lightpen.

Importantly, the panels of each can be full, each having whatever the contents were when you last left it.

File Web™

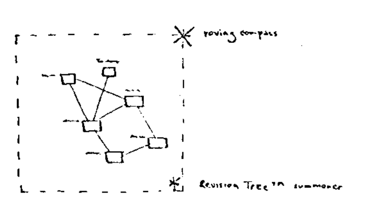

©1972 T. NELSON

The File Web™ is a map indicating what (labelled) files are present in the system, and which are collaterated.

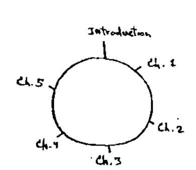

File Star™
(example)
©1972 T. Nelson

The File Star™ is a quick index into the contents of a file. It expands as long as you hold the lightpen to the dot in the center, with various levels of headings appearing as it expands. Naturally, you may jump to what you point at.

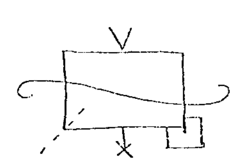

©1972 T. NELSON

EDITING

Rather than giving the user anything complicated to learn, the system is completely visual. All edit controls are comprised in this diagram, the Edit Rose™. Viz.:

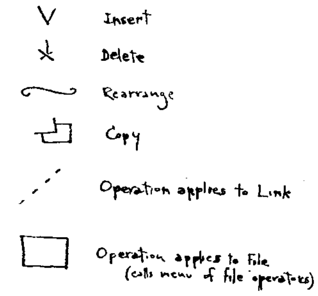

©1972 T. NELSON

Separate portions of the Edit Rose invoke various edit operations. (You must also point with the lightpen to the necessary points in the text: once for insert, twice for Delete, three or four times for Rearrange, three times for Copy.)

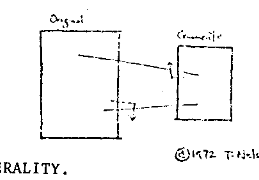

©1972 T. Nelson

GENERALITY.

The system may be used for comments on things,

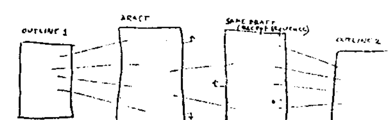

©1972 T. NELSON

for organizing by multiple outlines or tables of contents;

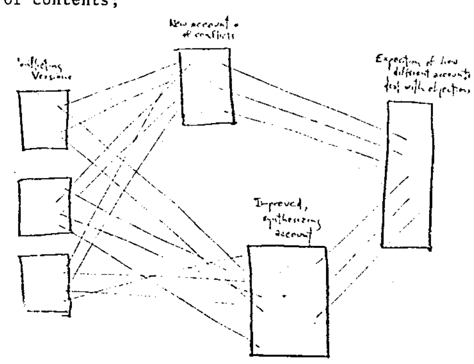

©1972 T. NELSON

and as a Thinkertoy, organizing complex alternatives. (The labels say: "Conflicting versions," "New account of conflicts," "Exposition of how different accounts deal with objections," "Improved, synthesizing account."

In other words, in this approach we annotate and label discrepancies, and verbally comment on differences in separate files or documents.

In ways this may seem somewhat obtuse. Yet above all it is orderly, and the complex of collateral files has a clarity that could be all-too-easily lost in systems which were programmed more specifically to each problem.

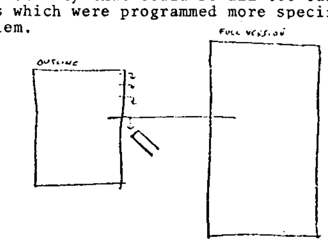

©1972 T. NELSON

The fundamental strength of collateration, seen here, is of course that any new structure collateral to another may be used as a table of contents or an outline, taking the user instantly to parts which are of interest in some new context.

* Oddly, this has the same logical structure as time-travel in science-fiction.

There are basically three alternate premises of of time-travel: 1) that the past cannot be changed, all events having preceded the backstep; 2) that the past can be changed; and 3) that while time-travelers may be deluded into thinking (2), that (1) is really the case-- leading to various appointment-in-Samarra plots.

Only possibility (2) is of interest here, but there are various alternative logics of mutability and time-line stepping. One of the best I have seen is in The Man Who Folded Himself by David Gerrold (Popular Library, 1973): logic expounded pp. 64-8. I am bemused by the parallel between Gerrold's time-controls and these, worked out independently.

Th3

A VERY ADVANCED (?) TEXT SYSTEM.
Read this at your peril. Multidimensional Concept-freaks only.

This design, Th3 (Thinkertoy in 3 dimensions), is one I have been working at while on the faculty of the University of Illinois. It is designed specifically for implementation using DeFanti's GRASS language (see p. DM 31), and the Vector General 3D display (see p. DM 30). Whether it will ever be actually programmed depends, of course, on numerous factors.

It is meant to be a very high-power thinkertoy, suitable for experimentation with creative processes, especially writing and three-dimensional design. (There is no room to discuss the latter here.) It is suited especially to the visualization of tentative structures in amorphous clusters. In some of its features it goes considerably beyond the more "commercial" thinkertoy system, the Parallel Textface (elsewhere in this spread).

Nevertheless, the same design criteria apply: a well-designed computer environment for any purpose should be learnable in ten minutes; otherwise the designer has not been doing his job. (I mean it would be learnable in ten minutes if you and I had it in front of us, working. This description will have to be weird and abstruse, I'm sorry to say.)

This system is designed around a three-dimensional display screen (the Vector General display, as manipulable by the GRASS language).

Now, most people do not think of text as three-dimensional. Laymen think of it as two-dimensional, since it's usually printed on rectangular pages. Computer people ordinarily think of it as one-dimensional, as a long string of characters and spaces--essentially what you'd get if you printed things in one line on a long, long ribbon. Well, frankly, I don't think of text as three-dimensional either; but like anything else, it has numerous qualities or dimensions, any three of which it's nice to be able to view at once (see "Dimensional Flip,") p. DM 31). And that's essentially the idea: the three dimensions we'll look at at any one time will be a particular view of a larger whole.

Now, the basic form of storage will be one of those Nelson-structures that drives computer people batty. Specifically, the basic data structure will be clusters of points.

Puns sometimes reflect a higher reality. Now it turns out that this structure in fact reflects a great Folk Truth: written discourse does in fact consist of "points" which you intend to get across. That we here intend to have them rotate as dots upon a screen reflects this structure.

Writing is, in fact, a projection from the intended "points" to a finished exposition which embraces them. Now, this is very like the view of language held in modern linguistics, namely, that a finished sentence is a "surface structure" constructed out of basic sentence kernels chewed up by certain transformations. Well, I am just pointing out here that writing is a surface structure of "points" which have been embedded and spliced in a structure of transitions, accordance-notes and so forth (see p. DM 43)*

The general idea of the Th3 system, then, is that the user may view the "points" he wishes to make, variously upon the 3D viewing surface. Successive drafts, then, will all be projections, geometrically, from this interior structure of points.

Finally, the unifying idea that gives the system simplicity is this: all views will be on faces of a cube.

