RCA Engineer

Vol. 31 No. 1 Jan./Feb. 1986

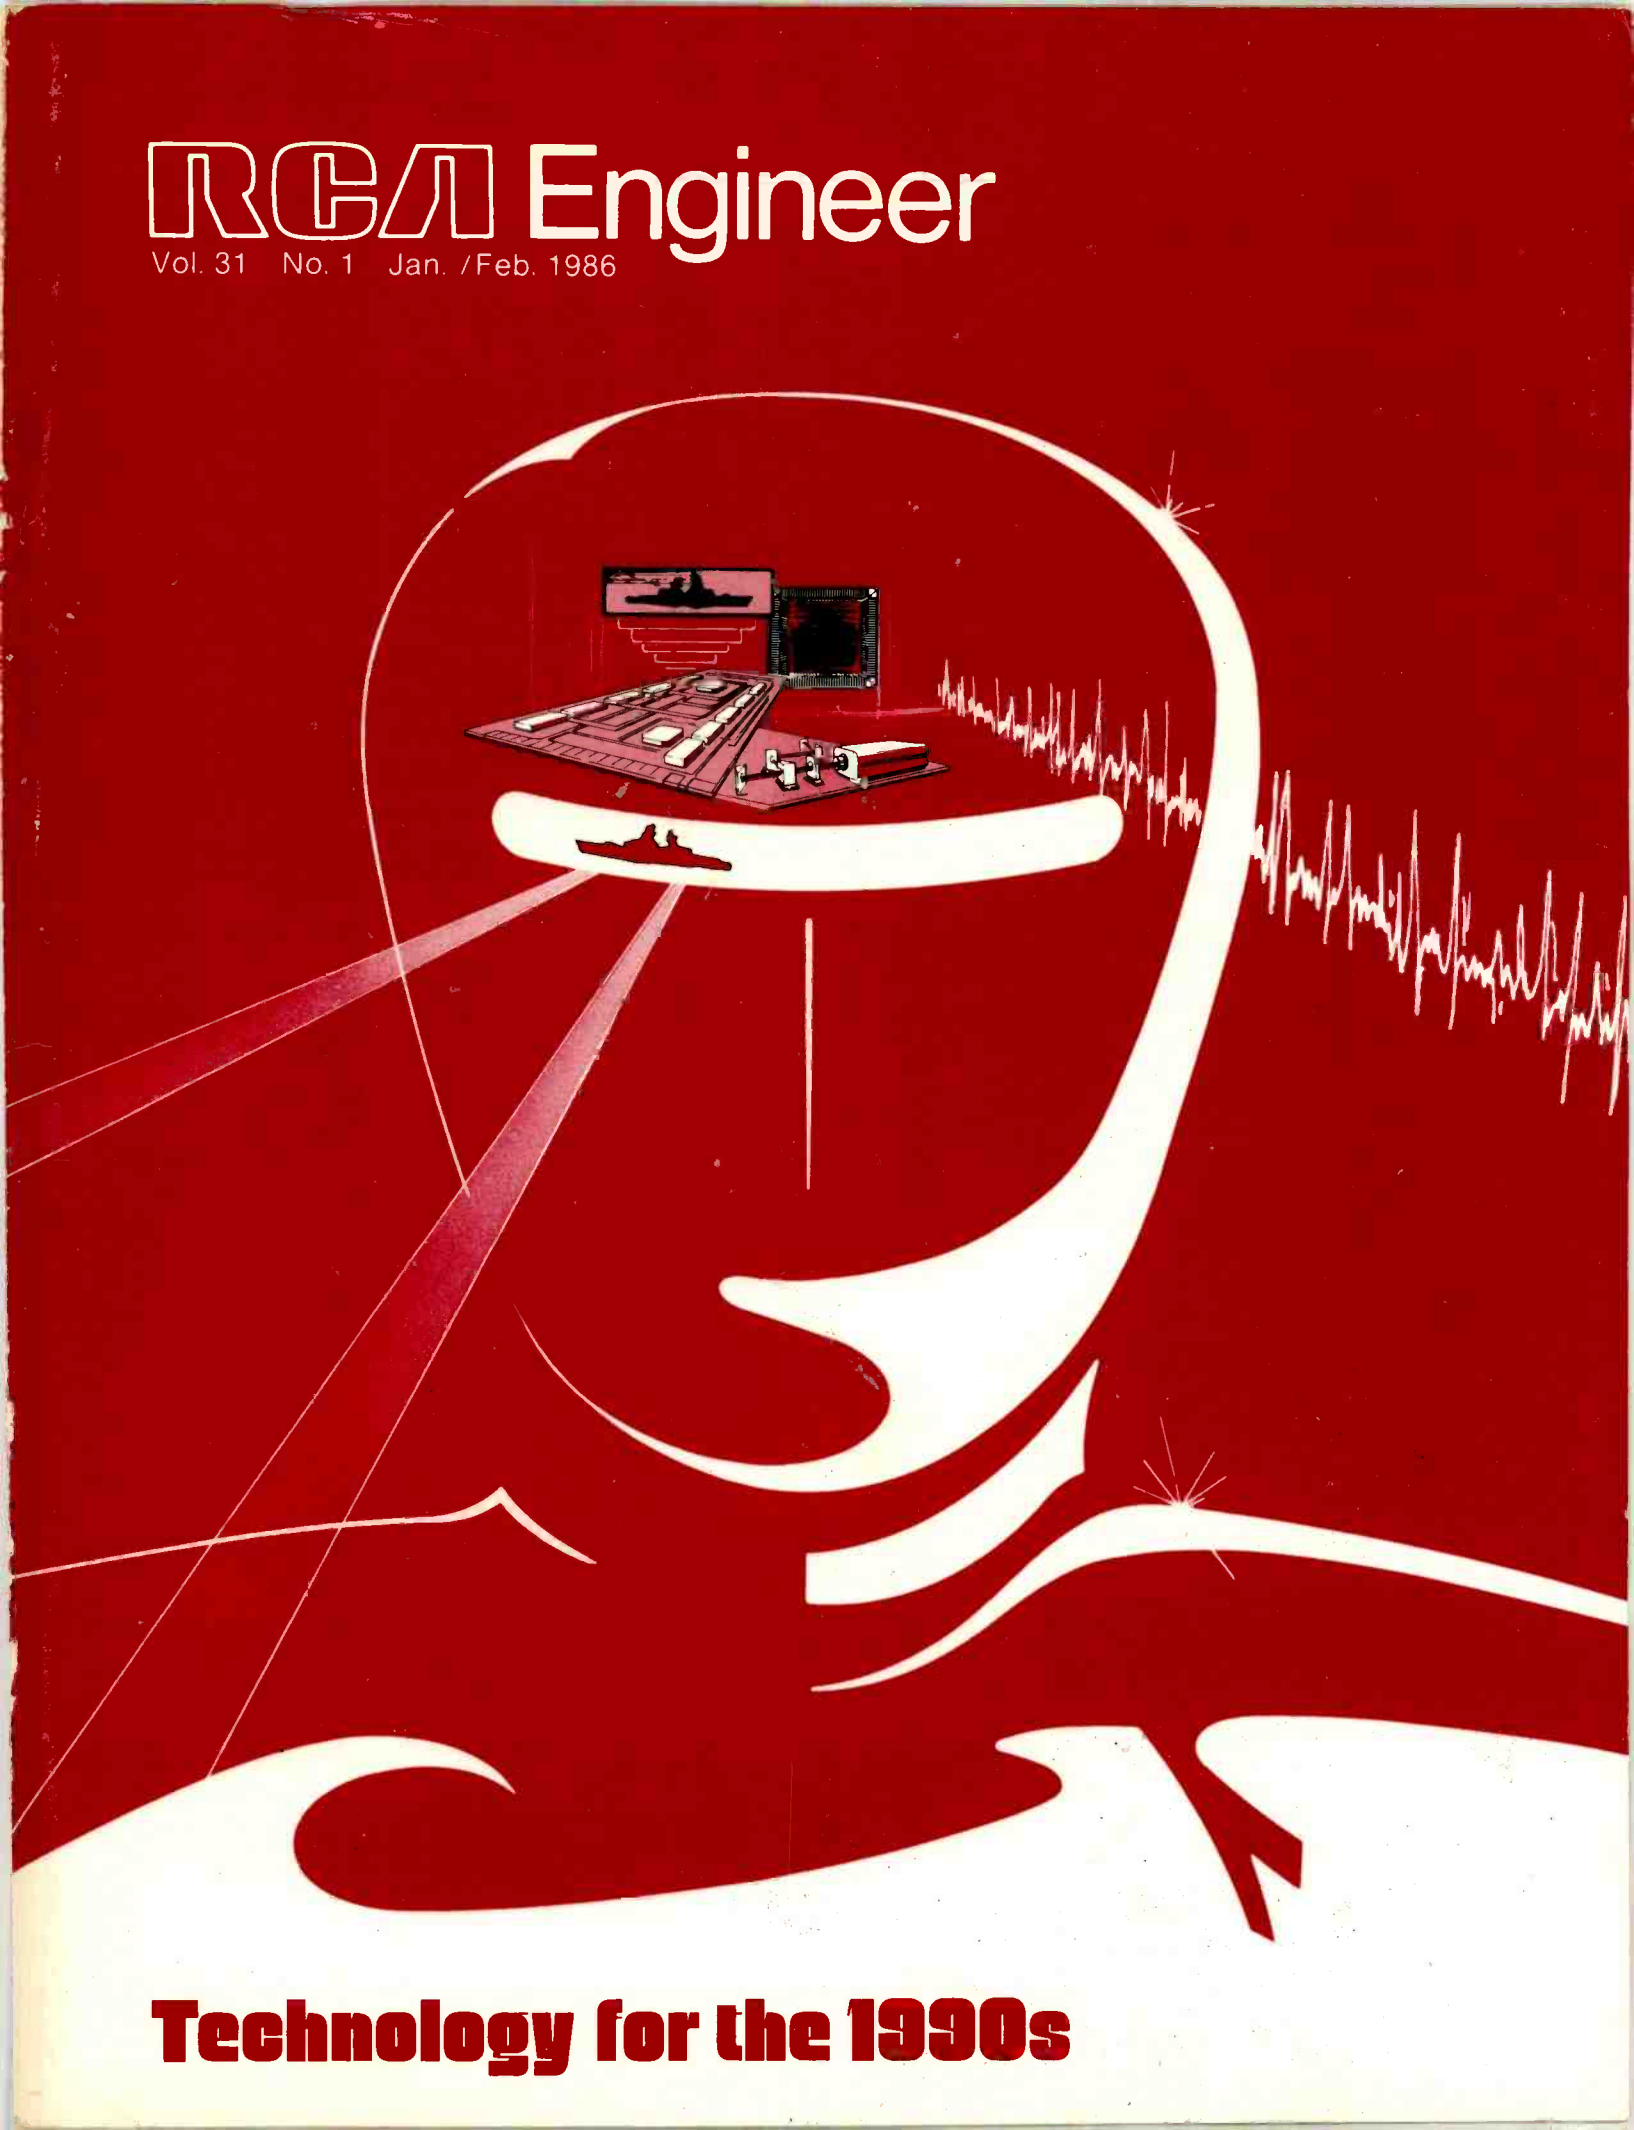

Technology for the 1990s

RCA Engineer

A technical journal published by the
RCA Technical Excellence Center □ 13 Roszel Road □ P. O. Box 432 □ Princeton, NJ 08540-0432 □ Tacnet: 226-3090

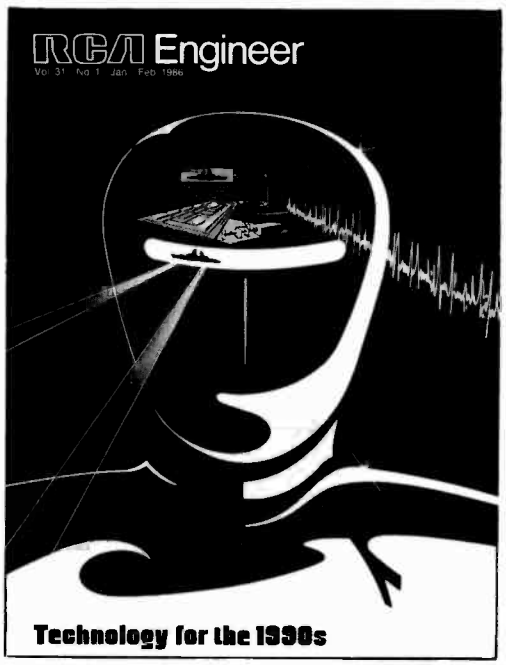

Our cover, we believe, aptly characterizes the subject matter of the articles in this issue—simplicity, power, and elegance.

Cover design and illustration by John Duffy, Advanced Technology Laboratories

□ To serve as a medium of interchange of technical information among various groups at RCA □ To create a community of engineering interest within the company by stressing the interrelated nature of all contributions □ To disseminate to RCA engineers technical information of professional value □ To publish in an appropriate manner important technical developments at RCA, and the role of the engineer □ To help publicize engineering achievements in a manner that will promote the interests and reputation of RCA in the engineering field □ To provide a convenient means by which the RCA engineer may review professional work before associates and engineering management □ To announce outstanding and unusual achievements of RCA engineers in a manner most likely to enhance their prestige and professional status.

Technology for future-generation defense systems

To a technologist concerned with electronics in its broadest sense, today's environment presents more challenges and is more exciting than has been the case for a long time. This situation is a result, in part, of recent major advances in computer power, speed, and cost reduction. In parallel with this, the DoD has taken on a technologically very aggressive strategy to accelerate the rate of advancement. The examples are clear.

The VHSIC (Very High Speed Integrated Circuits) program is providing computers on a single chip with clock cycle times of 40 nanoseconds or less. The DARPA (Defense Advanced Research Project Agency)-sponsored gallium arsenide (GaAs) program is attempting to lower this by another order of magnitude. The potential for this performance at the device level has made possible whole new concepts in computer systems. The Strategic Computer Initiative (SCI) routinely talks about computers that operate at about one GFLOPS (10^9 floating point operations per second). The Japanese Fifth Generation Computer program has similar objectives. Both SCI and Fifth Generation programs involve new architectures and computing concepts. The advanced signal and data processors from these programs will support new strategic systems requirements. The Strategic Defense Initiative, for example, will need computing power orders of magnitude beyond what is commonly available today. Radar systems that have to deal with Stealth technology will need super-robust signal processors. Spaceborne sensors at all wavelengths from radar through the visible spectrum bring with them massive system signal processing requirements. The imagery data from such sensors will require huge data link bandwidth capabilities and signal processors that have yet to be developed for near real-time operations.

The advanced computer capabilities also enable an entirely new class of "intelligent" systems. Computing machines and systems of the future will have the power to mimic human-like attributes to better interface with the human operator and to relieve the human of burdensome tasks. This general concept is viewed by many as "artificial intelligence." Current applications include vision systems for object/target recognition, expert systems that capture and exploit expert knowledge, speech recognition and natural language machine processing, and advanced tactile sensors for robotic systems.

The conclusion is clear: There are numerous high-potential opportunities for the application of electronics technology to defense systems over the next several years.

The Advanced Technology Laboratories (ATL) has a history of providing applications of new electronics technology to the Aerospace and Defense businesses of RCA. A few notable examples from the recent past are narrowband and wideband voice terminals for secure communications, high-speed and high-capacity optical storage systems in a jukebox configuration, and the ATMAC special-purpose microprocessor for signal processing. Similar activity is continuing today. In fact, the range of technologies being explored and developed in ATL today is broader than ever before.

ATL's current programs range from the development of advanced microelectronic devices to architectural concepts that are critical for high-level systems. The programs represent a good match with the challenges and opportunities presented by today's environment. From a system complexity perspective, the highest level is intelligent or "human-like" machine interfaces. To address this level, ATL programs exist in digital image processing for vision systems, image understanding and target classification; in speech recognition for speaker verification, word cueing, and improved man-machine interfaces; in artificial intelligence for expert systems, natural language processing, and data fusion in support of command and control requirements; and in robotics to integrate vision, speech, AI, and tactile sensors to support autonomous space-based maintenance requirements. One level down in complexity is advanced computer concepts for signal and data processing. To address this level, ATL programs exist in advanced signal processing architecture to support radar, sonar, and image processing; in software engineering to support parallel processing needs and software development tools; in microprocessor/chip architecture to support custom VLSI system needs; and in optical storage to provide online storage systems for supercomputer systems. Finally, at the lowest level of complexity, there are programs in device technology. ATL is involved with VLSI design and related CAD tools to enable the development of the large VHSIC-class devices; with IRCCDs to build sensor subsystems for trackers, physical security, and surveillance; with laser diodes to build optically-activated switches and high-efficiency spaceborne laser sources; and with both silicon (CMOS) and GaAs technology for high-speed and radiation-hardened applications.

The articles in this issue represent a cross section of ATL programs, from device technology to examples of intelligent systems. They reflect the confidence and competence with which ATL and other RCA business units are responding to the challenges of today's environment.

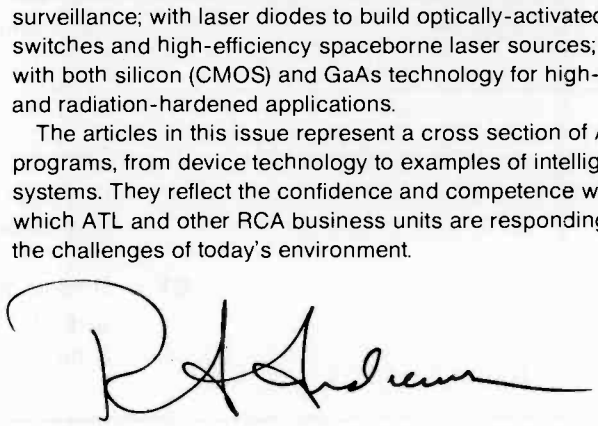

Ronald A. Andrews
Staff Vice President
Advanced Technology Laboratories

RCA Engineer
Vol. 31 | No. 1 Jan. | Feb. 1986

in this issue . . .
technology for the 1990s

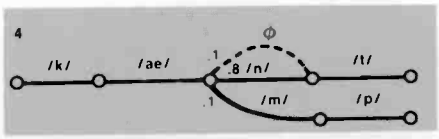

■ **Graff/Lynch:** *"Our primary effort is to develop acoustic-phonetic techniques that will permit the probabilistic segmentation of an input utterance into its constituent phones."*

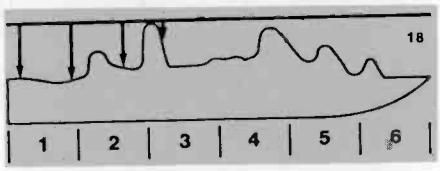

■ **Warren:** *"RCA has been a leader in CCD technology for the past decade, and this leadership has enabled unique contributions to both visible and IR sensors."*

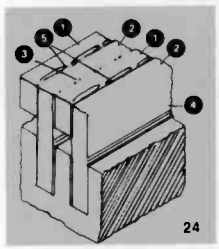

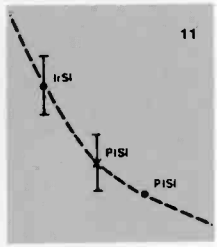

■ **Luce/Schaming:** *"The final goal of the project is to produce "flyable" hardware including an infrared sensor that will automatically detect and classify a target and make several tactical decisions in flight."*

■ **Wille/Rosen:** *"The pumping mechanism ensures that there are more atoms in the higher energy state than in the lower energy state, so gain occurs."*

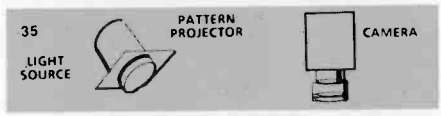

■ **Pierson:** *"To increase the capabilites of distant robotic servicing systems, we must develop predictive feedback techniques or give the robot increased autonomy."*

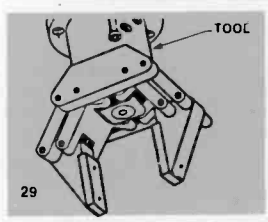

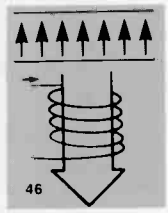

■ **Alexander et al.:** *"The apparent ease with which humans interpret the many elements of a scene conceals the great complexity of the processes involved."*

■ **Altman/Claffie/Levene:** *" . . . the Space Station, to be launched in the early 1990s, is expected to generate 1 terabyte of data per day at data rates that can exceed 1 gigabit per second."*

■ **Britton:** *"Verlangen is appropriate for many kinds of computer systems, including distributed systems, communications networks, and operating systems."*

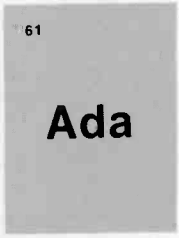

■ **Mebus/Armstrong/Rosenthal:** *"As far as the DoD is concerned, the Ada language is definitely the software systems language of the future."*

■ **Conklin:** *"Because a team of experts usually develops these systems, we will exploit the team approach to maximize the advantages it offers in design efficiency and quality."*

■ **Adams et al.:** *"Computer-based expert systems provide the means to capture such diagnostic expertise from a few human experts, and then to make this knowledge available to others who have less experience."*

■ **Helbig/Schellack/Zieger:** *"Improvements with GaAs processing technology are occurring at a rate that is three times what occurred in silicon processing during the 1970s and early 1980s."*

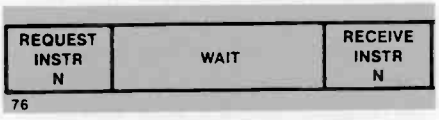

■ **Feller:** *"Current design VLSI methodologies, based on past LSI or even MSI design approaches, are responsible for excessively long design times."*

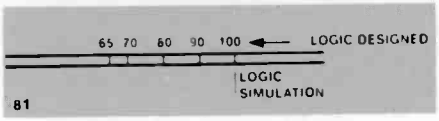

---

in future issues . . .
PC workstations on the job
laser technology

D.E. Graff | T.E. Lynch

Acoustic phonetic techniques for future automatic speech recognition systems

The closer we look at automatic speech recognition systems, the more we realize just how complex the human interpretation and filtering system is.

Ideally, an automatic speech recognition system (ASR) should respond to any spoken word from any speaker. However, despite active research in industry and academia over the past 30 years, this goal does not seem achievable in the near future.

Most current ASR systems, after being "trained" by the intended user, can recognize only a limited number of words (fewer than 300) spoken in isolation. Unfortunately, the techniques used in these primitive ASR systems have limited potential. We must develop new techniques before automatic speech recognition can become a viable technology.

Abstract: *At RCA's Advanced Technology Laboratories, we are pursuing fundamental research in automatic speech recognition (ASR), concentrating on acoustic phonetics. We are also particularly concerned with developing techniques that can be used in future ASR systems. Our research is performed in conjuntion with the University of Pennsylvania Linguistics Laboratory under an Air Force contract (F30602-84-C-0065). This effort parallels the current DARPA ASR work, although the two programs are not directly connected.*

Final Manuscript received November 20, 1985
Reprint RE-31-1-1

Scope of the ASR problem

To achieve an ideal ASR system from current technology, we need to extend the state-of-the-art in several directions. Some of them are depicted in Fig. 1.

First, the system must become speaker-independent. At present, most current systems can handle one speaker at a time, and that speaker must provide training models for each word that is to be recognized. This training process is very time consuming and tedious. When several thousand words are to be recognized, the task is prohibitive.

On the other hand, a speaker-independent system could handle any speaker, and thus we could bypass the training process for each individual speaker. To achieve good performance, the system should be made aware of the speaker's voice characteristics (average pitch, vowel patterns, etc.), but it can learn these characteristics much faster from estimates than by training each word model separately. Also, a speaker-independent system could adapt to a speaker over time.

Next, we must extend the vocabulary

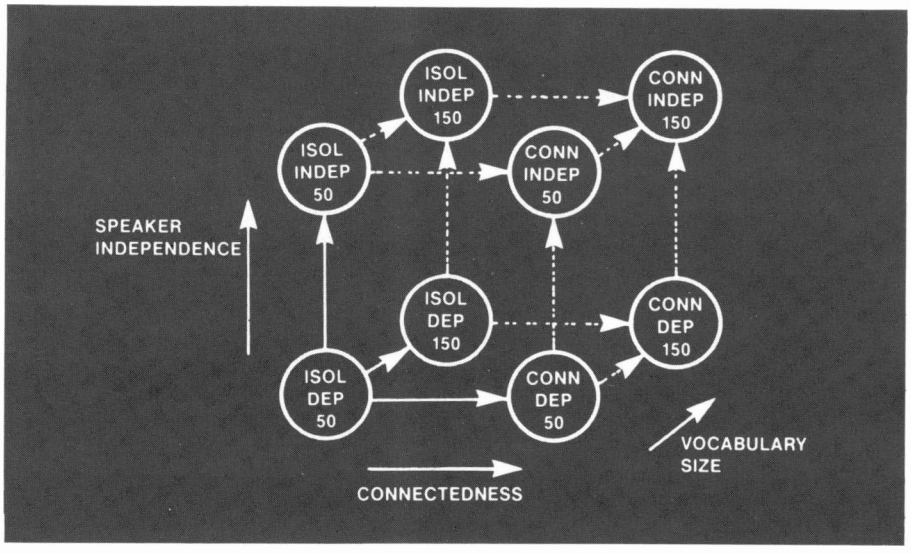

Fig. 1. *Three dimensions of automatic speech recognition complexity. The solid circles represent existing technology.*

from current limitations to a size suitable for the desired task. For instance, a voice input typewriter ("talking typewriter," or "talkwriter") would need a vocabulary of several thousand words, whereas a fighter cockpit function may require only a few hundred.

Most current ASR systems represent each word in the vocabulary as an acoustic model of the entire word. For vocabularies of several thousand words, the storage and pattern-matching computation requirements become prohibitive. English words can be separated into a finite number of smaller units, called phonemes (see Fig. 2). We could model each word as its phoneme string, with some rules for phoneme interactions both within words and at word boundaries (co-articulation).

Thus, a more reasonable approach is first to model the "phones" acoustically and then to model each word as the composite of its phones, with acoustic rules at phone boundaries. (A phone is the acoustic realization of the phoneme).

Third, we must extend the speech style that the ASR system requires from words spoken in isolation (discrete speech) to words spoken more naturally (connected speech). However, it may not be possible to handle totally "natural" speech (How 'bout dem Dawgs?), and we may require some degree of careful articulation by the speaker.

Restricted grammar is a related topic. Currently, most ASR systems require the user to speak within a constrained syntax (see Fig. 3). The alternative is to allow conventional spoken English grammar. However, such an unrestricted syntax requires an artificial intelligence (AI) interpreter to interpret the message, and is not a major focus of our effort. Figure 4 depicts some forms of speech that could be used as input to an ASR system. Other factors that affect ASR include noise, channel, and human factors, but these too are not of primary concern in our research.

PHONEMES

SONORANTS			OBSTRUENTS (TURBULENCE EXCITATION AT VOCAL TRACT OBSTRUCTION)			
VOWELS			VOICED			
FRONT	MID	BACK	STOPS	AFFRICATES	FRICATIVES	WHISPER
i (BEET) I (BIT) e (BAIT) ɛ (BET) ae (BAT)	ɜ (BIRD) ə (BUT)	u (BOOT) U (FOOT) o(BOAT) a (HOT) ɔ (BOUGHT)	b,d,g	dz (JOURNAL)	V (VOICE) ð (THEN) Z (ZOO) ʒ (AZURE)	h
DIPTHONGS			UNVOICED			
FRONT GLIDING		BACK GLIDING	STOPS	AFFRICATES	FRICATIVES	
aI (BUY) ɔI (BOY) eI (BAY)		aU (HOW) oU (BOW) jU (YOU)	p,t,k	tʃ (CATCH)	F (FINE) θ (THIN) S (SEE) ʃ (SHUN)	
CONSONANTAL SONORANTS						
LIQUIDS	GLIDES	NASALS				
R,l	y (YET) w (WERE)	m,n ŋ (RING)				

Fig. 2. *The phonemes of American English. Glides, liquids, and nasals are produced with some vocal tract obstruction, but are generated by vocal cord vibrations. Voiced fricatives combine vocal cord (glottalic) and obstruent (turbulence) excitations.*

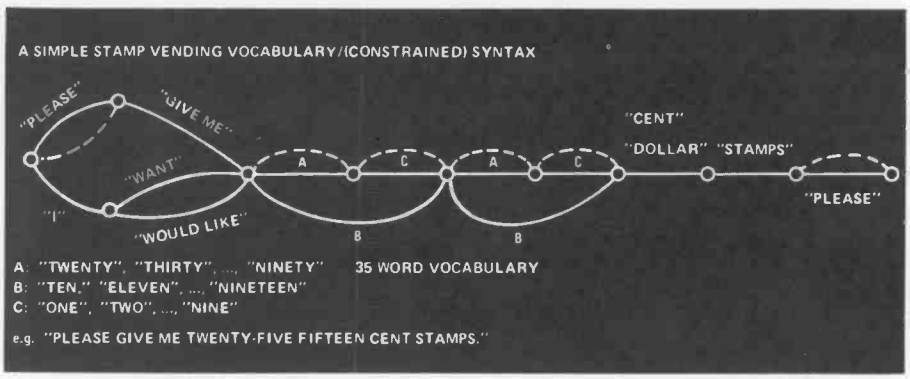

Fig. 3. *A simple stamp-vending vocabulary with constrained syntax.*

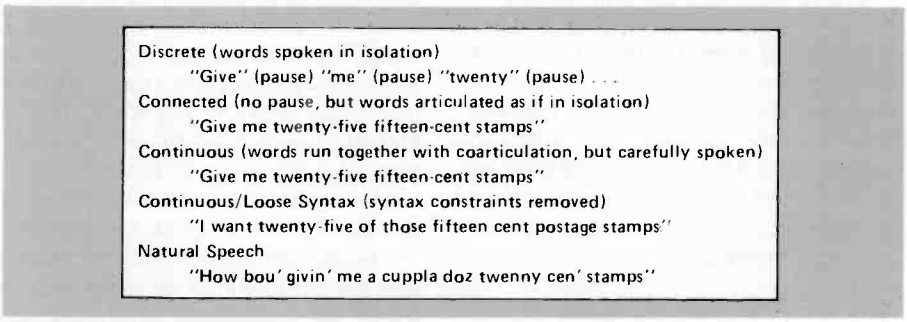

Fig. 4. *The progression from discrete speech to natural speech.*

ATL ASR research

RCA's Advanced Technology Laboratories (ATL) is engaged in fundamental research focused on developing future, speaker-independent connected-speech (SICS) ASR systems with a medium to large vocabulary (greater than 200 words). Our primary effort is to develop acoustic-phonetic techniques that will permit the probabilistic segmentation of an input utterance into its constituent phones. We have also directed some effort to developing recognition algorithms for determinating the original word sequence from the segmented phones. We are not concerned with how the resulting word sequence is interpreted, leaving such problems to future AI work.

Will true speaker-independent automatic speech recognition ever become a viable technology? Certainly that question clouds all work being done in this field. For instance, in some urban dialects of the Great Lakes region, the word "locks" is pronounced identically to a Philadelphia area "lax." Thus, no ASR system could distinguish between the two words absolutely without con-

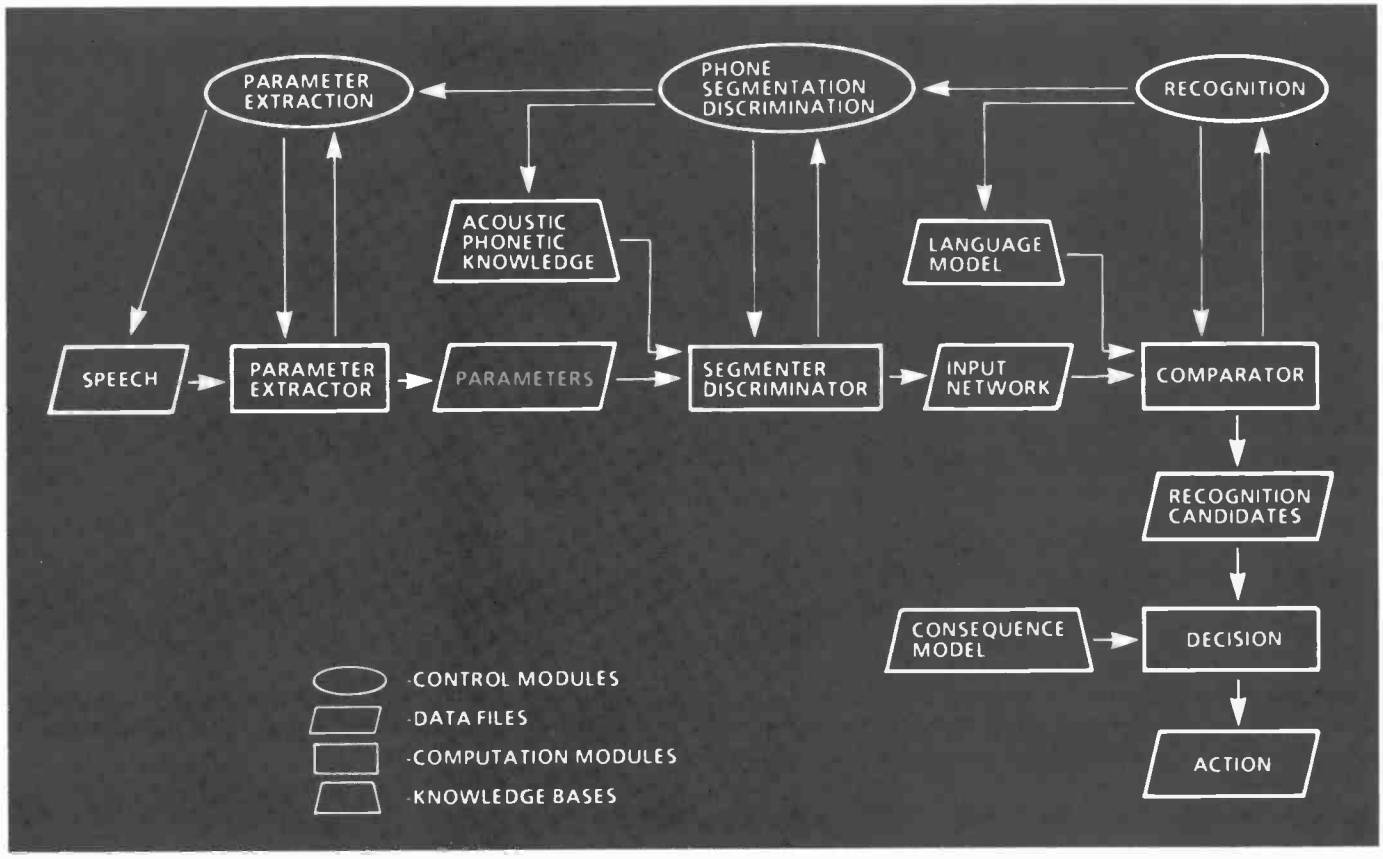

Fig. 5. *Phonetic recognition system.*

text or speaker-dependent information.

Furthermore, we do not find it desirable to handicap system performance so severely when it is relatively easy to inform the system about the speaker's vowel characteristics and other dialect (or ideolect) information.

Thus, we focused our acoustic-phonetic research on finding speaker-independent speech attributes and a small set of parameters that can be easily adjusted to an individual speaker. Whether we achieve true speaker-independent ASR or not, whatever we learn in this area can be applied to speaker-independent and speaker-adaptive systems.

ASR model

Figure 5 depicts our automatic speech recognition system model, which consists of four modules: parameter extraction, phone segmentation/discrimination, word string recognition, and recognition decision.

Parameter extraction processes input speech to produce parameters that locally characterize the incoming utterance. In human hearing, the cochlea (part of the inner ear) does the first stage of parameter extraction. This stage consists (roughly) of a time series of spectral amplitude information derived from a filter bank (the basilar membrane, which is part of the cochlea).

Our parameter extraction consists primarily of a time series of spectral amplitude information that we obtained from a "critical-band" filter bank (a rough approximation of the basilar membrane), and more direct information about peaks in the spectrum (formants). The novelty of our parameter extraction is that we do the analysis in a pitch-synchronous manner, which is described in the next section.

Phone segmentation/discrimination uses the parameter array to segment (in time) the incoming utterance into phones and then identify each phone (phone discrimination). Segmentation and discrimination are done on a probabilistic basis. For instance, if we find a stop consonant but cannot absolutely determine if it is to be a /p/, /t/, or /k/, then we maintain the three candidates as well as their relative likelihoods.

Word string recognition takes the probabilistic output of the phone segmenter/discriminator and, working from the language model (vocabulary and syntax), recovers the likely reconstructed word strings. We have developed an algorithm to do this recognition using the probabilistic output of the phone segmenter/discriminator as input.

Finally, the consequence model is used to decide if the resulting word string makes sense and if the likelihood of a correct recognition is enough to take action or should be verified. This falls in the realm of AI or operations research, and again is not one of our focuses.

Note that our model has three knowledge bases that contain:

- Acoustic/phonetic information (used in the phone segmentation/discrimination process).
- Language information (vocabulary and syntax models used in word string recognition).
- A consequence model (used in deciding what to do after recognition).

Further, most modules are interactive. To resolve a word-recognition ambiguity, the recognition module could request a more careful analysis from the segmentation/discrimination module that could, in turn, request finer param-

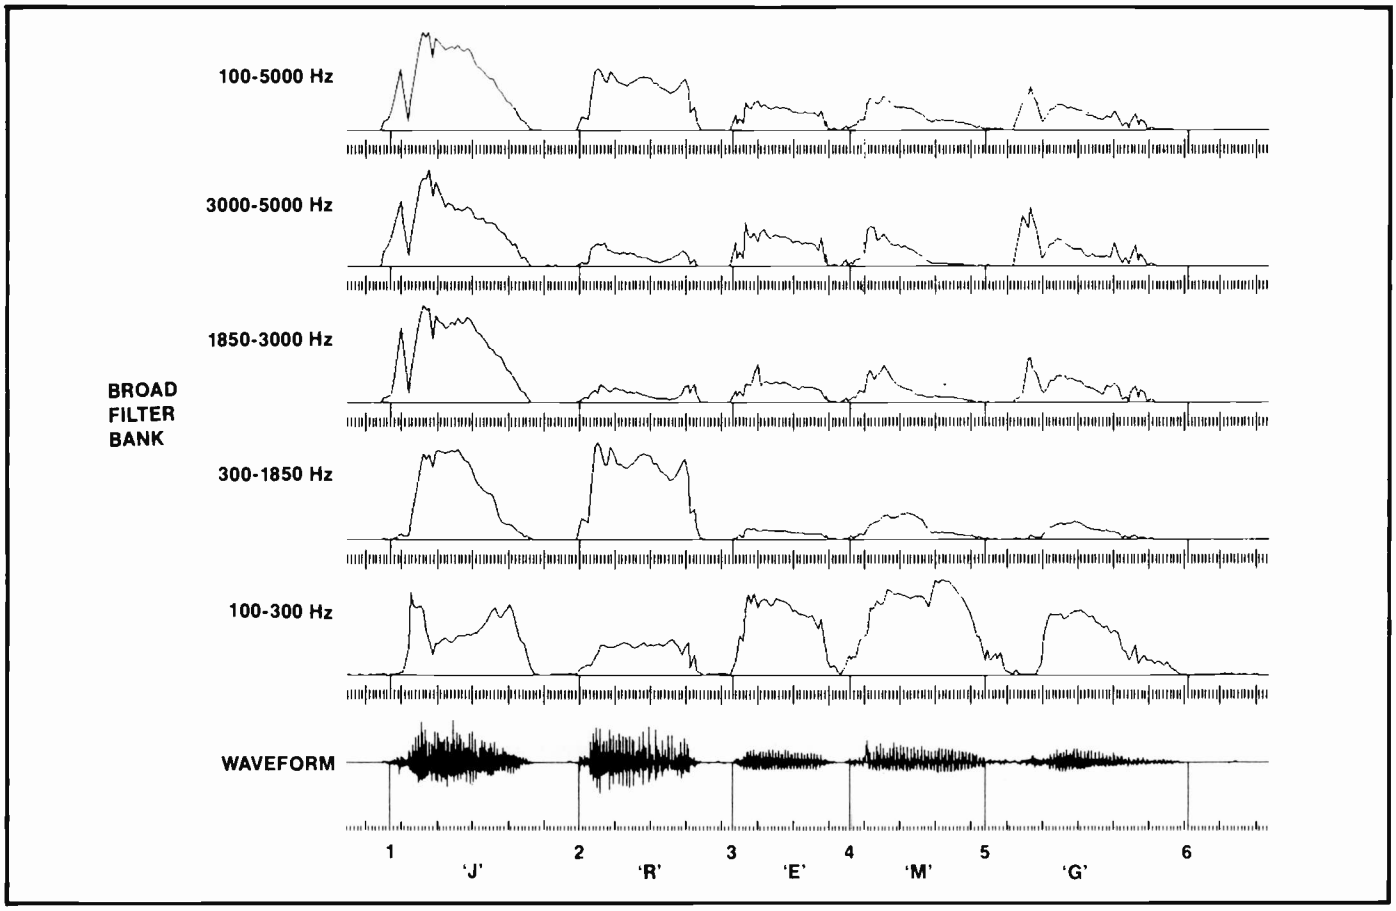

Fig. 6. *Filter bank structure of the utterance "J-R-E-M-G."*

eter extraction over a segment of the input speech. Thus, as Fig. 5 shows, lines of control lead right-to-left and lines of data flow left-to-right; so we can use both top-down and bottom-up processing. In particular, blackboard paradigms and other AI decision techniques fit in this model.

Parameter extraction

We generate speech sounds either by periodic vibrations of the vocal cords (sonorants: vowels, nasals, glides, etc.) or by turbulence near an obstruction in the vocal tract (obstruents: /s/, /f/, /p/, etc.). The sonorants include vowels (beet, bat), dipthongs (boy), glides (we), liquids (red, led), and nasals (bean, beam). The obstruents include stops (beet, peat), fricatives (seat), affricates (cheat), and whisper (heat). Figure 2 contains a complete list of these speech sounds.

The obstruents, created by turbulence, are noiselike and have their principal frequency components in the upper frequency range (greater than 2000 Hz). They can change rapidly in time (stop releases), and therefore are best analyzed with a small time window (less than 5 milliseconds) spectral analysis that is computed frequently (at least every 5 milliseconds). Thus, in obstruent regions, our spectral analysis is done using 5-millisecond effective time windows that are computed every 5 milliseconds.

The sonorants are more difficult to treat, because the results of the spectrum analysis depend on the position of the spectral-analysis window with respect to the pitch pulse. The classical way around this problem is to adopt a wide window (20 to 30 milliseconds) that will contain several pitch pulses and then smooth the effects of the individual pitch pulses.

Our approach is to adopt a very narrow window (5-millisecond effective width) and place it the window directly over the pitch pulse for each pitch period. This requires a reliable pitch tracker that can accurately locate the pitch-pulse for each vocal cord vibration. Our pitch tracker has proven to be remarkably robust and the resulting spectral analysis tracks the spectral movement very finely and with excellent stability, considering the short analysis windows involved.

Figures 6 through 8 illustrate our parameter extraction. The utterance in Figs. 6 and 7 is the sequence of letters, "J-R-E-M-G." Figure 6 presents the energy structure over several energy bands, and Fig. 7 indicates the formant structures, computed in a pitch-synchronous manner. Formants correspond to the passbands in the vocal tract (viewing the vocal tract as a bandpass filter). Notice that the formants occur in voiced regions and that the lowest two formants characterize the vowel being spoken. In particular, the formant patterns in "E" and "G" are roughly the same because they represent the "ee" sound. Figure 8 shows the formant pattern for the word "are". Notice the formant movements.

Phone segmentation/discrimination

Ideally, a phone segmentation/discrimination module would decode an input-utterance waveform into its underlying phonemes. Then, with some clever

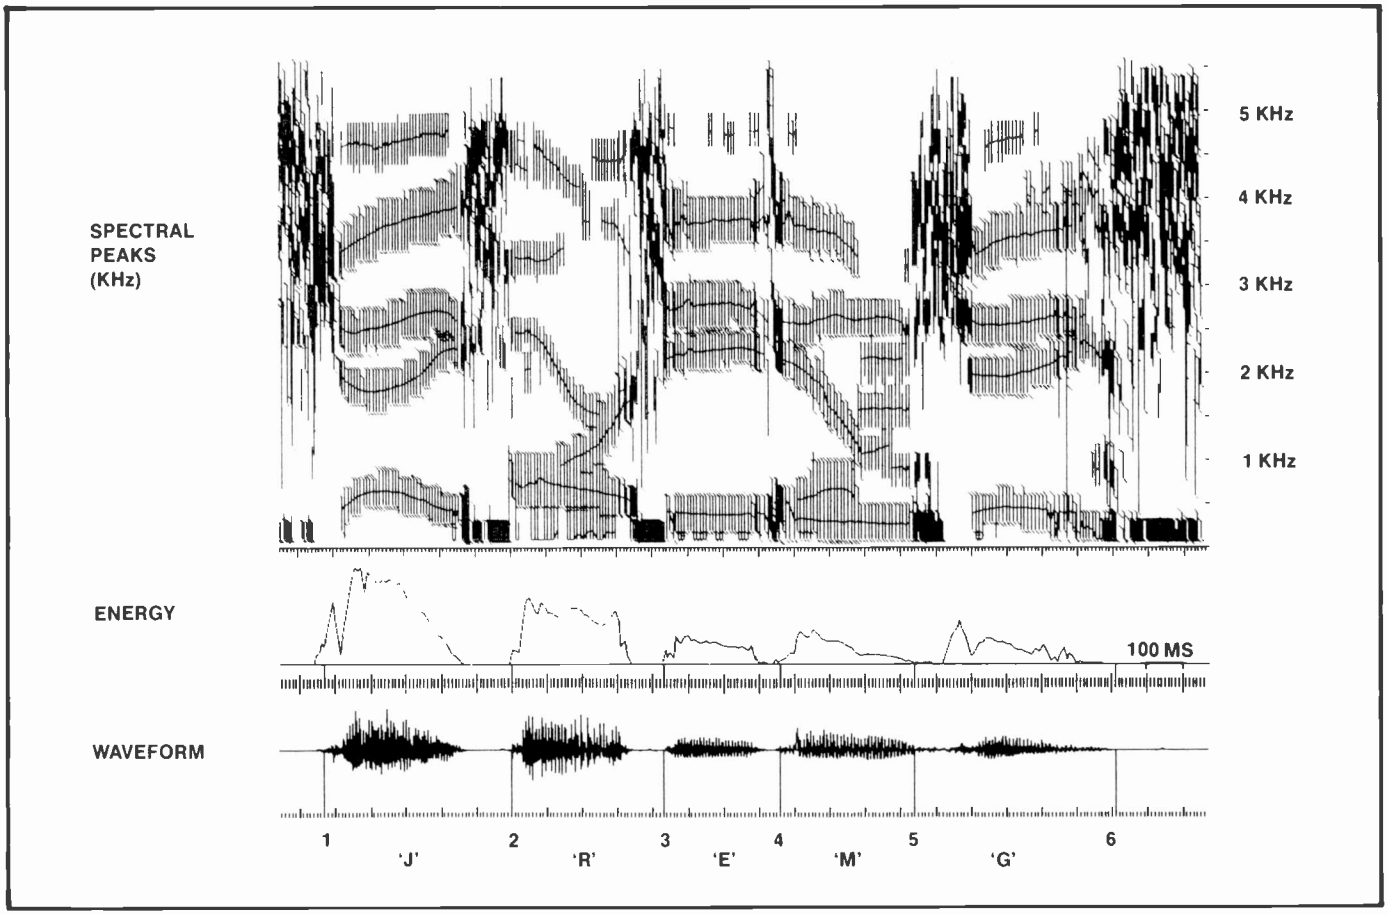

Fig. 7. *Formant structure of the utterance "J-R-E-M-G." Note the clear formant structure in the vocalic regions.*

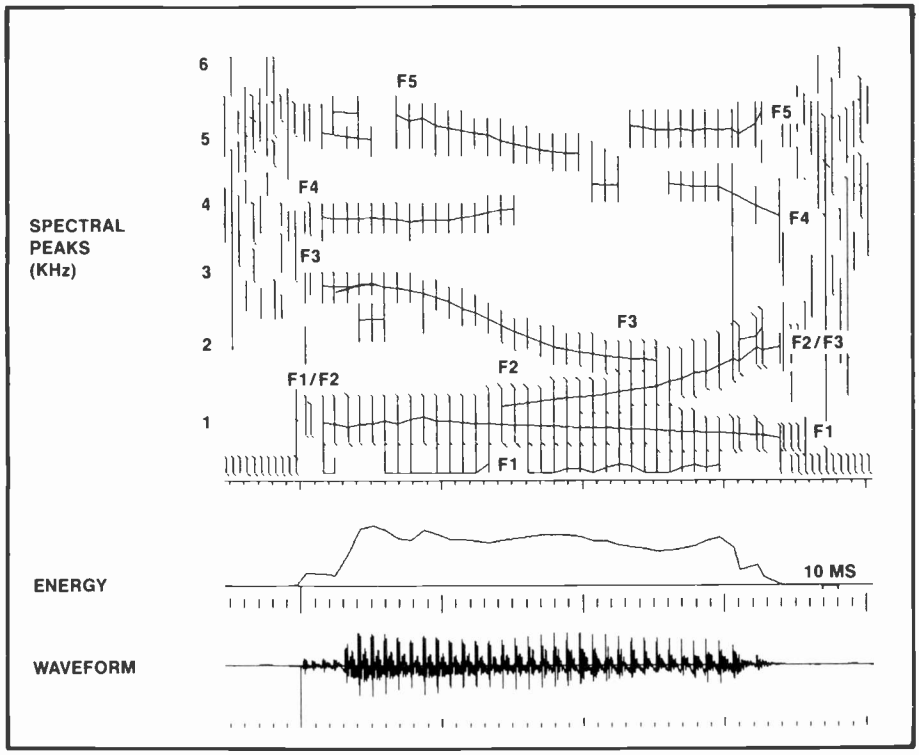

Fig. 8. *Formant structure for the word "are." Note the f1/f2 merger diverging into a f2/f3 merger as the low back vowel (ah) makes the transition to the retroflexed liquid /r/.*

string-matching and a good dictionary, it could recover the underlying message (although it might not be able to distinguish "six sheep" from "sick sheep").

However, even humans are not able to decipher absolutely the phonetic content of things they hear. Therefore, we needed a hierarchical, probabilistic model for a speech understanding/decoding technique. Some acoustic events, such as silence or strong vowels, can be distinguished reliably and can serve as "islands of certainty" around which segmentation proceeds.

In our system, we initially segment the utterance into regions: strongly voiced (sonorant), strongly unvoiced (silence, frication), and intermediate (devoiced sonorants, voiced fricatives, or transitions). In the strongly voiced regions, we find likely strong vowel regions. In the strongly unvoiced regions, we differentiate silence from frication. We then subdivide the regions until all segment candidates in the incoming utterance are labeled.

Discrimination then attempts to assign a phone label to each segment.

For instance, if we determined a segment to be a stop, then discrimination would try to assign (/b/, /d/, /g/, /p/, /t/, /k/, or "flap") to the segment.

As mentioned previously, humans are certainly not totally accurate in the phone segmentation/discrimination task. In fact, our colleagues at the University of Pennsylvania Linguistics Laboratory are currently measuring just how well humans do certain phone segmentation/discrimination tasks. This suggests that if we try phone segmentation/discrimination, the resulting output should not be a linear sequence of labels, but rather a transition network of phone labels and transition probabilities. Figure 9 depicts a typical local output.