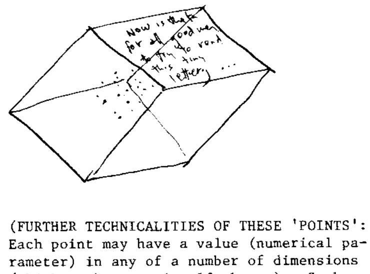

(FURTHER TECHNICALITIES OF THESE 'POINTS': Each point may have a value (numerical parameter) in any of a number of dimensions (which number may itself change). Such values may be null, as distinct from zero, showing that the point has no position on that particular scale.

Associated with each point may be one or more pieces and scraps or written material. Such scraps may be just phrases or single words. (Indeed, such scraps may be associated not just with a point, but with several specific values of a point.) Each scrap may also contain keywords.

Discrete relations between points may also be defined. There may be a variety of types of relation, which either exist between two points of don't.)

The crucial point here is that it's unified to the user: every version appears on a side of a box; and a typeset version is simply a magnified two-dimensional view in which the two dimensions are "position in overall text" (vertical) and "position on line" (horizontal).

Each side of a box may have a different view projected to it. This means that as many as three views of a specific cluster may be seen at once. However, for consistency these must have appropriately common dimensions.

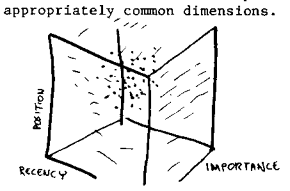

By rotation and zooming the user may focus on the original pieces, and work with them, writing and revising.

Moreover, by using a combination of zoom and hardware clipping (as available on this equipment), the user may restrict his work to a specific range of material on particular dimensions.

GALAXY AND BOX

There are basically two views of what you are working with: the Galaxy and the Box. They appear in various manifestations, allowing you to study discrete relations and structures in the material; various "dimensions" of the material; alternate versions and drafts to be made from the material; and the complex collateration (see under "Thinkertoys") of different structures.

In what follows we will discuss the screen functions but not the control structures, which have not firmed up particularly.

1. GALAXY VIEWS.

The points are seen as a cloud of dots on the screen. If no view coordinates are supplied, the dots will be randomly positioned.

A. "Star Trek" effect. Under a user's zoom control, the dots fly apart as if he is hurtling through space.

B. MAGNIFICATION. The user may "magnify" the dots, making each show its keywords, further text, and on up to the full Piece.

C. ROTATION. The 3D structure of the dots in space may be seen by the user at any time through short rotations.

D. Any relations that exist among the Points, insofar as they have been logged into the system, may be displayed among the points.

E. The user may sort the points by moving them with a lightpen.

F. The user may write within the individual pieces and splice them together, combining lightpen and keyboard operations.

2. BOX VIEWS

In the Galaxy Views, the individual Points simply swarm about with no definable position. "Box Views" allow you to order the points on any dimensions that have meaning to you, in an arbitrary coordinate-space.

The box is more than a mere measurement-frame. On request the user may see the points projected on a specific face of the box (orthographically); and on request he may also see projection lines between a box-face and its corresponding point in the point cluster.

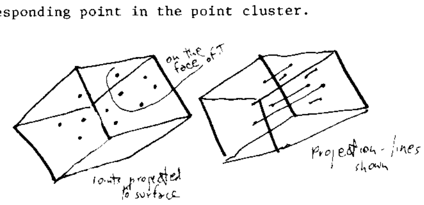

"Magnifying," as before, will create a view of the text: but in the box mode of viewing, the text appears on the side of the box. That is, the inner view will project to the outside, yielding a draft. Naturally, this is the current assembly of your pieces; if certain coordinates are selected it is even a "typeset" version.

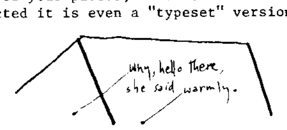

(Note: Vector General hardware does not allow character rotation; only keyword and headline rotation is possible, through software character generation. Thus text pieces on the side of a box show certain freaky movements if the side is not viewed square-on.)

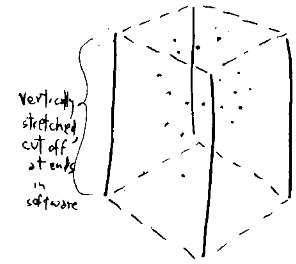

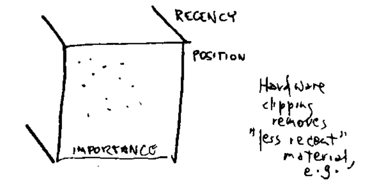

COLLATERAL GALAXIES AND BOXES.

Viewing of collateral structures works through the same mechanism. Galaxies and boxes may be collaterated:

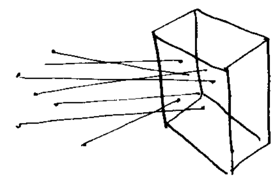

COMPLICATED NOTE: The extension of these mechanisms to pictorial graphics in two and three dimensions is straightforward, and to conceptual substructures (such as may exist) behind these graphics. The same goes for collateration and annotation of multidimensional cluster materials, e.g. in sociology: the system would allow, for instance, the viewing, annotation and collateration of sociometric clusterings.)

BOX FISSION. (The Beauty Part.)

For paired views of projections from the same cluster which do not share a common coordinate, a marvelous trick is possible: BOX FISSION. Starting with one box containing a galaxy, we pull it apart, making two boxes and two galaxies whose Points are linked.

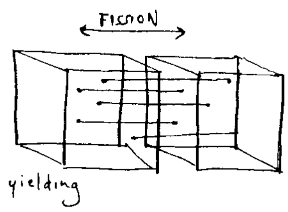

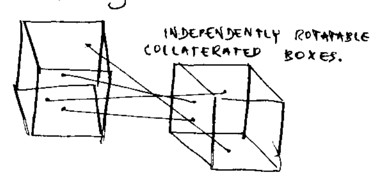

Now both boxes can be rotated independently, and any view considered; equivalence-linkages may now be viewed between any two views. (The eye must, however, turn two corners.)

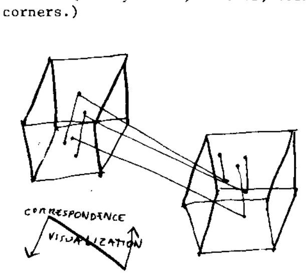

(It is interesting to note that the links in Box Fission are handled automatically, to an extent, by the hardware.)

WELCOME TO THE FUTURE. HUH?

This has summarized the development of some ideas for the viewing and manipulation of complex stuff. I offer this design, insofar as I have been able to present it here, as an example of fantic design (see p. DM 50). There is no logical necessity to it; it corresponds to the traditional structure of no technical system; it arises from no intrinsic or traditional data structures used for computer representation of these things.

But none of these considerations is to the point. This design has a certain stark logical simplicity; it extends itself plausibly from its basic outlook (or starting ideas, if you can isolate them) into a tool for truly intricate cross-consideration, without adding unnecessary and hard-to-remember "technicalities." At least that's how I think of it.