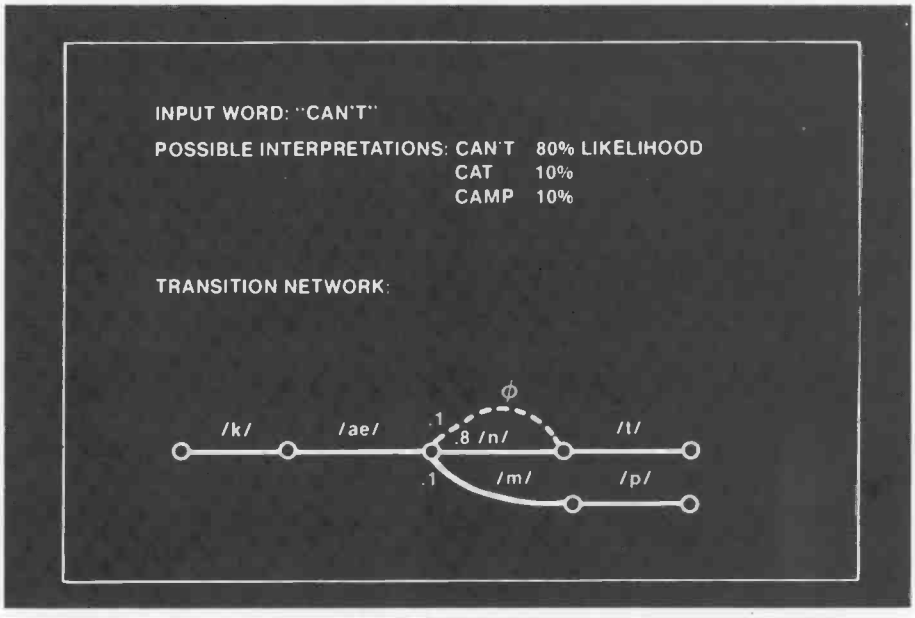

Fig. 9. *Construction of a probabilistic transition network from a probabilistic segmenter/discriminator.*

Word string recognition

Even when given a linear sequence of decoded phone segment labels, extracting the underlying word string sequence is difficult because of input errors in segmentation and labeling. Also, the same word may be produced in several different ways and adjacent words can interact acoustically at their boundaries (co-articulation).

For example, a southern speaker may say "farm" either in the standard manner (with an /r/), or as "fahm," or anything in between. A common example of co-articulation is the conjunction of the words "did" and "you" to form "dija."

In our research, we have developed a recognition algorithm that permits comparing probabilistic input transition networks with statistically derived reference transition networks. This algorithm incorporates compensations for segmentation/discrimination "errors," multiple pronunciations, and co-articulation. Like dynamic programming and string matching, our algorithm has a computational complexity that is roughly linear in relation to vocabulary size and grammatical complexity. Moreoever, our algorithm is an extension of dynamic programming and string matching; it incorporates these strategies as special cases. However, an exposition on these topics is not possible here, so anyone interested in further information should contact the authors.

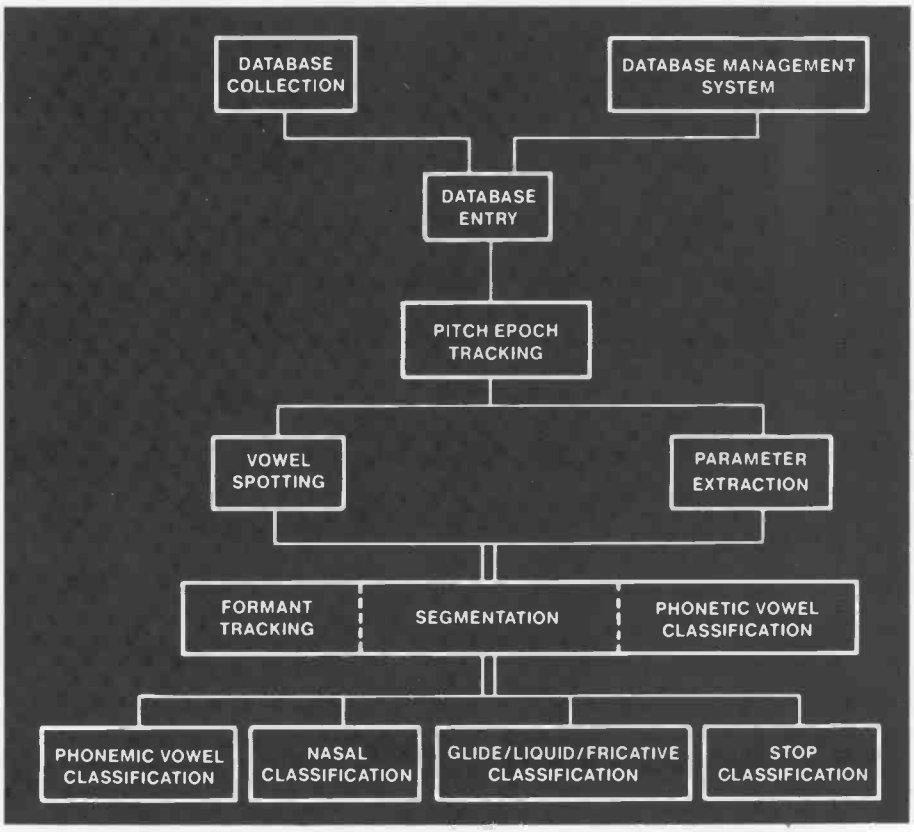

Fig. 10. *Acoustic/phonetic research modules.*

Current efforts

Figure 10 shows the acoustic/phonetic research modules being investigated under the SICS contract. In our current efforts we have identified stop discrimination and vowel assignment as the two most difficult acoustic-phonetic problems in speaker-independent automatic speech recognition.

Of these, stop discrimination is the most difficult acoustic problem because no known reliably computable statistics exist that will absolutely (or even reasonably) discriminate among stops. Through our robust, small time-window spectral analysis, we hope that new statistics will become computable

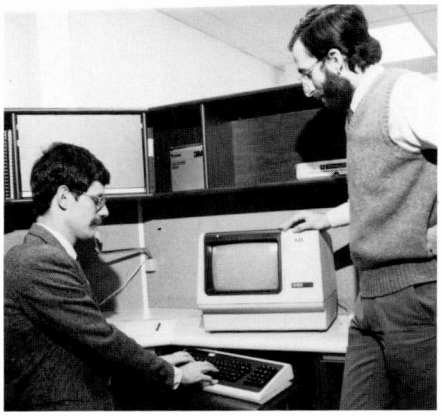

Lynch (left) and Graff.

Thomas Lynch is a Senior Member of the Engineering Staff in ATL's Speech Technology Laboratory. He is working on the development of advanced speech recognition systems. Previously, he was involved with secure speech transmission and modem software development. Prior to joining RCA in 1981, Dr. Lynch was at Verbex, where he concentrated his efforts on automated recognition of continuous speech. He received a PhD in Mathematics from Massachusetts Institute of Technology and a BS in Mathematics from the University of Georgia.

David Graff is a Member of the Engineering Staff at RCA's Advanced Technology Laboratories. He joined the Speech Processing group in 1983, participating immediately in the Speaker-Independent Connected-Speech (SICS) program. He is currently working toward a PhD in Linguistics at the University of Pennsylvania, with a focus on change and variation in American dialects. He received his BA in Linguistics from Pitzer College.
Contact him at:
Advanced Technology Laboratories
Moorestown, N.J.
Tacnet: 253-6613

and allow more robust discrimination than has been available.

Phonemic vowel assignment, which must be treated separately from the estimation of acoustic vowel quality, represents the most difficult task speaker independence faces. This results from two essential facts about dialect variation, already established by work at the University of Pennsylvania.

First, dialects differ about where particular vowels are located in the phonetic space defined by the positions of the first two formants. (Consider the previous example where "locks," as pronounced by some speakers in the Great Lakes region, sounds like "lax" as pronounced by a Midlands speaker, or like a Bostonian's "larks"). Second, dialects differ in how many vowel phonemes are pronounced differently, as well as how phonemes are distributed in the vocabulary.

Most southerners do not differentiate "pin" and "pen" or "mint" and "meant," while the country seems evenly divided between those who pronounce "cot" and "caught" or "hawk" and "hock" differently and those who do not. In addition, those who do have the latter distinction may differ in how they apply it. For Bostonians, "bomb" has the vowel of "hawk," but for New Yorkers, it has the vowel of "hock".

We will examine this problem with special emphasis on how to encode a person's vowel characteristics to permit better recogniton by an ASR system.

In parallel, the University of Pennsylvania Linguistics Laboratory will continue to do research in human perception, especially human capability relating to stop discrimination and dialects. Their work on the human capacity to do segmentation and discrimination provides useful information for developing of automatic techniques.

Summary

At ATL we are in the midst of a two-year project to do fundamental research on topics relating to automatic speech recognition. While we have done some work in recognition algorithms, our primary emphasis has been on the development of new acoustic/phonetic techniques. Progress to date has been very encouraging, although the ultimate usefulness of our research may be several years away.

F.B. Warren

Advanced technology for future imaging sensors

RCA's unique approach to building infrared sensors has a promising future.

A method for electronically shifting or multiplexing signals in two dimensions is essential for practical operation of a large array of individual detectors. At present, only silicon-based multiplexers, such as the silicon charge-coupled device (CCD), can do this.

Because silicon is sensitive to visible radiation, a multiplexed visible imager can be monolithic. In this case, silicon photosites produce charge as a result of incident photons, and we can transfer this charge down a line of CCD elements with little distortion. This creates a visible imager that is a mosaic of discretely addressable picture elements.

For the infrared bands, sensors previously consisted of sheets of IR detectors bonded onto CCD multiplexers so the CCD storage wells can collect the detector output charge. This creates a hybrid sensor chip. To store a signal, CCDs integrate the output current of a photosite or detector. Thus, all multiplexed imager arrays are integrating sensors. RCA now uses a monolithic approach to sensor construction, described in this article.

RCA has been a leader in CCD technology for the past decade, and this leadership has enabled unique contributions to both visible and IR sensors.

Abstract: *RCA has developed unique sensors that will solve future imaging problems in terrestrial and space applications. The major uniqueness and benefit of RCA's approach is the development of monolithic silicon sensors. This is the only technology that currently can achieve the large multidetector focal planes needed in the future.*

Final Manuscript received November 20, 1985
Reprint RE-31-1-2

Visible CCD technology

A recent RCA advance is the fabrication of a large, linear visible array (VIS/NIR) for the Multi-Spectral Linear Array (MLA) effort funded by NASA's Goddard Space Flight Center. MLA technology will be used in earth resource work, such as future LANDSAT missions. Present earth-satellite sensors for this purpose use a small number of detectors that are mechanically scanned orthogonal to the path of the satellite's motion.

RCA has constructed a large, hybrid, visible focal plane that contains four lines of 5120 picture elements (pixels) each. (The focal plane combines monolithic line-imager chips with hybrid electronic components that are placed directly on the focal plane.) This pixel density allows us to image the earth's surface from the "pushbroom" scan that the satellite's motion produces; we do not need an additional mechanical scan across track. As a result, each element dwells on a scene pixel 400 times longer than the sensors now in use.

This focal plane was fabricated as a combined effort by Advanced Technology Laboratories (ATL), RCA Laboratories, and New Products Division. RCA offered two advantages over competing CCD technologies: the high performance provided by a proprietary CCD thinning process, and integral spectral filters in the visible spectrum.

The VIS/NIR chip—a frame-transfer imager—uses buried-channel, three-phase polysilicon gate technology like the RCA 605 CCD imager.[1-3] To obtain exceptionally high photosensitivity and resolution (inherently limited by the pixel structure), we used a 10-micrometer silicon substrate illuminated through the unstructured rear side of the device. We cannot do this with the other CCD technologies that do not feature thinned devices, because visible light does not penetrate silicon that is thicker than a few micrometers. These technologies force us to use front-side illumination, which degrades the image partly because polysilicon gate features on the front side of the CCD scatter the signal.

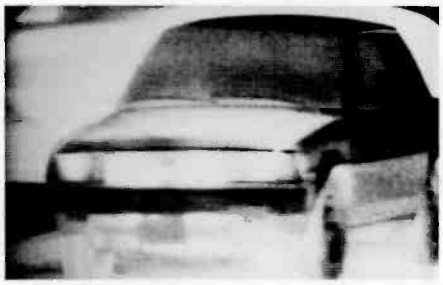

Night imagery from RCA's 160× 244 IRCCD array (reproduced from a television monitor). Note the heat pattern on the car's hood and the thermal reflection of the exhaust system on the road.

An important and unique feature of the VIS/NIR design concept is the integral sensor/filter approach. The set of MLA optical filters now being produced has 70, 80, 40, and 140-nanometer[2] passbands. Their design, which requires high in-band and low out-of-band transmissivity between 0.4 and 1.0 micrometers, was carried out at RCA Laboratories using a computerized modeling technology.[4]

To determine how many layers would provide the desired slope for each side of a wide bandpass filter, we started with an existing catalog of quarter-wavelength stacks. The stacks had alternating high and low indexes of refraction

Electro-optical sensors

The optical spectrum has four major atmospheric windows, and the atmosphere is opaque to electromagnetic radiation outside them. The visible spectrum lies within the most familiar window, followed by the short-wave-infrared (SWIR), mid-infrared (mid-IR), and thermal-infrared (TIR) windows. These mid-IR and TIR window regions are classically called the "infrared" bands, or the 3-5 micrometer and 8-12 micrometer bands, respectively. Fig. 1 shows atmospheric transmission over a 20-km horizontal path under typical conditions. We can see similar transmission for vertical paths from ground to space.

While all objects with temperatures above absolute zero radiate energy at all wavelengths, we cannot detect enough self-emission from a room-temperature object in the visible and SWIR bands. Therefore, we use these bands only to observe sources at elevated temperatures, or objects that reflect illumination from those sources. However, we can use the two infrared (IR) bands to detect self-emitted energy from bodies that are near room temperature. This is the basis for infrared sensors that "see in the dark."

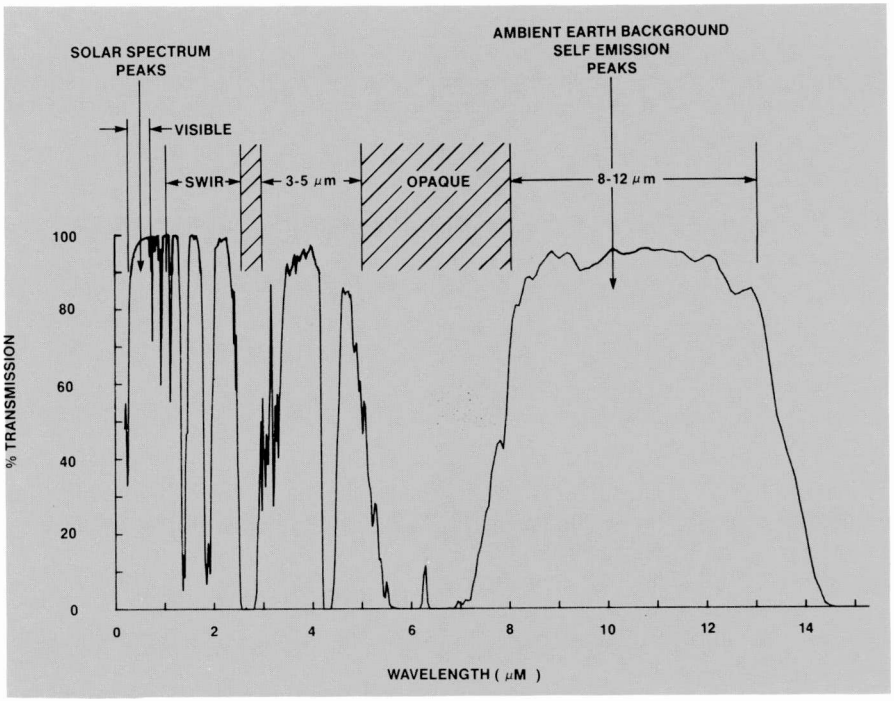

Fig. 1. *Typical atmospheric transmission at optical wavelengths, along a sea-level horizontal path.*

and were from one to 20 layers thick between $\lambda = 0.40$ to $\lambda = 1.0$ micrometers. To adjust the wavelength of the filter, we changed the thickness of the layers by successively altering the wavelength at which the layers are calculated. We added layers to match the stacks to the substrate and to each other.

Figure 2a illustrates a 12-layer stack at 0.38 micrometers; the first layer has an eighth-wavelength high index of refraction for matching to glass. Fig. 2b represents a 13-layer stack at 0.605 micrometers. When we superimpose the two stacks, we obtain a bandpass filter of our blue filter design.

However, from the diagram we can see that the transmission increases markedly beyond about 0.60 micrometers. If we need further blocking, we must add stacks that will transmit in the filter region and reject anything past 0.68 micrometers. Adding another stack (Fig. 2c) at 0.80 micrometers will then lower the transmissivity out to about 0.90 micrometers. In this way, we can get blocking to a desired wavelength on both sides of a filter. All this provided the starting point for the MLA design.

To complete the design, we used an RCA-developed computer program, Synthesis and Monitoring of Optical Interference Coatings. We made many iterations to minimize the number of layers required, to use primarily quarter-wavelength layers (for ease of monitoring), and to satisfy the final specifications. Where difficulties in direction arose, an RCA-developed optimization program guided us.

VIS/NIR chips were integrated into a hybrid focal plane that had five sensor chips, physically abutted, giving us four lines of 5120 pixels each[1]. The mechanical alignment and placement of these chips was very demanding. For example, no more than two 15-micrometer pixels could be lost in an abutment seam. SWIR devices were placed in a similar focal plane containing two lines of 2560 elements each.

Both focal planes (Fig. 3) are being delivered to NASA Goddard.

Infrared CCD technology

RCA has extended its expertise in visible CCD technology to the development of infrared CCDs (IRCCDs). We are using the metal-silicide Schottky barrier detector approach that was encouraged and funded by Dr. Freeman Shepherd of Rome Air Development Center (Hanscom AFB).[5,6] A detailed description of the physics involved with Schottky barrier detectors is beyond the scope of this article, but a brief explanation is possible.

When deposited directly onto silicon,[7] a metal, such as platinum, forms a metal-silicide (PtSi) layer. This layer is a photodiode that is sensitive to IR radiation.

If we arrange the architecture of a CCD sensor chip properly, we can configure many of these photodiodes to feed a signal directly into multiplexed

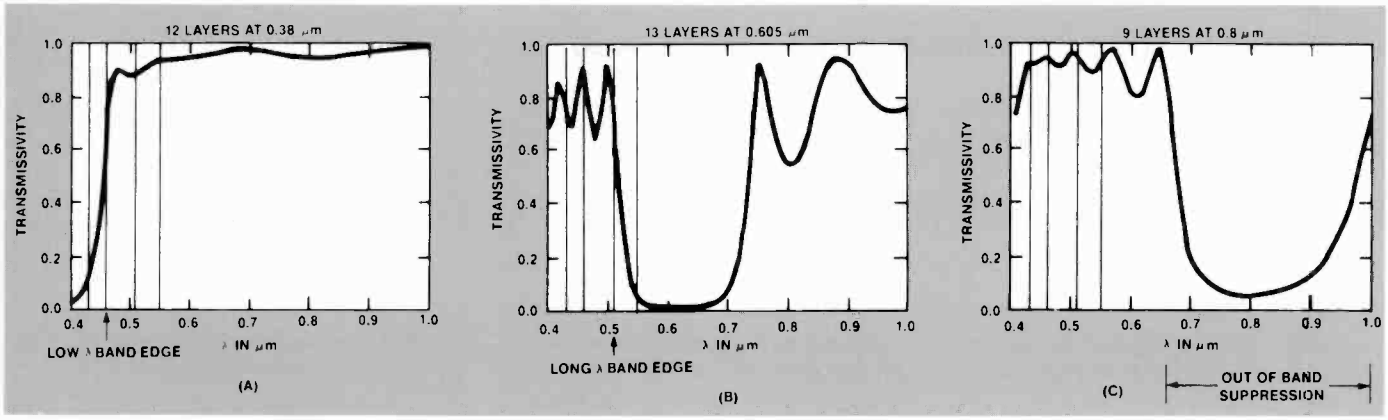

Fig. 2. *Filter stack technology: interference filter transmissivity versus wavelength for a number of alternating high and low refractive index layers.*

- 5 BUTTED DEVICES
- 4 SPECTRAL BANDS (INTEGRAL FILTERS)
- 5,120 PIXELS PER BAND
- ONLY 30μm (2 PIXEL) LOSS AT SEAM
- 3μm MISALIGNMENT
- 5μm PLANARITY DEVIATION (Z)

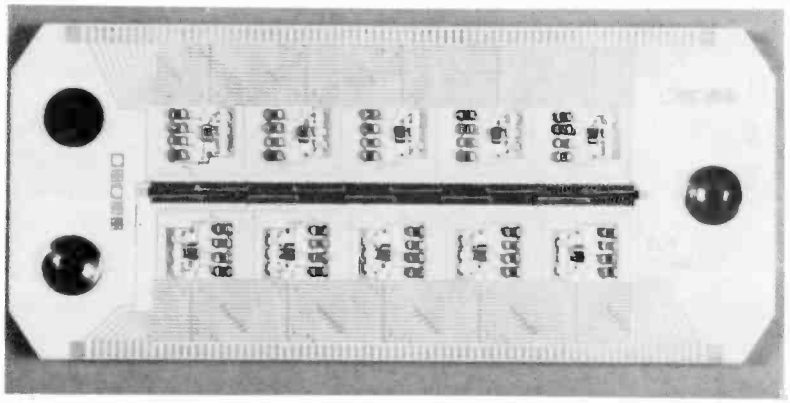

- 5 CCD's WITH FILTER
- 20 JFETS
- 30 RESISTORS
- 50 CAPACITORS
- 240 WIRE BONDS

BeO SUBSTRATE WITH

- KINEMATIC MOUNT
- PRECISION ULTRASONICALLY MACHINED OPTICAL APERTURE (f/2)

(a) VIS/NIR HYBRID

SWIR HYBRID

- 5 BUTTED DEVICES
- 2 SPECTRAL BANDS
- 2,560 DETECTORS PER BAND
- ONLY 60μm (2 DETECTOR) LOSS AT SEAM
- 5μm MISALIGNMENT
- 5μm PLANARITY DEVIATION (Z)

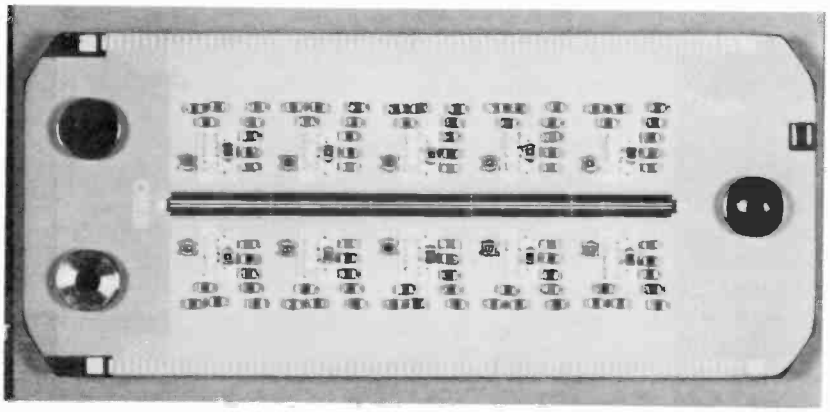

- 5 IRCCD's
- 20 JFETS
- 30 RESISTORS
- 70 CAPACITORS
- 245 WIREBONDS

THICK-FILM-ON-SILICON SUBSTRATE

- ALL MATERIALS MATCHED TO SILICON CTE

(b) SWIR HYBRID

Fig. 3. *Hybrid focal plane arrays.*

Important difference

The two IR bands have major differences that become important when we consider technologies that are competitive with RCA's. Background scenes near 300K (27°C) ambient temperature have peak emission at 10 micrometers, the center of the 8-12 micrometer band. As a consequence, the 8-12 micrometer band contains about 20 times more signal energy from an ambient scene than the 3-5 micrometer band. However, because of the physical advantages of 3-5 micrometer sensor technologies and the reduced scene clutter in the mid-IR region, the 3-5 micrometer region has equal or superior performance for some IR sensor applications.

The tradeoff between the 3-5 micrometer and 8-12 micrometer bands for a given application requires a complex analysis, and slight variations of basic assumptions strongly influence the results of that analysis.

Although the choice between the 3-5 micrometer or 8-12 micrometer bands is often controversial, some general-use trends have been established. We often use the 8-12 micrometer band to observe ambient scenes passively, and detect specific objects (targets) along lengthy horizontal paths (particularly where we expect high humidity). A typical application is an IR imager for piloting a vehicle.

The 3-5 micrometer band is often used to observe or detect scene objects, such as vehicles or aircraft, that have slightly elevated temperatures.

However, a good 3-5 micrometer imager can compete with an 8-12 micrometer imager in observing ambient scenes, and 3-5 micrometer imaging systems are usually less expensive than comparable 8-12 micrometer systems.

Sensor characteristics vary widely for different requirements at the four atmospheric windows.

Visible technology is the most mature for two reasons: need and availability. The need is obvious because our only source of natural illumination peaks in the visible region. Imaging sensors—such as the human eye, television, and film cameras— were developed earlier than IR imagers because detection of visible photons is easier.

In a relatively stable material—such as silicon—that can be operated at room temperature, a visible photon has enough energy to excite a signal electron across the energy gap (move it from the valence band into the conduction band). The weaker IR photon can only dislodge the more loosely held electrons from a less-stable material (lower energy gap). However, the material must then be cooled to prevent the excessive signal noise that arises because the thermal agitation releases random electrons.

Visible imagers such as television vidicons and film are, in effect, X-Y addressable because we can correlate the image signal stored in each physical location of the two-dimensional sensor surface with the location of the external scene object that is imaged onto it. These X-Y locations are analog in nature, with no crisp or quantized boundaries between adjacent scene locations. Such imaging sensors are feasible in the visible and near-infrared spectrum because photon energies exist in this region.

However, no mid-IR or TIR vidicon or film exists (except the pyroelectric vidicon, which is a nonquantum imaging IR sensor with low sensitivity and low general performance). The IR-sensitive materials can operate viably only as

CCD storage wells. The result is a one- or two-dimensional matrix of multiplexed detector elements that exhibit very high-response uniformity and excellent picture quality. For Schottky barrier IRCCDs we have measured a 0.5 percent total rms response nonuniformity over 39,000 array-detector cells (160×244 configuration). The hybrid imager technologies, indium antinomide (InSb) and mercury cadmium/telluride (HgCdTe), exhibit response nonuniformities of 5 to 50 percent over much smaller arrays.[8]

We have constructed multiplexed focal plane arrays (FPAs) with hybrid technologies and with IRCCDs. Our emphasis has been on area arrays rather than line arrays, although both have been successfully fabricated. While we have demonstrated the principle of large staring FPAs, affordability remains questionable. The leading technology in this area is HgCdTe.

About a decade ago, this country selected HgCdTe as its strategic detector material because HgCdTe detectors can be tailored to respond in the SWIR, 3-5 micrometer, or 8-12 micrometer spectral regions. They also exhibit high quantum efficiency. Since then, more than one-half billion dollars in developmental funding has been expended. This investment has created a marketing interest that has led to consistently optimistic predictions of future cost and performance. However, HgCdTe FPAs still are prohibitively expensive for two major reasons:

1. The hybrid (non-monolithic) nature of the FPA sensor chip adds cost and reduces yield.
2. The photovoltaic HgCdTe material necessary for FPAs has not yet been developed to its cost and performance potential.

Although the older photoconductive single-element HgCdTe detectors are now manufactured with only occasional production problems, photovoltaic HgCdTe has poor noise performance and does not respond across the entire 8-12 micrometer band. However, many military applications prefer operation at 8-12 micrometers, and HgCdTe remains the most promising material for this spectral region.

RCA's monolithic Schottky barrier technology can also operate in several IR regions. Figure 4 shows the cutoff wavelengths and operating temperature requirements for Schottky barrier detectors formed with various metals. While iridium silicide shows some promise of responding in the 8-12 micrometer region, not enough work has been done with this material to show overall feasibility.

RCA has concentrated on palladium silicide (PdSi) for the SWIR region,[9] and platinum silicide (PtSi) for the 3-5 micrometer band.[8] Figure 5 plots the

discrete detectors. Early IR imagers scanned a single small detector over the entire scene in a raster to form an image. This is analogous to the scanned, single electron beam that "reads" a television vidicon. However, each small area of the vidicon constantly collects scene energy, while the single-scanning detector sees each small element of the scene for a very short time.

When we deal with discrete detector elements, we can improve them proportionally to the square root of the number of detectors used to observe a given scene. (Each detector increases its elemental-dwell time as the overall number of detectors increases.)

The ultimate IR imager would be the discrete detector equivalent of a vidicon. Many small detector elements, in a two-dimensional X-Y matrix, are electronically multiplexed to form the image from a single serial-data stream. This eliminates mechanical scanning and excessive numbers of signal channels for handling the individual detectors. Many attempts are being made to achieve such an IR sensor, while also making it reliable and affordable.

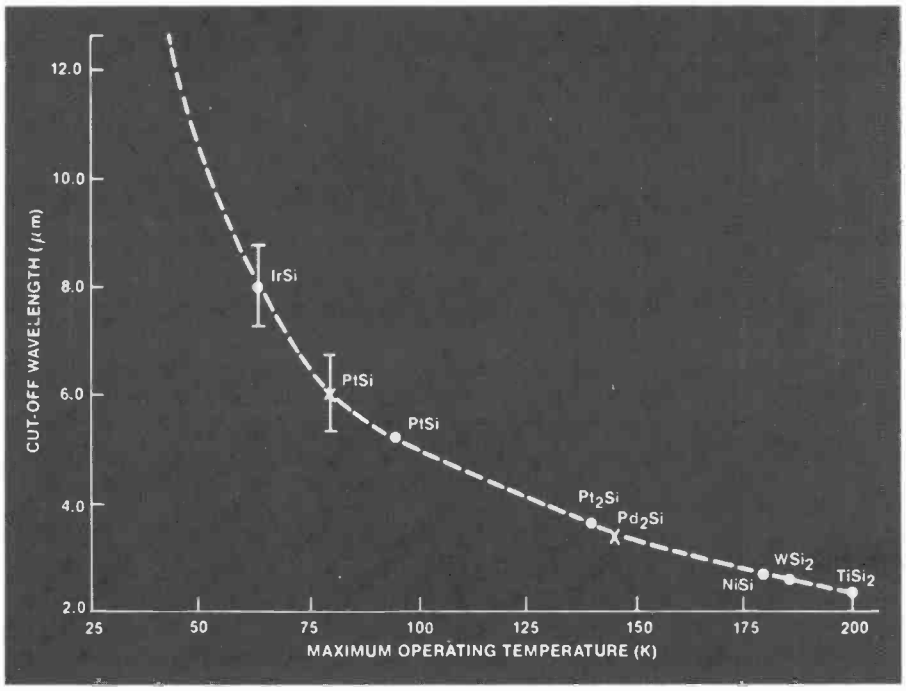

Fig. 4. *Schottky barrier detector cutoff wavelengths and operating temperature requirements.*

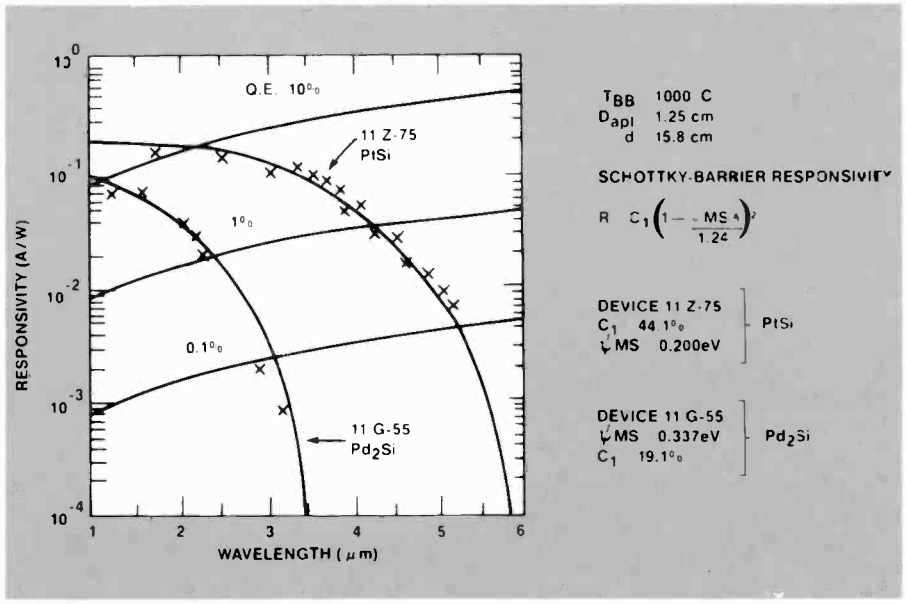

Fig. 5. *Responsivity of typical palladium and platinum devices.*

measured responsivity (amps per watt of incident scene radiation) versus wavelength for typical palladium and platinum devices. The graph also shows quantum efficiency (output electrons per input photon). Although platinum's performance is superior to palladium's at all wavelengths, palladium is sometimes preferred for the SWIR band because of its relaxed cooling requirements.

However, the responsivity and quantum efficiency shown in Fig. 5 do not entirely define detector performance. We must also consider detector noise, because the signal-to-noise ratio that we can obtain for a particular scene object is the real measure of a detector's performance.

Although the quantum efficiency of Schottky barrier detectors is roughly ten times less than that of the hybrid technologies, the Schottky barrier also contributes much less detector noise.

The most common parameter for evaluating an IR detector's signal-to-noise performance is specific detectivity, or D*, which is roughly equivalent to responsivity divided by internal detector noise. It is a figure of merit; high D* indicates good performance.

Figure 6 shows D* for Schottky barrier platinum silicide (PtSi) and indium antimonide (InSb) as a function of spectral wavelength. InSb, a classical hybrid-detector material used for FPAs that operate in the 3-5 micrometer region, is a better behaved material than HgCdTe and has good credibility as a 3-5 micrometer sensor. However, Fig. 6 indicates that PtSi is competitive with InSb.

RCA has fabricated monolithic FPAs that have more resolution than has been achieved with any other technology used for commercial or tactical military applications. Figure 7 displays a chronology of this development, culminating in the

present 160 × 244-pixel IRCCD.[8]

At this point, we should summarize the advantages that IRCCD technology offers.

- Performance—The IRCCD's 99.5-percent rms element-to-element response uniformity is unique. It contributes to an overall system performance that is competitive with, and sometimes superior to, other detector technologies.
- Producibility—The monolithic silicon IRCCD is the only existing IR FPA that offers real production economy. We can produce it on existing visible-CCD production lines. Some hybrid technologies could be produced in moderate quantities, but at ten to twenty times the cost. Hybrid technologies that operate in the 8-12 micrometer region cannot honestly be considered producible in high quantity (50,000 per year), or even in moderate quantities.
- Cost—The fundamental production cost of an IRCCD is not appreciably different from that of a visible CCD imager. The cost may even be less because military applications allow electronic-blemish correction and image compensation beyond what is usually feasible in a commercial camera. We incur a significant portion of production cost by testing to military specifications. Equivalent developmental hybrid FPAs are 100 times more expensive than the developmental IRCCDs that RCA is producing now.

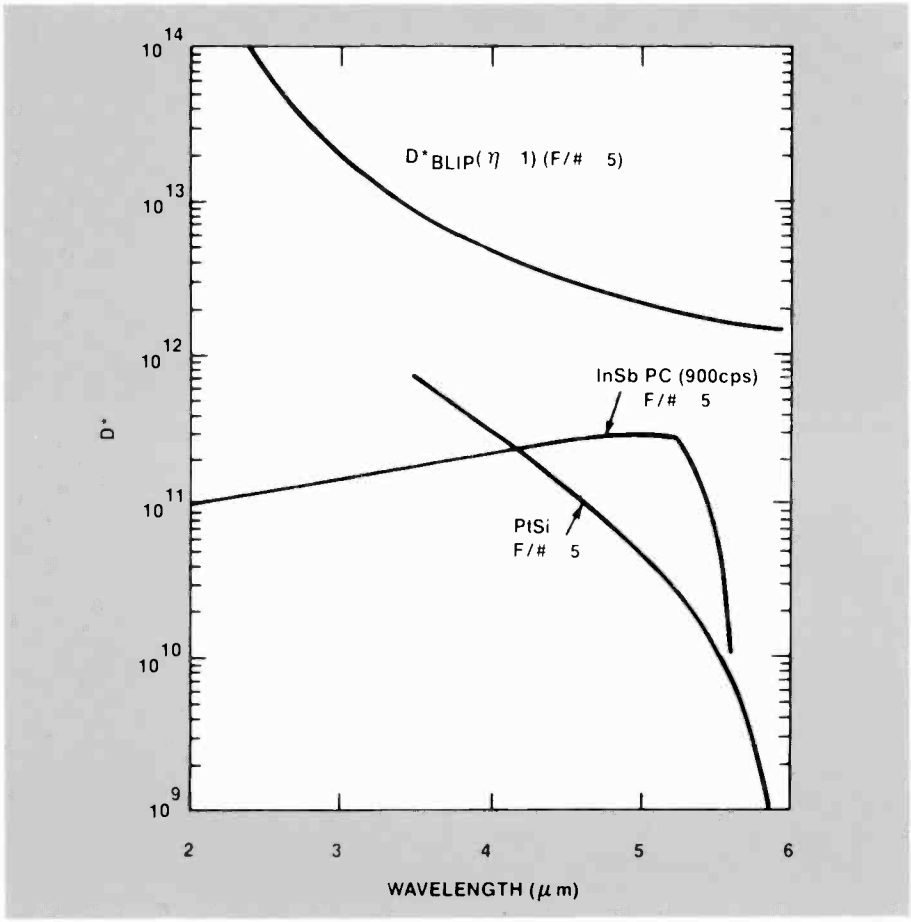

Fig. 6. *D* comparison between PtSi SBD and InSb detectors.*

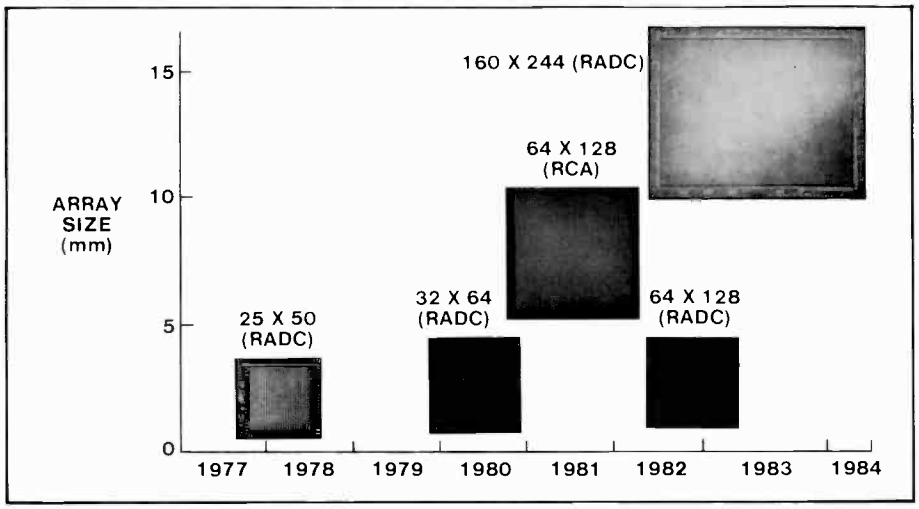

Fig. 7. *History of the development of Schottky barrier focal plane arrays.*

Although the government's prior technology commitment and resulting aerospace marketing thrust have favored the HgCdTe FPA, the government infrared business is a large area, and some significant niches are available to RCA.

RCA imager business development

RCA has a long history of developing imaging sensors, with emphasis on television technologies. The visible and infrared CCD devices encourage similiar development for future applications.

Table I outlines RCA's present contract involvement in imager technologies and describes Government and aerospace programs where ATL has a prime or major supporting role. The total value of the efforts represented in the table is approximately 7 million dollars. The table does not include programs that RCA's Astro-Electronics and Automated Systems Divisions and RCA Laboratories have pursued independently.

Recently, major ATL proposals were submitted to the Navy and other government agencies for extensive contracts based on IRCCD imagers. The general market for IRCCD and CCD imagers includes missile seekers, night-vision sensors, aircraft-tail warners (for approach of tactical anti-aircraft missiles), earth-resources sensors, strategic earth-observation sensors (especially for Strategic Defense Initiative pro-

grams), and night sights (for tactical anti-armor missiles).

RCA can anticipate more government and aerospace contracts than those listed in Table I, which will continue the evolution of an interesting and useful technology.

Table I. *Major ATL program*

PROGRAM	CUSTOMER	DESCRIPTION
Tankbreaker	*ASD/DARPA*	Applies area-array IRCCD to imaging military targets (tanks) for autonomous detection and guidance of man-portable anti-tank missile.
VIS/NIR	*NASA*	Use multispectral visible-CCD line imagers for earth-resources surveillance (LANDSAT).
SWIR	*NASA*	Use multispectral IRCCD line imagers for earth-resources surveillance (LANDSAT) in the SWIR spectral region.
IRCCD cameras	*various (aerospace prime contractors)*	Deliver 160x244 pixel IRCCD cameras and 64x128 pixel IRCCD sensors for testbed evaluation systems.
Shuttle camera	*Astro/NASA*	Deploy a 160x244 pixel IRCCD camera for a December 1985 shuttle experiment.
Laser beam control	*Lockheed*	Develop IRCCD sensing of a large space-based laser beam profile for feedback laser control.
FOG-M (fiber optic guided missile)	*Army (MICOM)*	Develop an IRCCD missile-seeker imager for target detection and guidance via a fiber-optic link.
Staring measurements sensor	*Navy (NRL)*	Assess the infrared clutter background for future IRCCD application in air defense.

References

1. J.R. Tower, et al., "Performance of 4 × 5120 Element Visible and 2 × 2560 Element Shortwave Infrared Multispectral Focal Planes," SPIE Conf., San Diego Ca., August 1985.
2. E.D. Savoye et al., "High Sensitivity CCD Imagers for Television," *SPIE* Vol. 501, pp. 32-39, 1984.
3. H. Elabd, A.D. Cope, F.V. Shallcross, W.F. Kosonocky, P.A. Levine, and D.M Hoffman, "4 × 1024 CCD Array for Remote Sensing Applications," *Journal of Imaging Technology*, Vol. II, No. 6, October 1985, pp. 198-204.
4. W.M. Kramer and D.M. Hoffman, "Design and Manufacture of a Filter Module for a Multispectral Detector Array," Electronic Imaging Conf., Boston, Mass., Sept. 1985.
5. F.D. Shepherd and A.C. Yang, "Silicon Schottky Retinas for Infrared Imaging," 1973 Int. Electron Devices Meeting, *Tech. Dig.*, pp. 310-313.
6. F.D. Shepherd, "Recent Advances in Platinum Silicide Infrared Focal Plane Arrays," 1984 IEDM, San Francisco, Ca., December 9-12, 1984, pp. 370-372.
7. H. Elabd and W.F. Kosonocky, "Theory and Measurements of Photoresponse of Thin Film Pd_2Si and PtSi Schottky-Barrier Detectors with Optical Cavity," *RCA Review*, Vol. 43, December 1982, pp. 569-589.
8. W.F. Kosonocky, F.V. Shallcross, T.S. Villani, and J.V. Groppe, "160 × 244 Element PtSi Schottky-Barrier IR-CCD Image Sensor," *IEEE Transactions on Electron Devices*, Vol. ED-32, No. 8, August 1985, pp. 1564-1573.
9. J.R. Tower, L.E. Pellon, B.M. McCarthy, H. Elabd, A.G. Moldovan, W.F. Kosonocky, J.E. Kalshoven, D. Tom, "Shortwave Infrared 512 × 2 Line Sensor for Earth Resources Applications," *IEEE Transactions on Electron Devices*, Vol. ED-32, No. 8, August 1985, pp. 1574-1583.

Frank Warren is Manager of ATL's Device Technology Laboratory. He is responsible for sensor system and infrared camera technology development, and solid-state laser system development. He has managed programs to develop visible and infrared sensor arrays for spaceborne applications, EO surveillance techniques for MX missile system physical security, and an IRCCD intrusion detection system for perimeter surveillance. He is currently directing the development and production of IRCCD cameras custom designed for a wide range of applications. Frank received a BSE degree from the University of Central Florida. Contact him at:

Advanced Technology Laboratories
Moorestown, N.J.
Tacnet: 253-6501

F.C. Luce | W.B. Schaming

Automatic classification of ship targets

Engineers at ATL have developed a user-friendly, flexible software system for classifying infrared images of ships.

The Advanced Anti-Ship Targeting Development (AATD) program started in 1983 under the auspices of the Naval Weapons Center (NWC), China Lake, California, and is divided into several phases. The current phase calls for developing software that automatically classifies thermal infrared images of ships. The final goal of the project is to produce "flyable" hardware including an infrared sensor that will automatically detect and classify a target and make several tactical decisions in flight.

Abstract: *Through IR&D projects and contracts with the Naval Weapons Center, the Image Technology Laboratory has developed a software system for automatic target classification of ship imagery. The system includes modules for target detection, segmentation, feature extraction, classification, and report/display products. Developed from a large infrared image database, the system can receive and process an image (or images) and produce classification results and associated confidence levels along with other optional outputs (e.g., total classification results and image display products). This article documents the development of the Advanced Anti-Ship Targeting Development (AATD) software and summarizes its current capabilities and the results achieved to date.*

Final Manuscript received November 20, 1985
Reprint RE-31-1-3

Background

Over the past two years, we have developed software for the current phase based on infrared image data that NWC supplied. The data were originally collected on videotape on-board an aircraft and later digitized by NWC and an independent contractor. During this phase, the procedure has been to receive one set of data for developing software and training a classifier, followed several months later by a set of testing data. We used the testing data to evaluate the development to date, as well as reveal problems with the software (or the data). The Naval Weapons Center is currently comparing various systems developed by several contractors.

The data from each successive evaluation was concatenated to form an image database. At the conclusion of this phase, the database contained over 10,000 images, consisting of 10 subclasses of ships at various distances and aspect angles. Each image of the database has a corresponding ground-truth record. The ground-truth information includes the correct ship classification for each image, the distance from the sensor to the target (the range), the aspect angle of the sensor with respect to the bow of the ship, the relative intensity of the target, and a sensor flag that identifies the particular sensor used in data collection.

The final delivered software system has a user-friendly interface that allows easy access and analysis. The interface prompts the user for the information required to run each module of the software and has logical default values. Although the system is designed to execute a standard algorithm set, it is also flexible enough to allow various processing options and combinations of data and classifier designs.

The AATD software

Although the system's modules were developed separately, they were integrated to run under a main driver routine. The detection module has two algorithm options, one based on histogram thresholding of the image data and the other on multifeature histograms and Bayesian decision rules. The segmentation module uses the detection output and the multifeature Bayesian approach to segment the target into a silhouette that is passed to the feature extraction routine. The feature extractor collects various features from the silhouette including geometry, texture, entropy, and height above water. The features for a given target are passed to the classification module that contains a statistical tree classifier. The report and display module reports the results for a given target and summarizes the results for a set of targets. As an option, it will display intermediate image products on a display processor.

The input to the software is one or more thermal infrared images that contain potential targets and the answers to a series of questions about the data and processing to be performed (Fig. 1). The image data and the corresponding ground-truth data can be on tape or disk, but must be in the specified format that NWC established. The software

prompts the user for the location of the data (is it on tape or disk), the names of the data and ground-truth files, and the name of a file that contains the classification tree. The remaining prompts are for processing and output options.

Detection—first algorithm

The detection algorithm locates an object in the image that is most likely the target. To do this, it uses a histogram thresholding process. The threshold used is based upon the shape of the histogram, the intensity of the target (is the target black or white) and the target/background contrast. Although selection of the threshold is somewhat adaptive, the adaptivity is limited.

To start the process, the detection module reads the black/white hot flag from the ground-truth label and uses it to control initialization. If the flag has a zero value (white is hot), the threshold parameters are set for the bright end of the histogram. If the flag has a value of one (black is hot), the parameters are set for the dark end of the histogram.

Next, the histogram storage area is cleared and the 8-bit intensity histogram is computed for the image. The intensity I_1 that corresponds to the peak of the histogram is located (see Fig. 2). Intensity I_2 that corresponds to the Nth percentile on the distribution function is also located; in Fig. 2, I_2 is at the 95th percentile. The threshold T is computed as the average of I_1 and I_2, and the input image is then divided into a binary image based on the threshold at intensity T.

If black is hot, all pixels that have an intensity less than T are set to one. If white is hot, all pixels that have an intensity greater than T are set to one. All clusters of pixels in this binary image are located and are then thinned based on the size of the expected target as dictated by the range estimate. (A cluster is a connected set of pixels.)

Next, the detection module evaluates the results, modifies the thresholds, if necessary, and constructs a rectangle around the target region.

If a target is found, the maximum and minimum x and y coordinates that define a rectangle around the target are stored along with the size of the area. If the detection process did not find an object, it sets the area to zero and sets the maximum and minimum x and y coordinates to out-of-bounds values.

```
Would you like the program to run a self test [y/N]: n
What is the test number (i3) : 4
What is the input image data-set ID number (i3) : 45


                DEFAULT OPTIONS
                ---------------

- Detection, Segmentation, and Classification will be
  performed
- Detection Algorithm I will be used
- Segmented output will NOT be generated
- Feature file will NOT be generated
- Classification results NOT displayed on terminal
- Input and Output assumed to be on DISK
- NWCPRG is NOT executed
- DEFTREE.dat is assumed to be the tree file
- The Debug option is NOT specified


Use default options [Y/n] : n

Is Detection and Segmentation required [Y/n] : y
Use Detection Alg. I or Detection Alg. II [ONE/two] (a1) : one
Save segmented output [N/y] : n
Save Feature file [N/y] : n
Is Classification Desired [Y/n] : y
Log Classifier results to terminal [y/N]: y
Is input image on tape or disk [D/t] : d
Output report file on disk or tape [D/t] : d
Is the Display program (NWCPRG) required [N/y] : n
Read in a new classifier tree [N/y] : y
Is Debug option required [N/y] : n


Sequence number of the 1st image to process (i4)[1]: 1
How many images should be processed (i4) [1]: 1
The interval between the images (i4) [1] : 1
Enter name of classifier tree [DEFTREE] (a8): newtree
```

Fig. 1. *Sample session of the AATD user interface. The system's requests also identify acceptable user responses.*

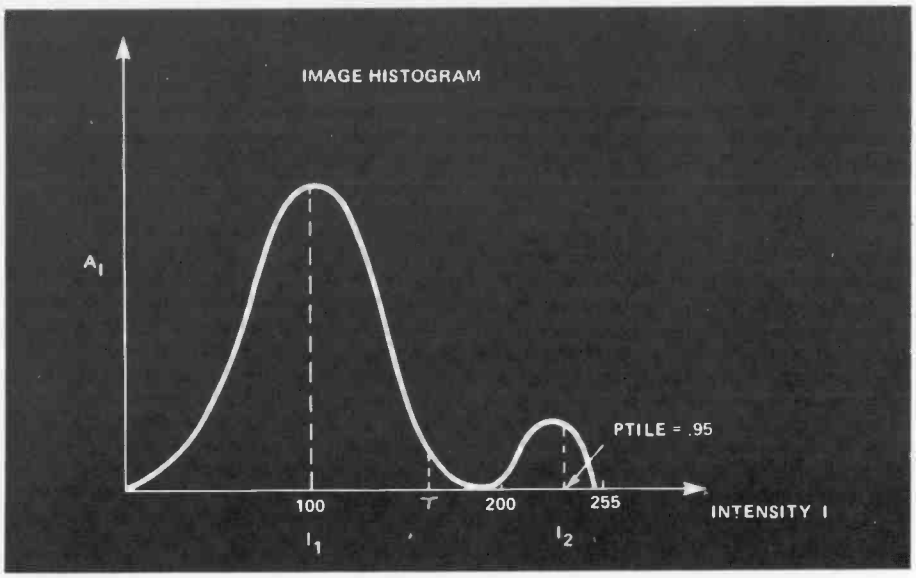

Fig. 2. *Example histogram of a hot target in a colder background. (White is hot.)*

Detection—second algorithm

The second detection algorithm was primarily devised for use with synthetic imagery where the performance of the first algorithm was questionable. Specifically, synthetic images that contained random noise in the background along with a gray target would fail in the

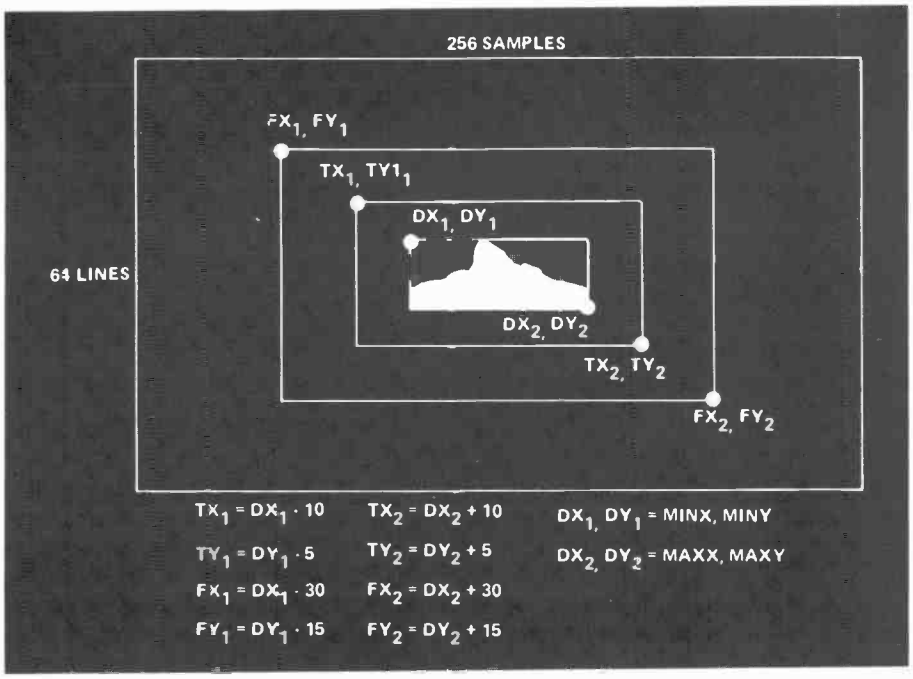

Fig. 3. *Target and frame regions defined for the segmentation process.*

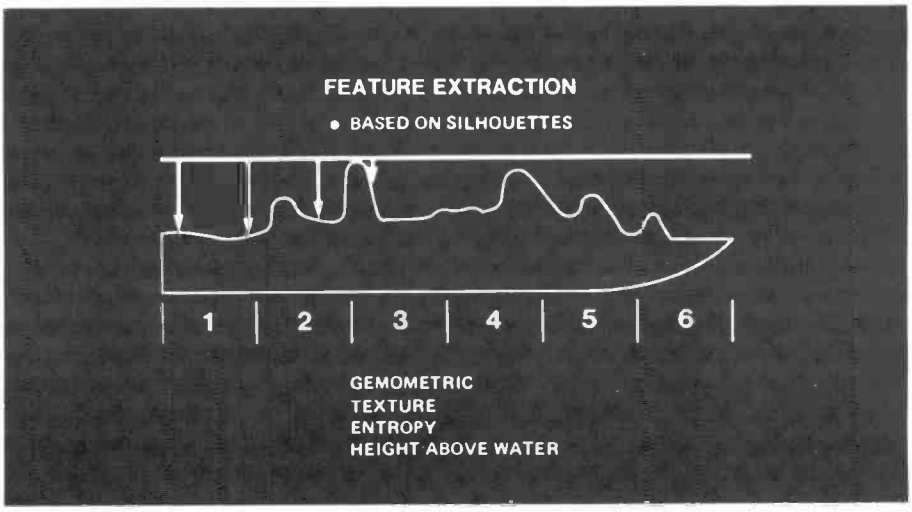

Fig. 4. *Sample silhouette that shows the six equal lengths along the orientation axis in which the classification features are computed.*

histogram-thresholding method. Therefore, we adapted the segmentation algorithm (described in the next section) to perform two-dimensional histogram thresholding, using intensity and edge information to detect targets.