Obviously the aesthetics of it are important to the designer. But a more final criterion of its goodness-- its usefulness--may depend on the same parsimony and organizational clarity.

*At the 1971 Spring Joint Computer Conference, I think it was, I was heckled by a linguist who accused me of being "unimaginative," insisting further that writing is merely an extension of speech and thus "merely" the application of further transformations; and he claimed further that what the user therefore needs is an input language to specify these transformations. This view, while interesting, is wrong.

A but/indeed control language might be interesting, however.

[Appended by the however-operation, a postfix "but." See "Writing," p. DM 43.]

XANADU™

(pronounced Zanna-Doo)

KUBLA KHAN : OR, A VISION IN A DREAM.

A FRAGMENT.

In the summer of the year 1797, the Author, then in ill health, had retired to a lonely farm house between Porlock and Linton, on the Exmoor confines of Somerset and Devonshire. In consequence of a slight indisposition, an anodyne had been prescribed, from the effect of which he fell asleep in his chair at the moment that he was reading the following sentence, or words of the same substance, in "Purchase's Pilgrimage:" "Here the Khan Kubla commanded a palace to be built, and a stately garden thereunto: and thus ten miles of fertile ground were inclosed with a wall. The Author continued for about three hours in a profound sleep, at least of the external senses, during which time he has the most vivid confidence, that he could not have composed less than from two to three hundred lines; if that indeed can be called composition in which all the images rose up before him as things, with a parallel production of the correspondent expressions, without any sensation or consciousness of effort. On awaking he appeared to himself have a distinct recollection of the whole, and taking his pen, ink, and paper, instantly and eagerly wrote down the lines that are here preserved. At this moment he was unfortunately called out by a person on business from Porlock, and detained by him above an hour, and on his return to the room, found, to his no small surprise and mortification, that though he still retained some vague and dim recollection of the general purport of the vision, yet, with the exception of some eight or ten scattered lines and images, all the rest had passed away like the images on the surface of a stream into which a stone had been cast, but alas! without the after restoration of the latter:

Then all the charm
Is broken—all that phantom-world so fair,
Vanishes and a thousand circlets spread,
And each mis-shape the other. Stay awhile,
Poor youth! who scarcely dar'st lift up thine eyes—
The stream will soon renew its smoothness, soon
The visions will return! And lo! he stays,
And soon the fragments dim of lovely forms
Come trembling back, unite, and now once more
The pool becomes a mirror.

As a contrast to this vision, I have annexed a fragment of a very different character, describing with equal fidelity the dream of pain and disease.—1816.

KUBLA KHAN.

In Xanadu did Kubla Khan
A stately pleasure-dome decree:
Where Alph, the sacred river, ran
Through caverns measureless to man
Down to a sunless sea.
So twice five miles of fertile ground
With walls and towers were girdled round:
And there were gardens bright with sinuous rills
Where blossomed many an incense-bearing tree:
And here were forests ancient as the hills,
Enfolding sunny spots of greenery.
But oh! that deep romantic chasm which slanted
Down the green hill athwart a cedarn cover!
A savage place! as holy and enchanted
As e'er beneath a waning moon was haunted
By woman wailing for her demon-lover!
And from this chasm, with ceaseless turmoil seething,
As if this earth in fast thick p ants were breathing,
A mighty fountain momently was forced:
Amid whose swift half-intermitted burst
Huge fragments vaulted like rebounding hail,
Or chaffy grain beneath the thresher's flail:
And 'mid these dancing rocks at once and ever
It flung up momently the sacred river.
Five miles meandering with a mazy motion
Through wood and dale the sacred river ran,
Then reached the caverns measureless to man,
And sank in tumult to a lifeless ocean:
And 'mid this tumult Kubla heard from far
Ancestral voices prophesying war!
The shadow of the dome of pleasure
Floated mid way on the waves;
Where was heard the mingled measure
From the fountain and the caves.
It was a miracle of rare device,
A sunny pleasure-dome with caves of ice!
A damsel with a dulcimer
In a vision once I saw:
It was an Abyssinian maid,
And on her dulcimer she played,
Singing of Mount Abora.
Could I revive within me
Her symphony and song,
To such a deep delight 'twould win me
That with music loud and long,
I would build that dome in air,
That sunny dome! those caves of ice!
And all who heard should see them there,
And all should cry, Beware! Beware!
His flashing eyes, his floating hair!
Weave a circle round him thrice,
And close your eyes with holy dread,
For he on honey-dew hath fed,
And drunk the milk of Paradise.

"Is that the river that runs down to the sea?"

James Stewart
in
"The FBI Story."

Everything is Deeply Intertwingled.

Patent work on Xanadu is in progress.

Xanadu, friend, is my dream.

The name comes from the poem (nearby); Coleridge's little story of the artistic trance (and the Person from Porlock) make it an appropriate name for the Pleasure Dome of the creative writer. The Citizen Kane connotations, and any other connotations you may find in the poem, are side benefits.

I have been working on Xanadu, under this and other names, for fourteen years now.

Originally it was going to be a super system for handling text by computer (see p. 12 and 13). But it grew: as I realized, level by level, how deep the problem was.

And the concept of what it was to be kept changing, as I saw more and more clearly that it had to be on a minicomputer for the home. (You can have one in your office too, if you want, but that's not what it's about.)

Now the idea is this:

To give you a screen in your home from which you can see into the world's hypertext libraries.

(The fact that the world doesn't have any hypertext libraries-- yet-- is a minor point.)

To give you a screen system that will offer high-performance computer graphics and text services at a price anyone can afford. To allow you to send and receive written messages at the Engelbart level (see p.DM46). To allow you to explore diagrams (see p. DM12 and P. DM51). To eliminate the absurd distinction between "teacher" and "pupil."

To make you a part of a new electronic literature and art, where you can get all your questions answered and nobody will put you down.

* * *

Originally Xanadu was programmed around the Parallel Textface (see p. DM53). But as the requirements of the Parallel Textface were better and better understood, Xanadu became a more general underlying system for all forms of interactive graphic environments. Its data structure has Virtual Blocklessness and is thus well related to the smooth motions needed by screen users. Thus in its final form, now being debugged, it will support not only the Parallel Textface (see p. DM53), the Walking Net (see p. DM51), Stretchtext (see p. DM19), Zoom Maps (see p. DM19) and so on, but indeed any data structure that needs to combine complex linkages with fast access and rapid changes. Because the data structure is recursively extensible, it will permit hypertext (see p. DM44) of any depth and complexity, and the collateral linkage (see p. DM52) of any objects of contemplation.

Xanadu is under private development and should be available, if the economy holds, in 1976. Regrettably, first prices will not be at the $3000 level necessary for the true Home System. Exact equipment for the production version has not been selected. A number of microprocessors (see p. 44) are in serious contention, notably the Lockheed SUE, but there's something to be said for a regular mini. The PDP-11 is of interest (see p. 42); (so especially is its Cal Data lookalike-- unless DEC would like to build us a PDP-11X with seven modes of indirect display addressing. Are you reading this, Ken Olsen?) And here's a laugh: a company called IBM may in fact make a suitable computer, except that they call it the "3740 Work Station." So for those customers who want IBM equipment, maintenance and prices, with Xanadu software, it's a definite possibility.

So, fans, that about wraps it up. I'll be interested in hearing from people who want this system; many hardheaded business people have told me nobody will. Prove 'em wrong, America!

Of course, if hyper-media aren't the greatest thing since the printing press, this whole project falls flat on its face. But it is hard for me to conceive that they will not be.

Xanadu:

BRASS FOX

WHAT IT IS: the heart of the Xanadu system, now being debugged, consists of a highly integrated program for use on minicomputers ("software"-- see p. 36) or microprocessors ("firmware"-- see p. 44). It is an operating system with two programs: a highly generalized data management system for handling extremely complex data in huge files, and a generalized display system, married to the other, for handling branching animation and retrieval and canned display programs. These ordain retrievals by the data system. The Parallel Textface (see p. DM53) and the Walking Net (see p. DM51) are two such canned programs.

These internal systems are intended to be sold with consoles of various types, as illustrated nearby, for stand-alone turnkey use (see p. 13). Xanadu is self-networking: two on the phone make a network, and more can join.

LANGUAGES: Xanadu programs will not be made available in any higher languages, mainly because of their proprietary character, but also because the display routines (and some of the retrieval routines) must be programmed in machine language.

The system has its own under-level language, XAP (Xanadu Assembly Program). While two higher-level display languages, DINGO (Display Lingo) and X^{ult} (the ultimate?) are contemplated, these will not ordinarily be accessible to the user. The purpose of Xanadu is to furnish the user with uncomputerish good-guy systems for specific purposes, not a chance to do his own programming.

Important features of the data system are huge addressability (in the trillions of elements) and Virtual Blocklessness. For advantages of this latter, see Zoom Map, p. DM19.

COMPATIBILITY: because of its highly compacted and unconventional structure, it is not compatible with other operating systems (including time-sharing). Anyway, to put it on a larger machine is like having your Mazda driven around in a truck. Because it uses a line-drawing display (see p. DM 22-3) and therefore draws individual arbitrary lines on the screen repeatedly, it is not compatible with television either-- unless you point a TV camera at it, or the equivalent. Sorry.

STANDARDIZATION. Taking a lesson from the integrated work of various people whose work has been described in this book, we see that if you want a thing done right, you have to do it yourself. (Great Ideas of Western Man: one of a series.) My good friend Calvin Mooers with his TRAC Language (see pp. 18-21) has discovered that trademark is one way to nail this as a right.

Several levels of standardization are important with Xanadu. One, all Xanadu systems must be able to work with all Xanadu files (except for possible variations in screen performance and size of local memory). Now, there are those who would not be concerned for this sort of universality, and who might even try to make sure systems were incompatible, so that you had to buy accessories and conversion kits up and down the line. That is one of the things that must be avoided: "partial" compatibility, subject to expensive options and conditions, a well-known technique in the field.

By stabilizing the "Xanadu" trademark, I hope to prevent such shenanigans. Thus every accredited Xanadu system will offer full compatibility with the data structure, and either full performance or substitutes as necessitated by the hardware. The "Xanadu" trademark can thus in principle be made available to manufacturers abiding by all design features of the system.

Second, all Xanadu systems should be able to work with outside systems either through or off the net, if they conform to the unusual data rules required by the unusual design of the system. This assures that Xanadu systems will be compatible with any other popular networks. It also assures others who want to offer Xanadu-class services to system owners (through, e.g., conventional time-sharing) that if they adhere to the rules (see "Canons," p.DM58) they can play the game on a certified basis.

AVAILABILITY. It is hoped that Xanadu will be available in 1975 for at least one machine (guess which). As a program it will be available only in absolute form, without source or comments.

AHEM. There is a lot to talk about, but a lot of time can be wasted talking. It is suggested that thoughtful computer firms, interested in some form of participation, study this book carefully-- at least enough so no one's time need be wasted.

BIBLIOGRAPHY

"Nelson's the Name, and What He Proposes Could Outdo Engelbart." Electronics, 24 Nov 69, 97.

A recent report by Arthur D. Little, a Boston firm that makes its money by seeming to be omniscient, commented on the considerable market potential for on-line data supply systems. The report cost $400 or $4000, I forget which. Big-time interests are aprowl.

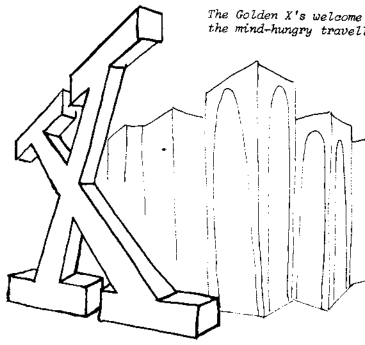

The Golden X's welcome the mind-hungry traveller.

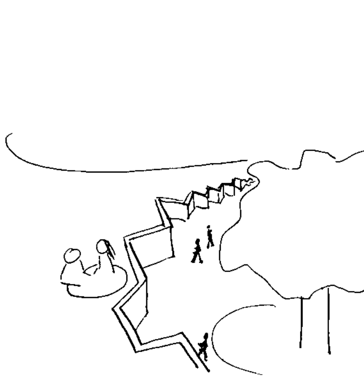

View from the snack balcony of a large Xanadu installation, overlooking the internal greenery. Hexagonal architecture permits physical expansion without interruption of services. (The mollusks have been telling us something about expansion.)

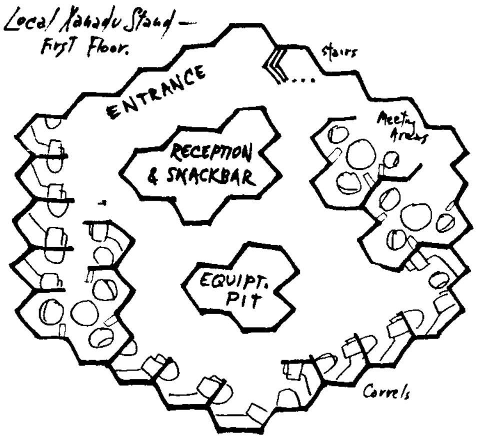

Porta-Xan. (Mockup by Tom Barnard.) Faceplate reflects CRT to user while he's abroad in the world. One-hand typewriting and pointing device frees the other. Can be built with available ruggedized components.

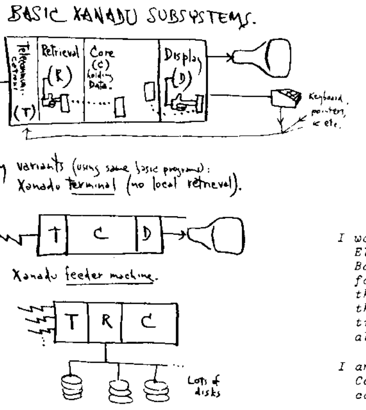

THE AUTHOR ANSWERS THE QUESTIONS HE IS MOST FREQUENTLY ASKED.