The output from the second detection algorithm is the same as the first. It passes target coordinates to the segmentation routine.

Segmentation

If the area of the target located in the detection process is zero (no object was detected), the system bypasses segmentation processing.

If a target had been found in detection, then the system defines target and frame regions (Fig. 3). To define the target region, it adds ten pixels to both ends of the detected region and five lines to the top and bottom of the region, to allow for inaccuracies in the detection output. To define the outside bounds of a frame region, the system adds 30 pixels to both ends of the detected region and 15 lines to the top and bottom of the detected region.

For segmentation purposes, then, the target lies inside the rectangle defined by coordinates TX_1, TY_1 and (TX_2, TY_2). The area between this rectangle and the frame rectangle defined by the coordinates (FX_1, FY_1) and (FX_2, FY_2) will be used to estimate the statistics of the local background.

The system generates an edge image using the absolute-value approximation in the Sobel-edge operator. The original intensity data and the edge magnitude data are used to generate two-dimensional feature histograms from the pixels contained within the target window and the pixels contained within the frame region. These two joint distributions are used to generate a Bayesian-decision rule that indicates whether a pixel is to be called target or background, depending on the joint occurrence of intensity and edge magnitude associated with the pixel.

Then, the system uses the decision rule to classify all pixels within the target window as either target or background. It sets target pixels to one and background pixels to zero. The target pixels are again grouped by clusters and the largest cluster is preserved as the target. All pixels that are not part of the largest cluster are set to zero. This represents the segmented target used for classification.

Feature extraction

The feature-extraction algorithm computes the set of features that the classifier module requires and uses the segmented target silhouette generated during segmentation. Figure 4 shows an idealized ship silhouette from which features are to be extracted and lists the general types of features that will be computed.

The silhouette is divided into six equal lengths along the orientation axis. Many features are computed individually within these six regions to provide relative-position information about the superstructure. In reality, superstructure features are computed with respect to a mensuration line that passes through the highest point in the silhouette rather than from the bottom of the object. This eliminates some of the discrepancies that may occur because of segmentation irregularities at the water line.

The original set of features evaluated in the classifier-design process included 71 different measures, but this set has been trimmed to 43. The general characteristics of the current feature set are listed in Table I.

Classification

To evaluate the features computed in the feature-extraction module, the classification algorithm passes them through a binary tree classifier. The exact structure of the binary-tree is based on the characteristics of the training dataset.

We developed an example of a binary tree (Fig. 5) by analyzing a training image-data set we received in the early stages of the program. At each node in the tree, a subset of the computed features will be used to decide whether to proceed down the left or right branch from the node being examined. To determine the decision rule at each node (as well as the tree structure), we analyzed the training data features by classes or subclasses, and found a model or function that will differentiate between two groups.

The function of this algorithm and its input/output requirements are well defined, even though the exact tree structure is not. The only required input is the feature set that is passed from the feature extraction program, and the output is the set of decision levels required to generate a report file.

To create the tree structure, we used a classification algorithm that yields a binary-decision rule. The first decision, or node, is selected by using available training data, separating the data into two groups using a classifier, finding the "best" binary split among the classes (or combinations of classes), and then applying the decision functions. These two groups are then analyzed for the next "best" split among the classes, and each is divided by the decision function. The process continues until terminal node conditions exist (we reach the bottom of the tree).

We selected linear discriminant analysis (LDA) as the classification scheme for this project because it is easy to implement and has relatively high performance. The goal in LDA is to find classification functions (linear combinations of the original variables) that best characterize the differences between groups. In this case, the groups are subclasses or combinations of subclasses. We can structure other classification schemes—such as minimum distance, nearest neighbor, quadratic, and Bayes—to produce similar binary trees that can be readily adapted for use in the classification algorithm.

To improve classification performance, we concentrated on the imagery

Table I. *General characteristics of classification features*

Feature number	Description
1—8	Features derived from length, area, height, perimeter
10—17	Eight directional probabilities
24—26	h7, h8, h9 = (h1+h6/2), (h2+h5/2), (h3+h4/2) respectively Six average height above water features, h1—h6 (The ship is divided equally into 6 vertical sections along the length of the silhouette)
27, 28	(h7/h8), (h8/h9)
30, 33, 34	Three superstructure entropy features
42—44	Texture features, analogous to Features 24—26
45, 46	Texture features, analogous to Features 27, 28
48, 50	Divide Features 42, 44 by total superstructure texture (To make scale and rotation invariant)
58, 59	Texture features: analogous to Features 43, 44 but are computed with different formula
62	Linear combinations of two texture features (Feature 57/Feature 59) Note: Feature 57 is analogous to Feature 42 but computed with different formula
65	Texture: analogous to Feature 50 but computed with different formula
66—68	Silhouette moments M10, M01, M11
69	Total superstructure entropy
71	Total superstructure textures for Features 51—56

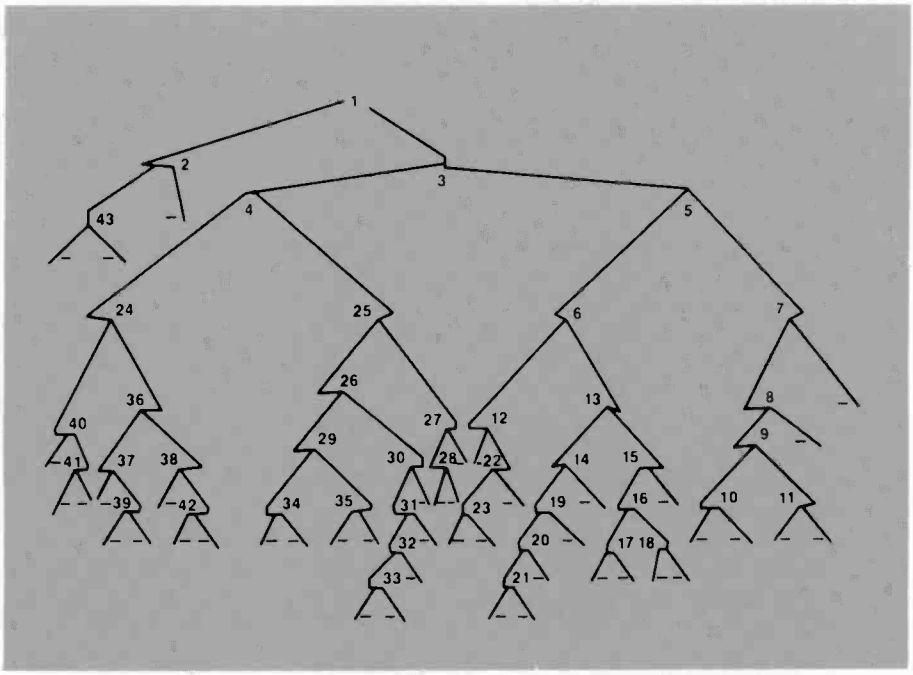

Fig. 5. *Sample decision-tree structure developed during the early stages of the AATD program.*

```
-------------------------------------------------------------
  IMAGE ID      1.750        SC = 205        RANGE      52 Kilofeet
-------------------------------------------------------------

  ASPECT INDICATOR      1

  DETECTION (1) / NO DETECTION (2) :    1   CONFIDENCE       93.78

  TARGET (1)     / NO TARGET (2)     :    1   CONFIDENCE      100.00
  FOE (1)        / FRIEND (2)        :    2   CONFIDENCE      100.00
  COMBATANT (1) / NON-COMBATANT (2):    1   CONFIDENCE      100.00


  1st CHOICE CLASS          2               CONFIDENCE       91.34
  2nd CHOICE CLASS          3               CONFIDENCE        6.69
  3rd CHOICE CLASS          4               CONFIDENCE        1.97

  1st CHOICE SUB-CLASS    205               CONFIDENCE       90.94
  2nd CHOICE SUB-CLASS    301               CONFIDENCE        6.69
  3rd CHOICE SUB-CLASS    401               CONFIDENCE        1.97

-------------------------------------------------------------
```

Fig. 6. *Sample output of the AATD software for a single image.*

that contained side-views of the target but refused to classify head-on targets. Therefore, we added one node to the top of the tree for determining favorable and unfavorable aspects. We divided the training data into two classes (head-on and side-view) and again used LDA to find the best classification function to separate the two groups.

Likewise, because many classification errors were caused by poor segmentation, we added nodes near the top of the tree to screen-out the poorly segmented samples. All segmented training data was viewed and labeled as good segmentation, over-segmented, or under-segmented. Again, we used LDA to find discriminating functions to seperate the poor segmentations from the good. The poorly segmented samples were classified as refusals, and the good samples continued down the tree to be classified into a class and subclass.

Report and display

The report module outputs the results for each image as it is classified and generates output files for calculating final classification matrices. Figure 6 is an example of the output for one image. The output contains various levels of decision and their associated confidence levels.

The display routine is simply an interface between the AATD software and a display-processing system, but the software is generic enough to be compatible with many systems. The detection, segmentation, and original images can all be displayed.

Performance assessment

The NWC held four classification evaluations and one segmentation evaluation to find deficiencies in algorithms and their development and to compare results from various contractors. Each evaluation consisted of a training data set that was sent for training a classifier, and a test data set of unknown images for testing the software. The results of the classification evaluations accuracy over all classes ranged from 63 to 90 percent total. Each classifier was self-tested and, as expected, performed very well. When testing the classifier on the data, the results always dropped, as was also expected. However, the magnitude of the drop in overall percent correct indicated problems in the algorithms, or possibly in the image data or corresponding ground-truth label information. Table II, a sample classification matrix, shows a self-test on the most recent classifier design. The classification is made to the ship subclass level, and the subclasses varied from 84.6 to 95.8 percent correct.

In the segmentation tests, our two detection algorithms performed as anticipated. The first algorithm did well on real imagery and produced good segmented output, but it performed poorly

Table II. *Self-test classification results.*

```
     4 Images have unfavorable segmentations
     0 Images have unfavorable aspects (i.e. head-on)
     0 Images have zero area (i.e. not detected)
     0 Images are false alarms (i.e. area less than 25)

Classification Matrix:
----------------------
```

Group	Percent Correct	Number of samples classified into group →									
		SC201	SC202	SC205	SC301	SC302	SC303	SC401	SC402	SC403	Total
SC201	90.91	190	0	4	5	0	0	6	0	4	209
SC202	0.00	0	0	0	0	0	0	0	0	0	0
SC205	92.99	4	0	292	12	0	0	1	0	5	314
SC301	90.54	3	0	23	287	0	0	3	0	1	317
SC302	92.72	1	0	3	0	191	0	2	0	9	206
SC303	0.00	0	0	0	0	0	0	0	0	0	0
SC401	84.55	6	0	5	2	3	0	104	0	3	123
SC402	0.00	0	0	0	0	0	0	0	0	0	0
SC403	95.76	4	0	2	1	0	0	7	0	316	330
Total	92.06	208	0	329	307	194	0	123	0	338	1499

```
1380 samples are correctly classified from 1499 samples
Total classification percentages =     92.0613708
```

on some synthetic images that contained random noise as a background. The second algorithm did reasonably well with both the real and the synthetic imagery and had good segmented output images. The only drawbacks with the second algorithm were its slow computation time to detect a target and its assumption that the targets were queued and near the center of the image.

The software system was also tested for robustness to range-estimate errors and to determine unfavorable aspects. The results indicated that severe errors in the range estimates could be the cause of large drops in classification percentages. The problems with favorable aspect determination are more manageable. If we use LDA and the current feature set, a single node in the classification tree can determine favorable or unfavorable aspects with a high degree of accuracy. Targets that have unfavorable aspects can then be called refusals until a favorable aspect is obtained.

Summary

A software system now exists that will automatically classify ship targets from thermal infrared image data. The system has been tested and produced favorable results. It has a user-friendly interface to allow easy processing under a variety of parameter combinations. This will allow analysis with various data classes and subclasses, and provides the ability to introduce new classification schemes without software revisions. If software revisions are deemed necessary, the system's modules adapt to change without major software revisions.

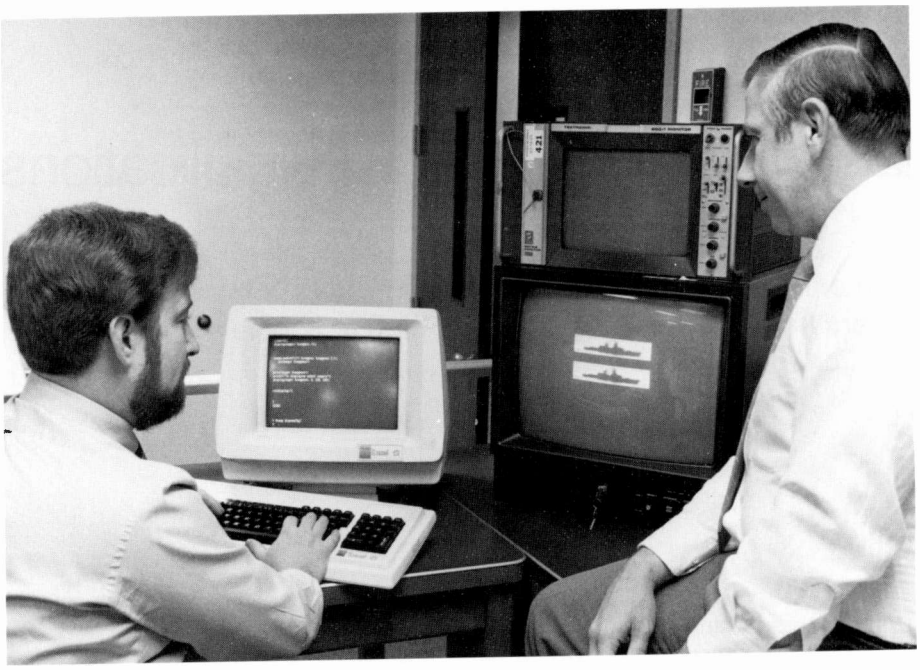

Frederick Luce was a Senior Member of the Engineering Staff in ATL's Image Technology Laboratory. He was a lead engineer on the Ship Target Classification program, responsible for design, development, testing, and implementation of the software system. He has also performed analysis of the entire system and its corresponding inputs and outputs. Prior to joining RCA in 1984, he worked on a variety of image processing, remote sensing, and database development projects at the Jet Propulsion Laboratory. Fred has a BS degree in Forest Science, and an MS in Forest Resources and Operations Research from Pennsylvania State University. He is currently at Corning Glass Co.

Bernie Schaming is Unit Manager of the Target Recognition Systems group in ATL's Image Technology Laboratory. His responsibilities include the development of automatic target recognition technology. He was the program manager for the Advanced Antiship Targeting Development Program for the Naval Weapons Center. He received a BSEE degree from Gannon College and an MSEE degree from the University of Pennsylvania.
Contact him at:
Advanced Technology Laboratories
Moorestown, N.J.
Tacnet: 253-6465

Acknowledgment

The authors would like to thank Richard Gould of the Naval Weapons Center, China Lake, California, for his ongoing support of the project and his assistance and ideas that improved our software system. In addition, we would like to thank ATL engineers Marybeth Kenig, for her professional software support, and Lou Gardiner, for his assistance in data analysis.

D. Wille | A. Rosen

Laser development and applications

The efficiency and high reliability of diode-pumped solid-state lasers has spurred interest in their use in a variety of applications.

Abstract: *Advanced Technology Laboratories (ATL) is currently involved in the development and application of lasers in conjunction with RCA Laboratories and Astro-Electronics Division. Future efforts will expand the area of interest to include other divisions. The current laser activity, which started in 1983, has grown significantly since then. The initial thrust of this activity has been in the area of diode-pumped solid-state lasers for space applications, but interest has expanded to include other areas, such as the application of lasers to the control, generation, and processing of radar signals.*

Final Manuscript received November 20, 1985
Reprint RE-31-1-4

Important applications for lasers are found in communications. One area that is being addressed is the Navy program for communications with submarines from satellite platforms (SLC-SAT). A second prominent application is LIDAR (Light Detection and Ranging), which uses space-based lasers to monitor the state and constituents of the atmosphere. This includes such measurements as wind velocity, cloud heights, and altitude of atmospheric constituents. For both of these applications, two important requirements must be met: the laser system must be long-lived (10 years is desirable), because repair is not usually feasible if the system fails; and the system must be power-efficient. The latter is important because the increased requirements for prime power means increased power generating capability. This, in turn, will increase the launch weight, and will mean a weight restriction someplace else.

Achieving high-power efficiency is an important goal, because solar panels must provide the prime power on a satellite. Also, reaching the efficiency goal can be the difference between making the mission feasible or not.

Diode-pumped solid-state lasers

The usual practice with solid-state lasers is to use a flash lamp to pump the laser material, but flash lamps are both inefficient and unreliable. We need another approach, such as the use of laser diodes as the pump source. For these applications, the important laser materials are Nd-doped YAG (yttrium aluminum garnet) or Nd-doped glass. With these materials, lasing action occurs with a rod or a slab of the material in an optical resonator formed by two mirrors. Light in the resonator travels back and forth between the mirrors and interacts with the laser material. The light either grows or decreases in intensity, depending on the relative number of atoms in the two energy levels involved in the laser transition. The pumping mechanism ensures that there are more atoms in the higher energy state than in the lower energy state, so gain occurs.

A typical resonator is shown in Fig. 1. It is a flash-lamp pumping configuration where both the lamp and the laser rod are at the foci of an elliptic cavity that has reflecting walls. Thus, all the light that the flash lamp emits will pass through the laser rod and have the best chance of being absorbed by the laser material.

Now, consider the question of efficiency. In Fig. 2, we show the absorption curve for Nd:YAG. As can be seen, absorption occurs only at discrete frequencies. A flash lamp, on the other hand, has a very wide emission curve, and this is the source of the efficiency problem. Most of the energy that the flash lamp emits is at the wrong wavelength, and is therefore wasted.

A laser diode, on the other hand, has a line emission (Fig. 3) that can be tail-

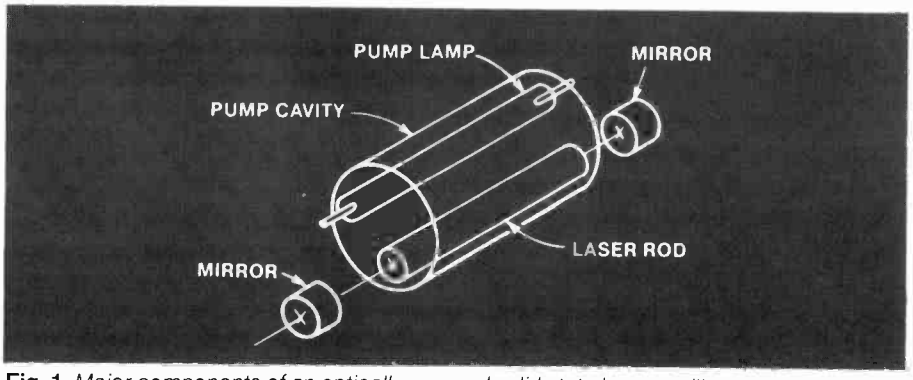

Fig. 1. *Major components of an optically pumped solid-state laser oscillator.*

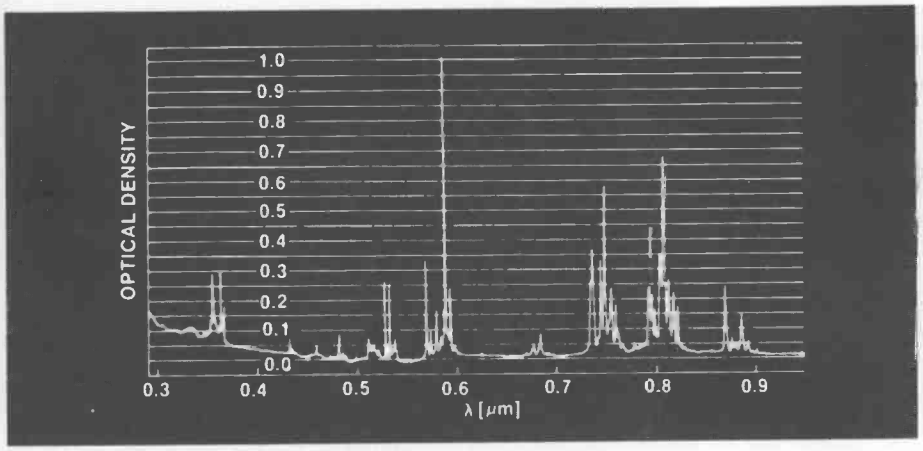

Fig. 2. *Absorption spectrum of Nd:YAG at 300K.*

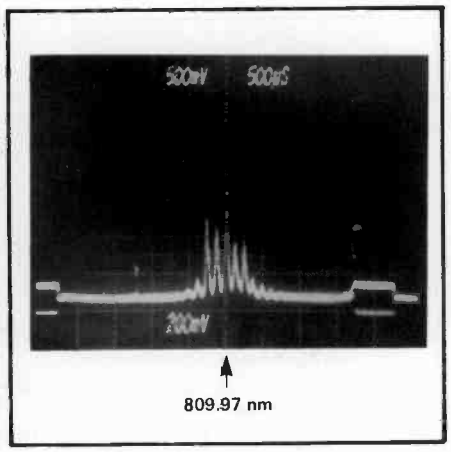

Fig. 3. *Laser diode emission spectrum.*

ored to occur over a range of wavelengths and can be matched exactly to an absorption line of Nd. The wavelength selected is at the strongest absorption line of Nd— 0.8085 microns. The laser diode is itself a semiconductor junction device and uses an electrical pumping mechanism; that is, electrons are injected into the device at the upper energy level and are ready to provide gain. Figure 4 shows the layer structure of the diode. These diodes are referred to as oxide-stripe, double-heterostructure devices, and the layers are grown epitaxially. The diodes have a wide junction (100 microns) and are at a spacing of 300 microns.

For these devices the output power, P, is given by

$$P = \frac{\eta h\nu}{Q}(I - I_{th})$$

where η is the external differential quantum efficiency, $h\nu$ is the photon energy, Q is the electron charge, I is the operating current and I_{th} is the threshold current for lasing to occur. As can be seen from this expression, one of the more important design characteristics for a diode laser is a low threshold current.

Diodes grown at RCA Laboratories put out as much as one watt of peak power, but even with this output we need an array of laser diodes to provide all the power necessary to pump the Nd laser. Figure 5 outlines the array concept that is now being developed.

The basic building block is a bar of six diodes that is cleaved from a wafer. These bars are attached to insulating substrates, and the units are assembled onto a grooved baseplate and soldered into place. Assembling the arrays using single diodes would be much too complicated, and very long bars would be hard to handle and would cause yield problems. The use of six-diode bars is a compromise.

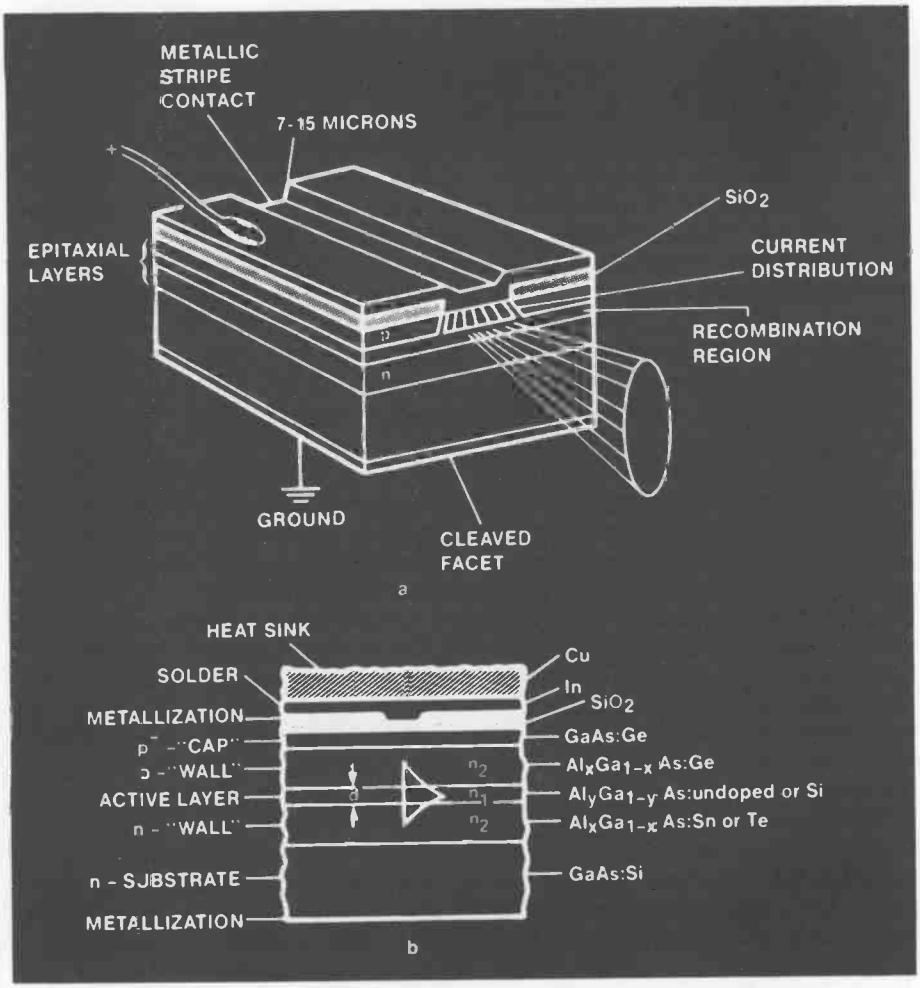

Fig. 4. *Simple continuous-wave lasers. (a) The layer structure of the oxide-stripe double-heterojunction laser; (b) The cross section of a heterojunction aluminum gallium arsenide laser.*

The laser diodes show a power conversion efficiency of greater than 25 percent, and because their emission spectrum overlaps with the Nd absorption, we gain orders of magnitude increase in overall efficiency compared to flash tubes.

A typical Nd laser using a diode pump is shown in Fig. 6. Here, the Nd:YAG slab is pumped by four units of diode arrays, each consisting of five subar-

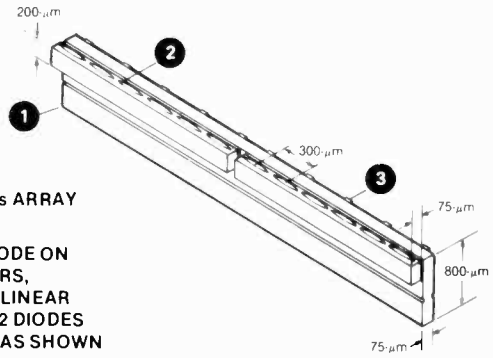

1. BERYILLIUM OXIDE LINEAR GaAs ARRAY SUBSTRATE
2. SINGLE AlGaAs WIDE-STRIPE DIODE ON APPROXIMATELY 300-μm CENTERS, APPROXIMATELY 6 DIODES PER LINEAR ARRAY, AND APPROXIMATELY 12 DIODES PER BERYILLIUM OXIDE WAFER AS SHOWN
3. PLATED GOLD BUMP/STRIPS APPROXIMATELY 25x25x200μm

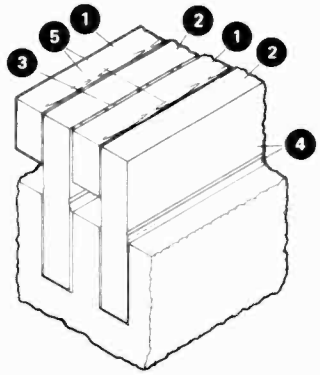

1. AlGaAs LINEAR LASER ARRAY
2. BERYILLIUM OXIDE
3. STRESS RELIEVING PLATED GOLD STRIP BUMPS (VERTICAL) APPROXIMATELY 25μm HIGH ON 150-μm CENTERS
4. BeO METALLIZATION SHOWING ELECTRICAL ISOLATION BREAK
5. TYPICAL DIODE EMITTING REGION

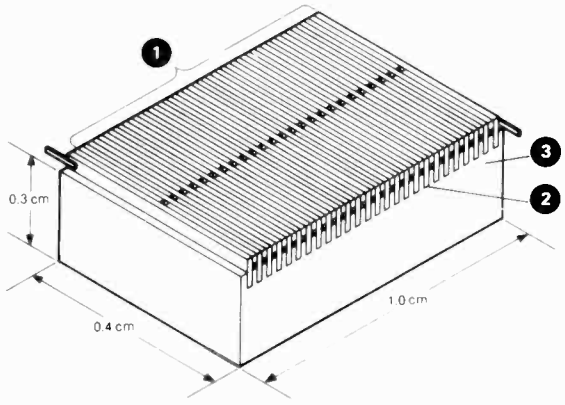

1. Approximately 57 Rows of Linear AlGaAs Diode Arrays on 175-μm Centers (114 Arrays, Total)
2. Heat Sink/Mechanical Mount for 57 Rows of Arrays to Submodule Holder by Mass Solder Joints in Slots in Holder
3. Submodule Holder — Copper or Beryillium Oxide

Fig. 5. *Aluminum-gallium-arsenide (AlGaAs) laser diode array. Schematics represent: (a) the AlGaAs laser diode mounted on beryllium oxide (BeO); (b) the end view of the diode array; and (c) the diode array submodule.*

rays. The subarrays, in turn, consist of units like those in Fig. 5, but each contains 684 diodes. The complete pumping scheme shown here has 13,680 diodes. We have developed techniques for assembling the diode arrays economically. Without such techniques, arrays with this many diodes would be prohibitively expensive.

Normally, we would want a short pulse from the laser, and this can be achieved with a Q-switch device (Fig. 6). The Q-switch is used to spoil the Q of the optical cavity and prevent it from building up the lasing resonance. While this prevents lasing, it allows a greater-than-normal population inversion to build up and, when the Q of the cavity is restored, a very short, high-power pulse is emitted. Typically, a non-Q-switched laser might have a pulse length of tens of microseconds, while a Q-switched laser would have a pulse length of tens of nanoseconds.

The laser diodes we are developing are highly reliable devices, and we expect them to last much longer than the average mission duration. There are various failure mechanisms for laser diodes, but after a burn-in period of approximately 100 hours, our testing has shown that we can expect a very long mean time between failures.

In Fig. 7, we show the results of life tests of laser diodes. As can be seen, after 7×10^3 hours only half the diodes have failed. These tests were part of an accelerated life test where we used a short period of operation at an elevated temperature to extrapolate to the results that would be obtained at room temperature. Diode failure follows a lognormal distribution function and its probability density function is given by

$$f(t) = 1/\sqrt{t^2\sigma^2 2\pi}\ \exp\left[-\frac{(\ln t-\mu)^2}{2\sigma}\right]$$

where $f(t)$ is the probability of failure at time, t. The symbols μ and σ refer to the mean and standard deviation of the related normal distribution, although they are not the mean and standard deviation of the resultant lognormal distribution.

Optical control of microwave devices

A new application area that we are exploring is the use of lasers in conjunction with microwave devices. This field

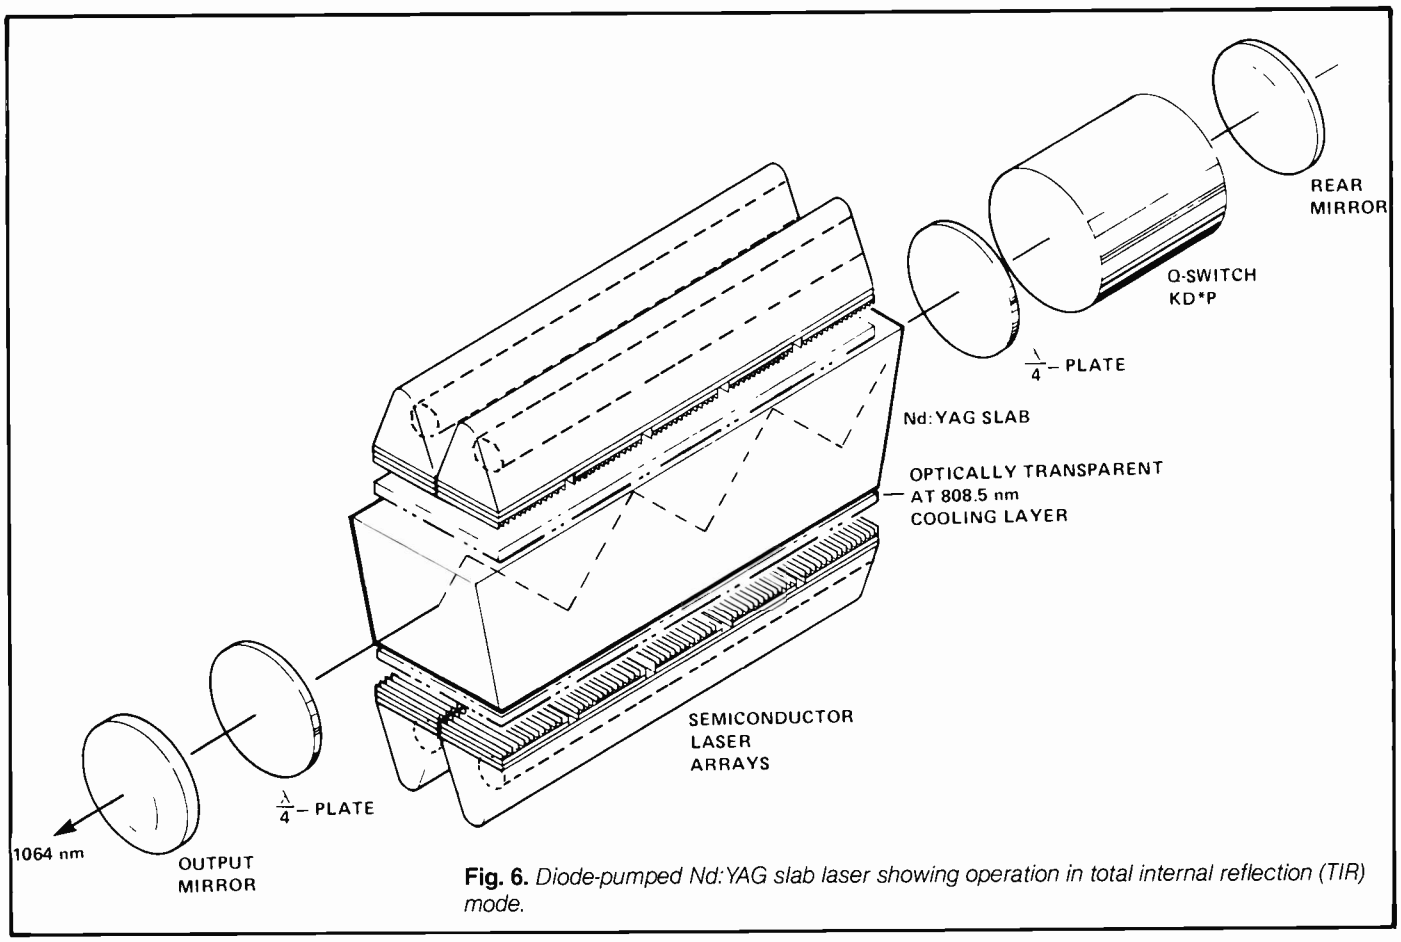

Fig. 6. *Diode-pumped Nd:YAG slab laser showing operation in total internal reflection (TIR) mode.*

is now a collection of separate techniques that, as their primary focus, use lasers to generate, control, or process microwave signals. Among the specific techniques being considered are:

- □ Fiber-optic rf transmission lines
- □ Optical control of phased-array antennas
- □ Optical delay lines
- □ Optically controlled switching of electronic signals
- □ Frequency locking of electronic devices using optical signals
- □ Optical data processing of electronic signals
- □ Optical switch networks for distributing microwave signals
- □ Microwave signal generation

Central to these ideas is the concept of using a laser source and modulating its output at microwave frequencies to send microwave signals on an optical fiber. Fiber-optic communications technology is proceeding in the direction of optical wavelengths in the 1.3- or 1.5-micron range. However, since long path transmission is not desired in this case, we prefer the 0.8-micron region because the technology of sources and detectors is more advanced. Laser diodes are the sources of choice because of their small size, high efficiency, and high output power.

We can modulate the laser diodes directly by varying the drive current, or we can pass the emitted light through an external modulator. Lasers have been directly modulated at frequencies up to 11.6 GHz, and have been externally modulated at frequencies up to 18 GHz. Because lasers can be modulated at microwave frequencies, it is possible to use optical fibers to transmit the microwave signals. With this approach there is significant freedom in choosing the length of the transmission path and complete absence of electromagnetic interference effects.

Today there is great interest in the use of phased-array antennas composed of individual emitters. Arrays are envisioned with thousands of elements, and the problem of controlling the array is formidable. Figure 8 shows how optical sources could be used for array control. The output of a laser diode modulated

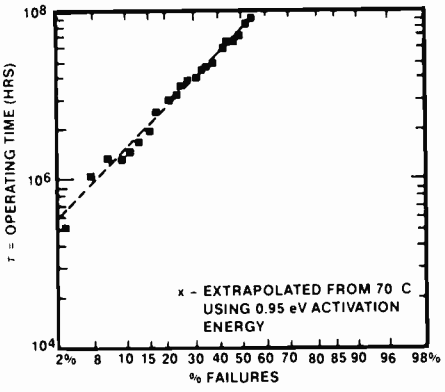

Fig. 7. *Results of life tests on laser diodes.*

at microwave frequency is coupled into a fiber, and an integrated optical device divides that energy into a number of fiber output ports. Each output is then fed to an integrated optical-phase shifter; it can be a switch matrix that simply switches in different lengths of coupling fiber to produce the required phase shift. The signal from the phase shifter is then fed to the radiating element that is turned on and frequency-locked by the incoming optical signal.

To provide optical delay lines, a switch array and a set of fibers with dif-

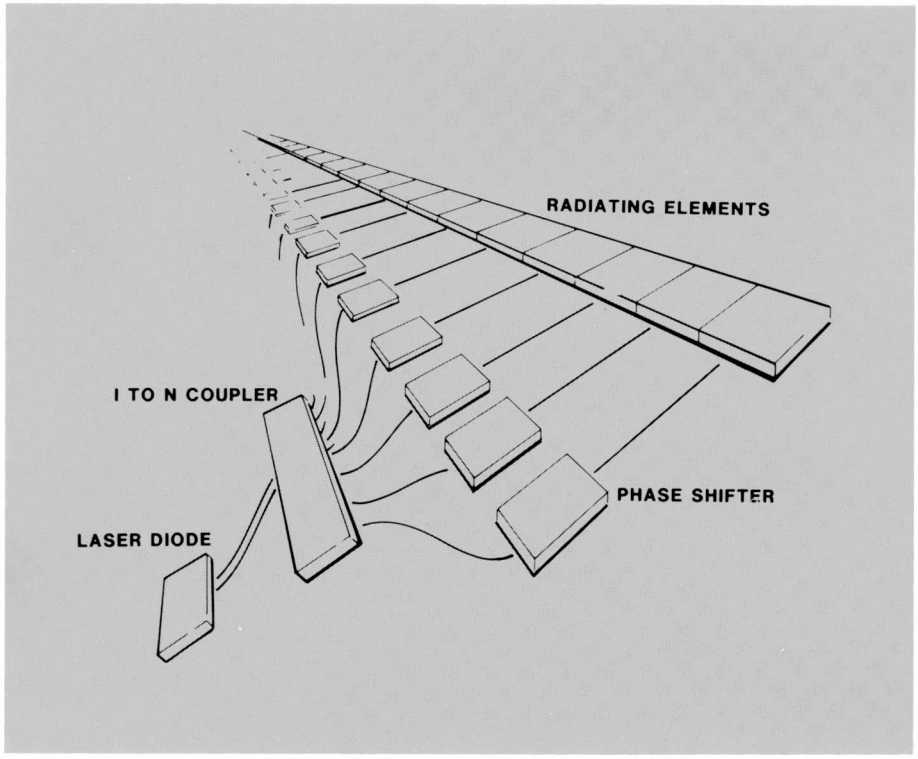

Fig. 8. *Concept for optical control of phased array microwave antenna.*

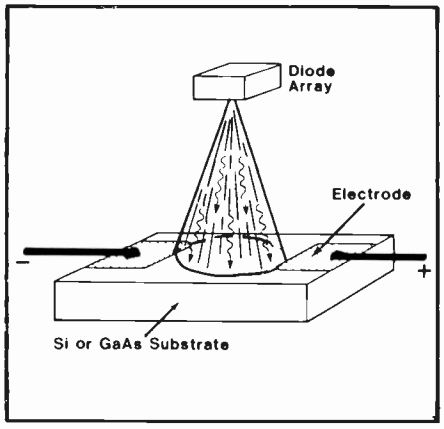

Fig. 9. *Optical switch concept.*

ferent lengths that can be switched in to provide the required delay. Delays of up to tens of nanoseconds are easily achieved, and delays in the microsecond range can be achieved using hybrid techniques.

In a more general sense, we can do data processing of microwave signals using optical techniques; functions, such as signal correlation and spectrum analysis, are prime examples. In addition, numerical processing can also be done optically.

Switching of signals at high speed is a difficult problem that is compounded when the signals are high-voltage and high-power. Optical techniques have significant advantages in speed, size, and reliability. In its simplest embodiment, an optical switch could be merely a slab of semiconductor material that is illuminated by a laser diode (Fig. 9). If we use high-resistivity material, the switch can hold off a high voltage and, when illuminated, is effectively shorted out by the optically generated carriers. With proper design, optically activated switches can withstand hundreds of volts, pass tens of amps, and operate at subnanosecond speeds. If we suitably connect groups of switches, they can operate in the kilovolt-kiloampere range. The same diodes that we use to pump YAG lasers find application in driving switches, but smaller arrays are generally needed, and there are no significant restrictions on the output wavelength.

Summary

Two development areas of laser technology with important practical uses have been discussed. They are currently receiving emphasis at ATL and will be actively pursued in the future. Other areas, such as the application of lasers in fiber-optic communications systems, are being reviewed and will receive emphasis in the future, if warranted.

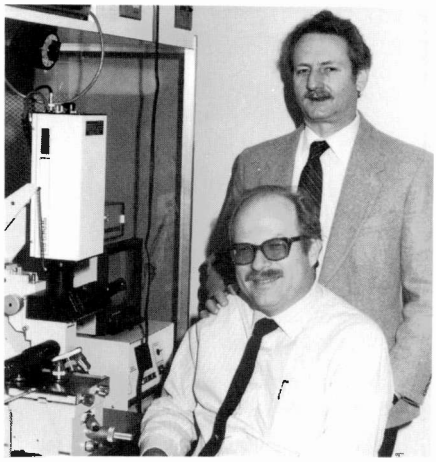

Wille (standing) and Rosen

Douglas Wille is Unit Manager of the Electro-Optics Applications group in ATL's Device Technology Laboratory. His responsibilites include development of laser systems and applications. He is currently managing the PILOT unit-cell definition program, the PILOT 2-D array development, and the Diode-Pumped Laser program. He recently supervised a study of techniques for using lasers for satellite communications with submarines. Prior to joining RCA, he was a Unit Work Leader in the Avionics Laboratory at Wright-Patterson Air Force Base, engaged in development of acousto-optic and electro-optic systems. Dr. Wille received his PhD in Physics from Ohio State University.
Contact him at:
Advanced Technology Laboratories
Moorestown, N.J.
Tacnet: 253-6498

Ayre Rosen received the BSEE cum laude from Howard University in 1963, the MScE from Johns Hopkins University (where he attended on a Gillman Fellowship) in 1965, and the MSc in Physiology from Jefferson Medical College in 1975.

He joined RCA Laboratories in 1967 as a Member of the Technical Staff in the Microwave Technology group. Mr. Rosen is responsible for the research and development of millimeter-wave devices and circuits. His work on integrated diode arrays paved the way for power combining in microwave and millimeter waves. He received a 1972 RCA Laboratories Outstanding Achievement Award for a team effort in the development of S-band TRAPPATT amplifiers and, in 1982, an individual RCA Laboratories Outstanding Achievement Award for developing high-performance silicon PIN diodes for application to communications systems.

Mr. Rosen now holds an appointment of Adjunct Professor at Drexel University, in the Department of Electrical and Computer Engineering. He is the author of more than 60 technical papers and presentations, and holds 21 U.S. patents in the microwave field.
Contact him at:
RCA Laboratories
Princeton, N.J.
Tacnet: 226-2927

P.B. Pierson

Advanced robotics projects at ATL for remote servicing

Robots can help repair satellites in space, and researchers at ATL are getting them ready for the job.

The use of robots in automating factory operations been rising steadily in recent years. Their precision and repeatability allow them to do high-volume repetitive tasks without relying on sensors and to perform at maximum speed and reliability. To achieve this precision in the workspace we use jigs and fixtures. They are costly, but we can amortize the costs over the large number of components the robot processes.

We are also using robots to perform operations in environments that are hazardous or beyond man's reach. For example, undersea robots recently located and retrieved valuable items at great depths. Robots have been used in nuclear-radiation contaminated areas to handle materials and clean up. The Army, Navy, and NASA are actively developing robotic systems to meet specific needs in remote or unsafe environments.

Robotic systems are also being developed to process non-uniform production pieces. With such systems, automation would be economically efficient for small production runs. Researchers are developing robotic systems that can machine just a few parts at low cost. Government agencies are interested in such systems for their potential to manufacture limited-quantity items through automation. The U.S. Postal Service wants robotic mail sorting systems that can handle a wide variety of nonuniform pieces with a minimum of human involvement.

In remote operations and low-quantity or nonuniform production operations we do not have precise data on the location and orientation of work objects. Also, exact motions usually are not repeated. Robot motions may require flexibility and freedom to reach around unanticipated obstacles in the workspace. Robots may need a variety of tools or dexterous "hands" to perform different tasks, but some tools may not be determined until the time of task execution. Sensor systems must provide the detailed information about the location, orientation, and status of work objects, and flexible software systems must command the robot to perform unique operations with minimal human involvement. Currently, these requirements exceed the capabilities of commercial robotics.

Advanced Technology Laboratories (ATL) decided to combine its skills and capabilities to meet a variety of automation and robotics needs in aerospace and defense, and formed the Robotics activity in January 1985. Under contract to the U.S. Postal Service, we are developing sensors capable of collecting three-dimensional information on irregular Parcel-Post items, and separating them using unique machine-vision techniques and special-purpose robot hands. (These programs are described in another article in this issue.)

Our current IR&D effort focuses on developing telerobotic systems and techniques for remote servicing of satellites in space. We are working closely with the Astro-Electronics Division (AED), where RCA is designing satellites and space station subsystems that astronauts or robots may service.

Abstract: *For many years, Advanced Technology Laboratories (ATL) has been involved in many technologies that are integral to advanced robotic systems: advanced sensor systems, television-signal processing systems, precision mechanism and control-system design techniques, machine-vision systems, computer architectures, artificial intelligence software techniques, word recognition systems, and very-large-scale integrated circuit designs capable of real-time signal processing. These are critically needed components for systems capable of servicing satellites in remote orbits in space.*

Final Manuscript received November 20, 1985
Reprint RE-31-1-5

Satellite servicing objectives

Recently, astronauts have successfully repaired several satellites during extravehicular activity (EVA) from the Space Shuttle. They repaired the Solar Maximum Mission by manually replacing the main electronics box of the Coronograph/Polarimeter experiment on Shuttle flight 41C in April 1984. On Shuttle flight 51I in August 1985, astronauts successfully modified faulty circuitry on LEASAT, a communication satellite owned by the Hughes Aircraft Co. They then performed manual spin-up prior to deploying the satellite into its required orbit. These successful repairs validate the concept of space-based servicing. However, EVA is expensive (estimated at about $75,000 per hour) and hazardous to the astronauts. Servicing by robots near the Shuttle or Space Station

would reduce risk to humans, and could reduce costs of servicing operations.

Besides repair, robotic systems could do planned or routine maintenance activities—such as exchanging payloads or orbital replacement units, refueling, and replacing batteries—thereby minimizing risks and operating costs.

Many satellites are in orbits beyond the 275 nautical mile range of the Shuttle. These include many of RCA's communication satellites at geosynchronous earth orbit (GEO), approximately 22,000 nautical miles, and our weather satellites in polar orbits, approximately 350 nautical miles. Servicing at GEO would be difficult for humans, even if the Shuttle could reach them. Because the earth's geostationary altitude lies in the Van Allen radiation belt, manned stations there would have to carry radiation shielding to protect the crew. Robotic servicers operating on site, therefore, offer important alternatives.

We can retrieve satellites at non-Shuttle orbits with an orbital-maneuvering vehical (OMV) and service them at the Shuttle or on the ground. Or, robotic devices attached to OMVs could service satellites on site, with minimum interruption to operating service. Time-consuming orbit change and re-establishment operations would not be required.

Initially, people on the ground, at the Space Station, or on the Shuttle would operate and control robotic servicing systems. (The Space Station is scheduled for assembly in space beginning in 1992.) We will modify techniques developed for undersea work and in areas of nuclear contamination. However, space introduces a problem that land or sea telerobots do not experience—the long communication delays (up to 5 seconds) between the robot and the operator. Velocity or position control techniques involve tedious and risky move and wait operations. Techniques that reflect forces experienced by the robot back to the operator are not suitable if there is any noticeable delay. To increase the capabilites of distant robotic servicing systems, we must develop predictive feedback techniques or give the robot increased autonomy.

Video bandwidth compression

TV cameras provide the primary means for observing and controlling operations of remote robotic systems. Conventional video transmission requires significant amounts of bandwidth, more than is practical for most space-based systems.

ATL is currently helping AED develop a video-processing system that can transmit high-resolution (768 pixels per line), real-time color video at 27 megabits per second. This system is planned for use on the Shuttle in the secure digital television (STV) system that AED is developing for the Air Force and NASA (see Fig. 1).

Video bandwidth compression will provide maximum resolution and frame-rate TV feedback to remote operators of robotic servicing systems in space that transmit over data links ranging from less than 500 kilobits to 27 megabits per second. Our STV design will have a wide range of transmission rates through variable reduction of frame rates and compression ratios between 1 and 2 bits per pixel.

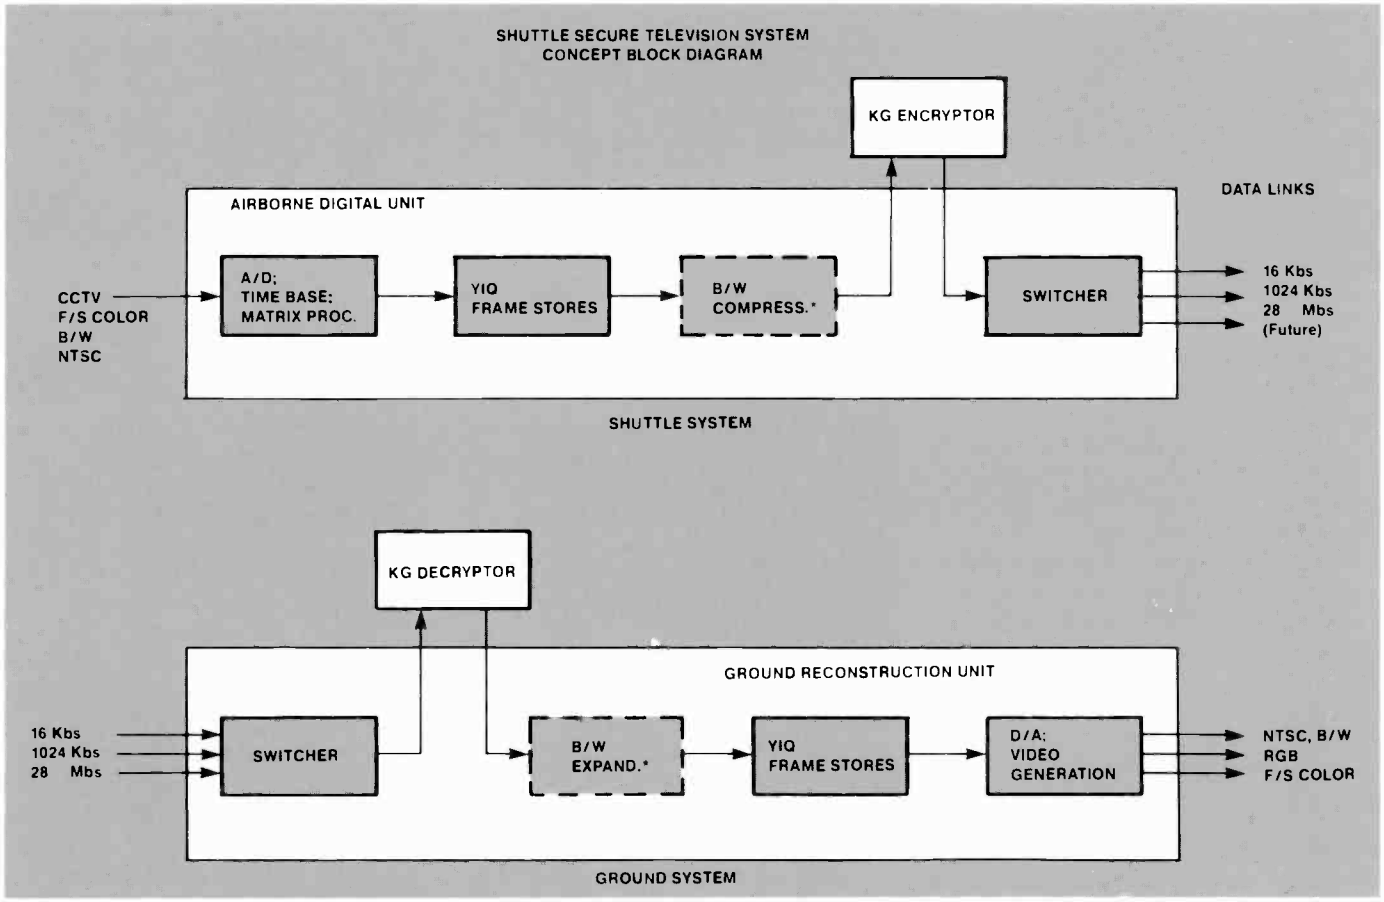

Fig. 1. *Shuttle secure digital television (STV) system.*

Satellite control techniques

We have three primary control techniques for satellite servicing: manual control, automatic control, and supervisory control as shown in the accompanying figure.

Manual control includes extra vehicular activity (EVA), teleoperation, and telepresence. In EVA, astronauts perform tasks directly from within their space suits. Teleoperation is remote manual control from a console where television views and simple sensor readouts assist the operator. Teleoperation commands control either position or force and torque. In telepresence, an advanced form of teleoperation, the teleoperator has the illusion of visual, tactile, auditory, or other sensory information.

Automatic control includes hard automation, flexible automation, adaptive robotics, and autonomous robotics. Hard automation uses preprogrammed sequences to execute a single task, such as module exchange. Flexible automation systems can execute a variety of preprogrammed tasks, with minimal retooling or reprogramming. Adaptive robotics uses sensors to provide fine control of robot motion where exact location of objects or servicer elements is not known. Autonomous robotics (or intelligent robotics) is adaptive robotics in which artificial-intelligence-based reasoning or planning programs develop the detailed control steps. The programs rely on goal-oriented commands from people and respond autonomously to unforeseen conditions and events during a mission.

Supervisory control mixes manual and automatic control. Although a human teleoperator has primary control, flexible automation or adaptive robotics may do segments of the operation. In this mode, planned task segments are performed on cooperatively designed subsystems under automatic control, while the operator handles contingencies. Supervisory control is the recommended technique for remote servicing on the initial operating configuration of the Space Station. As operator confidence and techniques of automatic control develop, more tasks will be performed automatically. When we add intelligence and the ability to resolve contingencies, supervisory control will gradually evolve toward autonomous robotics.

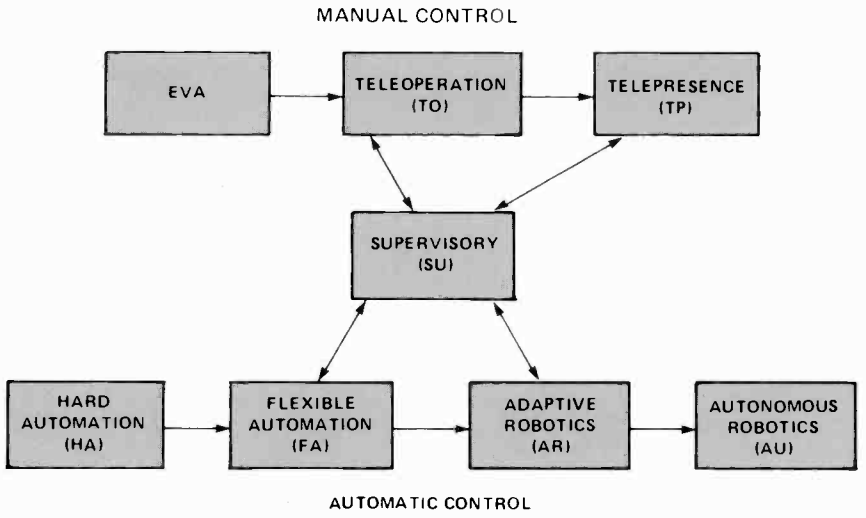

Control techniques for satellite robotic activity.

ATL provided the system architecture for the video-processing section of the STV system and designed the hardware system for the bandwidth-compression modules. It also did computer simulations of various bandwidth compression algorithms and computer hardware emulation of the selected algorithm and the rest of the color processing provided by the STV system. ATL's real-time, two-dimensional, cosine-transform processor occupies only two circuit boards of the STV system's airborne digital unit, and consumes less than 25 watts. It operates on composite Y, I, and Q fields separately, and processes the image in 16×16 blocks of pixels.

Satellite attitude determination

Remote servicing operations on satellites require docking to a specific fixture on the target satellite. Currently, the rendezvous and docking manuevers are done manually, and each docking operation requires about 16 hours. Automation of the rendezvous and docking task is identified as a target for early application on the Space Station because more and more docking manuevers will be performed with the increased number of servicing missions planned for the 1990s and beyond.

To manuever to the proper side of the target satellite we must determine the relative attitudes of the satellite and the approaching vehicle. NASA and aerospace contractors currently are developing various sensing techniques, including various laser-based schemes as well as methods that rely on camera-video processing. Under the direction of AED, we have developed a technique for determining the attitude of the target satellite, which uses camera video as its primary sensor signal. On the image plane of an RCA CCD camera we detect the locations of an array of retroreflective stickers attached to the surface of the target satellite. We search for the optimum match between mathematical projections of a geometric model of the reflector locations and the actual detected reflector positions.

Our test fixture (Fig. 2) supports demonstration of attitude computation over a range of angles of rotation about the x, y, and z axis under a variety of ambient lighting conditions. Current performance is accurate enough—within 2° on each axis—to automate rendezvous and docking to within several feet of contact. Other techniques that offer higher accuracy over close ranges may then provide the final guidance.

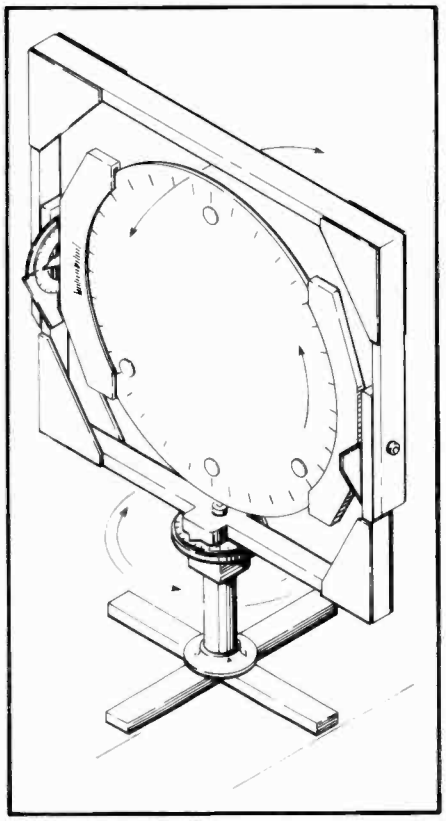

Fig. 2. *Attitude determination test fixture.*

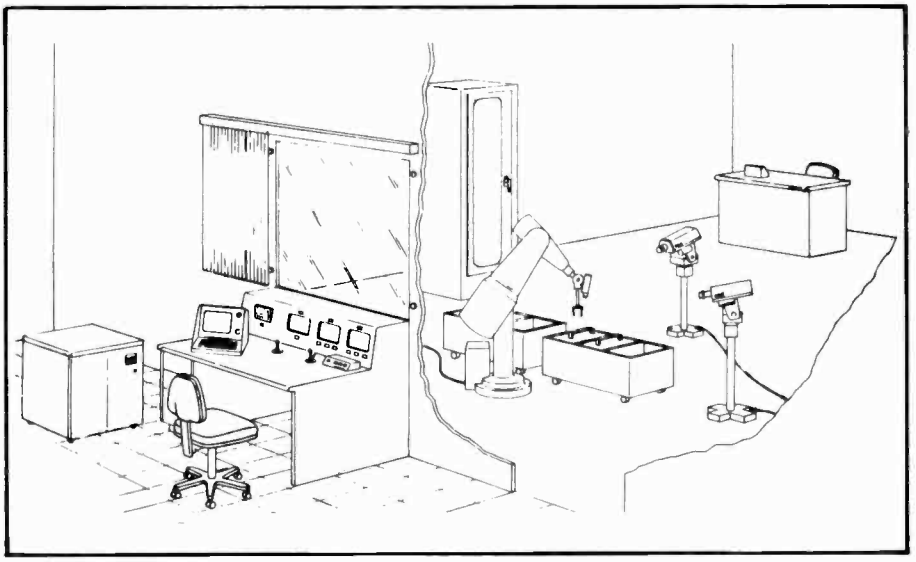

Fig. 3. *Teleoperations and robotics testbed.*

Telerobotic demonstration testbed

ATL's telerobotic demonstration testbed (Fig. 3) provides a physical environment for verifying new control techniques and mechanisms. The control room is next to the robot manipulator room. The curtain on the window may be left open. This simulates operation on the shuttle, where the astronaut controls the remote-manipulator system while looking out the aft windows into the cargo bay. To simulate remote operations, the curtain is closed, and the console provides all cues to the astronaut.

A camera mounted on the wrist of the robot manipulator provides primary feedback to the operator. Other cameras provide orthogonal peripheral views. Other feedback includes computer graphics of robot motions and forces, and synthesized speech. The operator uses joysticks to command robot motions or forces, and voice or keyboard entry for higher level sequence descriptors.

The robot has a set of interchangeable tools: a gripper, nut-runner, point probe, and location tool (for registering the coordinates of objects within the workspace of the robot coodinate system). Figure 4 is a sketch of the gripper, and how it connects to the quick-change wrist.

Prototype mechanisms—designed and built by ATL and AED engineers—include umbilical connectors, fluid connectors, panel covers, orbital replacement units, and complementary gripper systems. They are tested and assessed for remote robotic exchange using the telerobotic testbed.

Natural language voice interface between man and robotic systems

In future space operations, robots will work side-by-side with astronauts to accomplish tasks neither could do alone. Voice control of robot activity will be required for efficient operation, especially in the external environments where astronauts wear space suits.

ATL's Speech Recognition, Artificial Intelligence, and Robotics Laboratories are jointly developing a natural man-machine interface for robotics control. Their long-term objective is an AI natural language interface integrated with AI expert systems software. The robot will be able to respond to sentences spoken by anyone. This will make the human comfortable and allow for a conversation-like dialog between man and machine.

The initial development uses interactive, natural-language, voice recognition and synthesis to control the pan, tilt, zoom, and focus of two perimeter video cameras of the telerobotic testbed.

During execution of remotely controlled robotic tasks, the viewpoint of the camera systems may have to change frequently. Today, operator may control the robot motions, while a second operator controls camera positions. Or, one person could alternately control the camera and the robot. Both scenarios are inefficient. However, continuous robot and camera control is possible by a single operator, who could control the cameras by voice, while simultaneously manipulating the robot-control joysticks. The interface recently developed, Fig. 6, includes a continuous-speech recognizer, speech synthesizer, and natural-language parser integrated with expert-system software. It currently understands 100 words, which limits freedom to some extent, but demonstrates significant improvement over fixed-syntax voice control. A message generator provides spoken feedback to confirm that the system understood the operator's commands, or to ask the operator to repeat the command. The speech synthesizer provides a parallel means for the system and the operator to communicate in command input and system feedback.

Predictive graphic display

The problem of effectively controlling a robot through communication links with long time delays (such as would be experienced at GEO or polar orbits) has not been solved. The operator that provides velocity commands must know the real position of the robot; otherwise he might send the robot right through delicate satellite surfaces.

ATL is developing predictive feedback techniques that improve the opera-

tor's ability to control distant robotic activity. For example, real-time graphic displays of the commanded robot position will appear as graphic overlays on the real (delayed) camera views of the robot in the workspace. Before robot motion starts, the operator will see a stick figure of the robot superimposed on the camera view of the robot. As the operator provides velocity commands with the joystick, the stick figure will move away from its previous position, and indicate in real time the commanded positions of the robot. The camera view of the robot will catch-up to the stick-figure view after the round-trip time delay is completed. This technique will simulate real-time performance and overcome most problems of remote control of robot position with long time delays.

The graphic preview system will also show the operator what will occur during the next pre-programmed motion sequence. The operator may then see if such motions are appropriate in the real workspace, and observe interference by any obstacles that were not anticipated by the programmed motions. The operator could avoid the obstacle through joystick control and then resume programmed task segments.

Force/torque control

An industrial robot is controlled primarily by executing a sequence of motions described by changes in its position. This is satisfactory if everything in the workspace is controlled to tight tolerances, or if objects in the workspace can orient themselves to comply with the exact positions of the robot gripper or end-effector. Many satellite objects to be serviced will not do this. Many simple operations, such as opening hinged panels, require the robot to move along precise trajectories defined by the elements in the workspace, rather than by the robot. Operations may require control of terms of force to break work objects loose from their previous position, then in terms of position to move them to new locations, then again in terms of forces or torques to seat them in their final postion. Industrial robot controllers that only control robot positions are inadequate for many satellite servicing tasks.

ATL is currently integrating a robot controller, control language, and manipulator/end-effector drive electronics to sense and control sequences of forces and torques, as well as execute sequences defined by position changes. Planned demonstrations include opening hinged panels, unfastening bolted objects and refastening them to prescribed torques, and mating and unmating electrical connectors where one part is firmly mounted to a stationary panel and the robot gripper holds the other.

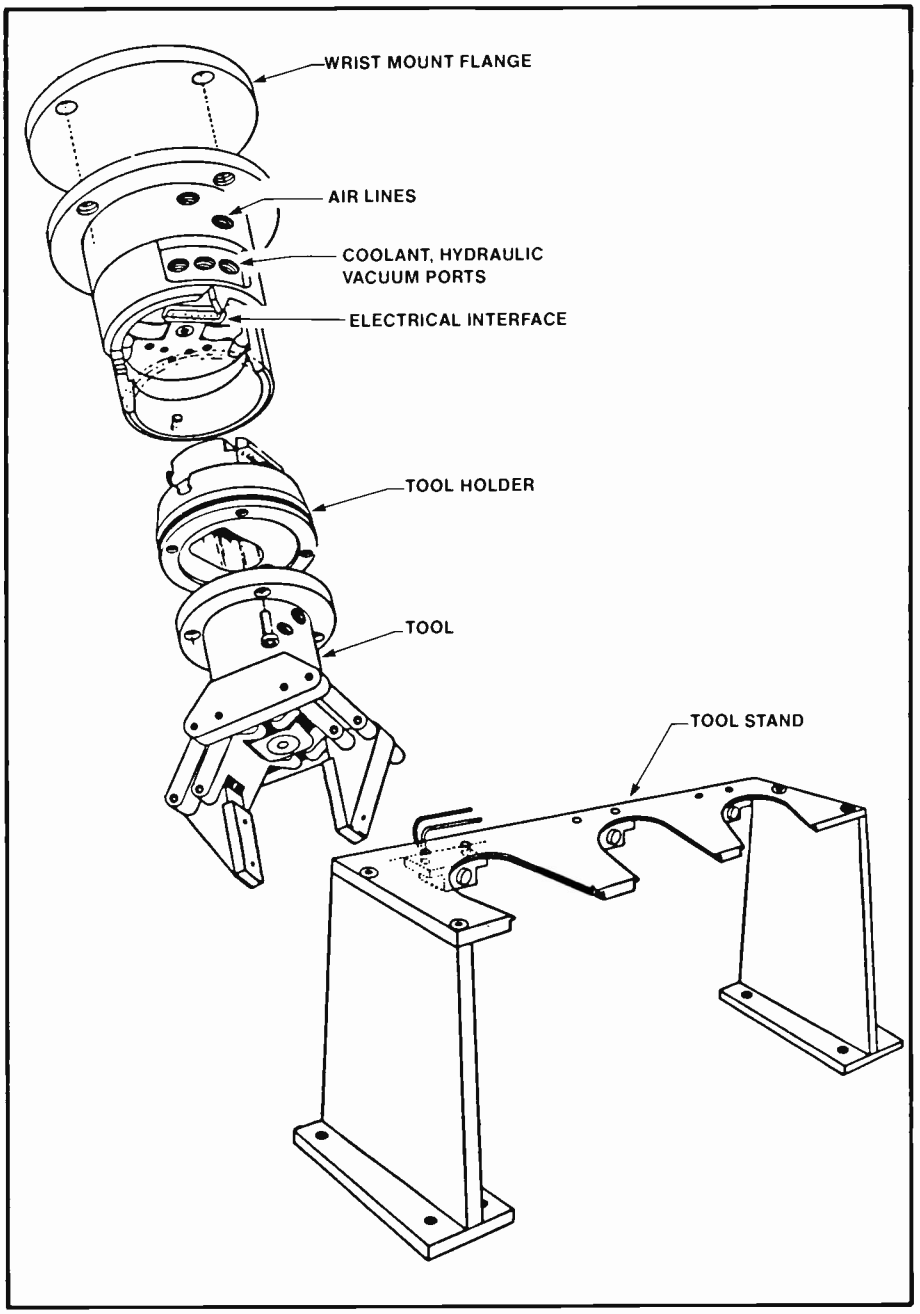

Fig. 4. *Quick-change robot gripper.*

3-D robot vision

Many teleoperation tasks will deal with the mating of two three-dimensional objects. Such tasks would be difficult to do with joystick control, even if the operator has stereo or orthogonal views of the work objects. We can automate such sub-tasks by suitable interpretation of 3-D range imagery and computer generation of required motions.

ATL has demonstrated a structured light range imaging system that can generate range maps of a robot workspace, with resolution of less than one tenth of an inch at rates up to four frames per second. We are developing techniques of interpreting those images to yield precise orientation and position information of

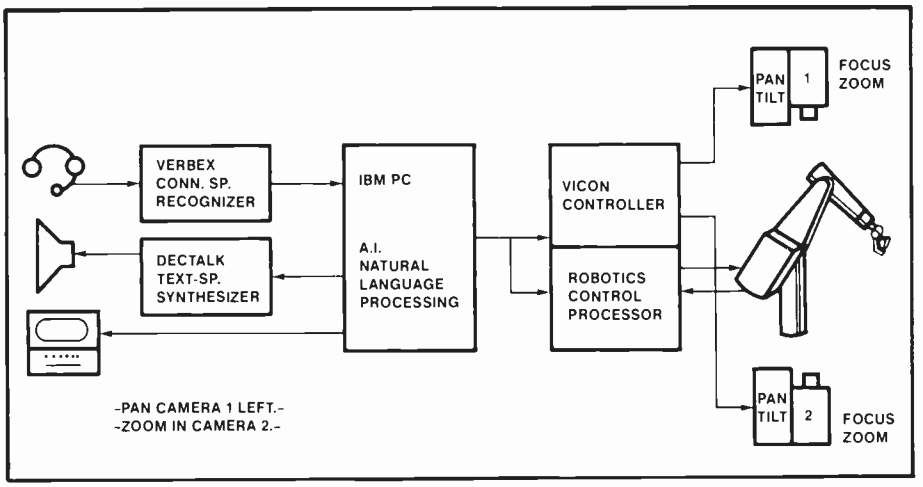

Fig. 5. *Natural-language interface for robot control.*

objects to be manipulated in the workspace. After the parts touch, mating of objects is completed through force/torque feedback and control.

Future development projects

Future development projects will include integrating ATL's VLSI capability to incorporate more intelligence at the remote site. Custom high-performance VLSI-based processors will be required to compute robot joint motions and forces in real time, in response to sensor inputs processed during robot actions. Real-time sensor processing, particularly vision, will also require new circuit technology for use in remote locations with minimum physical hardware and power.

Expert systems will be developed to integrate outputs from multiple sensors and intelligently process the appropriate signals for the task. Expert systems will also be required to plan and replan tasks and motions when unanticipated obstacles are detected, identified, and dealt with while accomplishing tasks that were defined as high-level goals rather than as specific trajectories or forces.

Improved word-recognition algorithms will accept speaker-independent connected speech with an increasing vocabulary. Natural language interfaces will interpret the astronaut's spoken words and respond in natural speech. Astronauts in space suits will then be able to communicate naturally and comfortably with robotic "astronaut apprentices" that help complete tasks, retrieve tools or parts, advise the astronaut of procedural steps, and inspect the quality of completed tasks.

Conclusion

ATL's advanced sensor systems, machine vision systems, computer architectures, artificial intelligence software techniques, word recognition systems, and VLSI designs all contribute significantly to answering emerging technical needs for servicing satellites in remote orbits in space. The desire for, and cost saving potential of, increasingly autonomous robotic servicing systems in space will provide an increasing need for these technologies. It should continue to do so through the 1990s and well into the twenty-first century.

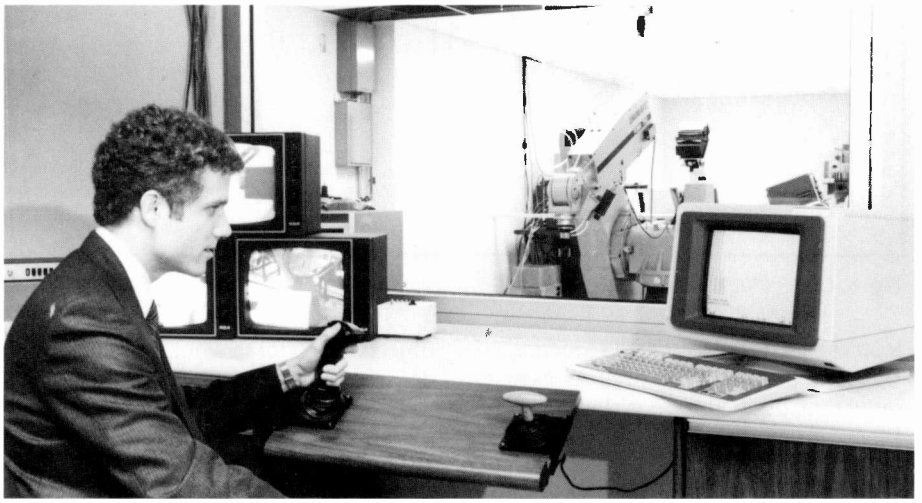

Paul Pierson is the Manager of ATL's Robotics Laboratory. He is responsible for the development of robotics technology and applications, and for associated disciplines. Prior to coming to ATL, he was Manager of Software Review and Evaluation in RCA's SelectaVision Videodisc Operations, where he was responsible for video and audio processing systems used in disc mastering, and for the quality of the tapes and the master discs used to produce discs. Before that he was a Program Manager in the Government Systems Division, responsible for high-resolution laser beam recorder development programs. Paul received a BSEE degree from Case Institute of Technology and an MSE degree in Systems Engineering from the University of Pennsylvania.
Contact him at:
Advanced Technology Laboratories
Moorestown, N.J.
Tacnet: 253-6502

E. Alexander | R. McClain | C. Tsikos | B. Siryj | R.M. Carrell

Advanced technology for mail processing in the 1990s

RCA has taken an adventurous approach to help the Postal Service handle the increasing volume of mail.

The U.S. Postal Service (USPS) currently processes and delivers over 120 billion pieces of mail per year, and this number doubles every 20 years. While letter mail processing has been highly automated, processing of other mail types is still very labor-intensive. To offset rising costs of labor and employee benefits that lead to increases in postage rates, the Postal Service has pledged itself to a long-range productivity improvement program, using the latest technology developments. Advisory committees and a Technology Resource Department have been set up within the Postal Service to conduct research in specific technologies, such as machine vision and electro-optics, robotics, artificial intelligence, and man-machine interfaces. Research in these technologies, viewed as critical in the quest for automation and worker productivity improvement, is being conducted under Postal research and development contracts by a variety of industrial and academic research organizations.

Abstract: *RCA is under contract with the U.S. Postal Service to research machine vision for robotic handling of irregular parcels and to construct an experimental robotic feeder mechanism that can "see" and handle large "flat" envelopes and magazines. This article discusses the postal requirements that give rise to these technology needs, the objectives of the research undertaken, and RCA's unique technical approaches to the automated postal system of the 1990s.*

Final Manuscript received November 20, 1985
Reprint RE-31-1-6

Evaluating robotics for future postal operations

Most mail in the postal system consists of ordinary business envelopes. They are handled by automatic machinery that accepts mailpieces within a specific range of sizes and shapes. There is a residue, however, that the existing machinery cannot handle. This residue comprises a large part of the immense mail volume and is, thus, a significant problem that costs over two billion dollars each year.

One category of this residual mail is known as irregular pieces and parcels (IPP), which includes small boxes, film mailers, rolled newspapers, hotel keys, and small bags. Another group, "flats," consists of generally flat mailpieces, such as tabloid weekly newspapers, magazines, oversized envelopes, packs of inquiry cards, and bank statements (stuffed envelopes). Open magazines, without sleeves, are especially difficult for machines to handle. This type of material includes much of the periodical and advertising mail, which brings little revenue but requires a large labor force for processing.

The fundamental tasks in mail handling are singulation, facing, Zip-code reading, and sorting. Singulation is the process of separating one mailpiece from a pile, a stack or a jumble on a conveyor. The many attempts to singulate residual mail have been variations of conventional material handling techniques such as differential motion or spreading. These techniques work well when the machines are designed for a uniform stream of a known product, but the techniques often fail for uncontrolled mixes of unpredictable products.

The USPS is sponsoring parallel work with several contractors in each aspect of non-letter-mail processing. Some of these contractors are trying to update the traditional technology in these fields. RCA chose the technically more adventurous path of investigation—robotic technology with emphasis on machine vision and programmable machines.

Our initial contract, "Machine Vision Research," required a search for machine vision technology that could analyze a pile of irregular parcels and pieces and direct a programmable machine to singulate it. After we analyzed 79 research papers and visited university and industrial laboratories, we concluded that a range image was the

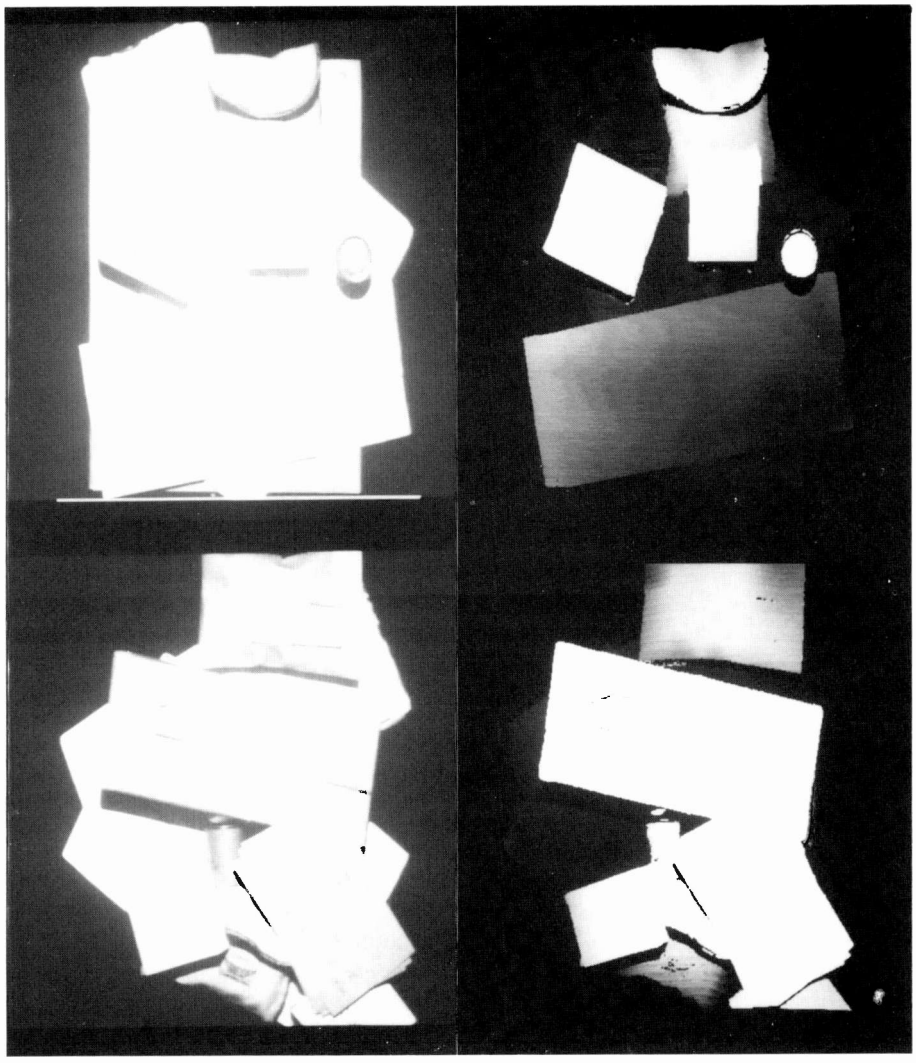

Fig. 1. *Intensity images (left) with corresponding range images.*

most promising technique for 3-D machine vision.

A range image of a scene replaces coordinates *X*, *Y*, and *Z* with *X*, *Y*, and intensity. To the eye, shapes will have abnormal shadings and no surface markings. A normal scene and its corresponding range image are shown in Fig. 1. There are two feasible ways to obtain a range image—structured light and laser ranging.

While RCA has built laser rangers in the past and several experimental models have been reported in the literature, these functioned at long ranges and were slow. In the last year or so, progress has been made in laser ranging at short distances, but for the postal application we recommended a structured light technique[1] that is detailed in the section on coded mask projection. We were subsequently awarded a contract extension to perform a laboratory demonstration of this method using the resources of ATL's Image Technology Laboratory.

A second contract was awarded to RCA for Advanced Feeder Research. This requires us to develop and construct a feeder that can take a wide range of generally flat mail, and singulate it at the rate of one piece per second with 97-percent efficiency. For this contract, ATL has provided a robotics research facility that consists of a PUMA 762 robot and an IRI D256 vision system. We chose this equipment not for speed, but for versatility. Besides being used to develop technology for the Postal Service, the equipment is serving military and space programs, such as those to develop robots for repair and maintenance of spacecraft.

A main requirement of the advanced feeder system is the ability to recognize and locate a single piece of flat mail (usually the top piece in a pile) automatically and then direct the robot to pick it up properly. The findings of the machine vision research on IPP problems were put to use to solve the "flats" problem.

Machine vision requirements

The apparent ease with which humans interpret the many elements of a scene conceals the great complexity of the processes involved. To emulate vision in machines we must carefully define just what is needed and find ways to reduce the complexity of the scene so it can be processed within the time and computing resources available.

Often, we simply need to know whether an object is there or not; simple reflective or transmissive sensors will do. For inspection, the shape and dimensions are needed (Are the holes all there? Are they the right size?). Silhouetted images are commonly used, because only binary (black and white) images with minimum data are produced. Usually, the job is to identify, measure, or locate known objects.

For the Postal application, the objects are not known; they vary widely in size, shape, surface markings, and attitude, and they are not separated. This closely resembles the classical bin-picking problem that has never been solved in the general sense. (The bin-picking problem is the robotic emulation of the human ability to reach into a bin of parts and select and orient one part. Assembly operators do this with ease, but vision-aided robots can only do it for known parts.)

Once a mailpiece is found and singulated, there are other tasks. The address label must be found, and then the Zip code. Finally, the Zip code must be read. Mailpieces, especially parcels and magazines, may be covered with distracting patterns that effectively camouflage the address label. If machine location and reading of the label fails, an image of the mailpiece is transmitted to an operator station where the Zip code is located, read and manually keyed.

Machine vision techniques

The development of machine vision systems to perform these functions is a challenge. If the human visual system had not been so successful at extracting meaningful information from a two-dimensional array of grayscale numbers, workers in the field might have

given up on the problem long ago. The human visual system processes images created by the IGRP factors: illumination, (surface) geometry, (surface) reflectivity, and (camera perspective) projection. The transformation from scene to image is many-to-one and, hence, noninvertible. One cannot take an image and "solve" for surface shapes and reflectivities. The interaction of IGR factors and the ambiguity of the perspective projection combine to present an extremely complex interpretation problem.

Most industrial machine-vision applications have relied on optical or mechanical ways to simplify the problem by selectively eliminating IGRP factors during the image generation process. For example, when a letter is imaged in an OCR machine, the surface geometry is flat and the illumination is uniform; hence, the images indicate surface reflectivity. Structured lighting techniques are often used to illuminate three-dimensional scenes with an angularly modulated light source, such as a planar projection of light that forms a stripe on the scene. A strong light source is used, so the stripe can be reliably detected despite ambient illumination (I) and varying scene reflectivity (R). Projection (P) ambiguity is eliminated because of the known geometry between the light source and the camera. Hence, structured light provides direct access to geometric information, such as range to a surface or surface shape.

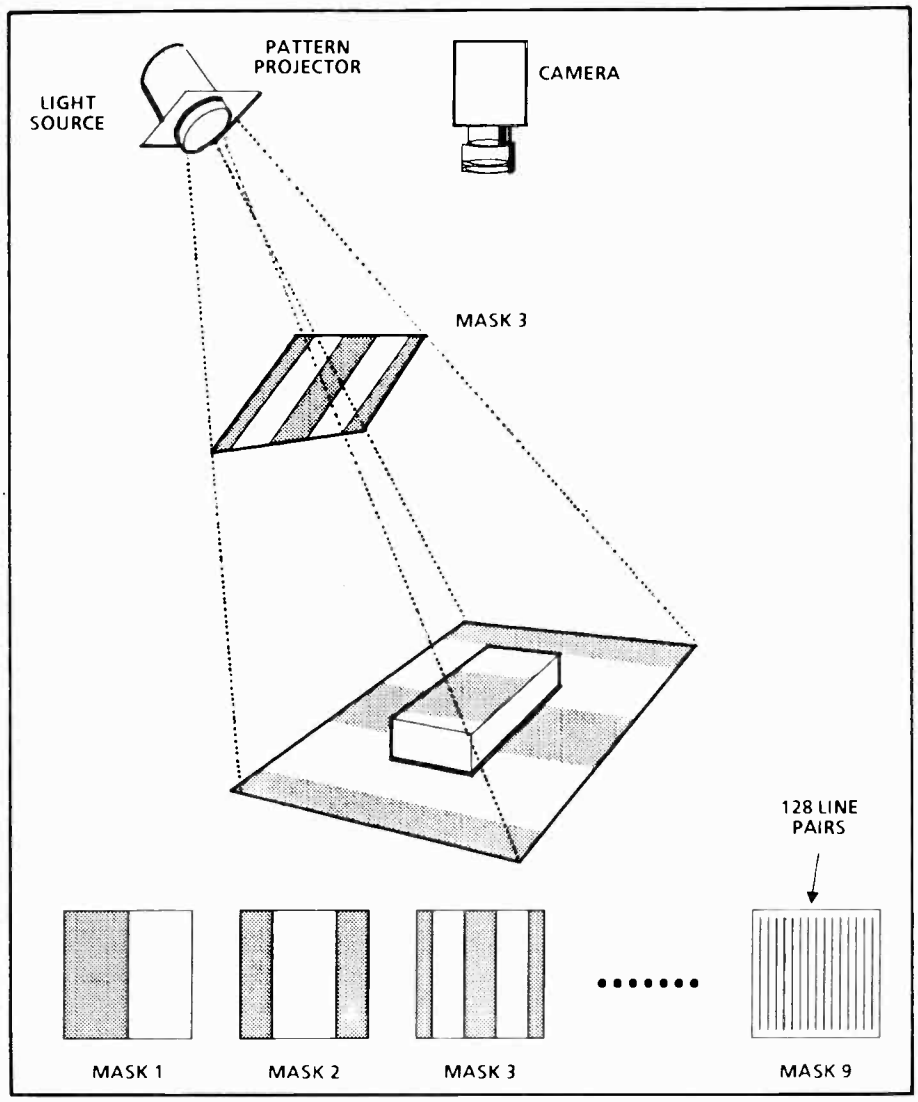

Fig. 2. *Inokuchi structured-light techniques for range imaging.*

Coded mask projection

ATL developed two structured light techniques for the USPS. The more general method, after Inokuchi,[1] is illustrated in Fig. 2. A point source of light illuminates the scene through a series of masks that project wedges of light with well-defined geometric boundries. The masks contain a binary- or gray-coded series of patterns, and a camera views the scene from a different angle. Each pixel on the camera sensor projects a narrow cone of perception that intercepts the projected wedges of illumination at geometrically defined points. These points are unambiguously identified by the mask code; they will be illuminated by some masks and not others.

For a voxel resolution of 256 by 256 by 256, eight masks are projected. (A voxel is a volume element; a three-dimensional pixel.) Each image is binary and is mapped to one level of an eight-bit memory. The code in each byte can be directly interpreted as a range. The basic triangulation scheme yields an image in polar coordinates that is corrected by a lookup table.

This process is straightforward for objects of uniform reflectivity viewed in a dark room. For mailpieces in room light and with surface markings, we cannot be certain whether a specific point in the scene is illuminated by a particular mask. To resolve this, we take two more images—one with the projector off; a second with the projector on, but with no mask— so all points are illuminated. For each point that the camera sees, we split the difference in intensity in the two images to form a threshold for comparing each masked image. The success of this technique is demonstrated in Fig. 1, where surface markings are removed from the range image.

This technique is fast and modest in computational requirements. The primary time required is to change the masks. We expect to use a custom liquid crystal cell that operates in the dispersion mode and should change mask patterns in 50 ms.

Next-piece location

In the advanced feeder research system, layers of scattered flat mailpieces are presented to the vision subsystem, which selects the next (usually top) mailpiece to be picked up. Here, the general shape and attitude are known. Therefore, it is possible to use a simpler structured light technique.

The vision subsystem's camera views the mailpieces on a conveyor from

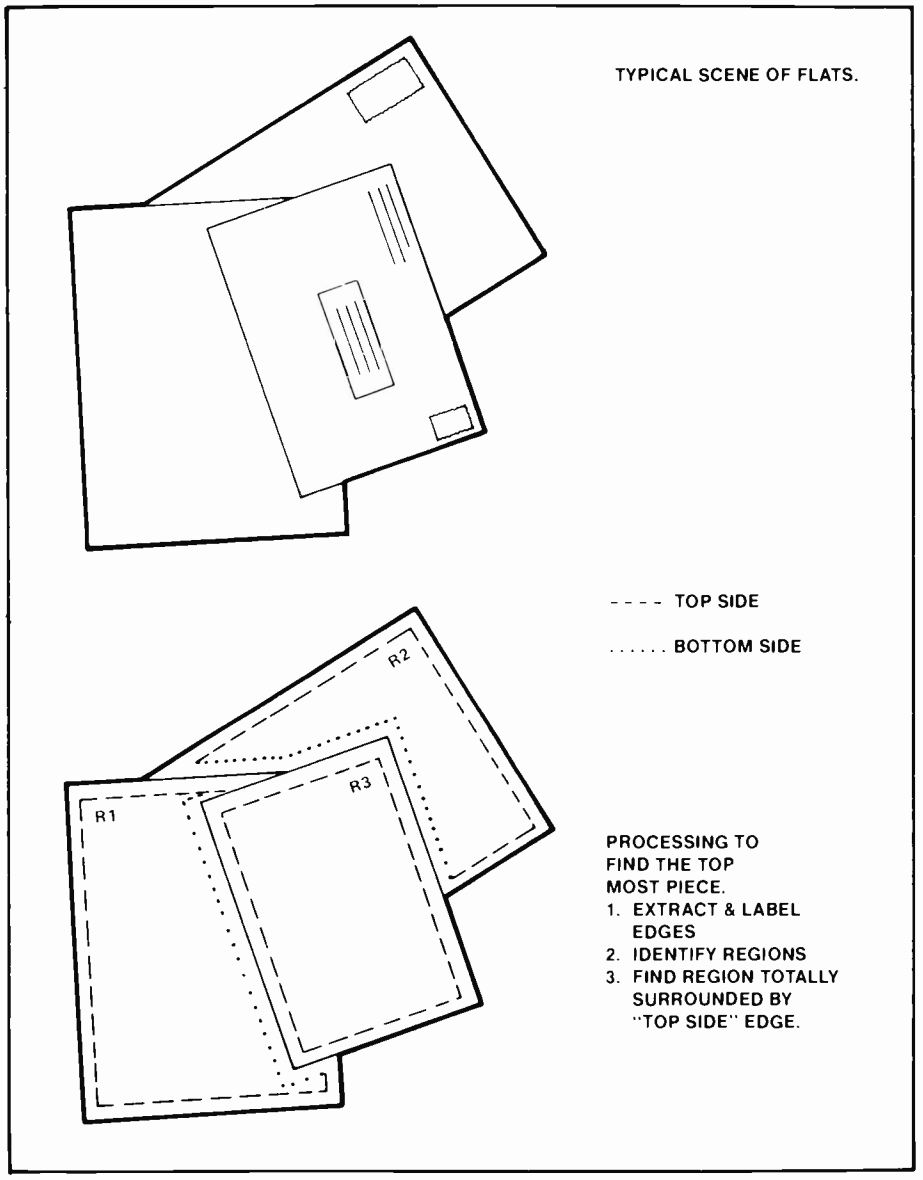

Fig. 3. *Method used for finding the top piece in a pile of mail.*

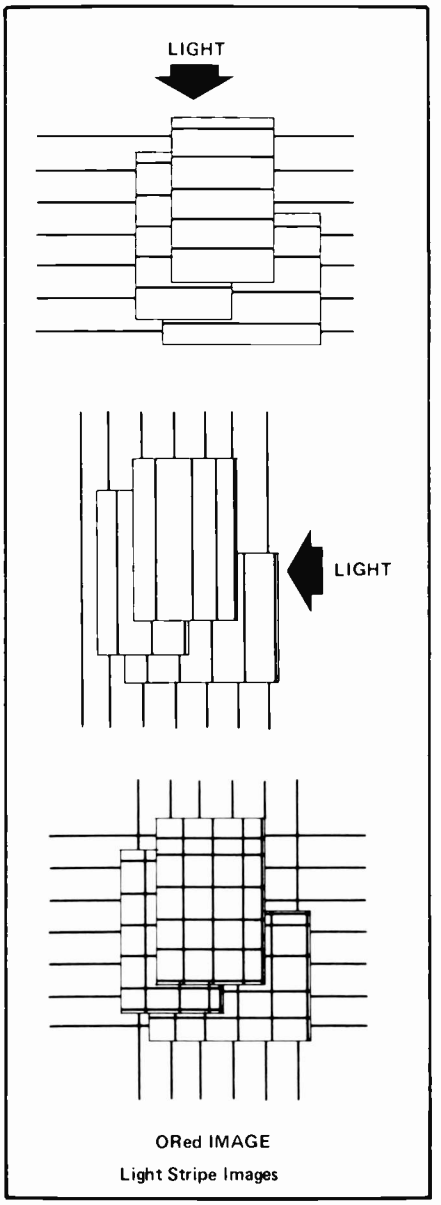

Fig. 4. *Use of structured lighting (stripes) to obtain geometric information on flats.*

above. The camera image is processed by an algorithm that identifies the top-most piece. The process consists of locating edges in the scene, using the edges to break up the scene into regions, and labeling the regions by examining the type of edge that surrounds each one.

Figure 3a shows a simple case consisting of three flats. In Fig. 3b, the edges have been identified, and they break up the non-background area into three regions. For each region, all points along the edges have been labeled as either the top side of the edge (denoted with a broken line) or the bottom side of the edge (denoted with a dotted line). Only the region associated with the top piece is totally surrounded by edges that have been labeled as the top side of the edge. To identify a top piece, we determine which regions are totally surrounded by these "drop" edges.

To find the drop edges, we project light stripes from an angle and note the direction of the discontinuity in the stripe as it passes over the border between two mailpieces. If the second piece is lower, the stripe "drops" to the next surface and is displaced along the path of projection (Fig. 4.) At each discontinuity, we know the direction of the drop because we can derive the height of the surface at each point along the stripe from the location of the stripes in the camera image. Two stripe images are acquired and analyzed, but the stripe directions in one image are rotated 90 degrees from the direction in the other image. This guarantees that all four edges of the top-most flat will intersect multiple stripes. The two stripe images are logically combined and analyzed into groups that correspond to surface regions (Fig. 3b). With the stripe implementation, edge sampling only occurs at the intersection of the stripes with the edges. After the stripes are grouped into regions, the identification of top-most flats involves typing the stripe discontinuities as "drop discontinuities."

We have investigated interpolation of pixel samples across the stripes as a way to increase the resolution of the imaging system and allow discontinuity detection for ever thinner flats. For this method to work, there must be no ambiguity as to which stripe is seen. Because

the stripes are projected from a 45-degree angle, the stripe image moves horizontally as a surface rises. If the surface is high enough, the movement can equal or exceed the nominal spacing between the stripes. Now there can be an ambiguity about which stripe is seen at any point in the scene. We can devise various rules for keeping track of the stripes, but various mailpiece positions can defeat them.

There are only two robust ways of identifying the stripes: spacing and coding. The coding method (described above) requires multiple images. For the feeder project, we spaced the lines 1.5 inches apart and restricted the mailpiece layer thickness to two inches. Wider spacing might miss small mailpieces; closer spacing would require us to restrict the layer thickness even more, or use a simple coding scheme.

In flat-mail processing, unbanded magazines present a particularly difficult problem. A suction cup pick-up device can adhere to the cover, but the magazine will fan apart if the cover is pulled away. The next section describes a gripping tool that lifts the bound edge and moves a support under the magazine. For this to work, the bound edge must be identified. One way of doing this is to get a closeup of the mailpiece edge through a telephoto lens, and then analyze the texture with two convolution filters. One filter emphasizes the texture of the cut paper edges; the other suppresses it. By comparing the filter outputs for the edge's image, we can determine if the edge is unbound (showing a paper stack) or bound (not showing the stack). The details of the whole system are beyond the scope of this article.

Grasping the mailpieces

In robotics technology, the functioning device at the end of the "arm" is known as an "end effector" or "end-of-arm tooling." There are some more-or-less standard versions on the market today. Most are two- or three-jawed clamps with very specific capability, but they are not versatile enough to handle parts with unknown or consistently varying geometries. In this article, we use the term "gripper" to refer to an end effector that is custom designed to handle a variety of flat mail.

Usually, the mailpieces are presented to the gripper in a disoriented pile. It is not good practice to slide the mailpieces together when manipulating the pile, because edges tend to catch on other mailpieces that have bent corners or partially detached labels. This is particularly true if the pile contains unbanded magazines.

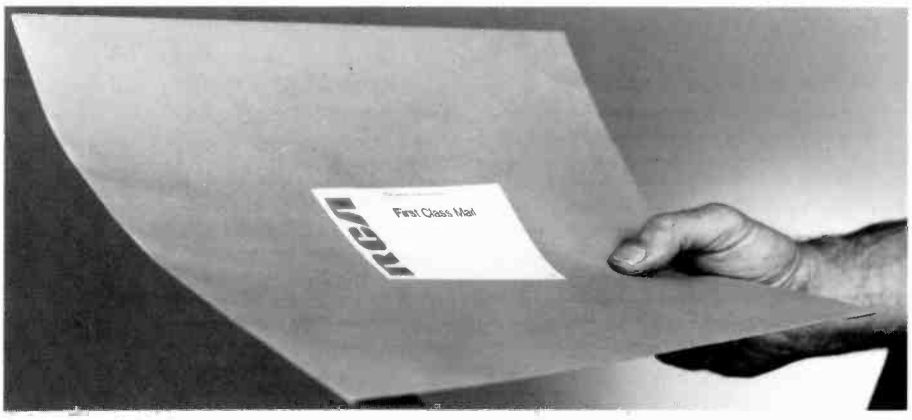

Fig. 5. *The mail gripper is designed to grasp a flat mail piece in a manner similar to the way a human hand clamps an object between the thumb and fingers.*

We concluded that the optimum way to depile and then singulate a stack is to remove the top piece first and then the next one, and so on, without sliding any piece. This approach closely approximates the way humans would remove the mailpieces, one by one, starting at the top of the pile.