Q. If you publish your ideas like this, aren't you afraid someone will steal them?

A. No.
(The Law of Intellectual Property is about the strongest backing the individual has in this society.)
Besides which, there is here no revelation of the Xanadu Sneakrets.

Q. Won't some big company sweep your Xanadu under if they imitate it?

A. Let 'em. If they come up with a system having equivalent scope, which seems unlikely (see Canons, p.DM58), I might even feel I had achieved enough. But in the meantime, like the tortoise, and like DEC, I am going to continue to try to do it right.

Q. Aren't you afraid that writing a flippant book will keep people from taking you seriously?

A. I do not want to be taken seriously in some quarters until it's too late.

I have heard rumors that someone else in the field calls a computer product "Xanadu."
I tend to doubt this; and even if they did, my usage goes back to 1966.

I would like to thank (in chronological order) Elliot Klugman, Nat ("Kubla") Kuhn, Glenn Babecki, Cal Daniels and John V.E. Ridgway for the considerable time and involvement they gave to the Xanadu program design sessions; thanks also to various others who sat in from time to time. For the final selection of algorithms, however, no one is to blame but me.

I am grateful to the good offices of Swarthmore College for the use of their equipment in the continuing efforts to debug the Xanadu programs.

the XANADU NETWORK

First of all, bear in mind that Xanadu is a unified system for complex data management and display. This basically means that the same system (without the displays) can serve as a feeder machine for the data network itself.

So far, so good. That means that we can have a minicomputer network handling the entire structure: sending out library materials to users on call, and storing any materials they want saved. This saves all kinds of hassles with big computers and big-computer-style programming.

But who will pay for it? To build the kind of capacity we're talking about-- all those disks, all those minicomputers in a network-- won't it take immense amounts of capital? How, people ask me, will any American company ever back such a Utopian scheme?

Aha.

One method of financing has proven itself in the postwar suburban era, this time of drive-ins and hamburger stands.

Franchising.

What I propose, then, is the Mom-and Pop Xanadu Shop. Or, more properly, the Xanadu stand. "Mom and Pop" are the owners of the individual stand. But the customers can be families, too.

From far away the children see the tall golden X's. "Oh, Daddy, can't we stop? I want to play Spacewar," says little Johnny. Big Sis adds, "You know, I have to check something for my paper on Roman politics." And Mom says, "Say, that would be a good place for lunch."

So they turn in past the sign that says "OVER 2 BILLION SCREEN HOURS," and pull into the lot. They park the car, and Dad shows the clerk his Xanadu credit card, and the kids run to screens. Dad and Mom wait for a big horizontal CRT, though, because there are some memories they'd like to share together...

Sis's paper, of course, goes to her teacher through his Xanadu console.

THE PLAN. IS IT AS CRAZY AS IT SEEMS?

Deep inside, the public wants it, but people who think of computers in clichés can't comprehend it. This means "the public" must somehow create it.

One way to go is to start a new corporation, register it with the SEC and try to raise a lot of money by selling stock publicly. Unfortunately there are all kinds of obstacles for that. ("Reg A" is about as far as it will go.)

Through the miracle of franchising, now, a lot of the difficulties of conventional backing can be bypassed. The franchisee has to put up the money for the computers, the scopes, the adorable purple enamel building, the johns and so on; as a Xanadu franchisee he gets the whole turnkey system and certain responsibilities in the OVERALL XANADU NETWORK-- of which he is a member. He is assigned permanent storage of certain classes of materials, on call from elsewhere in the net. (Naturally, everything is stored in more than one place).

The Xanadu subscriber, of course, gets what he requests at the screen as quickly as possible-- or on priority if he wants to pay for it-- and may store his own files, including linkages among other materials and marginal notations to other things that can be called. (See collateral structures, p.DM52; these can automatically bring forth anything they're linked to. (See "Nelson's Canons," p.DM58.)) A user's historical record will be stored to whatever degree he desires, but not (if he chooses) in ways that can be identified with him.

Home users need only dial a local phone number-- their nearest Xanadu stand-- to connect with the entire Xanadu network. (The cost of using something stored on the network has nothing to do with where it is stored.)

(Special high-capacity lines need not be installed between storage stations, as appropriate digital transmission services are becoming available commercially.)

Various security techniques prevent others from reading a subscriber's files, even if they sign on falsely; the Dartmouth technique of scrambling on non-stored keywords is a good one.

The Xanadu stand also has private rooms with multiple screens, which can be rented for parties, business meetings, design sessions, briefings, legal consultations, lectures, seances, musicales, and so on.

The choice locations for the Xanadu stands are somewhat different from hamburger spots. But that's probably not anything to go into here.

Within the Xanadu network, then, people may read, write, send messages, study and play.

****

XANADOODLES.™

fourteen years on a project and you get a little punchy.

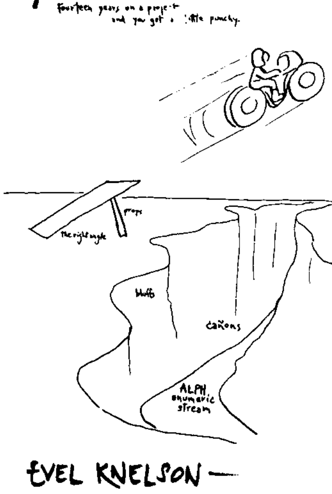

EVEL KNELSON—
WILL HE MAKE IT?

Thanks a lot, Sam Coleridge, for those two symbols.

Xanadu.
And the Albatross.

"Listen," Mr. Wonka said, "I'm an old man. I'm much older than you think. I can't go on forever. I've got no children of my own, no family at all. So who is going to run the factory when I get too old to do it myself? *Someone's* got to keep it going—if only for the sake of the Oompa-Loompas. Mind you, there are thousands of clever men who would give anything for the chance to come in and take over from me, but I don't *want* that sort of person. I don't want a grown-up person at all. A grownup won't listen to me; he won't learn. He will try to do things his own way and not mine. So I have to have a child. I want a good sensible loving child, one to whom I can tell all my most precious candy-making secrets—while I am still alive."

Roald Dahl, Charlie and the Chocolate Factory, p.157.

I am sorry I have not been able to reply to all those who have written to me saying they wish they could work for The Nelson Organization at even a low salary.

So do I, my friends, so do I.

[As heard at the National Joint Computer Conference, 1973.]

How are we going to sell the Home Computer?
Well if you want to sell computers, let me tell you what to do:
You've got to talk to the housewives, and the children, too;
No one wants to program, they want something they can view...

It's got to offer fun, and it's got to offer truth;
It's got to give you something that'll lift you from the booth;
It's got to be uplifting to the Lady from Duluth.
You've got to have a vision; you've got to have an angle;
You should maybe sing a jingle (in a way that doesn't jangle);
It's got to have a tingle, in a way their minds can't tangle--

So continuing under our guidance inertial,
Let's have the XANADU SINGING COMMERCIAL.

[strings] It's got everything to give.
It'll get you where you live.

[chimes] Realms of mind that you may roam:
Grasp them all within your home.

[brass flourish] The greatest things you've ever seen
Dance your wishes on the screen.

[bass hoot] All the things that man has known
Comin' on the telephone--

[TUBULAR BELLS!] Poems, books and pictures too
COMIN' ON THE XANADU --

XAN-A-DU, OO--

[kettledrums] THE-- WORLD-- OF-- YOUUUU!

Is Xanadu worth waiting for?
That depends, doesn't it, on the value of the hand-bush differential bird utility ratio.

fox box

quick brown

CRAZY LEICA FOX

WHAT NELSON IS REALLY SAYING

Told so that anybody can understand it without a Ph.D.

and maybe some wit.

From "Barnum-Tronics"
(citation p. DM 2.)

1) Knowledge, understanding and freedom can all be advanced by the promotion and deployment of computer display consoles (with the right programs behind them).

2) Computer presentational media, coming soon, will not be technically determined but rather will be new realms for human artistry. This point of view radically affects how we design man-machine systems of any kind, especially those for information retrieval, teaching, and general writing and reading. Some practitioners see such systems as narrowly technical, with the computer hoisting up little pieces of writing on some "scientific" basis and showing them to you one grunt at a time. A Metrecal banquet. I disagree. The systems should be opulent.

3) The problem in presentational systems of any kind is to make things look good, feel right, and come across clearly. The things that matter are the feel of the system, the user's state of mind, his possible confusion, boredom or enthusiasm, the problems of communicating *concepts*, and the very nature of concepts and their interconnection. There will never be a "science" of presentation, except as it relates to these things.

4) Not the nature of machines, but the nature of *ideas*, is what matters. It is incredibly hard to develop, organize and transmit ideas, and it always will be. But at least in the future we won't be booby-trapped by the nature of paper. We can design magic paper.

It is time to start using computers to hold information for the mind much as books have held this information in the past. Now information for the mind is very different from "information for the computer" as we have thought of it, hacked up and compressed into blocks. Instead we can stretch the computer.

I am proposing a curious kind of subversion. "Let us design," I say; and when people see the systems, everybody will want one. All I want to do is put Renaissance humanism in a multidimensional responsive console. And I am trying to work out the forms of writing of the future. Hypertexts.

Hypertexts: new forms of writing, appearing on computer screens, that will branch or perform at the reader's command. A hypertext is a non-sequential piece of writing; only the computer display makes it practical. Somewhere between a book, a TV show and a penny arcade, the hypertext can be a vast tapestry of information, all in plain English (spiced with a few magic tricks on the screen), which the reader may attack and play for the things he wants, branching and jumping on the screen, using simple controls as if he were driving a car. There can be specialized subparts for specialized interests, instant availability of relevancies in all directions, footnotes that are books themselves. Hypertexts will be so much better than ordinary writing that the printed word will wither away. *Real writing by people*, make no mistake, not data banks, robot summaries or other clank. A person is writing to other people, just as before, but on magical paper he can cut up and tie in knots and fly around on.

I believe in calling a spade a spade -- not a personalized earth-moving equipment module; and a multi-dimensional spade, by gum, a hyperspade-- not a personalized earth-moving equipment module with augmented dirt access, retrieval and display capability under individulaized control.

I want a world where we can read the world's literature from screens rather than personally searching out the physical books. A world without routine paperwork, because all copying operations take place automatically and formalized transactions occur through formalized ceremonies at consoles. A world where we can learn, study, create, and share our creations without having privately to schlepp and physically safeguard them. There is a familiar, all-embracing motto, the jingle we all know from the day school lets out, which I take quite seriously: "No more pencils; no more books; no more teachers' dirty looks." The Fantic Age.

From "Computopia and Cybercrud."
(Citation nearby.)

MINIFESTO

My work is concerned principally with the theory and execution of systems useful to the mind and the creative imagination. This has polemical and practical aspects: I claim that the precepts of designing systems that touch people's minds, or contents to be shown in them, are simple and universal: making things look good, feel right and come across clearly. I claim that to design systems that involve both machines and people's minds is art first, technology second, and in no way a derivative specialty off in some branch of computer science.

However, presentational systems will certainly involve computers from now on.

Since hundreds of such systems are now being built, many of them all wrong, we must teach designers (and certain others) the basics of computers, and give them some good examples to emulate (such as Sutherland's Sketchpad, Bitzer's PLATO, and, I hope, some of my own designs).

Further, the popular superstitions about computers must be fought-- the myths that they are mechanistic, scientific, objective or independent of human intent and contemplative involvement.

NELSON'S CANONS

A Bill of Information Rights

It is essential to state these firmly and publicly, because you are going to see a lot of systems in the near future that purport to be the last-word cat's-pajama systems to bring you "all the information you need, anytime, anywhere." Unless you have thought about it you may be snowed by systems which are inherently and deeply limiting. Here are some of the things which I think we will all want. (The salesman for the other system will say they are impossible, or "We don't know how to do that yet," the standard putdown. But these things are possible, if we design them in from the bottom up; and there are many different valid approaches which could bring these things into being.)

These are rules, derived from common sense and uncommon concern, about what people can and should have in general screen systems, systems to read from.

1. EASY AND ARBITRARY FRONT ENDS.

The "front end" of a system-- that is, the program that creates the presentations for the user and interacts with him-- must be clear and simple for people to use and understand.

THE TEN-MINUTE RULE. Any system which cannot be well taught to a layman in ten minutes, by a tutor in the presence of a responding setup, is too complicated. This may sound far too stringent; I think not. Rich and powerful systems may be given front ends which are nonetheless ridiculously clear; this is a design problem of the foremost importance.

TEXT MUST MOVE, that is, slide on the screen when the user steps forward or backward within the text he is reading. The alternative, to clear the screen and lay out a new presentation, is baffling to the eye and thoroughly disorienting, even with practice.

Many computer people do not yet understand the necessity of this. The problem is that if the screen is cleared, and something new then appears on it, there is no visual way to tell where the new thing came from: sequence and structure become baffling. Having it slide on the screen allows you to understand where you've been and where you're going; a feeling you also get from turning pages of a book. (Some close substitutes may be possible on some types of screen.)

On front ends supplied for normal users, there must be no explicit computer languages requiring input control strings, no visible esoteric symbols. Graphical control structures having clarity and safety, or very clear task-oriented keyboards, are among the prime alternatives.

All operations must be fail-safe.

Arbitrary front ends must be attachable: since we are talking about reading from text, or text-and-picture complexes, stored on a large data system, the presentational front end must be separable from the data services provided further down in the system, so the user may attach his own front-end system, having his own style of operation and his own private conveniences for roving, editing and other forms of work or play at the screen.

2. SMOOTH AND RAPID DATA ACCESS.

The system must be built to make possible fast and arbitrary access to a potentially huge data base, allowing extremely large files (at least into the billions of characters). However, the system should be contrived to allow you to read forward, back or across links without substantial hesitation. Such access must be implicit, not requiring knowledge of where things are physically stored or what the internal file names may happen to be. File divisions must be invisible to the user in all his roving operations (FREEDOM OF ROVING): boundaries must be invisible in the final presentations, and the user must not need to know about them.