The approach to singulation using a robot arm with a gripper is to lift the mailpiece in place and then move it to a singulation conveyor. For most envelopes, we can do this easily with the help of suction cups. In handling unbanded magazines, however, the task becomes much more difficult. The suction cups will attach themselves only to the top sheet(s), leaving the rest of the pages unsupported except at the bound edge. It is not usually possible to lift heavy magazines by just holding onto the cover, especially under high accelerations. One way to solve the magazine problem is to locate the bound edge (discussed in the preceding section), lift it up slightly, place a support under it, and then clamp the magazine in place, much like the human hand grasps an object between the fingers and thumb (see Fig. 5).

As previously mentioned, the USPS has specified that mailpieces shall be singulated at a rate of one per second. (This is the current rate for keyboard entry of Zip-code data by postal operators.) To achieve this throughput, we have allocated specific times to machine vision, robotic arm movements, and gripper action. The goal set for the gripper is to pick up a mailpiece in 200 ms. Associated with this very short pickup time are high accelerations and periodic requirements for high energy. To ensure that a mailpiece will not be damaged during the pickup cycle, we need to control accelerations by programming the pickup motion.

We chose the following design goals for mailpiece singulation:

- The gripper shall be capable of picking up all "nonmachinable" and "machinable" flat mailpieces, including unbanded magazines.
- To preserve mailpiece integrity, the gripper shall pick up all mailpieces in place without sliding them, especially when dealing with unbanded magazines.
- The gripper shall curve (crimp) each mailpiece to stiffen it during transport, if possible.
- Programmed accelerations shall be used during mailpiece pickup.
- Once the robotic arm and gripper are positioned at the top surface and an edge of a mailpiece, the gripper shall tolerate a thickness variation of one inch.

The performance goal for the gripper is to pick up a mailpiece in 200 ms.

The general use of vacuum or suction cups for lifting or holding is a well-known art. The unique aspect of the Postal application is that the suction cups have to attach themselves to different, nonuniform, and sometimes porous materials. Another problem is we might lose the vacuum when picking up a piece—such as a stack of papers in a very pliable envelope—and, as a result, the piece forms a bag (Fig. 6). A similar condition might occur when picking up

The postal technology program

In 1983, the U.S. Postal Service (USPS) made a significant change in its organization for technical research and development. In the past, such work was conducted by the USPS Laboratories in Rockville, Maryland, with the support of private sector contractors and other government research activities. RCA has performed several technical contracts for the USPS Laboratories, including experimental systems that Advanced Technology Laboratories (ATL) designed for parcel sorting and voice Zip coding. In addition, the Communication and Information Systems Division (CISD) completed a series of electronic mail studies and experimental system designs that culminated in the national implementation of the E-COM system for processing electronic computer-originated mail.

The reorganization in 1983 directed the USPS Laboratories' role to technical support of the Postal Operations Group. Responsibility for advanced research was assigned to a Technology Resource Department at postal headquarters, with the mandate that research be conducted by the best qualified industrial and academic research organizations. The Technology Resource Department operates with oversight by various USPS executive committees and with the direct support of a consulting engineering staff provided by the Arthur D. Little Corp. Further, it is subdivided into an Office of Operations Research and Systems Requirements and an Office of Advanced Technology.

The main objective of current postal research is productivity improvement through automation. Postal workers are generally well paid, both in salary and benefits. Unless their productivity can be improved by increased use of automation, normal wage-rate escalations will continue to produce corresponding postage rate increases. Because the amount of mail handled grows annually, automation can reduce the labor content of processed mail and produce desired cost savings without layoffs. To date, significant use of automation has been limited to letter mail that is uniform enough to permit rapid, consistent mechanical handling. The thrust of the current research that the Technology Resource Department sponsors is to identify and develop techniques for processing irregular parcels, "flats," and periodicals in production systems designed for the 1990s time frame. Toward this end, specific requirements have been identified for parallel research award in several technical areas.

To establish a community of qualified research organizations, the Office of Advanced Technology negotiated basic-ordering agreements with universities, research laboratories and manufacturers who have qualifying experience and facilities to work in one or more of the following areas:

- □ Materials handling
- □ Electro-optics and character reading
- □ Man/machine interfaces
- □ Code printing, scanning, and labeling

RCA has established its qualifications for the first three of these technical areas and is now under contract to conduct research in machine vision for sorting irregular parcels, and create a research model of a robotic feeder mechanism for obtaining a variety of flat mailpieces from a random pile and delivering the pieces in a single stream for address reading and coding.

RCA established its qualifications

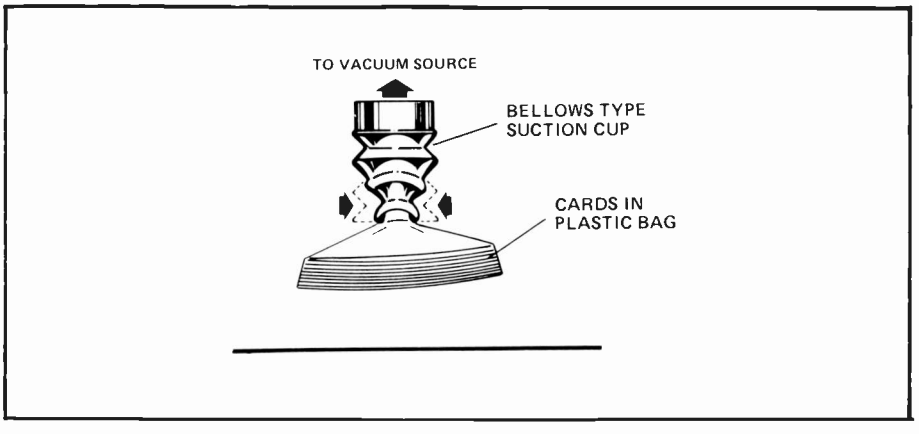

Fig. 6. *Suction cup gripping very pliable envelope.*

a heavy magazine near its bound edge (Fig. 7).

Single sheets of paper or very light magazines present yet another problem. If the suction cup is large, it tends to wrinkle the paper; the center of the cup is an apex. Quite often, air passage forms wrinkles that produce imperfect suction cup adhesion, venting the cup vacuum.

Results we observed during static testing led us to the following conclusions about suction cup design:

- □ The lip of the suction cup should be as flexible as practical restraints will permit.
- □ The diameter of the cup should be small. If additional lifting force is required, then several small cups should be used. An added advantage of the small cup is that the air volume is small, so the cup can be evacuated and then pressurized in a shorter time than a larger cup.
- □ Bellows type suction cups behave better than standard suction cups.

The proposed gripper design is shown in Fig. 8. It consists of two basic components: the pickup/pressure pad and the support platen. The energy source is a constantly rotating motor/flywheel that

for the basic ordering agreements on the strength of the broad technical capabilities found in our Corporate divisions. Current programs are staffed by engineers from Advanced Technology Laboratories, RCA Laboratories, and the Communications and Information Systems Division. The work is planned and directed by a Program Management Office that was recently transferred from Camden to ATL's new location in the Moorestown Corporate Center. Eugene M. Alexander is the RCA program manager responsible for all current contracts with the Postal Service. Mike Carrell of RCA Laboratories serves as senior technical consultant for the work in machine vision and robotics. The work currently in progress is staffed primarily by members of three of ATL's laboratories: the Robotics Lab under Paul B. Pierson, Gerry Claffie's Optical Devices Lab, and William Ealy's Image Processing Lab. Future projects in the man/machine interface area could include participation by ATL's Artificial Intelligence and Software Laboratory, the Missile and Surface Radar Division's Human Factors and Industrial Design organization, and RCA Laboratories' Video Systems specialists.

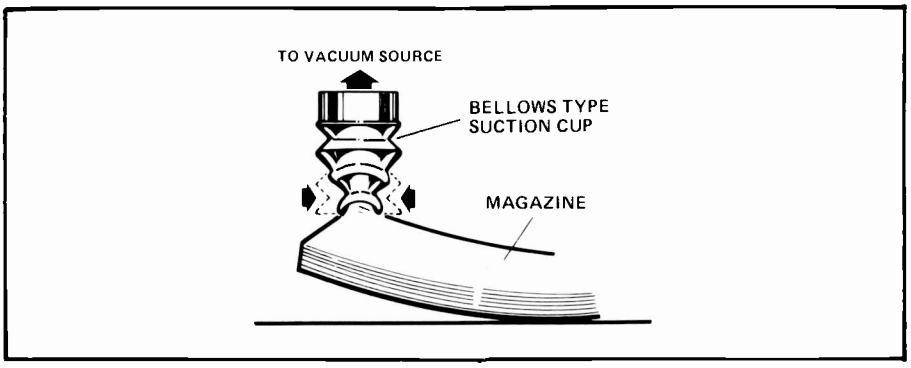

Fig. 7. *Gripping heavy magazine by bound edge.*

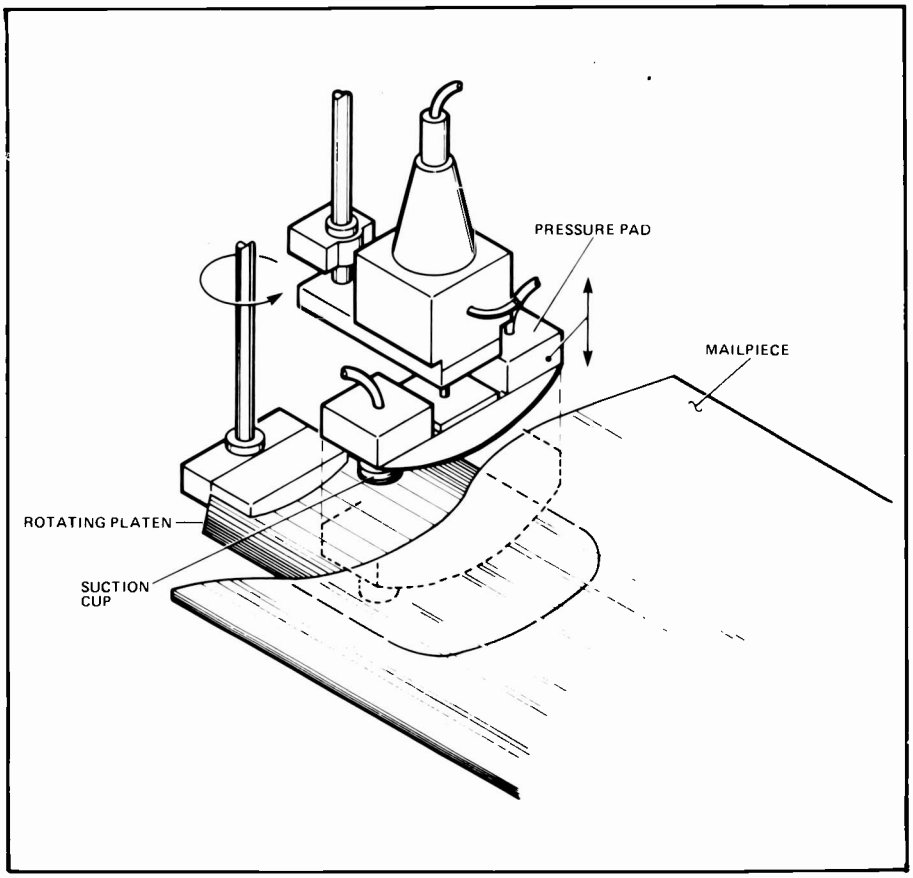

Fig. 8. *Design of gripper for flat mail handling.*

is coupled and decoupled from the two basic components by two spring clutch/brakes. A dc source drives the motor and provides versatility, easily varying the speed and cycle time of mechanism motions and accelerations. Timing belts transfer rotary motions, while a sinusoidal drive picks up mailpieces. Reflective detectors provide position and status information on the pick-up pressure pad and support platen to the controlling electronics.

Figure 9 describes the mailpiece pickup and drop-off cycles in detail, while Figs. 10a-10e depict the physical positions of the gripper and mailpiece.

For many flat mailpieces, principally large manila envelopes, it may not be necessary to hold the piece against the platen mechanically; the vacuum cups alone can support the load stably. The advantage of this approach is that we can reduce the 200-ms cycle to 133 ms by eliminating the platen-actuation cycle. The disadvantage is the additional burden placed on the mailpiece identification (vision) system.

Progress and projections

The PUMA robot was installed in the feeder testbed in July 1985 and, by November, we were able to demonstrate successful singulation from scattered layers of flats on a conveyor (Fig. 11). We are now integrating the components of the system into the testbed. Transfer-cycle times that average 4.7 seconds per mailpiece have been measured using a single suction cup attached to the general-purpose laboratory robot. This was achieved even though image processing presently requires 5 seconds and the robot cycle consumes 3 seconds. (Time is saved because two or more accessible "top pieces" are often found and fetched for each image processed.)

Future plans include modifying the testbed to handle IPPs. An upgraded vision system is planned that will obtain the contours of varying three-dimensional objects using available processing

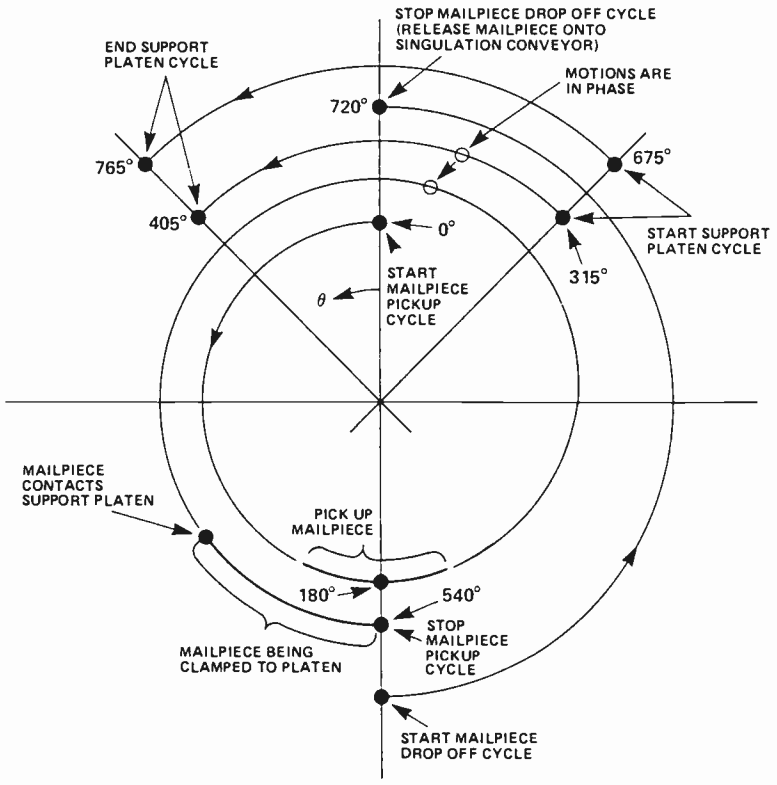

t	LINK (L_1) POSITION	PICKUP/PRESSURE PAD STATUS	SUPPORT PLATEN STATUS
0	$\theta = 0$	AT ITS HIGHEST POINT (CUPS ARE NOT COLLAPSED).	OFF*
1	$\theta = 180^\circ$	SUCTION CUPS ARE ATTACHED TO MAILPIECE.	OFF
2	$\theta = 315^\circ$	PAD & MAILPIECE ARE ALMOST AT THE HIGHEST POINT.	PLATEN STARTS TO MOVE INTO THE ON POSITION.
3	$\theta = 360^\circ$ $t_3 - t_0 = 133$ ms	THE VISION SYSTEM HAS DETERMINED THAT A PLATEN SUPPORT IS NOT NECESSARY. THE PRESSURE PAD (L_1) HAS STOPPED AFTER ONE COMPLETE REVOLUTION.	OFF

ALL END EFFECTOR DYNAMICS ARE STOPPED.

4	$\theta = 405^\circ$ $t_4 - t_2 = 34$ ms	PAD & MAILPIECE HAVE STARTED TO MOVE DOWNWARD. VERTICAL POSITION IS SAME AS IN t = 2.	PLATEN IS IN POSITION.
5	$\theta = 540^\circ$ $t_5 - t_0 = 200$ ms	MAILPIECE IS COMPRESSED ON THE PLATEN.	PLATEN IS IN POSITION – SAME ROTARY POSITION AS IN t = 4.

ALL END EFFECTOR DYNAMICS ARE STOPPED.

MAILPIECE IS READY TO BE SINGULATED BY THE ROBOTIC ARM.

6	$\theta = 675^\circ$	PAD & MAILPIECE ARE ALMOST AT THE HIGHEST POINT – SAME VERTICAL POSITION AS t = 2.	PLATEN JUST STARTED TO MOVE INTO THE OFF POSITION.
7	$\theta = 720^\circ$ $t_7 - t_5 = 67$ ms	PAD & MAILPIECE ARE STOPPED AT THE HIGHEST POINT. VACUUM CAN BE TURNED OFF IF THE ROBOTIC ARM IS OVER THE SINGULATION CONVEYOR.	PLATEN IS MIDWAY BETWEEN ON AND OFF.
8	$\theta = 765^\circ$ $t_8 - t_7 = 17$ ms	PAD IS STOPPED AT THE HIGHEST POINT. PRESSURE IS INTRODUCED INTO THE VACUUM CAVITY, AND THE MAILPIECE IS RELEASED ONTO THE SINGULATION CONVEYOR.	PLATEN IS IN THE OFF POSITION.

*SUPPORT PAD OFF MEANS THAT THERE IS NO SUPPORT FOR THE MAILPIECE UNDER THE PICKUP/PRESSURE PAD.

Fig. 9. *Mailpiece pick-up and drop-off cycles.*

An interesting ancedote

Range imaging by stripe projection has been used in some interesting ways. One pioneer company opened a portrait sculpture studio; the subject sat on a stool while a set of stripe projectors and cameras mapped his head and shoulders from all sides. Microdensitometer analysis of the films produced a database used to drive a special milling machine that created a bust needing only finishing touches by a professional sculptor. While technically successful, the company decided the robot business was more profitable, and it became Robot Vision Systems.

algorithms with a 3-second run-time. By incorporating a high-speed SCARA robot that alternately serves two input conveyors (each with its own vision system), mailpieces will be singulated to the output belt at a rate of nearly one per second. Such a system is depicted in Fig. 12.

Using available system components, we expect to achieve the one-per-second goal for enveloped flats. Extension of this capability to magazines at half this rate is within reach but, for IPPs, is a year or two further out.

After appropriate testing and postal evaluation, we can undertake specific prototype designs, with the goal of fielding production-level systems in the early 1990s. The productivity improvements envisioned can mean a potential savings of hundreds of millions of dollars annually.

References

1. Inokuchi S., Sato K., Matsuda F., "Range-Imaging System for 3-D Object Recognition," 7th International Conference on Pattern Recognition, Montreal, Canada, pp. 806-808, August 1984.

Side View of Mailpiece Gripper

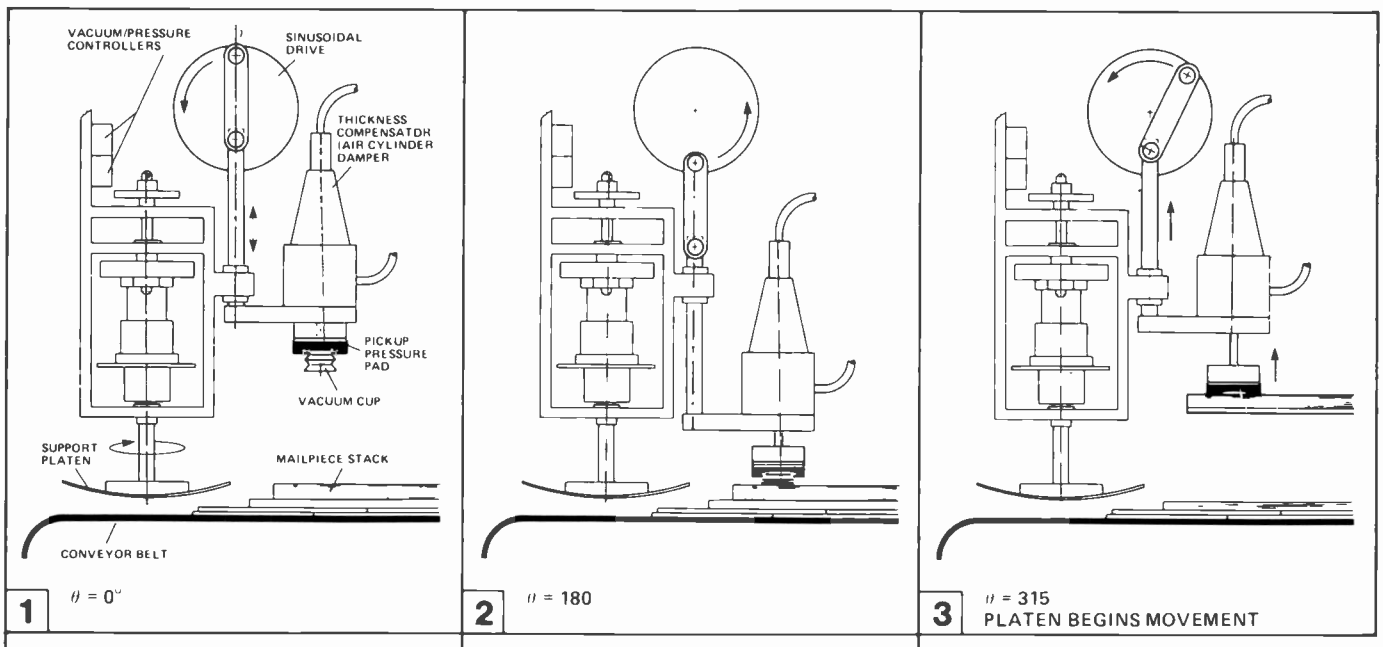

Fig. 10. *Positions of gripper in relation to mailpieces during phases of the cycle.*

Top View of Support Platen

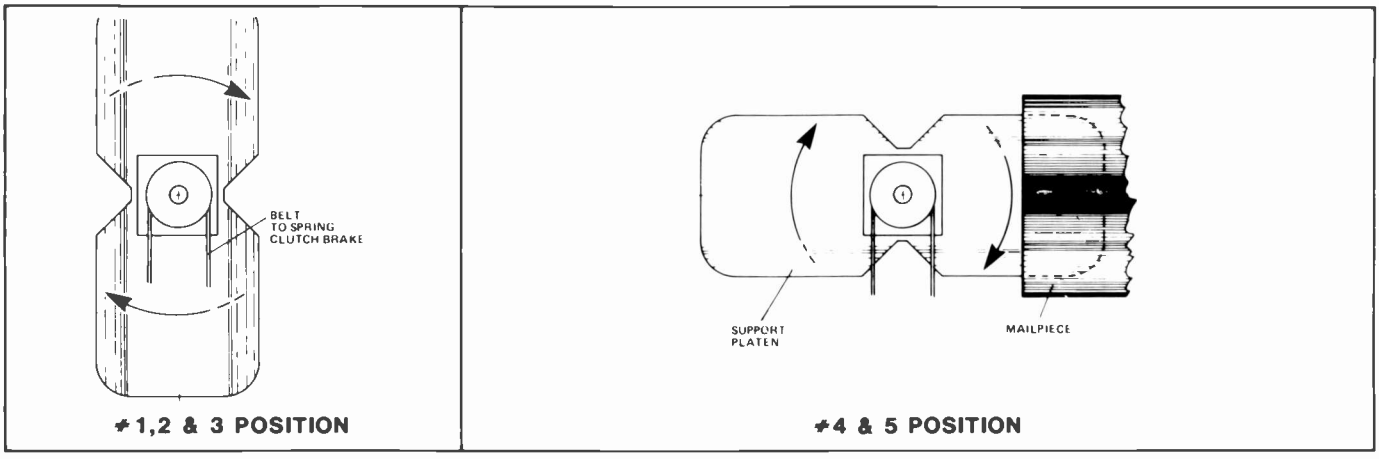

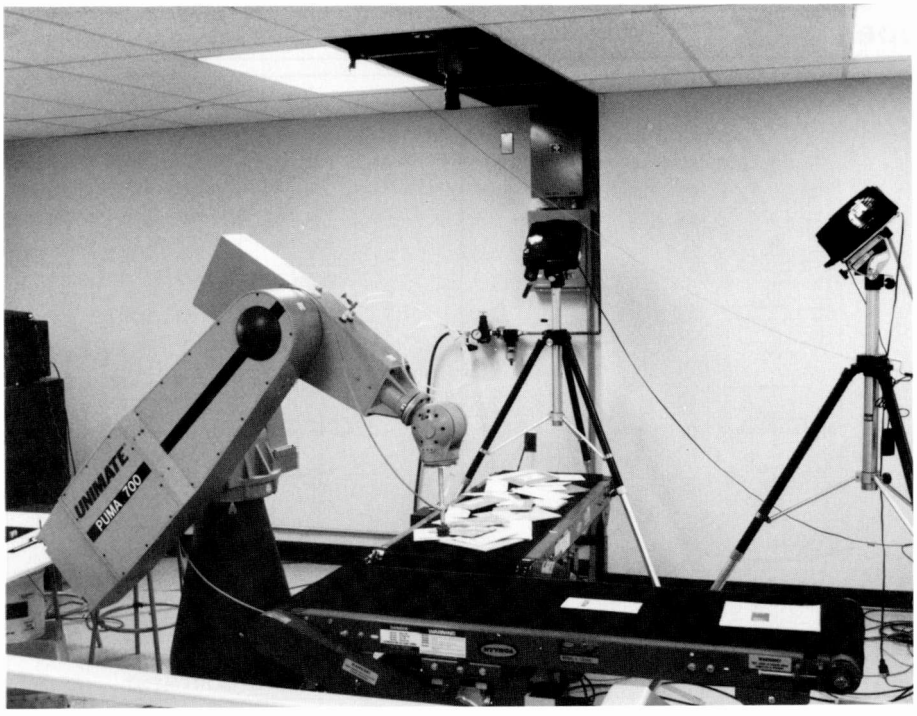

Fig. 11. *Flat mail in Advanced Feeder Research testbed.*

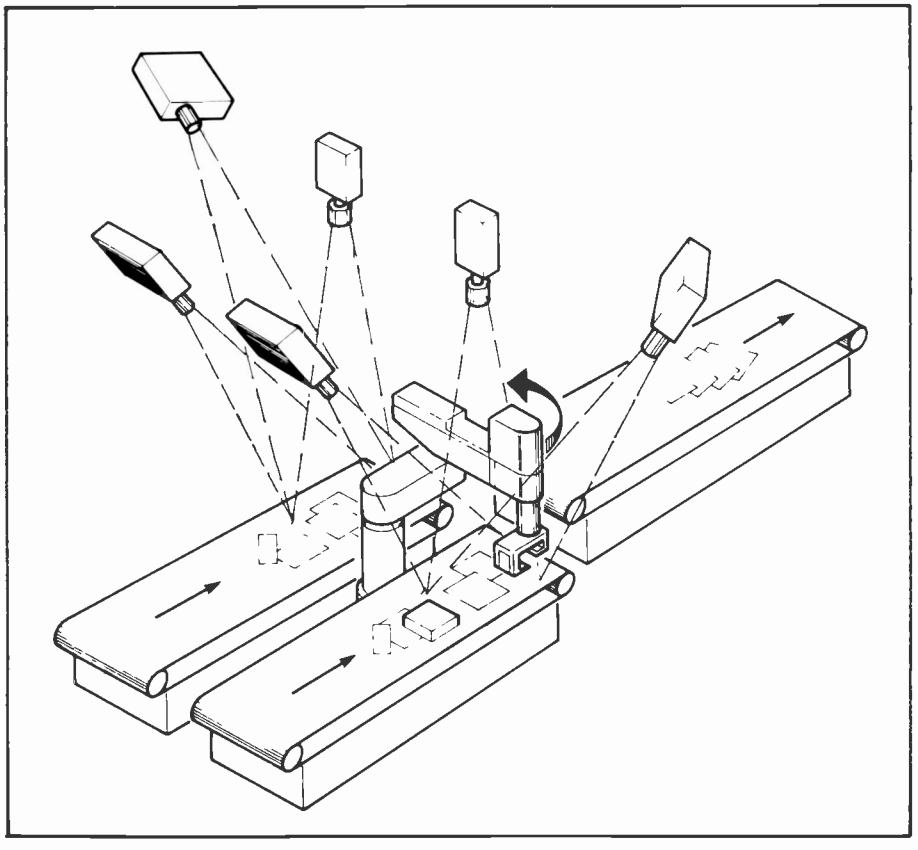

Fig. 12. *Proposed dual-line singulator.*

Gene Alexander is Program Manager for all USPS programs in RCA. During the past several years, he has been involved in system requirement studies, design projects, and implementation of systems intended to introduce the benefits of advanced technology to postal services. He is currently Program Manager for the USPS Advanced Research in Machine Vision and Advanced Feeder Research and Development contracts. Previously, he was responsible for overall management of RCA's contracts for development, implementation, and technical support of the USPS E-COM system. On the Electronic Message Service System (EMSS) study performed for the USPS, he was a key contributor to the modeling of centralized vs. decentralized mail processing systems, to the Parametric Analysis Model (PAM), and to the overall system economic analyses. He served as Program Manager for the Evaluation and Test System (ETS), which was designed by RCA and installed at USPS Laboratories as a working model of numerous electronic mail technologies. Gene received a BSEE and an MSEE from the University of Pennsylvania.
Contact him at:
Advanced Technology Laboratories
Moorestown, N.J.
Tacnet: 253-6420

Richard McClain is Unit Manager of the Machine Vision Systems group in ATL's Image Technology Laboratory. Recently, he was Technical Director of the Machine Vision contract for the USPS. He was responsible for the development of the structured light 3-D imaging system. On the Advanced Feeder research and development project he has been contributing to the development of special vision algorithms. In addition, he is directing an IR&D project for improving target recognition systems. He received BSEE and MEngEE degrees from the University of Pennsylvania, and he is presently a Doctoral Candidate at that university.
Contact him at:
Advanced Technology Laboratories
Moorestown, N.J.
Tacnet: 253-6457

Michael Carrell is a Member, Technical Staff at RCA Laboratories. He is a consultant on robotic automation throughout RCA, and recently has been working on the development of the Advanced Feeder for the USPS. Other projects include mixes of robots and material handling

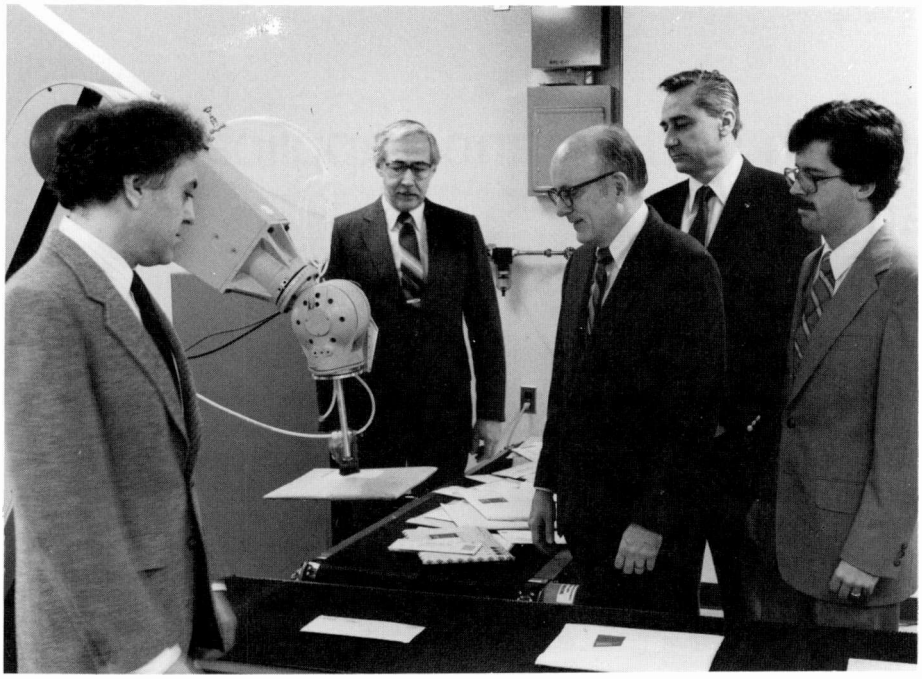

Grouped around the flat mail feeder test-bed are authors (left to right) Constantine Tsikos, Gene Alexander, Michael Carrell, Bohdan Siryj and Richard McClain.

systems involved with the production of TV picture tubes, satellites, phased-array radars, and printed circuit boards. Prior assignments have involved work with graphic systems for use in high-intensity noise. Mike received the BSEE degree from Iowa State University.
Contact him at:
RCA Laboratories
Princeton, N.J.
Tacnet: 226-2083

Constantine Tsikos is Unit Manger of the Video Processing/Robot Vision Systems group in ATL's Robotics Laboratory. He is currently directing machine vision research for the USPS Advanced Feeder project. Previously, in the Machine Vision project for USPS he recommended methods for automatic processing of irregular parcel post, and he was the principal investigator in the structured-light range imaging study. Prior to joining RCA, he was a Research Scientist at Siemens, investigating technology in the areas of robotics, image processing, intelligent sensors, and artificial intelligence. Constantine received a BSEE degree from National Technical University of Athens, Greece, and an MSEE degree from the University of Pennsylvania. He is currently working toward a PhD in Computer and Information Science at the University of Pennsylvania.
Contact him at:
Advanced Technology Laboratories
Moorestown, N.J.
Tacnet: 253-6464

Bohdan Siryj is Unit Manager of the Electro-Mechanical Systems and Thermal Systems group in ATL's Optical Storage Laboratory. Included in his responsibilities is direction of robotics and optical disk recording programs. At present, he is Technical Director of the electro-mechanical tasks in the USPS Advanced Feeder project. He is the recipient of a David Sarnoff Award for Outstanding Technical Achievement. Dan received the BSME and MSME degrees from Drexel University.
Contact him at:
Advanced Technology Laboratories
Moorestown, N.J.
Tacnet: 253-6488

W.P. Altman | G.M. Claffie | M.L. Levene

Optical storage for high performance applications in the late 1980s and beyond

RCA is developing high-data-rate, high-capacity optical recording systems for government applications that cannot be effectively addressed by alternate recording technologies.

The vastly increased amount of information that will be generated by advanced sensors and computing systems must be captured and stored by a new generation of high-data-rate, high-capacity recording and buffering systems that provide rapid, reliable access to the stored data. For example, the Space Telescope to be launched in July 1986 is expected to generate a terabyte of data per year for its 15-year planned life. A computer-compatible magnetic tape archive for this data would consist of approximately 100,000 tapes at 150 megabytes per tape. As a more dramatic example, the Space Station, to be launched in the early 1990s, is expected to generate 1 terabyte of data per day at data rates that can exceed 1 gigabit per second.

Abstract: *Over the last ten years, RCA has moved optical storage and retrieval systems from mere concept to reality. Two engineering models of "jukeboxes" have been delivered to the Air Force and NASA. These systems provide access to over 1200 gigabytes of stored information in six seconds.*

Current development effort is being directed toward jukebox systems with erasable disk media, optical buffers with extremely high speed data access on erasable media, and optical tape recorders with very high volumetric storage.

Final manuscript received November 20,1985
Reprint RE-31-1-7

Since 1975, the RCA Advanced Technology Laboratories (ATL) has been developing optical storage technology to address specialized user requirements, such as those of the Space Station, that cannot be satisfied on either technical merit or cost effectiveness by recorders intended for the more general purpose commercial market. Figure 1 shows a chronological progression of these technology developments—a progression to high data rate and high areal packing density via multiple channel recording. Rapid access to the stored data is achieved via simultaneous access to multiple disk surfaces, or automated disk handling mechanisms.

This paper briefly describes the most recent development work at ATL—optical disk "jukebox" recorder, optical disk data buffer and optical tape recorder—developments that will be applied to record at data rates above 1 gigabit per second and to achieve an on-line storage capacity above 1 terabyte.

Optical disk "jukebox" recorders[1]

Engineering models of optical disk "jukebox" recording systems were delivered to the NASA Marshall Space Flight Center and to the Air Force Rome Air Development Center in September 1984 and February 1985, respectively. The 14-inch diameter disks in these systems are arranged in an on-line library configuration and accessed via an automated handling mechanism similiar in concept to a phonograph record jukebox. Access is provided to any data in a store of 10^{13} bits (1250 gigabytes) within six seconds. The optical disks are housed in protective cartridges at all times to facilitate handling by both the operating personnel and the disk handling mechanisms. Cartridges are moved from the store to a load station by a precision belt-driven *X-Y* transport mechanism. The load station mounts the disks onto a precision turntable and the disks are spun up to speed so that data can be recorded or played back at user rates up to 25 Mpbs per channel using focused laser beams.

The NASA system has been interfaced to an image data base management system using a fiber optic bus. The Air Force system is interfaced to a Digital Equipment Corporation VAX 11/750 minicomputer to service a developmental image exploitation system.

NASA Optical Disk System (NODS)

The vast amounts of data from NASA deep-space probes prompted the development of an advanced data base management system. NODS, shown in Fig. 2, interfaces to this system using an advanced fiber optic data bus that can handle peak rates up to 100 Mbps. Packetized satellite sensor data will be archived within this system for rapid retrieval by an expand-

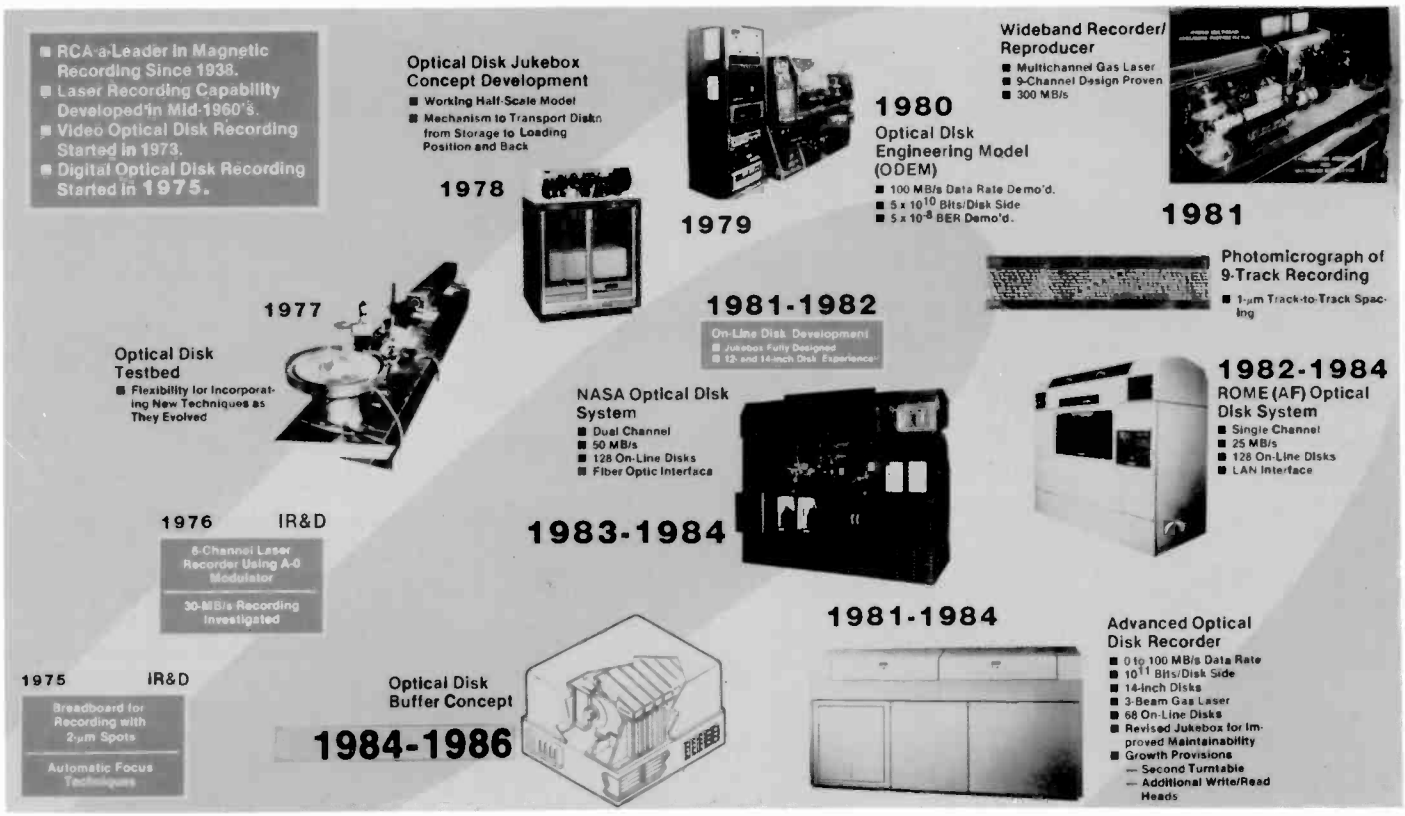

Fig. 1. *Chronological progression of optical storage development at RCA.*

ing number of users in the scientific community.

RADC Optical Disk System (RODS)

RODS will store and provide on-line access to large files for the Air Force Advanced Imagery Exploitation Program. This program is developing a testbed to evaluate techniques for rapidly manipulating and analyzing reconnaissance imagery and cartographic and geodetic data for mapping.

The VAX 11/750 computer interfaced with RODS provides the intelligence needed to manage the storage and retrieval of large files by file name in a very extensive data base. Each optical disk contains a directory listing every file on the disk and its location by track and sector number. When a command is received to record a file, the specified disk is loaded and its directory is input to a control computer in the RODS controller to determine where to locate the additional file. This file is recorded on the disk and its location is added to the disk directory. During file requests for playback, the specified disk is loaded onto the turntable and its directory is input to the control computer to determine the location of the requested file.

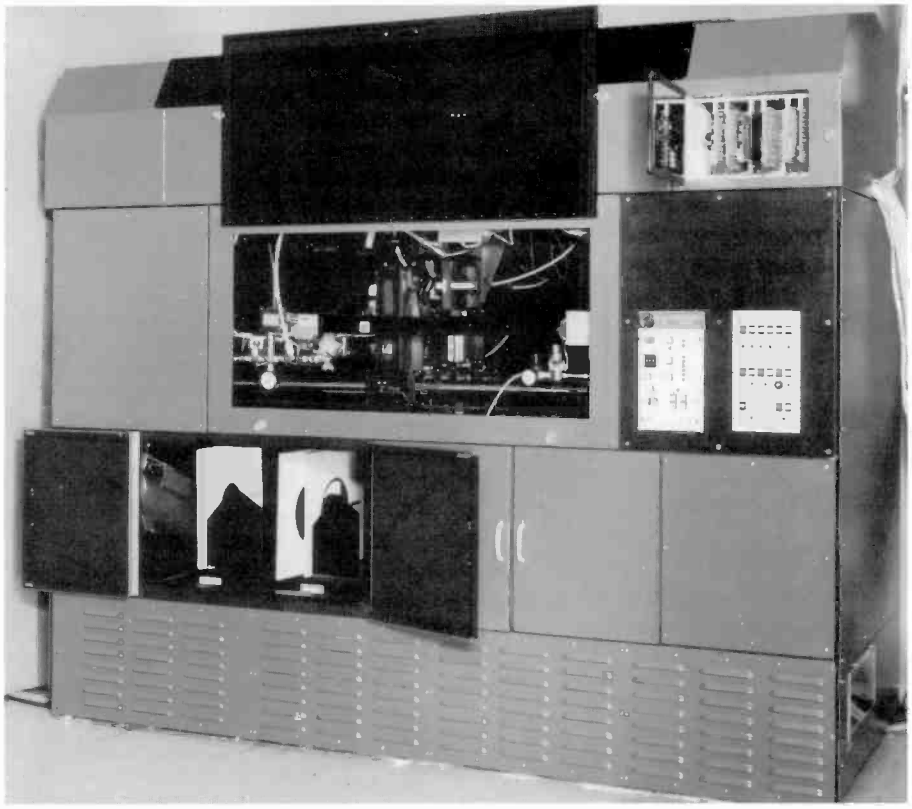

Fig. 2. *Photograph of NODS.*

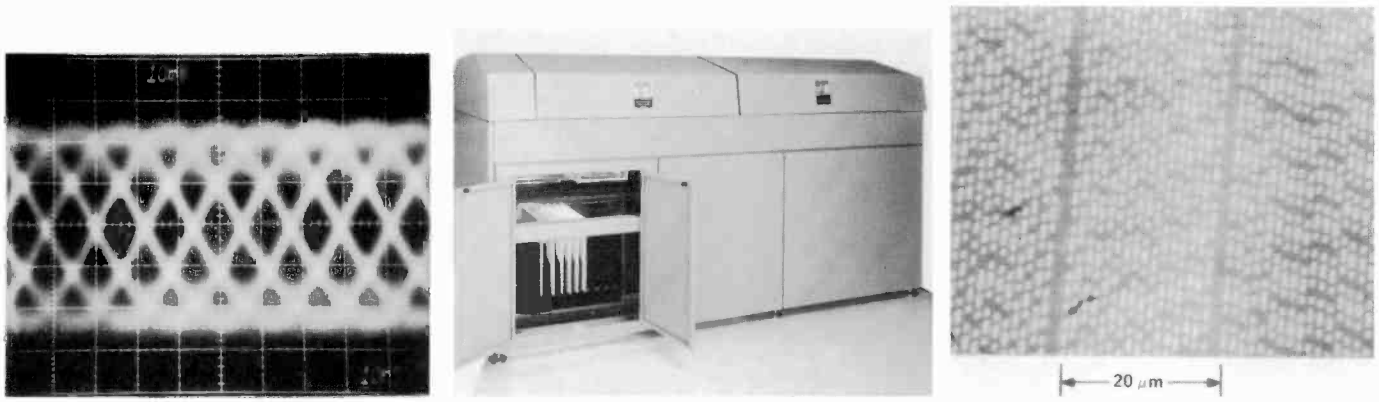

Fig. 3. *Collage of AODR photos (left to right: eye patterns, AODR, track-to-track spacing).*

Table I. *Adaptability to user requirements*

Use spiral or concentric data format on disk

- — Spiral
  - Long, continuous write/read files at fixed (high) data rate
  - Reduced data buffering
- — Concentric
  - Improved data access time
  - Simplified control

Provide higher continuous user data rate per optical head

- — Spiral (up to 9-beams/head) up to 450 Mbps
- — Concentric (up to 9-beams/head) up to 300 Mbps

Provide additional write/read and/or read-only heads

- — For double-sided disks
- — To increase data rate
- — To reduce data access time
- — For simultaneous write/read

Add second turntable

- — For improved access to on-line data
- — Improved capability for simultaneous write/read
- — For improved reliability

Implement continuous linear velocity (CLV) write/read

- — For increased disk capacity
- — Will increase data access time
- — Requires more complex recorder control

Implement user-specified interfaces

- — Control and data ports
- — Protocols
- — Data directory approaches

Provide additional on-line disk storage

- — More than 100 disks within recorder
- — More than 10^{13} on-line user bits
- — Automated disk exchange with external storage module
  - — less than 60-second access to 10^{14} on-line user bits

Advanced Optical Disk Recorder (AODR)[2]

The Advanced Optical Disk Recorder (AODR) shown in Fig. 3 is being developed as an evolutionary version of the delivered NODS and RODS "jukeboxes."

A key objective of the AODR program is to demonstrate the performance of critical optical, mechanical and electronic subsystems that could, if necessary, be combined in a modular fashion to provide a recording system with significantly increased capability. The potential design adaptations are summarized in Table I. This plan for growth in capability is demonstrated by the top view schematic of the optical table in Fig. 4, where it is shown that one of the two available turntable positions, and one of the four available translation stage positions at the single turntable position, are populated. A second key objective is to demonstrate operational reliability and maintainability enhancements with respect to the NODS and RODS by providing 1) more convenient access to the optics and the disk handling mechanisms, 2) optical system alignment control loops, 3) disk write power control loops, and 4) advanced subsystem status monitoring with software-controlled fault isolation.

The recording approach, demonstrated late in 1984, uses a three-channel argon laser-based optical system to input/output user data at rates up to 100 Mbps and to store up to 10^{11} user bits per side on a 14-inch diameter disk. This capability is achieved at a disk rotation rate under 30 rps and at a corrected read bit error rate less than 10^{-9} by using an efficient data encoding algorithm, read-while-write verification and a two-level error detection and correction architecture. The automated disk handling method achieves less than six-second access to the data from 68 on-line

disks with provision made for future interface to an external storage module to achieve automated access to several hundred disks in under 60 seconds. The recorder is configured as a Disk Drive Unit (DDU) containing the electronics, optics, and mechanisms required to write to, read from, and handle the optical disk media. The recorder also includes a controller that 1) controls and monitors operations, 2) buffers, formats and error protects the user data, and 3) interfaces the recorder with the host system.

Erasable/reusable recording

Another enhancement to optical disk technology being tested with the AODR is the use of erasable/reusable disk media with diode laser arrays as write/read/erase sources.

Recent advances in optical media technology have indicated that it is possible to erase and reuse optical storage media with one of several approaches: magneto-optics, phase change, and dyes encased in polymer films. The magneto-optics approach has been chosen for the current work because of its technical maturity and potential availability from multiple suppliers. The magneto-optic record/playback process is shown in Fig. 5. The sensitive, or recording layer is a vertically oriented magnetic material in which data is written by imposing a localized heated area, caused by focused laser illumination, in the presence of an external magnetic field. As the temperature of the region approaches a critical point, the coercivity of the magnetic material decreases such that the magnetic vectors within the media will align with a relatively weak external magnetic field in the heated region. Therefore, the direction of the resulting local magnetic vector in the sensitive layer is determined by the direction of the applied external field, and the size of this recorded feature is determined by the size of the focused spot. Data is read by reflecting a linearly polarized light from the surface of the sensitive layer. A very small polarization rotation based upon the Kerr effect is produced, dependent upon the direction of the magnetic field in the sensitive layer. The data readout signal is derived via differential detection of the polarization rotation. Erasure of the media to the original magnetic orientation is accomplished by the same process as recording, except that the imposed magnetic field is reversed. Because of the similiarity to magnetic recording, the magneto-optic technique has also been referred to as "thermo-magnetic recording" or "optically assisted magnetic recording." Being nondestructive, the record-read-erase cycle can be repeated millions of times.