3. RICH DATA FACILITIES.

Arbitrary linkages must be possible between portions of text, or text and pictures; annotation of anything must be provided for; collateration (see p. DM 50) should be a standard facility, between any pair of well-defined objects; PLACEMARK facilities must be allowed to drop anchor at, or in, anything. These features imply private annotations to publicly-accessible materials as a standard automatic service mode.

The AI people don't understand,
the IR people don't understand,
the CAI people don't understand,
and for God's sake don't tell IBM.

I believe that an introduction to any subject can be humorous, occasionally profound, exciting, vivid, and appealing even to experts on their separate levels.

Perhaps someday I can prove it.

4. RICH DATA SERVICES BASED ON THESE STRUCTURES.

The user must be allowed multiple rovers (movable placemarks at points of current activity); making possible, especially, multiple windows (to the location of each rover) with displays of collateral links.

The system should also have provision for high-level mooting (~~see p. —~~) and the automatic keeping of historical trails.

Then, a complex of certain very necessary and very powerful facilities based on these things, viz.:

A. ANTHOLOGICAL FREEDOM: the user must be able to combine easily anything he finds into an "anthology," a rovable collection of these materials having the structure he wants. The linkage information for such anthologies must be separately transportable and passable between users.

B. STEP-OUT WINDOWING: from a place in such an anthology, the user must be able to step out of the anthology and into the previous context of the material. For instance, if he has just read a quotation, he should be able to have the present anthological context dissolve around the quotation (while it stays on the screen), and the original context reappear around it. The need of this in scholarship should be obvious.

C. DISANTHOLOGICAL FREEDOM: the user must be able to step out of an anthology in such a way and not return if he chooses. (This has important implications for what must really be happening in the file structure.)

Earlier versions of public documents must be retained, as users will have linked to them.

However, where possible, linkages must also be able to survive revisions of one or both objects.

5. "FREEDOM FROM SPYING AND SABOTAGE."

The assumption must be made at the outset of a wicked and malevolent governmental authority. If such a situation does not develop, well and good; if it does, the system will have a few minimal safeguards built in.

FREEDOM FROM BEING MONITORED. The use of pseudonyms and dummy accounts by individuals, as well as the omission of certain record-keeping by the system program, are necessary here. File retention under dummy accounts is also required.

Because of the danger of file sabotage, and the private at-home retention by individuals of files that also exist on public systems, it is necessary to have FIDUCIAL SYSTEMS FOR TELLING WHICH VERSION IS AUTHENTIC. The doctoring of on-line documents, the rewriting of history-- cf. both Winston Smith's continuous revision of the encyclopedia in Nineteen Eighty-Four and H.L. Hunt's forging of historical telegrams for "The White House"-- is a constant danger. Thus our systems must have a number of complex provisions for verification of falsification, especially the creation of multilevel fiducials (parity systems), and their storage in a variety of places. These fiducials must be localizable and separate to small parts of files.

7. COPYRIGHT.

Copyright must of course be retained, but a universal flexible rule has to be worked out, permitting material to be transmitted and copied under specific circumstances for the payment of a royalty fee, surcharged on top of your other expenses in using the system.

For any individual section of material, such royalty should have a maximum: i.e., "by now you've bought it."

Varying royalty rates, however, should be the arbitrary choice of the copyright holder; except that royalties should not vary sharply locally within a tissue of material. On public screens, moving between areas of different royalty cost must be sharply marked.

BIBLIOGRAPHY

Theodor H. Nelson, "Computopia and Cybercrud." in Roger Levien (ed.), Computers in Instruction (Rand Corporation, 1971).

Theodor H. Nelson, "A Conceptual Framework for Man-Machine Everything." Proc. NCC 73.

"Rascal am I? Take that!"
Errol Flynn in "Swords of Valor" (?)

FLIP OUT.

> I have had a most rare vision. I have had a dream, past the wit of man to say what dream it was: man is but an ass, if he go about to expound this dream. Methought I was—there is no man can tell what. Methought I was,—and methought I had,—but man is but a patched fool, if he will offer to say what methought I had. The eye of man hath not heard, the ear of man hath not seen, man's hand is not able to taste, his tongue to conceive, nor his heart to report, what my dream was.
>
> *Bottom the Weaver*

Now you see why I brought you here. This Gem-maniacal book has, obviously, been created as a crossroad of several cross purposes: to furnish a needed, grabby layman's introduction to two vast but rather inaccessible realms; to present a coherent, if contentious, point of view, and unroll a particular sort of apocalyptic vision after preparing the vocabulary for it; to make bright friends and informed supporters for my outlook and projects; to get home to some of my friends the fact that what I am doing is at bottom not technical; and finally, if nothing else, to set forth some principles about the way things should be, which others will have to answer if they propose to do less. Thus, overall, this book is a message in a Klein bottle, waiting to see who's thirsty.

I suppose it all started in college. Swarthmore left me with an exaggerated notion of the extent to which ideas are valued in the academic world; it took two graduate schools to clear this up. After that, as far as I was concerned, Ph.D. stood for Poophead. But I still cared about ideas, and the deep necessity of finding their true structure and organization. From writing I knew the grueling difficulty of trying to make ideas get in order. I believed in the pure, white light of inspiration and the power of the naive but clever mind to figure out anything, if not obstructed but dumb dogmas and obtuse mental schemata fostered by the educational system.

When I finally got the idea of what computers were about, sometime in 1960, I took endless walks at night trying to hash these things out and see where they led. The text systems came clear to me, at least in their beginnings, in a few weeks; the realization that 3D halftone was possible came to me as a shock the following spring, I believe as I was walking across Radcliffe Common. Since then trying to build these systems for creation and the true ordering of intricate thought has been my driving dream.

My own life among these dream machines has been a nightmare, thoroughly unpleasant, and if people are right in telling me that nobody wants systems like the ones I am designing, I'll get the heck out of this and be a disk jockey or a toy salesman or something.

I first got into this as a writer; all I wanted was a decent writing system that would run on a computer. Little did I realize the immensity of what that entailed, or that for some reason my work and approach would engender indignation and anger wherever I went. There is a fiction that everybody in these fields is doing something fundamentally scientific and technical, and this fiction is usually upheld in carefully enacted mutual playlets. Trying to cut through that and say, "Let's build a home for mankind that will at last be shaped to fit man's mind," does not seem to generate immediate warmth and welcome.

But I'm glad for the friends I've made in this field, and of course there have been a lot of laughs. (I'd really have hated to miss being in this field, just for the thrilling madness of it all.) All in all my adventures have been a sort of participatory journalism, which I'd like to write up properly some time. Some highlights:

The days of madness in '68, trying to start an honest corporation to do all this stuff, and suffering endless lunches with Wall Street hangers-on who were looking for a vehicle to take public. They wanted another chicken-franchise type company, though, and certainly not ideas.

Being briefed by four different corporations, most of them major, on the fantastic powers their interactive-movie system was going to have. One of these briefings was in the board room of a famous skyscraper. And now, only one of those systems is left-- Kodak's.

Then there was the courtly gentleman who was going to be my Noah Dietrich, my Colonel Parker. He assured me that through his business connections all was going to go marvelously, and then later intimated that as a special favor he was going to put me in touch with other universes and the flying saucer people. I just didn't have time for other universes.

Then there was the suppression of my first book (this is my second). You might say it was a misunderstanding, at least on my part. My boss's understanding was evidently that the advancement of my ideas would be detrimental to his. If it had been a question of free speech in Yugoslavia it might have been different. Well, it takes a long time to get a book together, but here we go again.

Then there was the time I was called in as a consultant on a vast federal system, never mind what. Numerous computer programs were to be coordinated by a hypertext system they had created and they wanted to know if they'd designed it right. It took months to find out from the programmers exactly what the system was, so I ended up writing the manual; after which I explained what was wrong with the project and the whole hypertext system was scrapped. And my job with it. I never quite got the swing of consulting.

Flying coast-to-coast with the president of a large corporation, he and I planned the whole Xanadu budget for the following year at something like half a million dollars. Two years later, reduced in circumstances and driving a yellow cab in New York, the miserable vehicle breaks down in front of those same corporate headquarters. And the reason I had that bad taxi was that I was out of favor with the taxi dispatcher, on account of having been absent the previous week-- I had had to fly to California to give a banquet address at the Rand Corporation.

Then there were my adventures with the CIA.

I was sitting in my office at Vassar, sagely advising a student, when the phone rang and the caller identified himself as John W. Kuipers, head of computer research at the CIA. He told me I had been noticed as a new bright young man in the field, and would I like to work for them?

Now, there is something about being a cynic and a romantic. (They go together: the cynic deflates ideas, the romantic falls in love with them.) It is not impossible for the cynical romantic to surmise that because everything he has seen personally turned out to be so lousy, that the true hope may lie at the heart of the vortex, just where everybody thinks is impossible. Also the Kennedy aftermath, when sophisticated people had learned to laugh at simple idealism as a facade for the real wheel-and-dealing, slap-and-tickle, may have had something to do with it; anyway, I was enchanted. Thus began the Kuipers Caper.

YES, THERE IS A McLEAN, VIRGINIA

I was given a handler named Bob, a jolly fellow, who kept assuring me that much money was just around the corner. I was regaled with success stories of other people in the computer field who really, undercover Worked for Them. (They weren't doing anything very exciting.) I got to show my slides in the CIA office building in Arlington, and to see there very fancy display equipment behind shielded (!) double-doors in a shielded (!!) computer room-- shielded to keep any planted bugs from transmitting out the contents of the computers' working registers. I even got to visit the main CIA "campus" in McLean, Virginia, where the sign says Agricultural Research Station. It is an incredible feeling to walk across that big eagle in the terrazzo, and to be given the visitor's badge that says "United States Government" all in wiggly lines.

They told me that they would be glad to set me up in business as a hypertext company, but I would have to have a corporation, because that was the way they always did things. And so it came to pass that The Nelson Organization, Inc. was founded at the express request of the United States Central Intelligence Agency. I wouldn't have had it any other way. If life can't be pleasant it can at least be surrealistic.

... BUT NO SANTA CLAUS

I was encouraged to write proposals for them, and write proposals I did. (I happened to finish typing the first one during a lightning storm, and lightning crashed just as I was signing the page; I felt like Faust.) I explained how hypertext might have prevented the Bay of Pigs. After due consideration, I did not say what hypertexts might have done for the Warren Report. Numerous jolly phone calls assured me that my first $25,000 was just around the corner.

The break came when Bob called me and asked me to rewrite a proposal one more time. He had circulated it, he said, among various people "at the shop," who he reminded me were holders of advanced degrees, and it had been remarked that they found my proposal meaningless: "Every place you say 'hypertext' you could just as well put 'gobbledygook' instead; you'll have to clear that up a little."

That did it. They couldn't read either. Who turns out to be in charge of computer stuff in the heart of the CIA, the inner sanctum, the nest of vipers, but the same old poopy Ph.Ds. I decided to resuscitate my virtue.

As far as I know, there is still not a Decent Writing System anywhere in the world, although several things now come close. It seems a shame that grown men and women have to rustle around in piles of paper, like squirrels looking for acorns, in search of the phrases and ideas they themselves have generated. The decent writing system, as I see it, will actually be much more: it will help us create better things in a fraction of the time, but also keep track of everything in better and more subtle ways than we ever could before. But nobody sees this-- I suppose it's only writers and editors that know they're trying to "keep track of ideas"-- and I have been unable to get this across to anybody. (The professional writers, of course, won't talk to me either.)

So here I am after fourteen years with exactly two systems to show for it: the main one, Xanadu, the text-and-animated-picture network system, and Fantasm (I shouldn't have spent the time but it was a labor of love), the simulated-photography system. Actually, I don't have either of them to show, it's all just flow-charts, but it turns out that if I work on either of them with university equipment, my work of fourteen years gets confiscated. So much for that; the outside expedients for debugging continue.

And, to lighten the burden, I've finally given up on trying to reach professionals, who evidently need a thick gravy of technicalism to make the obvious palatable; with this bookity I am taking my case to The People. It is there, anyway, out in Consumerdom, that the real action is going to occur. So the important thing is for everybody to know what's really possible, and what they could have. That is why I have shot off my big canons (and this epistol).

To me, you see, this is really a holy crusade, whereas I know guys to whom it's just a living. It's no less than a question of freedom in our time. The cases of Solzhenitsyn and Ellsberg remind us that freedom is still not what it should be, anywhere. Computer display and storage can bring us a whole new literature, the uniting and the apotheosis of the old and the new; but there are many who would not necessarily want to see this come about. Deep and widespread computer systems would be tempting to two dangerous parties, "organized crime" and the Executive branch of the Federal government (assuming there is still a difference between the two). If we are to have the freedoms of information we deserve as a free people, the safeguards have to be built in at the bottom, now. And the opulence which is possible must be made clear to everyone before we settle on an inferior system-- as we did with television.

Some people have called my ideas and systems "Orwellian." This is annoying in two ways. In the first place it suggests the nightmare of Orwell's book Nineteen Eighty-Four, which obviously I want no part of. (But hey, do you remember what that world of 1984 was actually like? The cryptic wars against unseen enemies that kept shifting? The government spying? The use of language to twist and manipulate? To paraphrase Huey Long: "Of course we'll have 1984 in America. Only we'll call it 1972.")

The second reason the term "Orwellian" is offensive is that it somehow reduces the life of Orwell, the man, to the world of "1984." This is a shallow and shabby thing to do to a man who spent his life unmasking oppressiveness in human institutions everywhere.

In the larger sense, then-- in homage to that simple, honest, angry man, who cared about nothing more than human freedom-- I would be proud indeed if my systems could be called Orwellian.

That reminds me. Nowhere in the book have I defined the phrase "computer lib." By Computer Lib I mean simply: making people freer through computers. That's all.

Fantically-- or fanatically--
Yours for a better world,
Before we have to settle for Any--

Ted Nelson