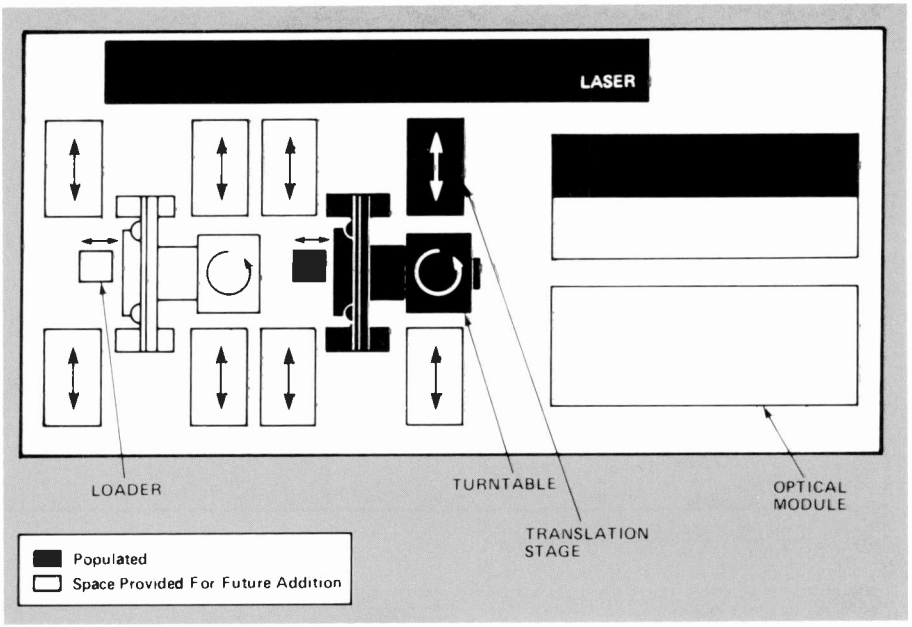

Fig. 4. *Top view of optical table (AODR).*

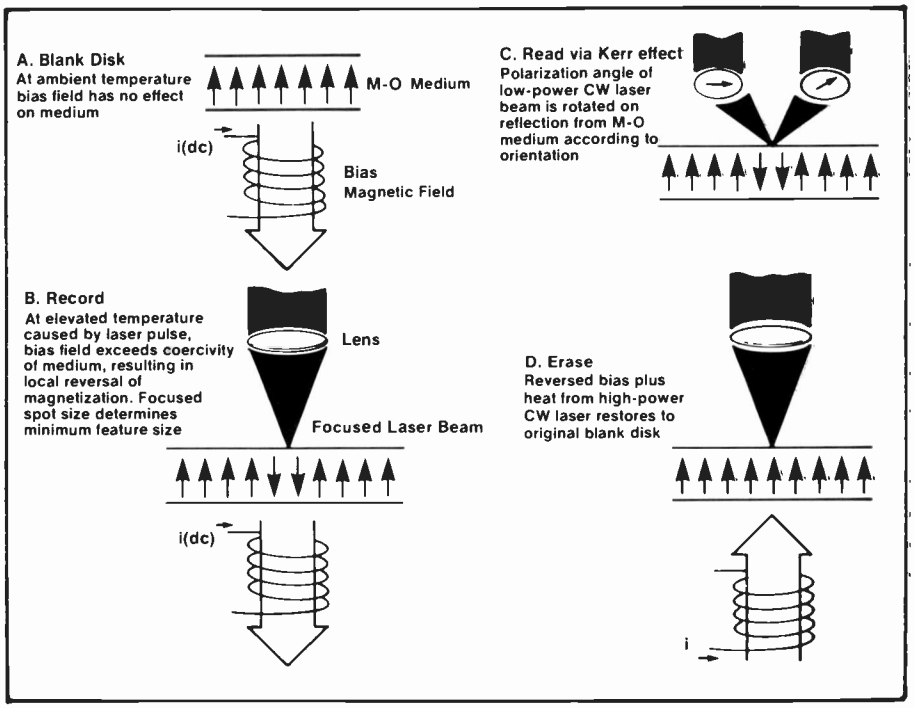

Fig. 5. *Magneto-optic record/playback/erase process.*

Recording with diode laser arrays

The very high data rate, high-capacity optical recorders developed to date have been based upon the use of argon lasers. These lasers require external components to generate and modulate the multiple beams that are applied to the disk media. The lasers are large, intrinsically inefficient, and require periodic maintenance. Laser diodes overcome these limitations and also provide the attractive capability for internal modulation at very high data rates.[3] However, they must be very closely spaced and precisely aligned to achieve a multiple channel source. This is a very difficult challenge that is not practical for produc-

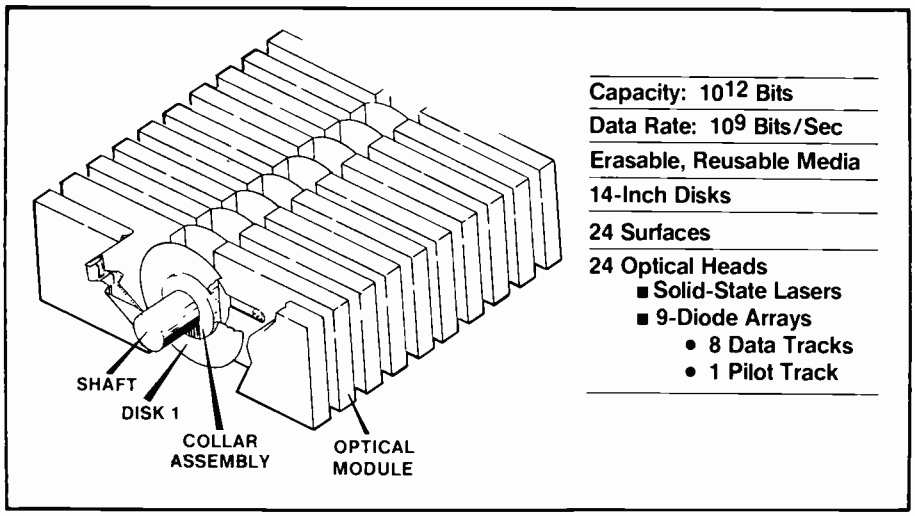

Fig. 6. *Optical disk buffer concept.*

Table II. *Projected performance of 12-disk system*

Capacity	
Number of user bits/ revolution/track	1,081,344
Track density	8 tracks/13.3 μm
Number of tracks/surface	37,813
Total buffer user data capacity (24 surfaces)	0.98×10^{12} bits
Data rate	
Disk rotation rate	15.413 r/s
User data rate/track	16.67 Mbps
User data rate/module	133.4 Mbps
User data rate/12 modules	1.6 Gbps
Access time	
Time to initial access	100 ms

tion recorders. To overcome this problem, the RCA Laboratories has been developing individually addressable, monolithic diode laser arrays and has demonstrated very encouraging performance from research samples of early design approaches.[3]

Early in 1986, a three-element array will be installed in the AODR to achieve three-channel write/read/erase operation with magneto-optic media. Since the feature sizes by which the optical recordings are made are linearly related to the wavelength of the write source, the larger wavelength of the laser diodes (830 nm) as compared to the argon laser (488 nm) implies a disk capacity smaller by approximately the wavelength ratio squared, a factor of four. In addition, diode laser power limitations (30 mW/diode CW), and the increased disk rotation rate required to achieve a given data rate per channel due to the increased minimum feature size, indicate that a maximum user data rate of 60-70 Mbps will be possible with the three-channel system, as compared to the current capability of 100 Mbps.

Optical Disk Buffer[4]

The Optical Disk Buffer, shown conceptually in Fig. 6, will meet the requirements of very large on-line data capacity (greater than 1 terabit) and very high input/output data rates (greater that 1.6 Gbps) with rapid access (less than 100 ms) to the stored data. Twelve double-sided 14-inch erasable magneto-optic disks are fixed to a common shaft in this system. Each disk surface has a data capacity of about 41 gigabits and is served by a dedicated electro-optic module that incorporates a nine-element diode laser array for record, playback and erasure of data at rates up to 133 Mbps. The projected performance of the twelve-disk drive is given in Table II. The relatively modest disk rotation rate (15.4 rps), which does not vary with the recording radius, and a moderate level of average drive power per laser diode (<30 mW-CW) both promise a favorable margin in system power and component life.

The projected linear data packing density on each disk surface is of the same order as that achievable with magnetic recorders. This system capacity advantage is gained with the Optical Disk Buffer by greatly increasing the number of recorded tracks per inch. Figure 7 shows the track geometry, in which eight data tracks and a central permanent pilot track used for position identification and tracking are placed on a 13.3-micron pitch.

Electro-optic module

Each electro-optic module is a fully independent subassembly containing a nine-element laser diode array, and optical systems for array focus, focus detection, data signal detection, and pilot track reading and tracking. The data signal is recovered in playback via differential signal detection using two nine-element photodetector arrays. Figure 8 illustrates the relative location of major subassembly components. Data access time is enhanced by translating only the low mass final focus objective and its two-axis positioning head. All electronic circuits for laser drive, focus and tracking control, data detection, and lens translation control are contained within the electro-optic module.

Development timetable

The Optical Disk Buffer is being developed as an exploratory development/brassboard model with a partial complement of electro-optic modules and active disk surfaces so that sufficient engineering performance data can be generated to predict full system performance. A major milestone will occur when the brassboard demonstration scheduled for early 1987 retires major technical risks. At this juncture, one or more prototype variants can be fully developed. These will include the high-capacity, 12-disk system that is currently planned, a less complex buffer with reduced capability, or a ruggedized buffer intended for a specific application such as the NASA Space Station. An artist's concept of the fully populated 12-disk terabit optical buffer is shown in Fig. 9. The less than one cubic meter volume will contain multiple, identical data processing electronics that are well suited

for LSI implementation. They will control the functions necessary for positioning the write/read heads and providing internal data synchronization. Table III compares the performance projected for this buffer with that of the well-known IBM 3380 magnetic disk system. As can be seen from this comparison, the emergent optical disk approach begins with a 25:1 advantage in capacity and over a 60:1 advantage in data rate compared to a relatively new commercial offering of the mature magnetic disk drive technology.

Data access time must include both file access time and data transfer time. Because the Optical Disk Buffer is configured to provide a very high data rate from a continuous spiral track, a data access time advantage is gained whenever data blocks greater than 1 megabyte are to be buffered. Figure 10 charts total data block access time against file size for the projected 14-inch magnetic disk drives of the future, and for the Optical Disk Buffer under development. Magnetic disk drives are available today with transfer rates of 3 Mbps, with projected growth to 10 Mbps.[5] Therefore, the total access time comparisons for the 1990 timeframe graphed in Fig. 10 apply for magnetic drives at 6 and 12 Mbps (to widely bracket the potential transfer rates), with initial data access in 33.3 ms (25 ms to track, plus a half-revolution latency of 8.3 ms). The Optical Disk Buffer total access time is graphed at the projected 200 Mbps (1.6 Gbps) transfer rate with 100 ms initial access time.

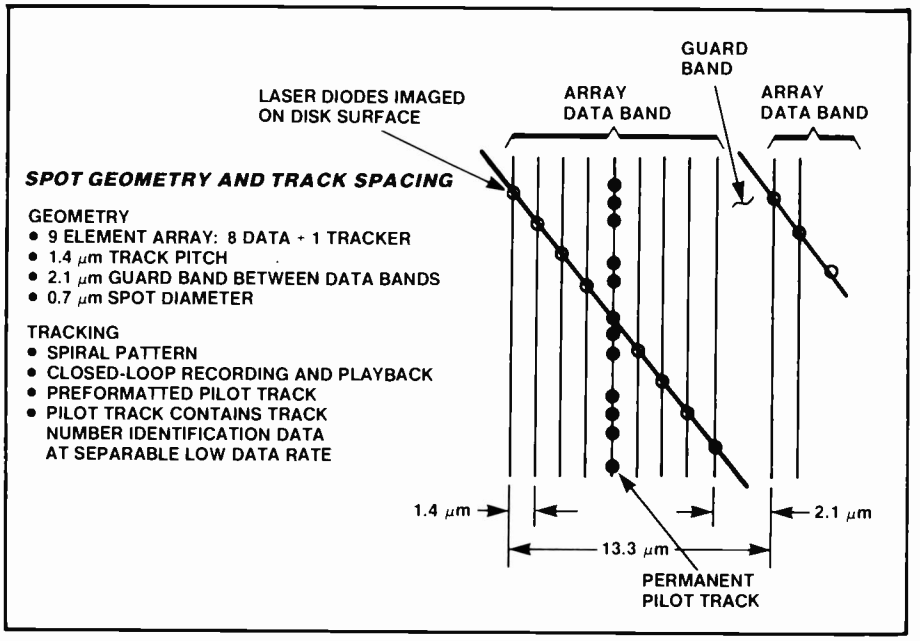

Fig. 7. *Buffer track geometry.*

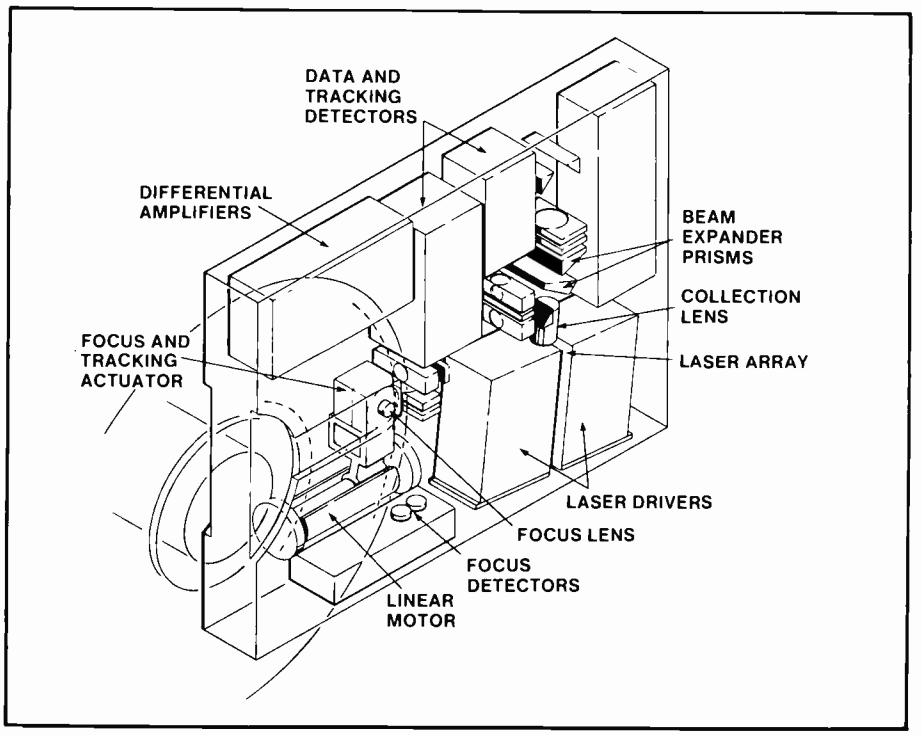

Fig. 8. *E-O module subassembly components.*

High-data-rate optical tape recorder

Development of a high-data-rate, high-capacity tape recorder with the specifications and performance goals shown in Table IV was initiated during 1985 for applications that do not require rapid, random access to the stored data. The approach is to combine the high areal bit packing density capability of optical tape recording techniques with the high volumetric bit packing density provided by the tape format in a recorder like that shown conceptually in Fig. 11. The basic recorder building block is a 100-Mbps electro-optic module that includes an eight-element diode laser array and an acousto-optic beam deflector. The developed recorder will, therefore, be applicable to data rates ranging from 100 Mbps to 1.6 Gbps in 100-Mbps increments.

One of the prime objectives of the recorder concept is to handle high data rates at a moderate tape speed. This is accomplished by demultiplexing the input data to as many as 16 electro-optic modules. The data is further demultiplexed to eight channels within each module to obtain up to 128 channels for a 16-module system. A tape speed of approximately 325 inches per second is required for the projected along-track data packing density if simple longitudinal recording is used to accommodate a 100-Mbps output from each module. Therefore, the output of each diode within the eight-element array is passed through an acousto-optic beam deflector to produce the needed tape scan velocity with the tape advanced at a maximum projected rate under 14 inches per second to position the tape from scan to scan.

Close contact is being maintained with potential tape suppliers to help assure that the recorder will be compatible with pro-

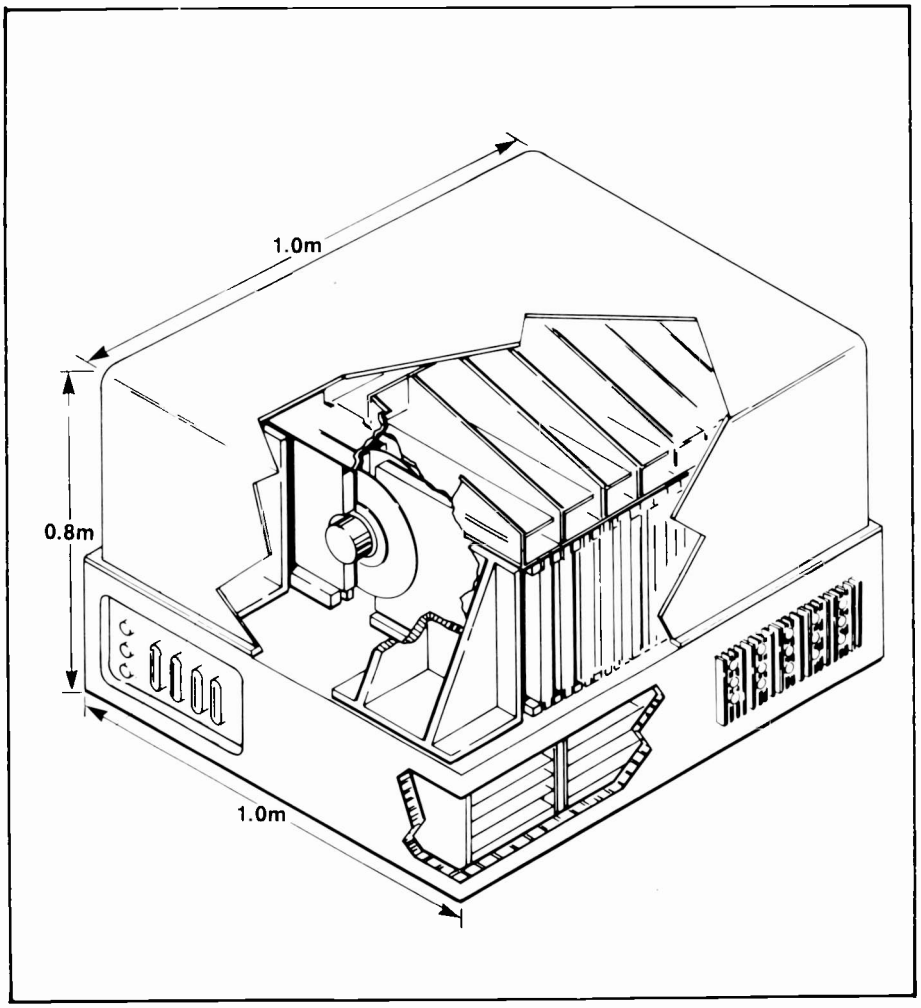

Fig. 9. *Fully populated optical disk buffer.*

Table III. *Comparison of projected optical disk buffer performance with commercial magnetic disk system*

User data density		
IBM 3380 (standard)	6.7×10^8 bits/surface	
IBM 3380 (double capacity)	1.34×10^9 bits/surface	
Optical disk buffer	4×10^{10} bits/surface	
System performance comparisons		
IBM 3380	*Optical disk buffer*	*Optical disk buffer advantage*
30 surfaces	24 surfaces	
2.5 Gbps=20×10^9 bits (standard)	1000×10^9 bits	50/1
5 Gbps=40×10^9 bits (double capacity)	1000×10^9	25/1
3 Mbps/sec=24×10^6 bits/sec	1600×10^6 bits/sec	67/1
379,808 bits/rev/track	1,081,344 user bits/rev/diode	3/1

duction media. Two critical parameters are tape thickness and dimensional stability. A tape base thickness of about 1 mil is required to meet the volumetric storage capacity goal. The tape must be dimensionally stable at this thickness to assure that data is recovered at an acceptable error rate after many wind/rewind cycles. The media development effort has concentrated on write-once media, but some consideration has been given to erasable media.

The optical tape is wrapped around a 16-inch diameter capstan supported by a hydrostatic air bearing to produce a very smooth, stable recording surface. Since there is a large tape-to-capstan surface contact area, virtually no slippage occurs. This minimizes tape mistracking as seen by the optical recording heads so that small guard bands may be used as an important step toward increasing the recorded data packing density. Since the capstan is very large in diameter and the tape speed is moderate, the angular velocity of the capstan is only 16.5 rpm. The capstan will be belt driven to permit a higher, more stable motor speed and to decouple motor speed variations from the capstan.

Dust and airborne particles are detrimental to successful recording and playback. Particles on the front side of the tape obscure the media from the laser beam since the media cannot have a thick overcoat to act as a dust defocus layer while maintaining an overall thickness of about 1 mil. Also, larger particles on the capstan surface will "print through" the thin media to the recording surface and cause a focus error. To eliminate these problems, the entire tape transport region will be sealed and continuously purged with filtered air.

The recording head for each module is composed of an objective lens, focus mechanism, and tracking mechanism. This assembly, which is about 1.5 inches in diameter, must be positioned very close to the tape due to the short working distance of the objective lens. Therefore, the optical modules fan out radially from the capstan, as shown in Fig. 11, with each module servicing a longitudinal swath of the tape about 16 mils wide. The lateral displacement of the adjacent modules and the use of acousto-optic beam deflectors with each electro-optic module results in the tape track data pattern shown in Fig. 12.

The output from a separate laser located on each electro-optic module is passed through the same objective lens as the

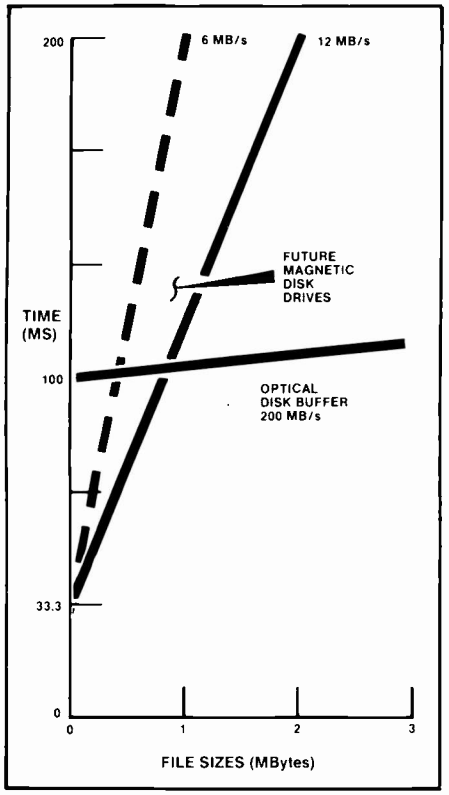

Fig. 10. *Data access and transfer time.*

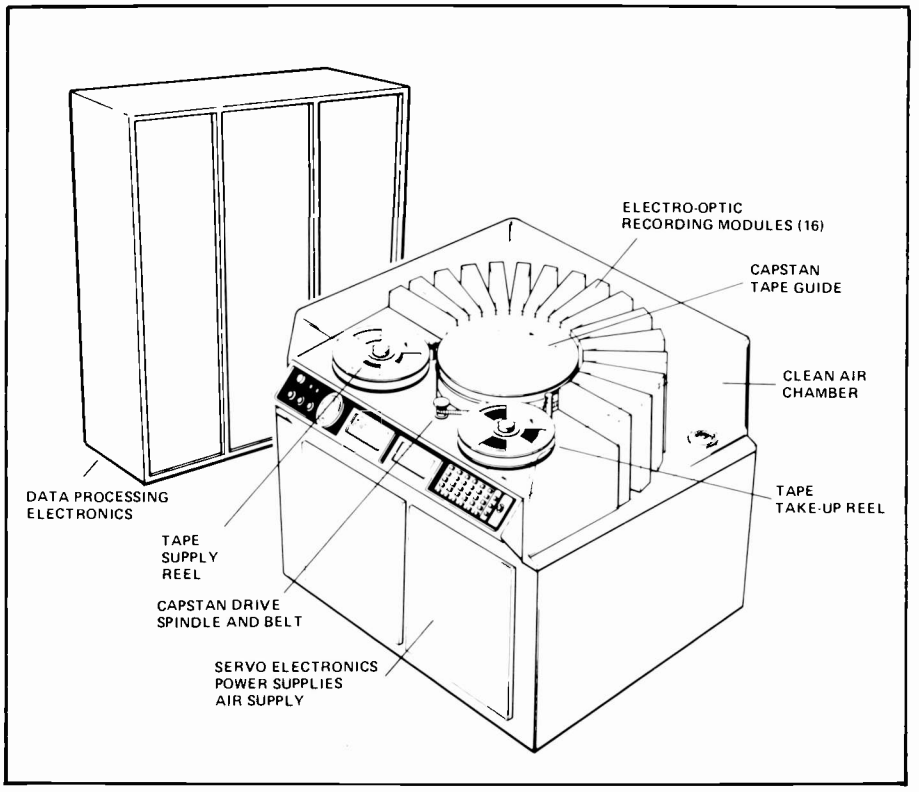

Fig. 11. *Artist's rendition of optical recorder transport.*

eight data channel beams to write a longitudinal control track. This is done to assure that the position of the control track with respect to the data tracks will be maintained during record and playback. During playback, the control track provides track position information to the tracking servo so that the data tracks can be properly scanned. Tracking is accomplished by moving the objective lens within the two-axis (track and focus) mounting assembly. The control track also provides timing information to the acousto-optic beam deflector in playback to keep the deflector scan properly phased to the recorded tracks.

The high-data-rate optical tape recorder brassboard will be developed over a two-year period, followed by the two-year development of an engineering prototype. The brassboard will include a single electro-optic module to demonstrate system feasibility using several hundred feet of tape media provided by one or more suppliers. The prototype will include multiple electro-optic modules to determine guard band requirements and the level of interaction between module swaths. Under the current schedule, successful demonstration of the prototype can be achieved in the 1989 timeframe, with production units available for Space Station ground support and other applications in the early 1990s.

Table IV. *Optical tape recorder specifications and performance goals*

User data rate	1.6 Gbps
Data capacity	2×10^{13} bits
Packing density	4.75×10^{8} bits/in^2
EO module data rate	100 Mbps/sec
Number of EO modules	16
Tape speed	13.8 inch/sec
Tape width	1/2 in
Tape length	7,000 feet
Tape thickness	1 mil
Reel diameter	11.5 in

Conclusions

By building on the developments of the last ten years, RCA is achieving very high performance optical storage systems that will capture, store and disseminate the vast amounts of data that will be generated by advanced sensors and computing systems in the late 1980s and beyond. The projected need for even higher areal-storage capacity using shorter wavelength solid-state sources and multiple layer recordings, and faster access times, is also being addressed to assure that RCA will continue to be the principal developer of advanced optical storage technology in the following decades.

Acknowledgment

The work presented in this paper has been sponsored for several years by the Rome Air Development Command, Griffiss Air Force Base, Rome, New York; NASA Marshall Space Flight Center, Huntsville, Alabama; and other Government organizations. Key Government contributors to the development of the optical recording technology have included A. Jamberdino and J. Petruzelli of RADC, and K. Wallgren and D. Thomas of NASA.

The optical recording technology developments represent the dedicated efforts of many talented RCA personnel from the Advanced Technology Laboratories, the

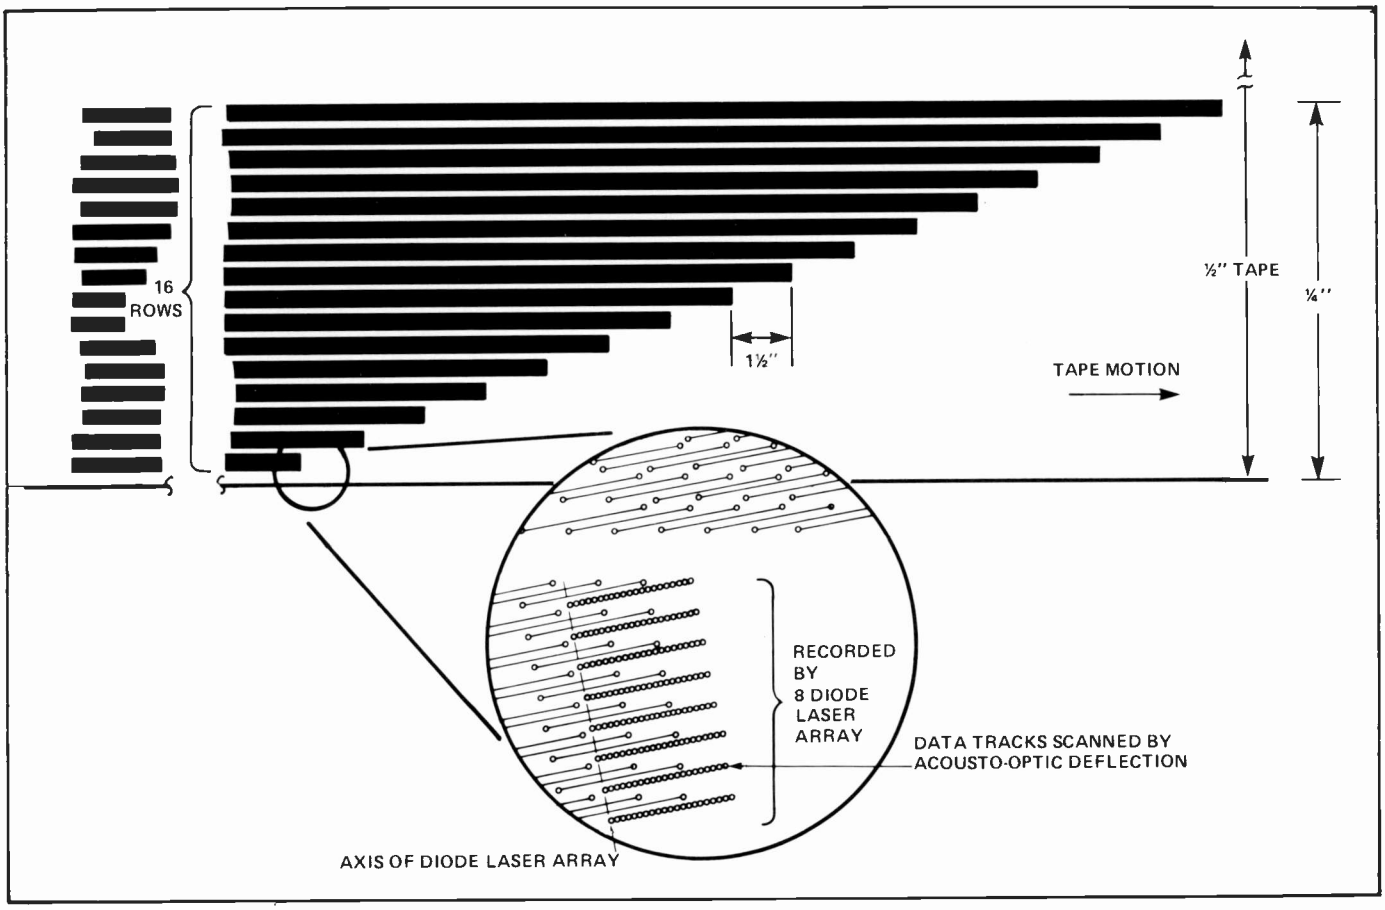

Fig. 12. *Tape track data pattern.*

Communication and Information Systems Division, and RCA Laboratories.

References

1. G.J. Ammon and J.A. Calabria, "Operational Performance of Optical Disk System," Third International Conference on Optical Mass Storage, *SPIE* Vol. 529, Los Angeles, Cal., January 22-24, 1985.
2. G.M. Claffie, "High Performance Optical Disk Jukebox," ibid.
3. D.B. Carlin et al., "Multichannel Diode Laser Array Optical Recording," *Proceedings,* OSA Topical Meeting on Optical Data Storage, Monterey, Cal., April 1984.
4. M.L. Levene, "A High Data Rate, High Capacity Optical Disk Buffer," Seventh IEEE Symposium on Mass Storage Systems, Tuscon, Ariz., November 4-7, 1985.
5. C. Panasuk, "Technology Report, Advanced Computer Technology: Mass Storage." *Electronic Design,* May 3, 1984, pp. 263 and 266.

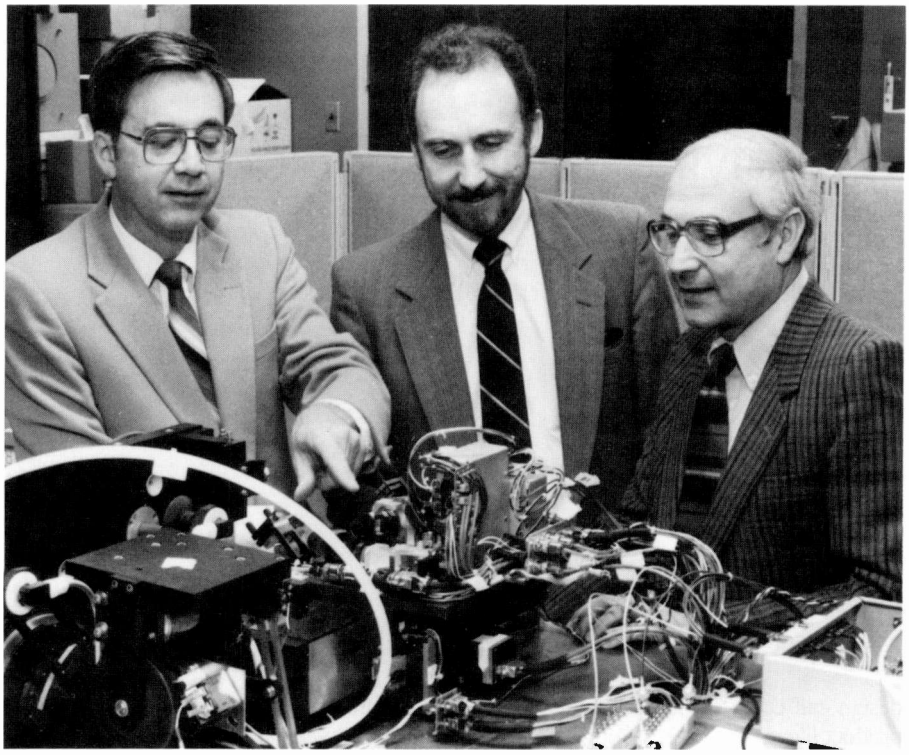

Authors, left to right: Claffie, Altman, Levene.

Wolf Altman is a Manager of Marketing Development at ATL. He is responsible for marketing planning and follow through in the area of optical data storage systems. Prior to joining ATL in 1984, he was a Principal Member of the Technical Staff at Astro-Electronics Division. He worked on laser remote sensing systems and solid-state and CO_2 laser technology. He has a BS in Physics from Drexel University and an MSEE from the University of Pennsylvania. He has completed graduate work toward a PhD at the University of Maryland.
Contact him at:
Advanced Technology Laboratories
Moorestown, N.J.
Tacnet: 253-6423

Gerald Claffie is the Manager of ATL's Optical Storage Laboratory, where he is responsible for the development of optical disk technology. Prior to this assignment, he held marketing, engineering management, and engineering design positions that addressed the concept design and/or advanced development of electro-optical and electro-mechanical systems for data collection, data processing, and mass data storage applications, He has a BSEE from Clarkson University and an MSEE from the University of Pennsylvania.
Contact him at:
Advanced Technology Laboratories
Moorestown, N.J.
Tacnet: 253-6484

Martin Levene is a Program Manager in ATL's Optical Storage Laboratory. He is currently responsible for the Optical Disk Buffer and Optical Tape Recorder contracts. Previously, he was responsible for programs in the areas of advanced optical disk recording, computer-aided mechanical analysis, and thermal management. He received BSME and MSME degrees from the University of Pennsylvania, where he has also completed the course work toward the PhD in Mechanical Engineering. He is the recipient of a David Sarnoff Award for Outstanding Technical Achievement.
Contact him at:
Advanced Technology Laboratories
Moorestown, N.J.
Tacnet: 253-6489

D.E. Britton

Verlangen: Formally verifying system designs

This set of tools can be used to prove that designs of computer systems meet their requirements.

Do we really need to invest our effort in formally verifying system designs? In part, the answer relates to how important it is for a system to function correctly. For critical systems, verified design can increase our confidence that a system will meet its requirements. Another benefit of verified design relates to the often enormous cost of system development. Verifying a system design can catch design flaws early in the product life cycle, while the cost of correcting them is relatively low.

For RCA Aerospace and Defense and other defense contractors there is another motivation, too. In August 1983, verified design was thrust into a position of special importance when the Department of Defense published its "Trusted Computer System Evaluation Criteria",[1] a document that defines criteria for evaluating how much trust to give to secure computer systems. Depending on security policies and assurances, a secure computer system can be rated as low as "D," minimal protection, to as high as "A1," verified design. The Department of Defense is currently using these criteria to rate existing systems, such as the Honeywell SCOMP,[2] and to provide rating requirements for systems yet to be implemented.

We have developed Verlangen for producing verified designs as required in the A1 category, but have not restricted Verlangen to verifying security requirements. Although several other languages and systems can specify and verify system designs, Verlangen has special strengths not found elsewhere.[3,4,5,6,7]

Verlangen encourages separating system designs and their verifications into tractable units. To do this, the language provides "classes" to support object-oriented design, and "levels" to support levels of refinement. These language features make Verlangen specifications readable and keep verification manageable. When a change in system design or requirements necessitates a change in a Verlangen specification, most theorems and proofs will remain unchanged, requiring relatively little added effort to verify the changed specification.

Verlangen is perhaps unique in its applicability to truly concurrent systems, such as distributed systems and networks. The model for communications between subsystems is extremely flexible, allowing a great variety of synchronization schemes to be specified. This contrasts Verlangen with the Gypsy Verification System.[3] Although Gypsy is commonly chosen for verifying network and distributed applications, it imposes a very restrictive message-based communications model.

The desire to verify designs of secure communication networks and distributed systems has played an important role in the development of Verlangen. We have used Verlangen to specify and verify the design of a communications network with end-to-end encryption (Fig. 1). All messages that travel over the network are encrypted to prevent their being tapped, and only hosts that are authorized to communicate exchange unencrypted messages. This application is a simplification of an example described in another paper.[8]

In another application (Fig. 2), a multilevel secure local area network (Secure LAN) connects several workstations that do not deal with or understand security levels. Each workstation is assigned a fixed security level and operates only at that level. Between the workstations and the local area network are guards, one for each workstation. To enforce multilevel security, the guards restrict the flow of messages between workstations. Examples

Abstract: *"Verlangen" is a German verb that means "to require." It is also, appropriately, the name of a set of tools for formally specifying system designs and verifying that they meet requirements:*

- *Specification language—provides high-level language features for specifying system designs and requirements.*
- *Compiler—translates a specification into definitions and theorems in first-order logic.*
- *Theorem prover—proves theorems to verify that a design satisfies requirements.*

Verlangen is appropriate for many kinds of computer systems, including distributed systems, communications networks, and operating systems. Eventually, a translator will be added for implementing verified designs in software. Verlangen is being developed by the Software Engineering Laboratory of RCA Advanced Technology Laboratories.

Final manuscript received November 20, 1985
Reprint RE-31-1-8

in this paper use the Verlangen specification for this application.

Although Verlangen is not a programming language, it includes many programming language features, such as block structure, identifier scope and visibility rules, and user-defined data types. These features are equally valuable for expressing program and design specifications.

Classes

Verlangen uses a language construct called CLASS to support object-oriented design and verification. Its concept of class combines the concept of abstract data type from programming languages with the concept of state machine from specification languages. An abstract data type defines a set of values and provides functions (or operations) that yield new values in the set from old ones. A state machine goes through states or cycles to do its job. For example, a washing machine cycles through fill, wash, spin and rinse. A Verlangen class is an abstract data type whose values represent the states of a state machine.

Since most system entities can be modeled as state machines, the class is a very useful and general construct. We use classes to represent data structures, processes, hosts, front-ends, communication links, communication networks, etc. We define a new class whenever we want to represent a new kind of system entity. If a system includes several similar entities that operate more or less independently—for example, the guards in the Secure LAN example—we define a class for that kind of entity and specify that there are several instances of that kind of entity in the system. A class instance appears in a Verlangen specification as a variable whose data type is given by the name of the class.

Each class instance represents a particular state machine and is called an object. Associated with an object is a sequence of values (states), called a history. The first value in the history represents the state machines's initial state, and the other values represent each subsequent state. To define the possible histories, a class definition specifies initial values for its objects (or a condition on initial values) and some operations that yield new values from old ones.

How do we use the Verlangen class construct? Consider the guard unit from the secure LAN. Each guard, which is assigned the security level of its workstation, has two functions. The guard uses its security level to label all data that goes

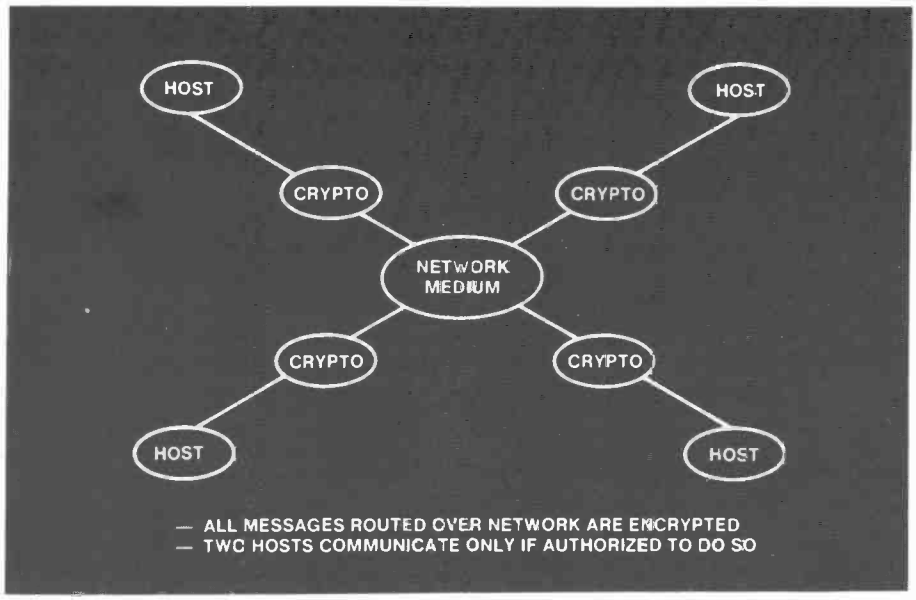

Fig. 1. *A network with end-to-end encryption. Only encrupted messages traverse the network. Two hosts communicate only if they are authorized to do so.*

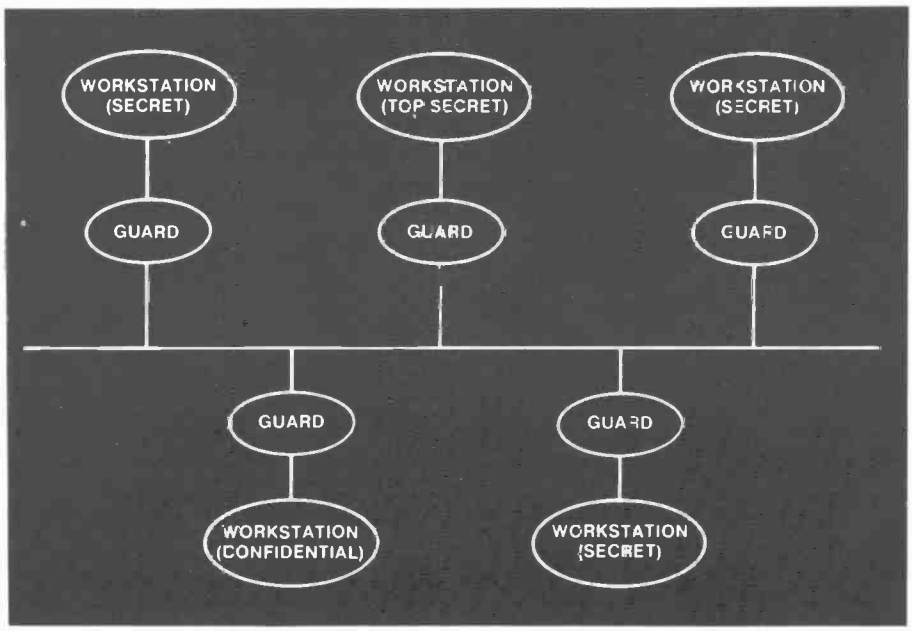

Fig. 2. *A secure local-area network. Information does not flow from a workstation with a low classification to one with a higher level.*

from its workstation to the network. The guard also prevents data with a higher security level from reaching the workstation.

We model each guard as a state machine and define a Verlangen class called *Guard* (Fig. 3) that specifies a class of state machines—all the guard units. The guards are not identical, since their security levels may be different. So the *Guard* class definition has a parameter: a constant *max* of type *Lev*. Each *VAR* declaration for a variable of type *Guard* will assign a value to *max*.

Each *Guard* value represents a state. The variables, declared by *VAR*, represent the components of the state. These components have types *Lev*, *Subj*, and *Obj* representing, respectively, security levels, names of individual workstations, and classified data objects.

The procedures in the *Guard* class definition represent operations that take the guard from one state to the next. Verlangen procedures include a precondition and an effect. The precondition and the effect are both expressed as one or more logical formulas. The precondition selects those

```
CLASS Guard(CONST max:Lev) IS
  VAR netin,netout:Obj;
  VAR dest:Subj;
  VAR netinfull,netoutfull:BOOLEAN
        INITIALLY FALSE;

  PROCEDURE fromnet(o:Obj; l:Lev) IS
    PRECONDITION NOT netinfull;
    EFFECT
      IF l <= max
        THEN (netinfull' & netin' = o)
        ELSE NOT netinfull';
      SAME netoutfull, netout, dest;
  END fromnet;

  PROCEDURE touser(o:Obj) IS ...
  PROCEDURE fromuser(d:Subj; o:Obj) IS ...
  PROCEDURE tonet(d:Subj; o:Obj; l:Lev) IS..

  INVARIANT netinok IS
    netinfull =>
      EXISTS g:Guard EXISTS l:Lev
      ( g << THIS &
        g.fromnet(netin,l) &
        l <= max );

  INVARIANT filter USING netinok IS
    FORALL o:Obj
    ( touser(o) =>
        EXISTS g:Guard EXISTS l:Lev
        ( g << THIS &
          g.fromnet(o,l) &
          l <= max )
    );
  ...

END Guard;
```

Fig. 3. *Guard class declaration from the Secure LAN specification.*

states in which the operation may occur; that is, the operation may occur in a state only if the precondition is met. The effect specifies the state after the operation in terms of the state before the operation. Primed identifiers, (e.g., *netin'*) refer to values after the operation; unprimed identifiers refer to values before the operation.

To express requirements on a class, we include assertions in the class definition. There are two kinds of assertions: invariants and constraints. An invariant is a condition that we want every state of every object in the class to satisfy. A constraint is a condition that we want to hold between every two subsequent states in every object history.

The *Guard* class defines two invariants, *filter* and *netinok*. The *filter* invariant states that the guard passes data to its workstation only if the data came over the network from a workstation with the same or a lower security level. We may regard the *filter* invariant as a requirement placed on the class, a property that will be maintained regardless of the class implementation. It refers only to the procedures and constant parameters of the class, not to the internal variables.

For every assertion that appears in a class definition, the Verlangen compiler produces verification theorems. Proving the verification theorems verifies that the design does indeed satisfy the assertions. For example, we verify an invariant by induction. So for an invariant the compiler produces theorems that correspond to the basis and induction steps of an inductive proof:

Basis: The initial state satisfies the invariant.

Induction: If an arbitrary state satisfies the invariant, then the next state does also.

Often an invariant is not inductive—not strong enough for the inductive proof to succeed. Then, to obtain a verification, we determine additional supporting invariants and include them in the specification. When these invariants appear in a USING clause of a non-inductive invariant, the compiler adds them as hypotheses to that invariant's verification theorems. For example, *netinok* supports *filter,* which is not inductive.

Concurrency

Verlangen allows us to decompose a system into (or compose a system from) simpler subsystems. This approach to system design is generally accepted as effective for operating systems. For distributed systems and communications networks, the approach is also a natural one. The system naturally decomposes into a set of concurrent, interacting subsystems—the host computers, front-end processors, gateways, etc.

We model a system composed of subsystems by a collection of state machines, and define a class for each different kind of state machine. In the class definition for the overall system, we declare variables that represent the component subsystems. The data types of these variables are the classes for the corresponding state machines. This specifies that the state of the overall system is composed of the subsystem states.

For example, the Secure LAN specification defines three classes: *WorkStation, Guard,* and *System* (Fig. 4). The class *System* represents the overall system, composed of several workstations and guards. In the definition for *System*, the workstations and guards appear as variables of type *WorkStation* and *Guard*, respectively.

Verlangen uses SYNC statements to specify how concurrent subsystems are to be synchronized. These statements correlate events (operations) that occur in the subsystems. A SYNC statement says that certain events in the subsystems cannot occur unless they occur together.

The *System* class definition in our example includes SYNC statements that state how the workstations interact with the guards, and how the guards interact with each other. For example, a user workstation sends data only if its guard receives it, and vice versa.

When a specification consists of several classes, the verification of each class is carried out independently. Supporting invariants may come from outside of a class, however.

For example, in the Secure LAN specification we included an invariant called *Origination* in the *System* class definition. This invariant asserts that any data object a user workstation knows was created by a user workstation on the network. This rules out, for example, a design where the guards spontaneously create data objects of their own. The *Origination* invariant is supported by invariants of the *WorkStation* and *Guard* classes.

Levels of refinement

Verlangen uses successive levels of refinement to support design. That means we can write a Verlangen specification as one or more ordered levels (Fig. 5), each a complete specification of the whole system. The first (or "top") level presents the most abstract view of the system. Each successive level presents a more concrete specification than the preceding one, and includes a map that specifies how it relates to its predecessor. To verify a specification that consists of more than one level we show that the individual levels are self-consistent and that neighboring levels are consistent with each other.

Usually, we organize a two-level specification so the top level represents a set of requirements placed on the system and the bottom level represents the system design. This approach ideally allows us to use a set of requirements—for example, a model of multilevel security—over and over again with different system designs.

The Secure LAN example is a two-level Verlangen specification. The top level specifies a model of multilevel security, and the bottom level specifies the design of the secure local-area network. The Verlangen fragments in Figs. 3 and 4 came from the bottom level.

Verification

The Verlangen compiler translates a Verlangen specification into a collection of definitions and theorems in typed first-order logic. The theorems state that classes satisfy their assertions, and neighboring levels are consistent. By proving the theorems from the definitions, we verify that the specified system design satisfies the specified system requirements. Figure 6 shows theorems the Verlangen compiler produced to verify the filter invariant of the *Guard* class.

The translation of a class definition declares a type that has the same name as the class. The type's values represent the states of the state machines that the class represents. The translation also defines a relation *precedes* and a function *next*. *Precedes* defines a partial ordering on values of the type, and *next* is a successor function that satisfies *precedes(x,next(x))*. When we give *next* a state (in the history of a state machine), *next* yields the next state in the history.

A variable declared within a class definition translates into a state function. Given a state, the function returns the value of the variable for that state. A constant also translates into a state function. An invariant translates into a state predicate. Given a state, the predicate determines whether the state satisfies the invariant. Initial conditions, constraints, and procedures all translate into state predicates. A SYNC statement translates into an axiom stating

```
CLASS WorkStation(CONST myself:Subj) IS ...

CLASS Guard(CONST max:Lev) IS ...

CLASS System IS
  CONST Clearance(s:Subj):Lev;
  VAR   user(s:Subj):WorkStation(s);
  VAR   guard(s:Subj):Guard(Clearance(s));

  FORALL s,d:Subj FORALL o:Obj
    SYNC user(s).send(d,o),
         guard(s).fromuser(d,o);
  ...

  INVARIANT Origination
        USING WorkStation.knowledge,
              Guard.filter,
              Guard.transport
  IS FORALL s:Subj FORALL o:Obj
     ( user(s).knows(o) =>
         EXISTS sys:LANSystem
         ( sys << THIS &
           sys.user(Originator(o)).write(o)
         )
     );
  ...
END System;
```

Fig. 4. *Excerpt from the Secure LAN specification defines concurrent classes.*

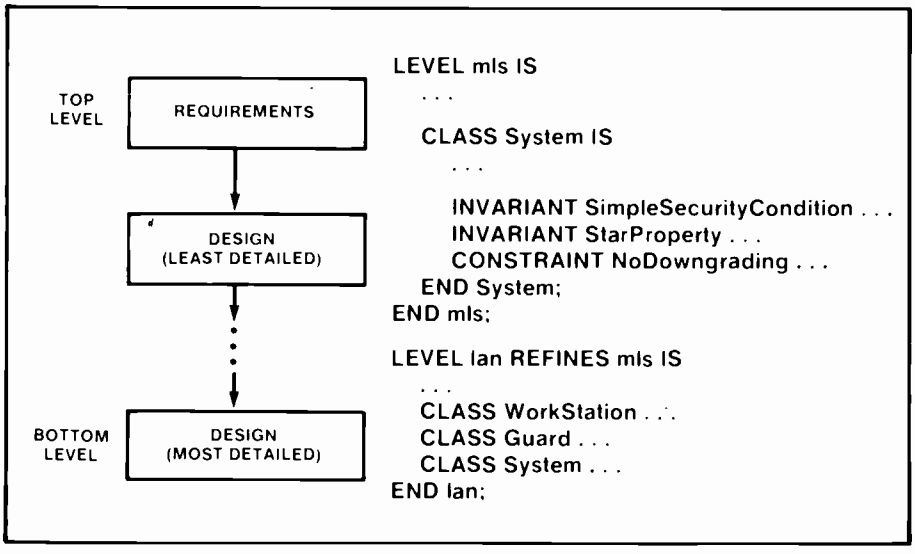

Fig. 5. *Levels of refinement.*

```
PROVE FORALL this:Guard FORALL max:Lev
      ( initial(this,max) &
        FORALL old:Guard netinok(old)
          => filter(this) )

PROVE FORALL this:Guard FORALL o:Obj FORALL l:Lev
      ( filter(this) &
        FORALL old:Guard netinok(old) &
        fromnet(this,o,l)
          => filter(NEXT(this)) )

PROVE FORALL this:Guard FORALL o:Obj
      ( filter(this) &
        FORALL old:Guard netinok(old) &
        touser(this,o)
          => filter(NEXT(this)) )

PROVE FORALL this:Guard FORALL d:Subj FORALL o:Obj
      ( filter(this) &
        FORALL old:Guard netinok(old) &
        fromuser(this,d,o)
          => filter(NEXT(this)) )

PROVE FORALL this:Guard FORALL d:Subj FORALL o:Obj FORALL l:Lev
      ( filter(this) &
        FORALL old:Guard netinok(old) &
        tonet(this,d,o,l)
          => filter(NEXT(this)) )
```

Fig. 6. *Theorems for verifying filter invariant, taken from the Secure LAN specification.*

the equivalence of the state predicates that represent the synchronized procedures.

Mappings between levels are generally represented by axioms that relate the entities of one level to another.

Formal verification at ATL

To meet the verification needs of RCA Aerospace and Defense, RCA Advanced Technology Laboratories supports a formal verification skill center in its Software Engineering Laboratory (SEL). Verification systems available to the skill center include the Gypsy Verification System[3] (from the University of Texas) and Verlangen.

Verlangen development is an ongoing effort of the skill center. The Verlangen compiler has been implemented in the Pascal programming language and runs on a Digital Equipment Corporation VAX under the VMS operating system. The Verlangen theorem prover comes from the commercially-available Verus verification system,[9] and runs in the same environment as the compiler. We are presently adding features to the specification language that improve its expressive power and extend the range of properties that can be verified.

Eventually, a translator will be added to the Verlangen toolset that will produce a separate translation from a Verlangen specification into the Department of Defense's Ada programming language code. This code will be a "skeleton" of a software implementation of the specified system. A skeleton is an incomplete implementation; it substitutes assertions (imbedded in formal comments, as in Anna[10]) for omitted code. To obtain a complete implementation that meets the system requirements, a programmer adds Ada code that satisfies the imbedded assertions.

Dianne Britton is a Senior Member of the Engineering Staff in ATL's Software Engineering Laboratory. She is working in the area of formal software and system verification, with emphasis on design and multilevel security verification. She developed Verlangen for specifying and verifying designs of systems and networks. At RCA Laboratories, she designed and built a protoytpe operating system kernel for distributed applications, and developed algorithms for maintaining consistent, fully-replicated databases in a distributed system. Dianne received an MS and PhD in Computer Science from the University of Arizona.
Contact her at:
Advanced Technology Laboratories
Moorestown, N.J.
Tacnet: 253-6654

References

1. Department of Defense Computer Security Center, *DoD Trusted Computer System Evaluation Criteria*, CSC-STD-001-83, August 1983.
2. Fraim, "SCOMP: A Solution to the Multilevel Security Problem," *Computer*, Vol. 16, No. 7, July 1983.
3. D.I. Good, R.M. Cohen, C.H. Hoch, L.W. Hunter, and D.F.Hare, "Report on the Language Gypsy Version 2.0," *ICSCA-CMP-10*, The University of Texas at Austin, Revised September 1978.
4. K.N. Levitt, L. Robinson, and B.A.Silverberg, *The HDM Handbook*, Vols 1-3, Computer Science Lab., SRI International, Menlo Park, California, June 1979.
5. R. Kemmerer, "FDM—A Specification and Verification Methodology," Proceedings of the Third Seminar on the DoD Computer Security Initiative Program, National Bureau of Standards, Gaithersburg, Maryland, November 1980.
6. R.R. Musser, "Abstract Data Type Specification in the AFFIRM System," *IEEE Trans. on Software Engineering*, SE-6,1 January 1980.
7. D. Craigen, "Ottawa Euclid and EVES: A Status Report," Proceedings of the 1984 Symposium on Security and Privacy, Oakland, California, May 1984.
8. D.E. Britton, "Formal Verification of a Secure Network with End-to-End Encryption," Proceedings of the 1984 Symposium on Security and Privacy, Oakland, California, May 1984.
9. Verus User Documents, Compion Corporation, a subsidiary of Gould Inc., 1984.
10. D.C. Luckham, F.W. von Henke, B. Krieg-Brueckner, O. Owe, "ANNA: A Language for Annotating Ada Programs," Preliminary Reference Manual, TR 84-261, Computer Systems Laboratory, Stanford University, July 1984.

G.W. Mebus | L.T. Armstrong | H.Rosenthal

Ada® software development tools

Ada, the DoD's programming language of the future, solves many software development problems, but also creates some problems of its own. The ATL Software Engineering Laboratory is developing some tools to handle them.

Abstract: *The Department of Defense's Ada programming language easily manages the complexities of large military software development. But its own complexity prevents people from understanding or maintaining Ada code. To solve these problems, the Advanced Technology Laboratories (ATL) is developing a set of software development tools for the Ada language. They consist of a frontend processor that analyzes Ada source code for handling by a cadre of specialized backend processors as well as a debugger/tracer. The tools are highly portable and easily targeted to different computer architectures.*

® Ada is a registered trademark of the U.S. Department of Defense, Ada Joint Program Office.

Final Manuscript received November 20, 1985
Reprint RE-31-1-9

The Ada language is a high-order computer language (HOL) of many firsts. It is the first major language named after a woman—the first programmer, Augusta Ada Byron (1815-1851), Countess of Lovelace, and Lord Byron's daughter. It is also the first language with a name that the U.S. Department of Defense (DoD) has trademarked, the first language for which DoD tests must be used to validate compilers, and the first that the DoD has mandated for use in all major implementations.

As far as the DoD is concerned, the Ada language is definitely the software systems language of the future. Indeed, it is quickly becoming the language of today as more sophisticated Ada compilers are now beginning to appear.

Ada language

The DoD mandated the Ada language for their programs to reduce many of the development and maintenance costs as well as other problems of large military software development. (Characteristics of the Ada language, its associated technology, and its place among other programming languages were discussed previously in the *RCA Engineer*.[1,2])

The Ada language is primarily designed to support real-time, parallel processing and hardware-intimate, large, complex software development. Several principles for handling the complexity of software are the concepts of program structuring, data abstraction, information hiding, reusability of existing proven code, and strong typing. The Ada language supports these concepts, allows separate compilation of pieces of a large program, and provides for parallel processing of special software pieces (tasks).

Program structuring allows us to avoid describing "jumps" or "go tos" in the program logic. Instead, statements like

if <condition> then <action1>
else <action2>

and

while <condition> do <action>

allow us to better understand and convey the intent of our programs to other people and to the computers. The Ada language fully supports the normal set of program structures.

With *data abstraction* we can emphasize the important details without being concerned with the unimportant ones, while *information hiding* prevents users of the software from knowing program details that they should not be able to access.

The Ada language supports these capabilities in the Ada package concept: having a specification that tells users (people or programs) what information and processing is available, and a separate "body" that tells how data and processing are actually implemented.

Strong typing is the ability to specify your own data types (such as "complex" for complex numbers, or "weekday" for a list of day names from Monday through Friday) and the operations allowed to process data of those types. The Ada compiler ensures proper use of the data and operations, based on the defined data types. The package concept allows this checking, even when program pieces are compiled separately and when body parts of packages have not yet been written.

All these beneficial features of the Ada language are there to support program reliability. But, like all solutions to problems, they create some problems of their own.

For large programming systems, the Ada language supports developing the various pieces separately; so, one person does not need to know all names used in the overall program. In some instances, multiple use of the same name is allowed, a feature called overloading of names. The names may be "qualified" by the program pieces in which they occur, or by the type of data that involves them. Although the compiler

keeps straight the various meanings of a name, people who must maintain the program source code can become confused.

Abstraction and hiding can also confuse things for code maintainers. For example, a record structure's definition could be distributed over several code-compilation units (some parts may be hidden while others are not).

Another problem exists with the current compilers themselves. The major effort in developing Ada compilers has been to conform to the Ada Programming Language[1] standard, not with providing extra features to help programmers. One glaring deficiency is that most Ada compilers do not provide even simple cross-reference listings to tell where a name is declared or used in a program, so debugging Ada software is often a difficult job.

To attack these problems, the Software Engineering Laboratory (SEL) in the Advanced Technology Laboratories (ATL) is developing new Ada software tools. They include a frontend for the compiler, backend tools, and a machine-independent Ada language-level debugger/tracer. The SEL had been using the TeleSoft Ada compiler, and has installed its second generation Ada compiler, namely Digital Equipment Corporation's VAX®/VMS Ada package.

Compiler frontend

The SEL is currently developing a compiler frontend for the Ada/M language, which is a large subset of the Ada language. The Ada/M language excludes generics, tasking, fixed-point numbers, and separate compilation of program pieces. In the future, the computer frontend will be upgraded to the full Ada language.

A compiler frontend is a software tool that reads source code written in a high-level language (for example, the Ada language) and produces an intermediate-language representation of the code that is more easily processed by subsequent software tools. We designed the frontend so it can be used not only as a frontend for a compiler, but also as a frontend for various backend tools that need or use information about the input Ada source program.

Most code for the frontend is written in a special, limited language called PAL (Parser Assembly Language), that has less than twenty instructions. PAL is specifically designed for implementing recursive-descent, table driven compilers. The advantage of the PAL code is it adapts easily to another host computer. Because the language is so small, writing an assembler program to process the language is almost trivial. Hence, frontends written in PAL, such as the one for the Ada/M language, are highly portable to different computers.

The rest of the code for the frontend is currently written in Pascal, but we intend to convert it to the Ada language in 1986 using our second-generation compiler. This will enable us to "bootstrap" the compiler when we put the frontend on another machine. The non-PAL code encompasses what are known as the "resolvers" (they determine the precise use of a name), and additional utility routines (for example, allocation routines and error routines).

As an example of resolving, consider the following expression in the Ada language:

$$(x+y)\times 2$$

When processed, the input expression will contain only block pointers (addresses stored on a symbol table) for x and y. But, the Ada language permits us to reuse identifiers. Therefore, the resolvers will determine exactly what versions of x and y we mean and return pointers to blocks that contain the particular contexts x and y represent.

The rest of the non-PAL code is primarily for machine-dependent memory allocation and error handling routines. To promote portability, we decided to allocate space via an HOL rather than use the VAX/VMS operating system space allocation routines. Also, we wrote the error routines in an HOL because PAL does not have any input or output instructions.

The frontend produces a data structure graph that represents the parsed input program and is a completely resolved graph, both for context (names of entities) and semantics. We initially designed the data structures to permit a compiler backend to take the frontend's output and produce object code. Thus, the graph keeps intact all semantic information about the input source.

Therefore, this graph provides a good generic representation for a variety of backend tools needed to evaluate an Ada program. The backends currently under development are:

- An Ada/M source code regenerator (that will be upgraded to the full Ada language in the future).
- A database tool with query/report capability. It is specifically designed to give the user semantic and contextual information about the input Ada source program.
- An interactive Ada debugger/tracer that does not depend on which Ada compiler is used.

Source-code regenerator and "pretty printing"

The Ada/M regenerator will provide the basis for all future compiler backends. As input, it takes the graph produced by the Ada/M frontend and reproduces the source program. This is valuable because it helps us verify that we have captured the semantic information of the input program.

Because the source code regenerator has all the information needed to reconstruct the original source code, it also has enough information to reformat and standardize the source code as well as enforce the coding guidelines. The reformatting is done in such a way that indenting source code lines clearly displays the structure of the code. This "pretty printing" enhances code readability, an important feature when people have to understand or document a program. The regenerator will also be the basis for both generating target-machine code and developing the debugger/tracer described below.

Database and query/report capabilities

Because the Ada language is primarily intended for large-scale software systems, properly managing the data becomes more difficult as the system's size increases. Proper modularization restricts as much data as possible to a few routines but, in some cases, data must be available to a large percentage of routines.

As data must be tracked over many routines, we chose to incorporate an Ada database into our Ada tool set. This database will be integrated into a user-friendly database management system (DBMS). When used properly, this facility will enable a software designer to

® VAX is a trademark of the Digital Equipment Corporation.

manage large streams of information quickly and easily. For example, if you want to list all procedures in a package that reference a specific variable *X*, then you can do it with a simple command such as:

Print procedures with variable *X*

This information may not seem very useful. But suppose we suspected that an unwanted change to *X* was causing a problem we were having with a huge software system. Then, we would simply alter our command to:

Print procedures where variable *X* is changed

Moreover, this tool can be used for much more than simple variable information. It will permit users to generate properly formatted hardcopies (reports), and thus quick-reference information about their large-scale software systems. This can be particularly helpful when dealing with program characteristics that are unique to the Ada language, such as user-defined overloading and generics.

To implement this concept, two approaches were considered. The first was to incorporate a DBMS package that already exists into the interactive part of the Ada tool set. This allows a great degree of generality and a user-friendly environment because these packages are designed for an end user to work in an English language-like environment with speed and efficiency.

The second alternative was to write a set of specific database routines for the existing backend code. This would allow the user to work in an environment specifically tailored to his or her needs. But, it would probably require the generation of huge amounts of code for each application.

For the Ada tool set, the SEL chose the first alternative because it satisfied requirements while minimizing risk. We are currently using the DATATRIEVE[4] package that allows a software designer to integrate database manipulation facilities with programs written in higher-level programming languages. This is achieved through the use of the DATATRIEVE access block (DAB), which is a series of type, variable, and external procedure declarations to which a program must have access. Once communication has been established, either the DATATRIEVE software or the program may handle the flow of control.

Authors, left to right: Mebus, Armstrong, and Rosenthal.

George Mebus is Unit Manager of the Languages and Compilers group in ATL's Software Engineering Laboratory. The work of this group includes creation of Ada development tools, software support for advanced pipelined processing architectures, parallel processing, and formal verification techniques and tools for verifying security and correctness of distributed communications systems and secure operating systems. Prior to coming to RCA, he worked with computer and software development systems for 24 years at the Naval Air Development Center NADC. George received a BSEE degree from the University of Pennsylvania, and an MSEE degree from the University of Michigan. He completed course work toward a PhD in Computer and Information Science at the University of Pennsylvania.
Contact him at:
Advanced Technology Laboratories
Moorestown, N.J.
Tacnet: 253-6634

Leonard Armstrong is an Associate Member of the Engineering Staff in ATL's Software Engineering Laboratory. For the Ada tool set development program, he has updated existing parser software to increase compilation speed, and he has developed both the Ada/M source code regenerator and a set of resolver routines needed for the Ada/M compiler to determine the context of names. For the CMOS MIPS (silicon advanced pipeline processor) program, he developed an assembly code translator to aid in cache memory studies. He received his BS in Computer Science from Saint Joseph's University, and is currently working on an MS in Computer Science at Drexel University.
Contact him at:
Advanced Technology Laboratories
Moorestown, N.J.
Tacnet: 253-6635

Harry Rosenthal is Manager of ATL's Software Engineering Laboratory, and is responsible for the development of advanced software engineering tools, transfer of software technology, and computer operations and administration. Areas of technical focus include validation/verification (both formal and informal), modeling/simulation, compilers, operating systems, project/product management, database management, and systems analysis. Before joining RCA, he spent twenty years in DoD consulting, primarily in the research and development sector, and also worked for two computer manufacturers. Mr. Rosenthal holds MA and BS degrees in Mathematics from New York University.
Contact him at:
Advanced Technology Laboratories
Moorestown, N.J.
Tacnet: 253-6451

Debugger/tracer

The Ada debugger/tracer, another member of the SEL's Ada tool set, will provide a debugging and tracing facility that can be used with any Ada compiler.

The features of the debugger/tracer will include, but not be limited to:

- Breakpoints at specified Ada statements. (A breakpoint is an instruction that temporarily stops program execution, so the user can check status.) Here, they involve one or more value- or range-dependent breakpoints that, optionally, permit the user to access variables by name in order to inspect or alter them.
- Noninteractive display of named variables.
- Noninteractive modification of named variables with defined values.

The user will select debugger/tracer features via an interactive menu. The debugger/tracer backend processes the data structure graph that the frontend produced from the input Ada source program, along with the debugging or tracing information that the user specifies. The tool then produces an extended Ada source program (using the source code regenerator). Although the new source program resembles the original Ada source program, additional code has been inserted for the debugger/tracer functions at program execution time. The user can compile the new version of the program using any Ada compiler.

Other current and future Ada efforts

As part of the SEL's formal verification (Verlangen) effort—also described in this issue—we will automatically produce an Ada code skeleton from the Verlangen design/specification statements. Furthermore, as another activity, ATL will generate Ada code that represents the contents of an expert shell (the knowledge/rule base as well as the control software). Future Ada tools are planned and will include program architecture diagrams that graphically display the various structures and interrelationships in the Ada code of large software systems.

Conclusions

The Ada language is becoming more accepted as the worldwide military language (the North Atlantic Treaty Organization—NATO—has decided to use it). While the language can manage the complexity of large software projects, its own complexity creates problems, particularly for managers, quality assurance personnel, and code maintainers. The SEL is developing Ada tools that will greatly alleviate growing pains of the new language and allow RCA development of Ada software to be truly productive and cost-effective.

References

1. R.M. Blasewitz, "Ada—Not just another programming language," *RCA Engineer*, Vol. 29, No. 1, pp. 23-31 (Jan./Feb. 1984).
2. G.W. Mebus and E.J. Braude, "Computer languages—A view from the top," *RCA Engineer*, Vol. 29, No. 1, pp. 14-22 (Jan./Feb. 1984).
3. *Ada Programming Language*, Department of Defense, Washington, D.C. NASI/MIL-STD-1815A-1983 (January 1983).
4. *VAX DATATRIEVE—Guide to Programming and Customizing*, Digital Equipment Corporation, Maynard, Massachusetts (September 1984).

Camden engineers:

- ☐ ***Where can you find the latest information on research in artificial intelligence?***
- ☐ ***Where can you obtain a current bibliography on antennas? On management practices? Parallel processing? Computer science?***
- ☐ ***Which are the most respected and cited journals for research in a specific field?***
- ☐ ***What RFPs and contracts are currently being placed for bids by the government?***

The answers to these and other questions can be found in the CISD Library, located in building 10-6-5. The scope of the periodical and book collection encompasses subject fields such as engineering, electronics, computers and control, management, mathematics, and physics. The library maintains a subscription to the Visual Search Microfilm system, which includes military documents, industry standards, and vendor catalogs on 16mm microfilm.

For information requests not readily filled on-site, the library has access to a wide range of sources. First and foremost is access to the collections held by other RCA libraries. Secondary sources of supply include, but are not limited to, the British Library Document Supply Center, the Engineering Societies Library, and the New Jersey State Library. Contacts are maintained with college and university libraries to obtain books and articles on an interlibrary loan basis.

Bibliographic citations and abstracts of articles and conference papers on a wide range of subjects are available on-line from more than 200 different databases through the DIALOG Information Retrieval Service. The library staff can access both the NASA RECON database and the Defense Technical Information Center database.

The Camden library was founded in 1927, and is the oldest library at RCA. It has an active books collection of more than 14,000 titles, and subscribes to approximately 300 periodicals.

For more information, contact:
Nina Arrowood
Tacnet: 222-4046

E. J. Conklin

Leonardo: Design environment of the 1990s

RCA is one of 21 members of the Microelectronics and Computer Technology Corporation (MCC) consortium, and one of ten members of MCC's Software Technology Program (STP). In that capacity, RCA is contributing to the development of a team design environment for very large/complex systems.

The mission of the Software Technology Program (STP) at the Microelectronics and Computer Technology Corporation (MCC) is to introduce, via prototypes, new software technologies—coherent techniques and tools—that will bring about extraordinary increases in productivity and quality in the laboratories and development groups of the program participants, such as RCA, in about seven to ten years.

The approach must be focused for the fulfillment of the mission. Accordingly, the STP is aimed at improving activities that precede programming, to effect profound changes in the way large, complex software is specified and designed, especially the *realtime* constrained and *distributed* systems. Because a *team* of experts usually develops these systems, we will exploit the team approach to maximize the advantages it offers in design efficiency and quality. Support of complex digital *system* design is also a prime consideration, because large portions of future applications will be implemented in microcircuits, not in software.

Research outside of MCC will continue to be about software development "in the small," and the aspects of program design that can be automated, such as executable specification and machine verification. We will incorporate the best results of these efforts and the STP results into a comprehensive software development environment prototype called Leonardo.

As the results of research on programming "in the small" become translated into technology, development costs will shift more toward the earlier, more creative aspects of design. During design, alternatives are explored and their impact assessed. We believe that we can eliminate or significantly reduce much of today's costly design efforts by using the computer as a powerful tool, for example, to offer alternative graphical representations. The computer will also be a reservoir for design knowledge that is available during project development and for maintaining and enhancing the product system.

Briefly put, the STP approach is to use computers to aid the collective innovation that must occur in large-scale software development done by a team of professional programmers. If participants rely solely on their own advanced technology efforts and on academic research, which is often fragmented and not committed to technology transfer, then essential parts of complex design will remain unaided by computers for more than a decade. Furthermore, there is a serious shortage of qualified researchers, which hampers the participants' ability to staff a comparable research effort independently.

One caution is necessary, however. In this rapidly changing field, anyone who tries to look more than a few years ahead must abandon strict logic. The picture we paint here will almost certainly be replaced by others as the research proceeds. No claims about scientific rigor, nor promises about specific deliverables, are being made.

Abstract: *Much has been written and said about the official research plans of the Microelectronics and Computer Technology Corporation (MCC) Software Technology Program (STP), but these plans do not adequately convey the long-term vision that shapes the research. This article describes Leonardo, one anticipated result of these research efforts. The intention of the STP is that program participants—for example, engineers and researchers in RCA—find this vision imaginative, stimulating, and compelling.*

Final Manuscript received November 20, 1985
Reprint RE-31-1-10

Leonardo

In general, Leonardo is a system that ties together a team of experts. It provides appropriate tools to each expert and has its own body of knowledge, set of skills, and expertise. In addition, each expert has a customized console located in a control room environment.

Leonardo is not just a collection of specialized tools. It is an integrated system for designing and developing systems— from cradle to grave. It is also a communication vehicle among the members of the design team and all others who are part of the system-development process.

A central feature of Leonardo is the Unified Product Model (UPM). This unified representation of all aspects of the system under development, at all

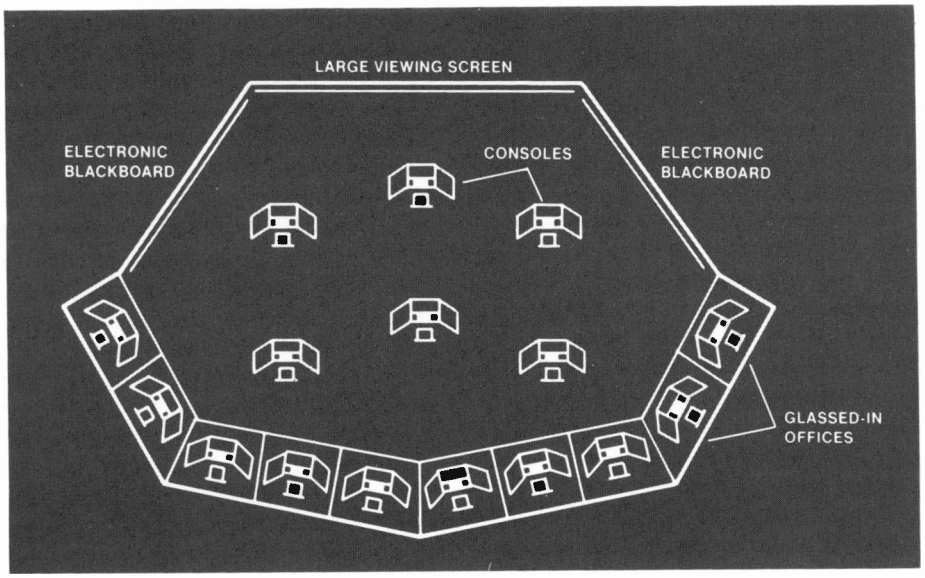

Fig. 1. *Plan view of the Leonardo control room.*

levels, is rich in pointers to related information in the requirements, specifications, design, implementation, performance and testing data, and documentation.

The hardware

Leonardo's hardware will be a collection of high resolution graphics workstations with a 10k by 10k pixel display, 100 megabytes of random-access memory, and a multi-gigabyte address space.

These workstations will be linked by a gigabaud fiber optic network. They will have limited voice input (not multiuser connected speech) and reasonable voice output.

Most workstations will be in offices that are clustered around a large, central "war room," and will have glass windows facing it (see Fig. 1). Several large electronic blackboards in the war room will allow a group to share an input/output surface and will provide them with color output as well as ("chalk") input facilities.

The people

The design team consists of user representatives (including problem domain experts) and a core group of Leonardo specialists. The size and structure of both groups are somewhat flexible, depending on the type and size of the problem. The Leonardo core group includes:

- Project manager—orchestrates the team's work; maintains the balance and flow among team members (analogous to an orchestra conductor); also is responsible for schedules, milestones, monitoring and reporting progress, and managing the design team.
- Unified Product Model (UPM) engineer—maintains all UPM subsystems, (analogous to an airline flight engineer), looking particularly for inconsistent points that Leonardo cannot spot among team members; tunes and maintains Leonardo; supports the other team members in their use of the Leonardo tools.
- Reuse Engineer—recovers and re-engineers any system elements for which previous work is available and appropriate, assures that work products being developed are reusable, and catalogs them (analogous to a librarian).
- Specification analyst—maintains the requirements and specifications for the system under design; ensures they are current, relevant, complete, and consistent; develops new representations of the problem as aids for creative problem solving; acts as mediator between user representatives and implementation specialists (analogous to a system analyst or system engineer).
- Design specialist—serves as an expert in functional decomposition and exploration; guides work distribution among team members and monitors the bushiness of the design tree to optimize the number of options specified, if not actually explored, at every decision point; develops new representations of the system to aid the team in problem solving; sees that the design works.
- Implementation engineer—runs the metacompilers and prototyping subsystem, monitors integration of the modules and subsystems being developed, and ensures they can be integrated; spots implementation snags (analogous to a programmer).
- Performance/reliability specialist—analyzes modules in any development phase for performance, reliability, and implications to the rest of the design.
- User interface specialist—serves as an expert in user-cordial, high-bandwidth interfaces; guides development of the interface; constructs prototype interfaces for user representatives to use and evaluate; is responsible for documentation (both on-line and printed) for the system.

Among the user representatives are:

- User management representative(s)—management people who are concerned with development cost and delivery schedules. Because development is extremely fast and uses rapid prototyping, the cycle for negotiating requirements is very short. Therefore, user representatives with decision-making power must be present.
- Problem domain expert(s)—users or independent contractors who understand the context and environment in which the developing system will be used and the constraints that that environment places on the system.
- End user representative(s)—representatives of the users who will install, use, and maintain the product system in the field. This role is similar to the problem domain experts.

The approach

Each role has a specialized view or perspective of Leonardo, which consists of a set of tools that automate the role's tasks (especially the mundane ones), plus a customized graphical interface.

Sometimes, all the roles "meet" together, focused on a single problem or activity. At other times each role works alone, or perhaps confers with one or two other roles. Many styles of meetings are supported: from team members

Software Technology Program

The Software Technology Program (STP) is one of the two research programs at the Microelectronics and Computer Technology Corporation (MCC) in which RCA participates—the other is the computer-aided design/very-large-scale integration (CAD/VLSI) program. Since September 1984, when Les Belady became the MCC STP Director, 32 researchers in software engineering and related fields have been hired. The major areas of research efforts are exploratory design, teamwork, and early phases of system development.

Four subgroups work on the MCC STP program:

- The Design Process Group seeks to create models of both individual and group design, as well as technologies to support these models.
- The Design Information Group focuses on ways to abstract the behavior of programs and on reusability issues.
- The Design Interface Group develops ways for designers to view the emerging product.
- The Design Environment Group works on the architecture of Leonardo and the integrated environment that it will present to its users.

ATL's Software Engineering Laboratory (SEL) is responsible for MCC STP technology transfer within RCA.

meeting around an electronic blackboard, to voice, video, and graphics teleconferencing via Leonardo; to a shared data base where team members can post and exchange the "objects" of their work, while graphically viewing the "space" of shared objects.

However, the team is not divided according to modules or subsystems. In this new paradigm, all the team members work on and apply their expertise to all parts of the system. Leonardo automatically does the mundane work—much of which is now done manually. The people are left with the analysis, creative decisions, and communication. Furthermore, the work no longer must proceed in a lockstep, waterfall development style. Leonardo supports many design methods, including top down, bottom up, most critical component first, and rapid prototyping. It is the UPM and its specialized views that permit this team approach. On a football team or a surgical team each member not only has a specific job and a set of skills for doing that job, but also has a single coherent view of the field of action. All members can see the problem, and they can see each other's problem-solving activities.

At present, no such single coherent view of the field of action exists in software engineering. Instead, the evolving software system exists as a loose, possibly incomplete and inconsistent collection of documents, code, and ideas in people's heads. By making the evolving system explicit, and by preserving its consistency, the UPM allows a team of experts to use their specialized views of the product, and to add their expertise in a coherent and synergistic way.

Conclusion

This is one vision of Leonardo. The specific roles described are only a guess at the right division of labor for a Leonardo-like environment. We believe that coordinating a team of experts through a graphical computer network will revolutionize the way large systems, especially software systems, are developed. A future article will discuss some of the tools and views that different roles have, and how these roles interact with each other.

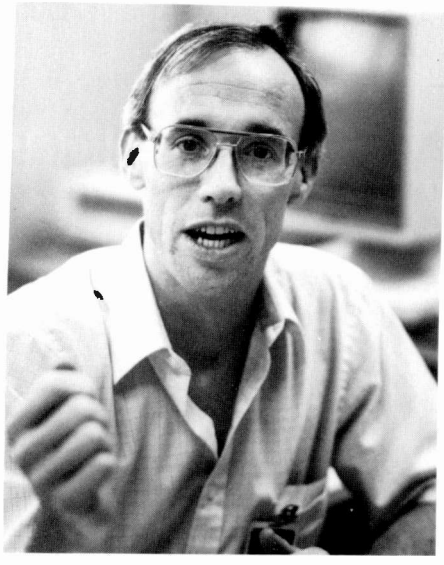

Jeff Conklin, a Unit Manager within ATL's Software Engineering Laboratory, is serving as RCA's liaison representative to the MCC STP. He is focusing on the area of software engineering—the scientific study of the engineering of large, complex computer programs. Prior to relocating to MCC in Austin, Texas, Dr. Conklin was a Unit Manager within ATL's Artificial Intelligence Laboratory. Jeff received his PhD and MA degrees in Computer Science from the University of Massachusetts. Contact him at:
MCC
Austin, Tex.
(512) 343-0860

References

1. *MCC STP Financial and Operating Plan (F&O), 1986*, August 1985. (MCC confidential and proprietary document; available upon request only to employees of MCC and MCC shareholder companies).
2. "MCC: Investment in the Future," *RCA Engineer*, Vol. 29, No. 4, pp 31-32 (July/August 1984).

J.A. Adams | M. Gale | J.W. Dempsey | G.W. Kaizar | N. Straguzzi

Artificial intelligence applications

ATL is investigating a number of applications for practical, deployable AI systems for government and industry.

This article begins with a discussion of the AEGIS Expert System—one of the more mature AI applications. The AEGIS system combines a knowledge database developed from experts in the field with information synthesis and classification software programs. Other applications discussed include automatic test program generation, an expert system to help engineers with computer-aided design of integrated circuits, voice-computer interfacing, and battlefield information fusion. The emphasis in all programs is on building practical, deployable AI systems for government and industry.

The AEGIS EXPERT system

The first AEGIS cruiser, *USS Ticonderoga,* demonstrated its extensive capabilities off the coast of Lebanon as it tracked and reported on all traffic in the area. It was so well received by the Navy AAW Task Group commander that there was an understandable reluctance to take the AEGIS system off-line for any reason—even to perform scheduled maintenance! The decision to take equipment off-line requires extensive knowledge of which faults have minor consequences (because of built-in redundant capabilities), and which faults have serious operational consequences and therefore must be corrected as soon as possible. Further, the decision to perform maintenance depends on detailed knowledge of the maintenance options that allow repairs with minimum impact on operational capabilities. These needs for extensive knowledge to support maintenance decisions led to consideration of a computerized expert system.

Abstract: *The Artificial Intelligence Laboratory at ATL is focusing most of its efforts on three broad application areas in aerospace and defense: information synthesis, classification, and signal interpretation. A long-term goal of this group is to develop systematic approaches to the analysis of new problems in order to select and apply appropriate techniques in a cost-effective manner. Identifying classes of algorithms effective for these applications is a principal approach toward achieving this goal.*

Final manuscript received November 20, 1985
Reprint RE-31-1-11

The AEGIS Operational Readiness Test System (ORTS) can be a major component of such a system. However, it is primarily oriented to the detection and isolation of hardware failures. It does not deal with system set-up problems and varying environmental conditions that must be considered when trying to diagnose fault symptoms, nor does it take advantage of the observations that a person can make at the various tactical displays within the system.

Figure 1 shows the problem of diagnosing faults in systems with deeply nested equipment interdependencies. Testing and monitoring equipment generally can identify, through a readiness reporting system, which of several equipment groups are not performing their functions. However, the monitoring system may not be extensive enough to identify a particular assembly as a common cause of several identified symptoms (see Fig. 1). These faulty assemblies must be quickly identified and repaired to assure uninterrupted availability of system capabilities.

In some cases, it takes considerable training and experience to infer the actual cause of faults from the observed fault indications. Computer-based expert systems provide the means to capture such diagnostic expertise from a few human experts, and then to make this knowledge available to others who have less experience. Such knowledge, based on extended shipboard experience, can help less experienced system users distinguish

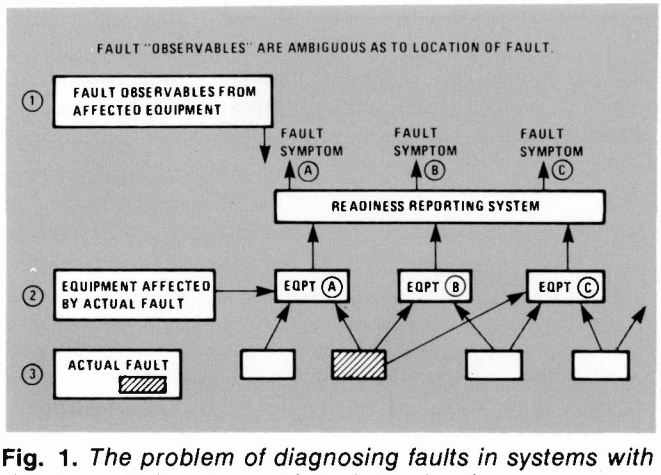

Fig. 1. *The problem of diagnosing faults in systems with deeply nested equipment interdependencies.*

anomalous observations caused by adverse effects of environment (sea returns, clutter, countermeasures, etc.) from the observations associated with system malfunctions or radar performance degradations. One objective of the AEGIS expert system is to capture and pass on just such shipboard experience.

How the AEGIS EXPERT system works

MSRD engaged the ATL AI Lab in a joint effort to build a computerized expert system for AEGIS. It was decided to use an existing program called EXPERT that was originally developed at Rutgers University for medical diagnosis. The program was designed so that it could accept a human expert's knowledge of facts and rules used for making diagnoses of problems—not just for medical problems. Briefly, here is how EXPERT works. The EXPERT system is a computer program that allows the user to build a database consisting of facts and rules (see Fig. 2). The facts provide descriptions of the problem domain in terms of observable symptoms, and the conclusions to be drawn or actions to be taken based on those symptoms. The rules describe the relationship between the symptoms and the conclusions. The relationships are in the form of IF-THEN rules, as shown in Fig. 2. These rules can show relationships between different observations, or between observations and conclusions, or between different conclusions.

An important advantage of the EXPERT system is that building the database does not require any computer programming skills. It is no more difficult than using microcomputer spreadsheet programs. After several hours of instruction, a database can be prepared by typing the facts and rules into a standard text file (a word-processor file is fine). The file is accepted by EXPERT, which then compiles it for further processing. Changes to the database are made by simply editing the text file and resubmitting it to the EXPERT compiler. This makes it convenient to upgrade the database with new facts or rules.

Once the database is prepared, the EXPERT program allows the user to consult it for a diagnosis. In the diagnostic session, the user responds to a sequence of questions presented on a computer terminal in a menu format. Essentially, the questions seek to determine which facts are relevant to the problem at hand. As each fact is identified, the computer program searches the list of rules, selecting those that are affected by that fact. When a rule is "fired" (i.e., when it makes use of a particular fact), the program will present the user with questions to determine the status of the other facts used by that rule. In this way, the EXPERT system will work its way through the rules, collecting the additional information it needs to finally come to the conclusions (diagnoses) specified by the rules. The portion of the EXPERT system that "navigates" through the database is called an "inference engine"—it makes inferences based on the facts provided by the user. Figure 3 shows a sample session sequence at an AEGIS EXPERT terminal.

Constructing the AEGIS EXPERT system database

The first attempt to develop a database for AEGIS followed the traditional methods of building expert systems. AI Laboratory "knowledge engineers" conducted interviews with an acknowledged AEGIS radar expert. The knowledge engineers had to determine if the expert's knowledge could be represented in the formats needed by EXPERT's IF-THEN rule system. The objective was to capture the expertise that an individual had accumulated at AEGIS test sites and aboard the *USS Ticonderoga*. These included problems with symptoms observed by radar operators and other equipment users. No effort was made to diagnose problems that ORTS was designed to detect, although the radar expert did make use of ORTS tests to confirm diagnoses.

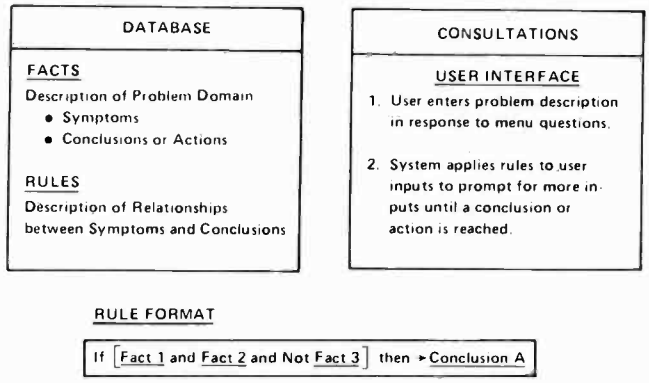

Fig. 2. *The EXPERT system shell.*

The knowledge engineers recorded five two-hour interviews with the radar expert. They then studied the transcripts to identify the relevant problem symptoms and the logical links leading from those symptoms to their diagnostic conclusions. These were written into the database, and the "human expert" was invited to consult the "computer expert" to determine how well the system could diagnose a problem from its symptoms. Other knowledgeable people were invited to consult the system and comment on its effectiveness during this debugging phase.

The resulting database, restricted to SPY-1A radar problems, was demonstrated to MSRD management; they authorized an extension of the database to include problems with computers and disks. A second human expert, acknowledged for his expertise in solving AEGIS computer and computer peripheral problems, was made available for a number of two-hour interviews. His database was developed, tested, and merged with the radar expert's database.

At this juncture, the AI Laboratory provided instructions to MSRD to allow the experts themselves to enter their knowledge directly into the system. A database development course was given to several MSRD staff members to describe how facts are entered, how rules are written, and how facts and rules are arranged so that the system makes efficient use of the information provided.

This new approach was tried by a former Navy Maintenance Chief who served on the second AEGIS cruiser, *USS Yorktown*. His objective was to enlarge the database with problem symptoms as detected by detailed ORTS tests, called Operational Performance Tests (OPTs). Since there are over 170 of these tests, and since they can fail in various combinations, it takes considerable expertise to interpret how to use them in the fault isolation process. It was decided that this would be an excellent demonstration of the shipboard capabilities of the AEGIS EXPERT system.

An additional objective was to include diagnostic interpretations that identified the operational consequences of OPT failures, as well as to identify the causes of those failures. This could be done with AEGIS EXPERT by using the same facts in association with different rules that identify the fault consequences.

Figure 4 shows the current structure of the AEGIS EXPERT

database. The system allows merging of separate database sections to permit the accumulation of the knowledge of several human experts.

Figure 5 shows the database designers at a terminal during a typical consultation session with the system.

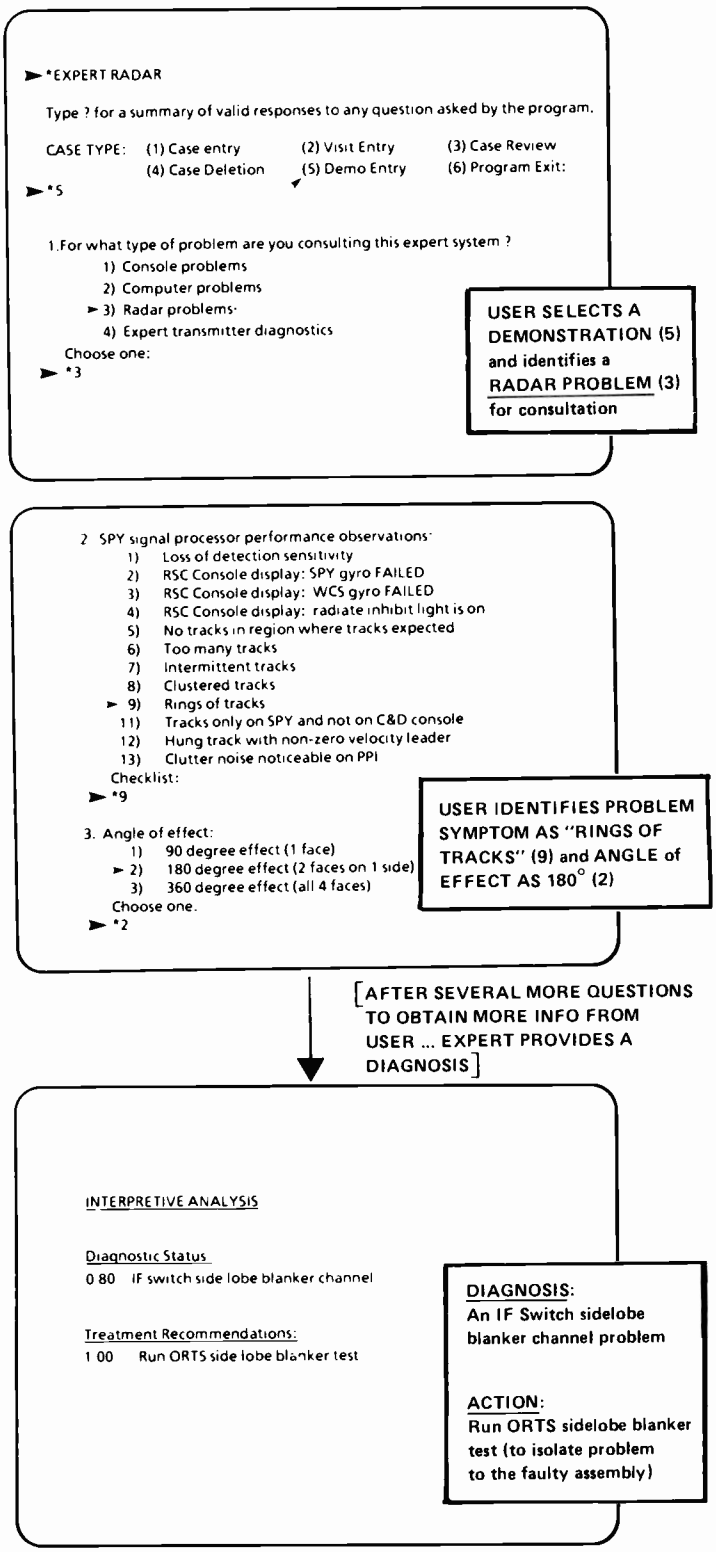

Fig. 3. *Portions of a diagnostic session at an AEGIS EXPERT system terminal.*

Current and future plans

Currently, the AEGIS EXPERT system is implemented on a VAX with access through local computer terminals at ATL and MSRD. Project plans call for transfer of the system from the VAX to a microcomputer that will permit shipboard evaluations and use. An invitation has already been extended to RCA to evaluate the system aboard an AEGIS cruiser at sea.

Implementation on a microcomputer will be based on current work at ATL: writing the essential user consultation features of the EXPERT system in C language. This will be implemented together with a program called ETC (EXPERT to C). The ETC compiler translates an arbitrary expert shell model file into an equivalent expert system in the C programming language. The Rutgers EXPERT model was selected as the shell because of its numerous knowledge engineering and debugging tools, its simple model file format, its proven effectiveness in several domains, and its case-saving mechanism. The programming language C was chosen because its compilers are available on microprocessors and it is capable of accessing low-level assembly-like constructs. When completed, it will be loaded into an IBM PC AT for evaluation at several user sites.

We also plan to include database development capabilities in future IBM PC AT implementations. They will serve as database building tools at various land sites linked to shipboard users.

EQUATE self-test expert system

EQUATE is a testing and troubleshooting program used by engineers at ASD in the development of automatic test equipment (ATE). The prototype expert system enhances the AN/USM 410 EQUATE self-test program and interfaces with software currently in use to provide additional capabilities in fault isolation and identification. An expanded knowledge base provides an improved capability for troubleshooting the ATE without the intervention of service engineers. The expert system reduces the need for expert help and decreases system downtime in the field. Written in Prolog, the system will run on a microcomputer with 400k of memory.

Automatic test program generation

When a new device is developed, a set of valid test programs for the device also must be developed before the unit can be fielded. Without these test programs, the device might as well be broken. As devices become more complex, ATE test programming is consuming a growing portion of the cost and time involved in designing a new device. An automatic system to write ATE test software could reduce both the cost and time involved in the generation of these software programs.

While ATE has been an RCA military business for some time, particularly for ASD, new hybrid circuits, including hypercomplex VHSIC chips, are rapidly outstripping the ability of humans to quickly develop valid tests that isolate faulty components of a given unit. Automatic test program generation (ATPG) software is needed to produce the programs required to test these systems and isolate faults on analog and digital circuits. The long-term objective of ATPG is to produce ATE program sets, given only the specifications of the ATE and a unit under test (UUT). ATL has advanced to the point of including fault isolation of a UUT with analog components or similar digital components with complex stored-state behavior.

In 1985, ATL enhanced the prototype ATPG system to

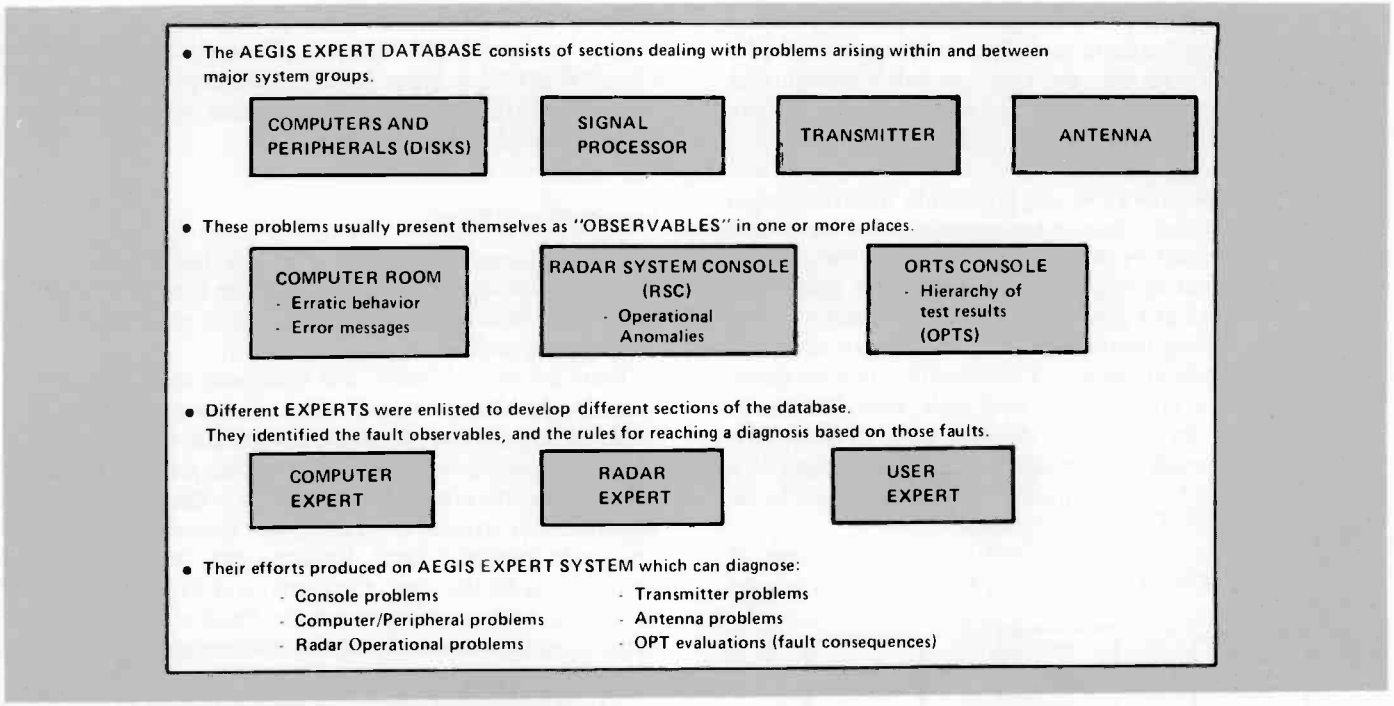

Fig. 4. *The current AEGIS EXPERT database structure.*

reason with resource constraints and test sequencing preferences acquired from the experience of expert test designers. Also, ATL developed a prototype architecture capable of combining resource planning with diagnostic decision-making, using a test problem from AEGIS combat system maintenance. The expert knowledge of personnel at ASD and MSRD provided the model for both of these advances.

Fault isolation for UUTs currently operates at two extremes: 1) a technician sits at a station with a collection of manuals and tools and tries to use them in a highly creative way, or 2) ATE is used to attempt to replace the creative technician with hand-coded UUT-specific software that generates an exhaustive sequence of tests. Both approaches have drawbacks. Technicians are scarce, slow, and of varying abilities. ATE software is very expensive, varies in quality, is very sensitive to engineering changes, and often does not exist when the UUT is fielded.

Recent research has been aimed at integrating various system architectures and knowledge sources, but there is no convincing evidence to date that any group has solved any major part of the analog ATPG problem. The lack of progress is an indication of the difficulty of this problem. The identification of a solvable subpart, therefore, would be a useful and significant contribution.

The test program set (TPS) development process begins with the study of the UUT. Inputs usually include schematics of the UUT and a functional block diagram of its operation. An ATPG system could change the current TPS development process by shifting the test designer's function toward maintenance of the knowledge base of the ATPG system. Work completed this year confirms that it is possible for a knowledge-based system to reason about functions of a UUT of realistic complexity and to generate an intelligent test strategy for such a UUT.

At some point, the decomposition of the functional block diagram results in functional components that are not divisible, and for which testing strategies are known. Individual components that can be tested by straightforward methods correspond

Fig. 5. *Database designers at a terminal session with the AEGIS EXPERT system.*

to primitive statements in the functional flow chart (FFC). Individual components that require tests that are not directly implementable by the ATE, or that require critical resources of the ATE, correspond to non-primitive instructions of the FFC.

ATL's current approach is characterized as a knowledge-based, depth-first traversal of the graph corresponding to the functional block diagram of the UUT. Each node in the graph represents a component of the functional diagram, and each edge a link between these components. Attached to each node is a frame of testing strategy and constraints expressed as axioms and rules. During the traversal, this frame is accessed, and subject to constraints, the strategy it expresses is implemented. A frame can be viewed as a micro knowledge base specific to its component. The algorithm is a depth-first traversal because a strategy for any output is generated completely before another output is considered.

An important feature of this design is that it explicitly mimics the problem-solving behavior of ATE software engineers in the process of writing ATE software. This approach is distinguished from most others, which use analytic methods or an abstract system design to attempt to solve the problem.

Basically, the expansion of the FFC to the detailed flow chart (DFC) can be viewed as a planning problem in which expansion of a non-primitive instruction to a sequence of primitive instructions is accomplished by an operator, subject to the physical resource constraints of the ATE and its interface equipment. The plan is viewed as a graph with a partial ordering on nodes that represent actions (instructions). The ordering on actions is extended if one action achieves a precondition of a successor action, or if two actions interfere with each other. The system develops a hierarchy of plans at different levels of detail. The hierarchy conceptually corresponds to the decomposition of a non-primitive instruction to primitive instructions, subject to the limitations of the ATE/interface equipment.

As an example, suppose that part of the current state of expansion of the FFC is represented by the following subgraph:

This is read as, "perform instruction 'a' or 'b' in either order, then 'c', 'd', or 'e' in any order, followed by 'f' or 'g' in either order." Examining the plan, a critic finds that the physical resource of the ATE used to test 'g' in fact must be allocated only after the physical resource used by test 'f' has been allocated The ordering on the plan can be refined to:

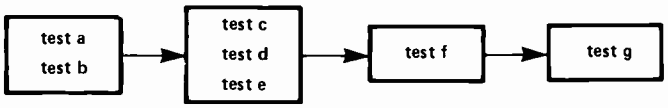

where tests f(1) and f(2) are primitive, but test f(3) is not. Examining the plan, a critic finds that test f(3) is a precondition for test f(1). The plan is reordered to:

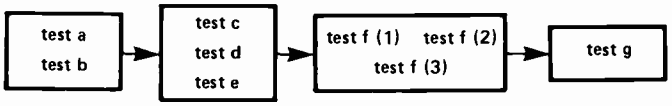

Suppose now that test 'f' is chosen for expansion as follows:

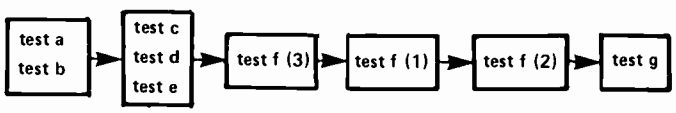

The process of expansion is completed when all non-primitive instructions have been reduced to primitive instructions. The set of components that can be evaluated in this manner has expanded to include decoders, sample-and-hold amplifiers, integrators, and flip-flops.

Test design experts at ASD and MSRD are providing criticism to ATL on the ATPG system's performance as part of the formulation of new heuristics for efficient test design. Implementation of these new heuristics will drive successive experimental cycles of revision and validation by human experts. In this manner, the scope and complexity of components and systems subject to ATPG is being improved and expanded. Currently, ATL believes this joint effort by RCA units is two years ahead of the published work in the ATE field.

Information fusion

Information or message fusion is a general class of interpretation problems that involves assessing evolving situations based on information from a variety of sources whose output may not be directly comparable.

There are several factors that contribute to the difficulty in designing an efficient and effective information fusion system. Different sources may contribute information in differing levels of detail, at varied sampling rates, and with different degrees of reliability and relevance. Furthermore, the volume of information available may exceed the information processing capacity of any single intelligent agent. Understanding the computational requirements for this class of problems is crucial for the rational distribution of tasks among agents, the design of decision support systems, and the structuring of effective command, control, and communication networks among the interpreting agents.

ATL is conducting computational experiments to investigate the quality, efficiency, and accuracy of a knowledge-based interpretation process for assessing battlefield situations. A variety of combinations of information source types and several control structure variations are being explored. This research is directed toward meeting the technology requirements of the battlefield of the future. In addition to basic research on information fusion, ATL is pursuing several advanced development efforts in which an AI system solves a portion of a particular fusion task and supports the decision making of a human expert.

One such project at ATL is the Cooperative Anti-Submarine Warfare Tracker. The goal of this project is to develop and demonstrate a methodology for designing "cooperative" man-machine problem-solving systems. In these systems men and intellignet machines are organized into coordinated teams to solve complex problems in a timely and effective manner. The strategy employed has been to redesign the existing operational tracking system, implement the prototype cooperative system, and compare the prototype to the existing system. The unit targeted for redesign has been a Navy airborne system that locates and tracks submarines using a variety of sensors.

The initial focus for this project was on the tracking subtask involving passive acoustic sensors, since this is an extremely difficult subproblem for the human expert. The man-machine system deploys sensors and interprets sensor contacts to determine a submarine's current location, course, speed, predicted location in the near future, and its tactical plan. A real-time emulation of the expert's current (operational) tactical workstation was developed to drive the design and evaluation steps.

AI and cognitive science tools and techniques are being used to design an intelligent machine component, to design the man-machine communication, and to evaluate the man-machine system. The design methodology includes specific steps to account for the user's capabilities and limitations, and the skills and limitations of the human expert are the basis for assigning a supportive role to an intelligent machine component. In this approach, we attempt to extend or enhance the user's expertise and performance beyond the limits imposed by the human information processing architecture, while introducing only min-

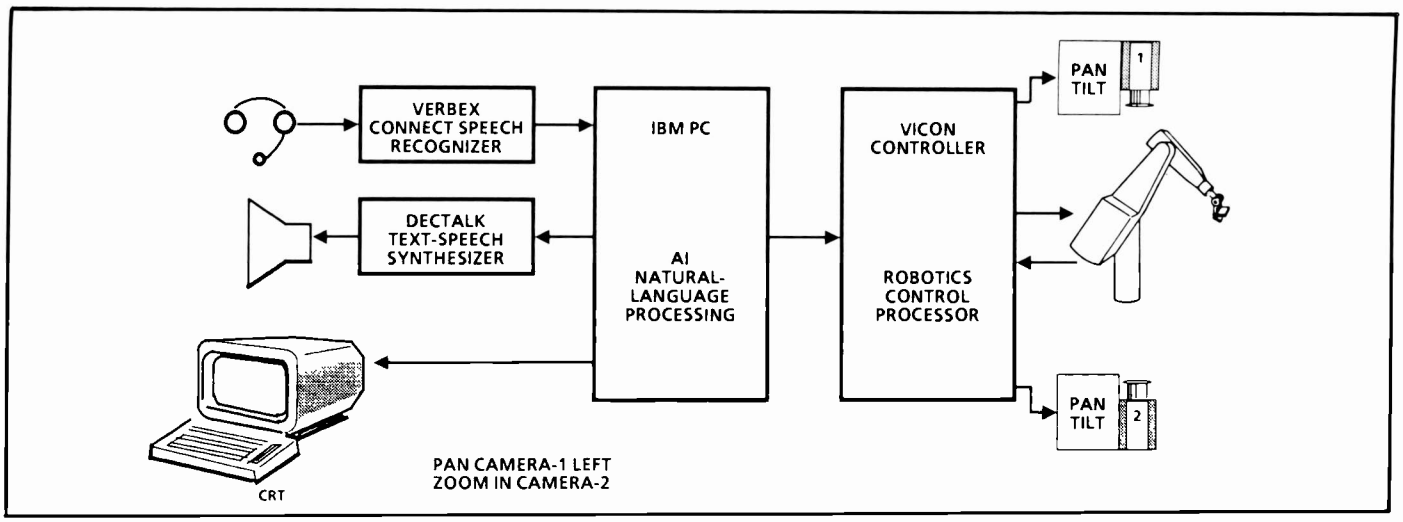

Fig. 6. *Natural-language interface demonstration system.*

imal overhead for the additional interaction/communication tasks.

An explicit model of the expert's problem-solving strategy and capacities is developed and used to identify a difficult subproblem. This subproblem is then distributed to the man and the machine. Cooperative interaction between man and machine is embedded in the tactical workstation by including AI-based components that give the machine the capability to model the expert's knowledge state and to initiate and carry on intelligent dialogue. The joint problem-solving activity of man and machine is driven by and affects a real-time simulation of the tactical environment.

Data collection/recording capabilities built into the simulation testbed (emulated tactical workstation) support the design, knowledge engineering, and system evaluation experiments.

Researchers at Rutgers University are collaborating with ATL on the information fusion work. Basic research at Rutgers on design and planning problems and expertise acquisition provides the tools, techniques, and results that make feasible the application of artificial intelligence to these complex and ill-defined problem-solving tasks.

Natural language processing

ATL is developing natural language interfaces that accept English queries and commands and translate them into directives that can be understood by the computer. The interfaces are being designed with maximum modularity to enable them to accept varied inputs (such as keyboard, voice, and batch text) and to be adapted to many different kinds of applications.

Natural language interfaces can ease the burden of remote control of multiple simultaneous functions by a human. User problems can arise when a sophisticated electronic workstation has many more controls, displays, and information sources than any one user can manage, especially during peak periods, or when intense concentration on one particular resource is required to complete a task. "Hands free" operation of remote devices allows teleoperators to concentrate on other tasks. The development of a natural-language voice recognition interface would benefit many existing government programs, including Space Station Automation (AED), Astronaut Apprentice (AED), Ground Site for Satellite Control (AED), Vetronics Program (ASD), TCAC Workstation (ASD), Integrated Interfaces (DARPA), AEGIS Control Room (MSR), Aircraft Communications (CISD), Hands-Busy Cockpit Control (NADC), and Robotics Command and Control. The major focus of natural language processing within the AI Laboratory is voice control of a robotic environment. In conjuction with the Speech Recognition and Robotics Laboratories of ATL, we are currently working to develop an interface that consists of a continuous-speech recognizer and a natural-language parser integrated with AI expert system software and interfaced to camera and robot controllers. The initial development goal is to demonstrate the use of natural-language voice recognition to control the pan, tilt, zoom, and focus of two video cameras mounted at a remote location, while manually controlling the operation of a robot arm mounted at the same location. At this time, the system is for a demonstration in which an operator performs a task via video monitors, using a Puma 762 industrial robot arm and two video cameras at a remote location, with natural-language voice commands controlling the cameras.

The working demonstration system developed late in 1985 (shown in Fig. 6) consists of:

- The operator interface—a head-mounted microphone, speaker, two joy sticks, two video monitors, and the IBM PC terminal.
- The processing hardware—a Verbex speech recognizer and the DECtalk speech synthesizer interfaced to the IBM PC as the central processor.
- The robotics hardware—two video cameras mounted on two separate pan/tilt units, the Vicon camera controller interfaced to the IBM PC with a bank of relays, and the PUMA robot arm.

Current technology, in the form of template-matching recognizers, can support only speaker-dependent, constrained-syntax, corrected-speech, and can recognize up to 100 words.

The current system allows the operator to enter voice commands that take the form of full or partial sentences. An Augmented Transition Network (ATN) parser determines the structure of the command, and discourse analysis software then extracts the meaning of the command on the basis of the situational context in which it occurs. If the command is determined to be sufficiently unambiguous, the desired low-level commands are sent to the camera controller. For example, if Camera 1 was most recently referred to, the command "move right a little" would cause Camera 1 to pan right a small,

VERBS:	PAN, TILT, ZOOM, FOCUS, MOVE, GO, STOP, CONTINUE, KEEP
GERUNDS:	PANNING, TILTING, ZOOMING, FOCUSING, MOVING, GOING
NOUNS:	CAMERA-1, CAMERA-2, BOTH-CAMERAS, THE-OTHER-ONE
ADVERBS:	RIGHT, LEFT, UP, DOWN, IN, OUT, BACK, MORE, AGAIN
LIMITERS:	A-LITTLE, A-LITTLE-BIT, SOME
ISOLATED:	GOOD, GREAT, OKAY, FINE, WHOA, YES, NO

Fig. 7. *Subset of words for demonstration.*

predetermined increment. On the other hand, if the previous command had been "Both cameras move right," a subsequent command of "The other camera zoom in" would cause a synthesized speech message to be sent to the operator requesting him to clarify which message was meant. The operator might say "Camera 1" and then Camera 1 would then start zooming in.

The Verbex speech recognizer permits a maximum dictionary size of 100 words. Figure 7 shows some of the words that are in use in the current configuration.

During 1986 the concept of a voice-controlled camera system will be expanded to include voice control of an entire robotics environment in which the human operator is attempting to manipulate an integrated work environment in which a robot, cameras, and a video imaging system are tied together. The operator will be able to issue relatively high-level goal commands via free-form spoken English. Typical goal commands might be "Remember this location as toolbox," "Camera 1 look at the toolbox," "Open the hatch," and "Replace unit 1 with unit 2."

The fields of speech recognition, speech synthesis, robotics, natural language, and artificial intelligence are each rapidly advancing on a day-to-day basis. The Voice-Controlled Camera project provides an excellent example of plans for further development representative of many multidisciplinary projects at ATL, and where rapid prototyping and advanced technologies are brought to bear on outstanding problems in aerospace and defense.

MP2D Advisor

To gain experience with chip design problems and their solutions, the AI Lab has developed a prototype expert system to support RCA's multiport, two-dimensional (MP2D) automated standard-cell placement and routing program. This program is part of ATL's computer-aided design/design automation system (CADDAS). The purpose of the MP2D Advisor expert system is to help designers who lack expertise in the use of MP2D to perform input parameter manipulation when iterating MP2D execution during chip design. The knowledge engineering process focuses on classification rules to determine which sets of parameters to use and how to use them. Experts from the Microelectronics Laboratory of ATL collaborated on the knowledge base and evaluation of the MP2D Advisor.

While MP2D has received high industry marks for being a large, comprehensive software system helping chip designers produce efficient 100-percent-connected chips, it is necessarily a very complex piece of software that offers hundreds of user-selectable options, features, and parameters. Less experienced MP2D users often find it difficult to use the program to full advantage, and may be unaware of the wide range of options and capabilities available.

To assist novice and intermediate level users of MP2D, ATL's AI Laboratory, in cooperation with the Microelectronics Laboratory, developed MPECS—the MP2D Expert Consultation System. MPECS is an AI-based expert advisor that assists users in optimizing their chips by identifying and resetting the appropriate input (IP) parameters to MP2D. The system was implemented using EXPERT.

Knowledge engineers from the AI lab interviewed a number of MP2D experts and determined the manner in which they applied their knowledge to a chip design problem. This information then was encoded into the MPECS knowledge base, which currently includes approximately 230 rules and 190 hypotheses. Because MPECS runs independently of MP2D, it allows AI technology to be applied immediately to the chip design process without interfering with the proven reliability of an existing VLSI tool.

Because of the size and complexity of the MP2D system, it was decided to confine the initial MPECS implementation to IP parameter identification and manipulation. A designer who has made at least one full run of MP2D on his current chip, including artwork, can thereafter call up the expert system.

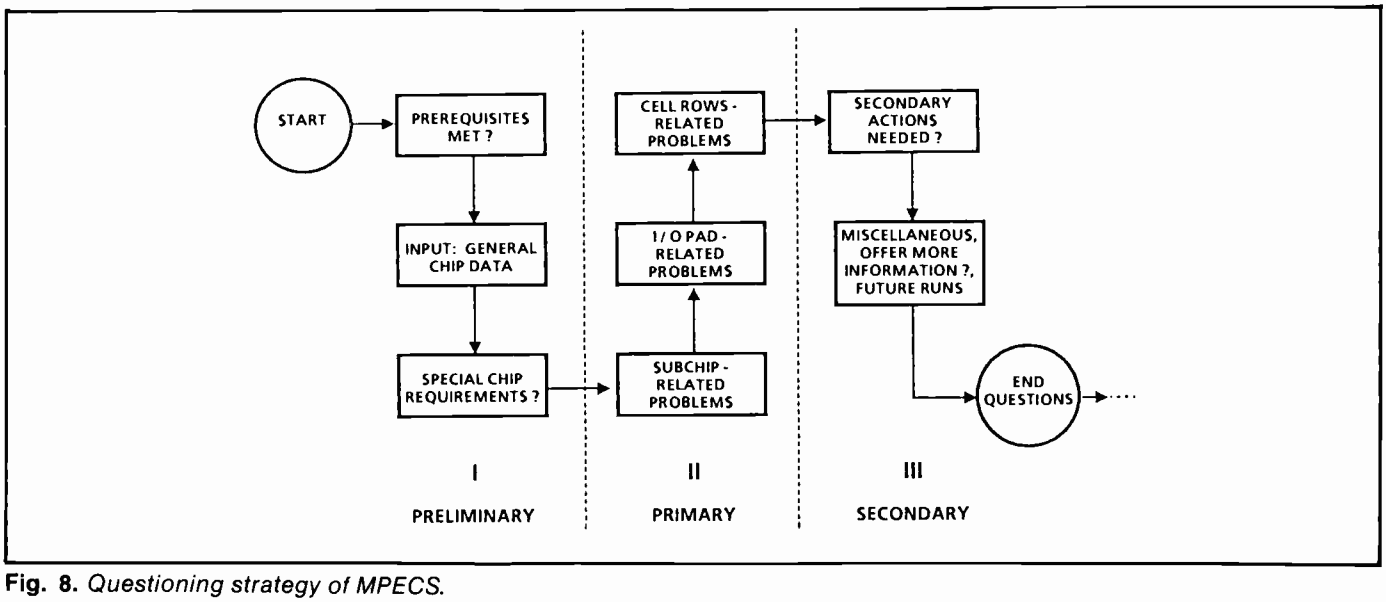

Fig. 8. *Questioning strategy of MPECS.*

MPECS asks the user a number of questions about the current state of his chip, and then recommends the parameters and settings that may be helpful in decreasing the size of the chip and increasing the efficiency of the layout. The procedure may be iterated as each improvement reveals additional areas for investigation. MPECS helps users meet certain special chip design requirements, such as non-square shape or special I/O pad placement. MPECS also offers a limited on-line help facility that allows users to access detailed information on certain topics of MP2D (See Fig. 8).

Future work in MPECS will include expansion of the depth and breadth of the knowledge base, work on improving the MPECS-user interface, increasing the on-line help facility, and field testing of Version 1.0 of the system.

Jacob A. Adams is a Staff Engineer at RCA's Advanced Technology Laboratories. He is responsible for coordinating all ATL Internal Research and Development Programs. Mr. Adams joined RCA in 1980 and served as Manager, Marketing Planning and Analysis for Command, Control, Communications and Intelligence Systems. He was appointed Manager, Marketing Development and Planning for ATL's Software Engineering and Artificial Intelligence Laboratories in 1983. Before joining RCA, Mr. Adams served as the New Products Manager, TU77/78 Tape Drive/Mass Storage System and Manufacturing Planning Manager with the Digital Equipment Corporation. His early career involved systems engineering and program management in aeronautical systems with the United States Air Force. Mr. Adams has a BA degree in Mathematical Engineering from Washington University and an MA degree in Business Administration from the Stanford Graduate School of Business.
Contact him at:
Advanced Technology Laboratories
Moorestown, N.J.
Tacnet: 253-6411

Joseph Dempsey has been a Senior Member of the Engineering Staff of ATL's Artificial Intelligence (AI) Laboratory since 1984. He is involved in the development of AI-based computer-aided instruction and the evolution and application of techniques for evidential reasoning. He also has devised a set of criteria for evaluating the applicability of AI to contemplated problem areas. From 1980 until 1983, as Assistant Professor of Computer Science at Baruch College, Dr. Dempsey taught both undergraduate and graduate courses in his field. Other experience includes the development of a control and analysis system for psychological experiments, the development of an AI model of infant intellectual growth, the design of a computer-based system for the development of high-quality tactual graphics for the blind, and system and numerical analysis programming in support of biomedical research. Dr. Dempsey holds a BS in Mathematics from Fordham University, an MS in Mathematics from Purdue University, and a PhD in Computer Science from the University of Pennsylvania.
Contact him at:
Advanced Technology Laboratories
Moorestown, N.J.
Tacnet: 253-6643

George W. Kaizar is Unit Manager, Readiness Engineering, at the Missile and Surface Radar Division. His most recent activities are concerned with the development of an integrated AEGIS Combat System Readiness Assessment and Reporting System that builds on existing AEGIS test capabilities. His earlier activities were involved with the development and testing of shipboard automatic test systems such as the AEGIS Operational Readiness and Test System (ORTS), and the Versatile Avionics Shop Tester (VAST) to support Navy carrier avionic systems. Mr. Kaizar holds a BS in Electronic Physics from St. Josephs University.
Contact him at:
Missile and Surface Radar Division
Moorestown, N.J.
Tacnet: 224-2157

Authors, left to right: Kaizar, Gale, Dempsey, and Straguzzi (Adams not available).

Morten Gale is a member of the Readiness Engineering unit at the Missile and Surface Radar Division. He is currently involved in efforts to enhance AEGIS readiness evaluation and reporting systems by the introduction of EXPERT systems and other artificial intelligence tools. His earlier activities included analysis of fault detection and fault isolation capabilities of AEGIS elements. Prior to that he was engaged in systems analysis in such programs as guided missile fusing, Ballistic Missile Early Warning System (BMEWS), postal automation, computer-aided education, etc. Mr. Gale holds a BS and an MS in Electrical Engineering from New York University.
Contact him at:
Missile and Surface Radar Division
Moorestown, N.J.
Tacnet: 224-7718

Nicholas Straguzzi has been a Member of the Engineering Staff of the Artificial Intelligence Laboratory at RCA's Advanced Technology Laboratories since 1983. He participated in expert system development related to the RCA/MSRD AEGIS prototype on-line expert diagnostic consultation program for SPY radar and computer systems. He is currently working on the MP2D Expert Consultation System (MPECS), which applies AI to VLSI design. Past projects included research into intelligent man/machine interface, including graphics programming. His professional competence encompasses LISP, PROLOG, and VAX/VMS programming; systems programming; and computer-aided instruction. Mr. Straguzzi has a BSE in Computer Science from the University of Pennsylvania.
Contact him at:
Advanced Technology Laboratories
Moorestown, N.J.
Tacnet: 253-6646

W.A. Helbig | R.H. Schellack | R.M. Zieger

The design and construction of a GaAs technology demonstration microprocessor

Is gallium-arsenide technology catching up to silicon for integrated circuits?

The technology maturity of gallium arsenide (GaAs) in processing digital integrated circuits is equivalent to that of silicon in the early 1970s. Today we can integrate between 1000 and 4000 GaAs gates on a single circuit. Yields for these medium-scale integration circuits are generally between 10 and 60 percent.

The major factor that currently constrains larger integration levels with better yields is the variation in pinch-off voltage of the GaAs transistors. Improvements with GaAs processing technology are occurring at a rate that is three times what occurred in silicon processing during the 1970s and early 1980s. Many silicon processing techniques that have been refined over the years apply directly to the processing of GaAs. As a result, integration levels approaching 10,000 gates are now being seen.[1]

Abstract: *Because technology has reached the point where we may integrate a few thousand gates on a single chip, RCA's Advanced Technology Laboratories (ATL) is designing and constructing an 8-bit, reduced instruction set microprocessor as a test case for technology development and demonstration. We expect to produce the first operational system by January 1986. This article discusses the technology status, design choices, and project progress and schedule.*

Final Manuscript received November 20, 1985
Reprint RE-31-1-12

With the low gate delays that the technology provides, a GaAs 100-million instructions-per-second central processing unit (CPU) was easy to design. However, construction of a complete system to operate at this speed presents some unique challenges in the designs of the system architecture, demonstration board, and support software. For example, chip-to-chip communication along conduction paths on the board (an insignificant problem in the silicon environment) is greatly complicated because propagation delays are about 5 percent of the CPU cycle per inch of conductor length, which makes control of data skew and delay a significant factor in board layout.

We are using the double-level-metal enhancement/depletion (E/D) process provided by TriQuint Corporation to construct the CPU. The chip is partially handcrafted and partially designed using standard cells, and was laid out using RCA CADDAS tools. A system that consists of the CPU, 512 words of memory, and a minimum I/O capability is being constructed using leadless ceramic chip carriers (LCCCs) and emitter-coupled logic (ECL) glue chips. It will be mounted on a single board made of alumina.

TriQuint's E/D foundry services

RCA is currently using the GaAs enhancement/depletion E/D process provided by TriQuint Semiconductor of Beaverton, Oregon. TriQuint is a newly formed subsidiary of Tektronix, also in Beaverton, Oregon. Research in GaAs technology began at Tektronix in 1978. Two processes are available from TriQuint: depletion-mode and enhancement/depletion-mode (E/D) processing. RCA is using the low-power, high-speed E/D process to build an 8-bit, reduced-instruction set CPU (RISC) microprocessor on a single chip.

The E/D process requires a single, 2-volt power supply. The E/D logic family is designed to have a 0.6-volt logic swing. Logic 0 and 1 levels are represented by 0.2 and 0.8 volts, respectively. The E/D logic family has a typical delay time of 200 ps that is varied based on the gate's output loading. Figure 1 illustrates output delay times with respect to loading for the E/D gate family.[2]

The E/D logic family generated for the 8-bit microprocessor design was created for use in a standard cell layout implementation. In this approach, the layouts for all gates in the family have the same cell height, and the length varies as a function of the gate fan-in. This permits automatic placement and routing of all standard cells.

The E/D logic family consists of 28 gate types for logic design, ranging from an inverter to a 10-input AND/OR logic

gate. Because we used standard cells to implement the logic, the 8-bit microprocessor—except for its general register file and four program counters—requires 1150 gates. We estimate that this level of standard-cell integration, combined with the register file and program counters, should dissipate 427 mW of power for the entire microprocessor circuit.

Alumina substrate packaging

We will use an 84-pin leadless-ceramic chip carrier (LCCC) to package the 8-bit microprocessor circuit. System-level packaging of the microprocessor with GaAs memory chips (in 40-pin LCCC packages) and emitter-coupled logic/transistor-to-transistor logic (ECL/TTL) translators will be done on an alumina substrate. The alumina substrate provides the support platform for the system and interconnection (on multiple levels) of signals between CPU, memory, and ECL input/output (I/O). Typical signal propagation delay time for transmission line interconnects built on alumina substrates is 250 ps/inch. This translates into a delay time per inch that is 5 percent of the microprocessor cycle time.

We considered alternative materials for packaging the system—beryllia, epoxy glass board, and Teflon—but problems with them dictated the use of alumina.

All the interconnects among the microprocessor, memory, and ECL interface parts are uniform transmission lines that have characteristic impedances of 50 to 100 ohms with adequate load terminations. This ensures that the rising- and falling-edge times of 1 ns do not produce unacceptable noise levels that would be caused by mismatched transmission lines.

Architecture

The GaAs technology demonstration CPU's low level of integration limits it to an 8-bit machine. A 16-bit address allows for memory growth to 64k bytes. The instruction set contains 23 instructions; 19 are 8 bits long and 4 are 16 bits long. Our goal of 100 million instructions per second (MIPS) requires that an instruction be fetched every 10 ns. To reach the 100-MIPS goal, we used a 400-MHz clock, a RISC philosophy[3] of simple, fast-executing, hardware-frugal instructions, and a pipeline[4] that permits the concurrent processing of several instructions. Table I lists the CPU processes.

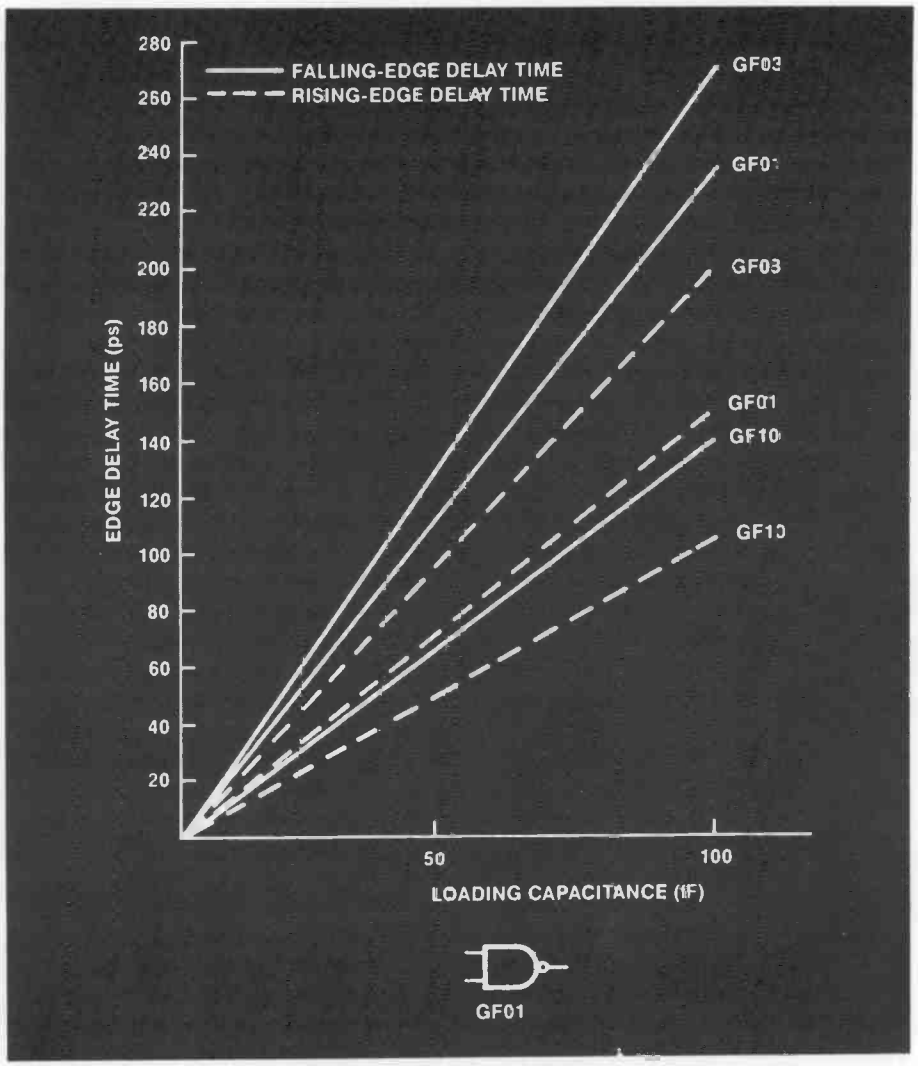

Fig. 1. *Delay versus loading for GaAs E/D standard cell family.*

A single chip-select bit divides the 512 available 8-bit memory locations into two sets of 256 by 4-bit memories. We did this so the lower 256 addresses fall into two of the four memory chips and the upper half of the addresses fall into the other two. The reason for this stems from the nature of instruction and data fetches from memory.

At 100 MIPS, an instruction fetch occurs every 10 ns. The data fetches occur between the instruction fetches. While future memory parts will be able to handle requests every 5 ns, we were uncertain that the available memory parts could do so. With this memory-chip-selection method, we can run two experiments on the system.

A program shorter than 256 (8 bits) instructions can reside in the lower half and access data in the upper half of memory. In this experiment, no memory part is accessed more often than every 10 ns. This experiment uses a form of interleaved memories.

A program that is longer than 256 instructions can request data from the same memory chip as current instructions. This will cause memory chip accesses 5 ns apart. If the system does not permit access at 5-ns intervals, then

Table I. *CPU processes*

CPU Pipestage	Duration (ns)	Phase
Request Instruction	2.5	3
Wait	15.0	4, 1
Receive Instruction	2.5	2
Decode Instruction	5.0	3, 4
Execute Instruction	2.5	1
Write-back Result	5.0	2, 3

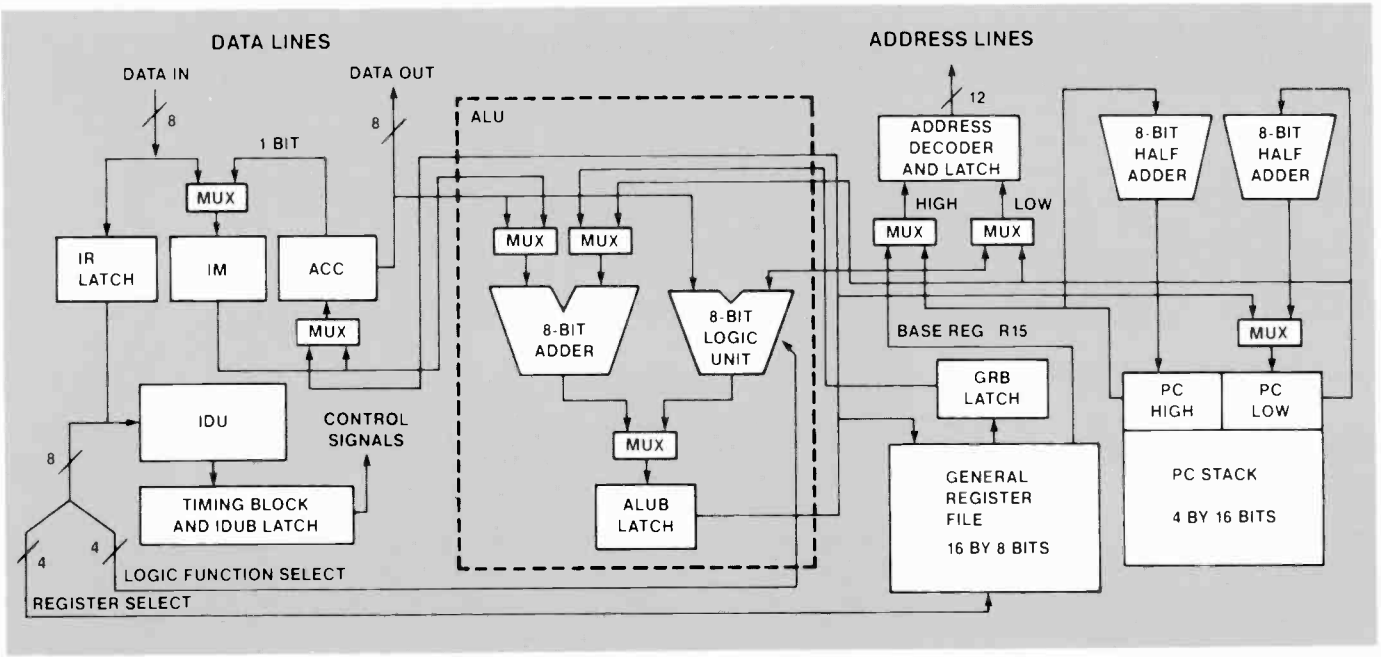

Fig. 2. *CPU block diagram.*

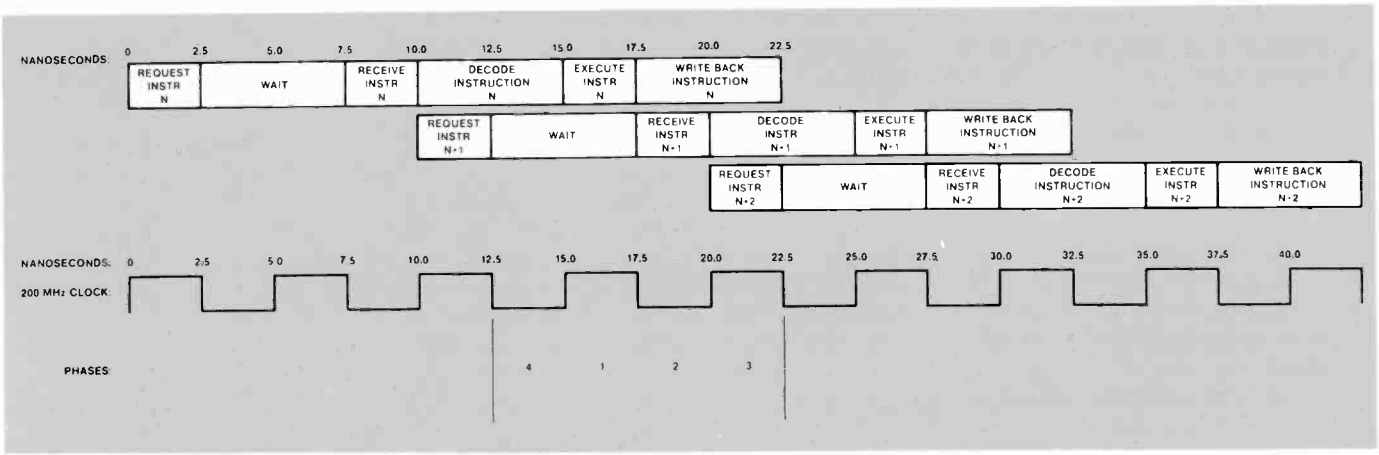

Fig. 3. *Typical instruction sequence.*

we can lower the clock frequency until the system works, and thus test the cycle time of the memory parts as they function in the system.

Figure 2 is a block diagram of the CPU developed from these goals and limitations. The general register stack consists of sixteen 8-bit registers, and the program counter (PC) stack consists of four 16-bit registers. The arithmetic logic unit (ALU) consists of an 8-bit ripplecarry adder, an 8-bit logic unit, and an 8-bit output latch (ALUB). All two-operand arithmetic and logic function instructions take one operand from the accumulator and the other operand from a register in the general register stack. All branching is relative to the value in the PC; the 8-bit offset is found in the 16-bit instruction.

All loops are implemented with a conditional branch. All subroutines require a push, a branch, and a pop. (The current address in the PC is saved, which is the push; the subroutine is executed, the branch; and the original address is retrieved, the pop, and loaded into the PC.) The PC stack exists only to save the return addresses that make subroutines possible. With four PC registers in the PC stack, subroutine nesting is limited to a depth of three.

The load and store instructions use addresses that reside in the general register stack. The most-significant 8 bits of the address are taken from register 15, and the least-significant 8 bits are taken from a general register that is specified in the instruction.

The typical instruction sequence given in Fig. 3 shows that an instruction request occurs every 10 ns and a pipeline depth of two to three instructions is processed at once. The instructions depicted in Fig. 3 represent 8-bit instructions, or the first half of a 16-bit instruction.

The instruction register sends the lower four bits of every instruction to the general register stack as a register address. It sends the upper four bits of every instruction to the logic unit to choose one of the 16 possible logic functions of two operands. The logic unit's result is disregarded for all instructions that are not among the chosen six logic function instructions. The general register stack's output will be disregarded for instructions that do not require an operand from the general register file.

This method assigns opcodes (operation codes) to the two-operand logic function instructions without choice by the designer. The two-operand nonlogic function instructions are chosen next from the remaining unused combinations. The four 16-bit instructions fall into the non-two-operand, nonlogic function category. Their first eight bits are decoded in the same manner as the 8-bit instructions. Their second eight bits are a constant or a branch offset; the constant or offset is anticipated by the decoding of the first eight bits (the opcode) of the instructions. The second eight bits of a 16-bit instruction are destined for the IM register instead of the instruction-register (IR). These constants or offsets are allowed into the IR latch register, but prevented from executing by resetting the IR latch register to zero, which is one of the opcodes for a no-operation instruction.

Chip and system construction

The GaAs technology demonstration system includes the CPU chip, four 256 by 4-bit GaAs memories, and ECL I/O parts. The system is shown in Fig. 4.

The GaAs technology demonstration CPU chip includes custom-designed macrocells for the general register file and the program counter stack. The rest of the chip is implemented with 1150 standard cells. If we had designed the entire chip with standard cells, there would have been about 2000 gates.

RCA's MP2D layout tool allows the engineer to specify the corners in which the macrocells are to be placed. Then, MP2D will fit the standard cells into the remaining chip space as efficiently as it can. Figure 5 shows a floor plan of the CPU chip.

The CPU will be packaged on an 84-pin leadless ceramic chip carrier (LCCC). There will be ten address lines, eight data-in lines, eight data-out lines, one 400-MHz clock input, one reset input, and one read/not-write output. This totals 29 signals. There will be 16 power and ground pins and 39 unused pins.

We have simplified the I/O of the technology demonstration CPU chip. The CPU is kept in a dormant state by the reset pin until the memories have been loaded. Then, the CPU goes into a no-operation instruction loop when the program is finished. The system's I/O involves loading the GaAs memories through the ECL interface parts from a Digital Equipment Corporation VAX 11/780 computer.

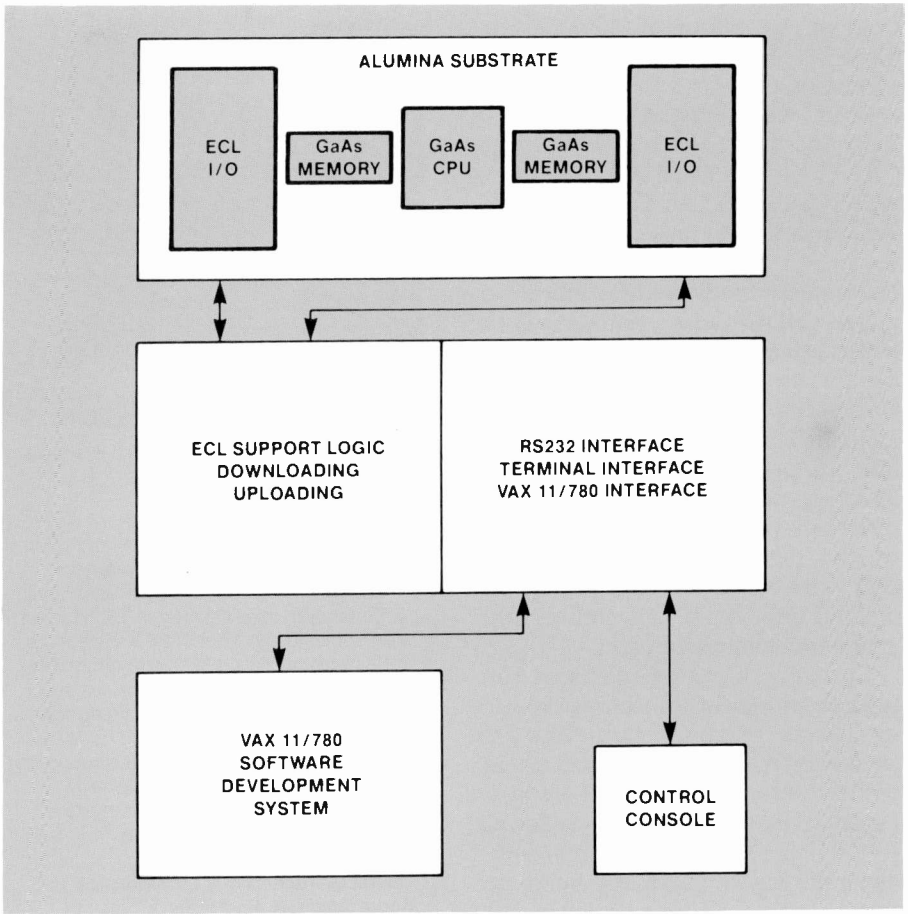

Fig. 4. *GaAs Technology Demonstration System block diagram.*

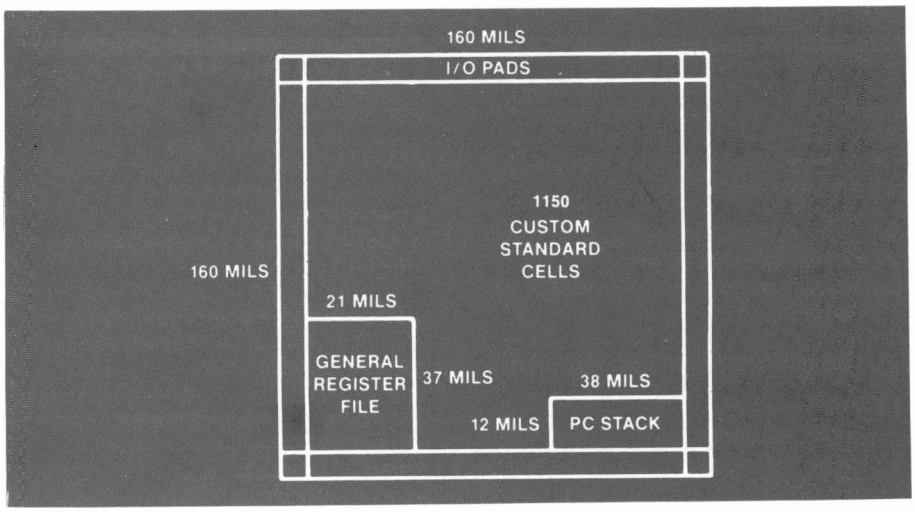

Fig. 5. *CPU chip floor plan.*

Support software/documentation

In the CPU development, we have used several approaches that are consistent with RCA's commitment to use formal hardware-descriptive languages. First, we wrote a behavorial model using ISP′. With this model, the assembler and linker in the ISP′ tool set are being used for software production. Initially, we used the N.mPc simulator to test the instruction set and the functional interrelationships of the system components.

Later on, we will use it to debug software.

Next, as the structural design progressed, we expressed the design in RCA's hardware descriptive language, CADL, and simulated the design using RCA's gate-level simulator, MIMIC. Subsequently, we transferred the design to a hierarchical description using Texas Instrument's hardware description language (HDL), and resimulated the design using Texas Instruments' simulator INTSIM. A translator has been written that fetches the connection list from the HDL database and converts it into the proper format for the RCA MP2D layout program. This provides a direct link from a high-level description of the design—in RCA's standard cells—to a complete artwork file (mask-making drive tapes.) Of course, other programs handle macrocells.

Currently, three other related software development projects are in various stages of completion. First, an assembler/reorganizer program is being written to translate the DARPA core set assembly language into inputs for the ISP' linker, reorganizing (optimizing) it along the way to obtain maximum performance from a pipeline machine.[5,6,7,8,9] Second, a set of diagnostic/demonstration routines is being produced to exercise all aspects of the system finally produced. This will help us debug the system and study construction aspects, with an eye toward expansion. And, third, a loader routine is being produced that, after hand-loading into the system, will enable us to load programs and data directly from the support software host, a VAX 11/780 computer.

Status and plans

On December 31, 1985, we were ready to begin layout of the 8-bit GaAs technology demonstration microprocessor.

Both the 16 by 8-bit general register stack and the 4 by 16-bit program counter stack have been hand-crafted and are complete. Work on the standard cell layouts for all 20 cell types is completed.

Final layout of the complete microprocessor circuit was done after we completed some final circuit simulations. These simulations address the timing and control of all major elements of the microprocessor.

We delivered a complete design tape to TriQuint Semiconductor for generation of masks and fabrication.

Work on the development of the demonstration system and its support software is scheduled to continue into the beginning of 1986.

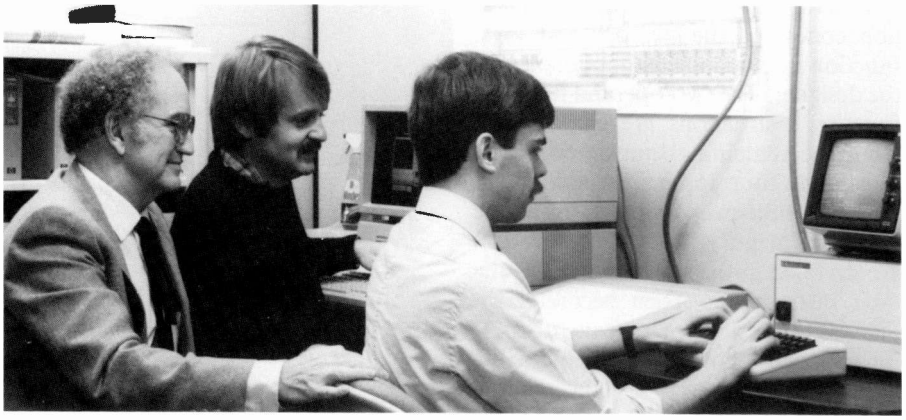

Authors, left to right: Helbig, Schellack, and Zieger.

Walter Helbig is Unit Manager of the Advanced Computation Systems Group in ATL's Systems Engineering Laboratory. He is responsible for the design and implementation of advanced computer system architectures. He directed the design of the GaAs 32-bit RISC architecture, and is participating in the construction of the 8-bit GaAs RISC Technology Demonstration System. Previously, he directed the design of two versions of the RCA ATMAC microprocessor, including all of the supporting software, as well as systems utilizing an ATMAC. Walt has a BSEE degree from the Milwaukee School of Engineering.
Contact him at:
Advanced Technology Laboratories
Moorestown, N.J.
Tacnet: 253-6530

Robert Schellack is a Senior Member of the Engineering Staff in ATL's Microelectronics Laboratory. His work involves GaAs device modeling and circuit analysis for high-speed digital applications. He is incorporating automated design capability for high-speed digital GaAs ICs into RCA's CADDA system. He is also Principal Investigator for a study of advanced GaAs IC technology. Bob has BSEE and MSEE degrees from Drexel University.
Contact him at:
Advanced Technology Laboratories
Moorestown, N.J.
Tacnet: 253-6519

Richard M. Zieger is a member of the engineering staff of ATL's Systems Engineering Laboratory, where he is involved in designing and developing computer architecture. He has worked with a System Architecture Simulator (SARSIM), modeling and simulating a high-speed signal processing system. Recently, he completed the architecture and logic design of an 8-bit gallium arsenide (GaAs) microprocessor, a pipelined, RISC architecture to be implemented in 1-micron double-level-metal enhancement/depletion MESFET technology. The microprocessor is designed to operate at a 200-MHz rate and perform 100 MIPS. Mr. Zieger holds a BS in Computer Engineering from Lehigh University, and has recently begun work on an MSEE at Villanova University.
Contact him at:
Advanced Technology Laboratories
Moorestown, N.J.
Tacnet: 253-6518

References

1. Roosild, Sven A., "Survey of Manufacturing Capabilities for GaAs Digital Integrated Circuits," *1984 IC Symposium Technical Digest*.
2. Flegal, T., LaRue, G., Rode, A., "A High-Yield GaAs Gate Array Technology and Applications," *1983 GaAs IC Symposium Technical Digest*.
3. Hennessy, J.L., Jouppi, N., Baskett, F., Gross, T.R., Gill, J., and Przybylski, S., "Hardware/Software Tradeoffs for Increased Performance," *Proc.*, SIGARCH/SIGPLAN Symposium on Architectural Support for Programming Languages and Operating Systems, ACM, Palo Alto, Cal., pp 2-11, (March 1982).
4. Hennessy, J.L., Jouppi, N., Przybylski, S., Rowen, C., Gross, T.R., Baskett, F., and Gill, J., "MIPS: A Microprocessor Architecture," *Proceedings*, Micro-15, IEEE, pp. 17-22, (October 1982).
5. Gross, T.R., "Code Optimization of Pipeline Constraints," PhD thesis, Stanford University (August 1983).
6. Gross, T.R., and Hennessy, J.L., "Optimizing Delayed Branches." *Proceedings*, Micro-15, IEEE, pp. 114-120, (October 1982.)
7. Gross, T.R., "Code Optimization Techniques for Pipelined Architectures," *Proc. Compcon Spring '83*, IEEE Computer Society, San Francisco, Cal., pp. 278-285 (March 1983).
8. Hennessy, J.L. and Gross, T.R., "Code Generation and Reorganization in the Presence of Pipeline Constraints," *Proc. Ninth POPL Conference*, ACM, pp. 120-127 (January 1982).
9. Hennessy, J.L. and Gross, T.R., "Postpass Code Optimization of Pipeline Constraints," *ACM Trans. on Programming Languages and Systems*, Vol. 5, No. 3, (July 1983).

A. Feller

A new methodology reduces design time for very large application-specific integrated circuits (ASICs)

Doing tasks in parallel and using modern tools shortens the design cycle for new denser chips.

As a reflection of competitive pressures, both the gate counts and complexity of custom VLSI devices continue to increase dramatically, along with the requirement of having these devices designed in a short period of time. Thus, custom chips with three, four, or five times the device count of chips several years ago must still be designed in the same amount of time as chips of the past.

Current design VLSI methodologies, based on past LSI or even MSI design approaches, are responsible for excessively long design times. However, one design methodology that uses current design aids and design tools does reduce the design time (by at least half) of custom LSI and VLSI devices having 30,000 transistors or more; still, the design time does increase as the gate count increases. This goal is achieved with little loss in performance and an area premium of 5 percent to 15 percent depending on the ratio of functional to random logic.

Abstract: *RCA Advanced Technology Laboratories has developed a new design methodology for very large application-specific integrated circuits (ASICs). Unlike conventional design methodologies that implement the various design steps in sequence, this new design approach does many of them in parallel. In addition, chip layouts are initiated with incomplete and unverified logic. This clearly accelerates the chip design schedule because generating the complete logic of the test vectors and verifying the logic is frequently the most schedule-limiting design step.*

Final manuscript received November 20, 1985
Reprint RE-31-1-13

Design methodology should not be confused with design tools. For example, standard cells, handcrafted design, macrocells, automatic placement, and automatic routing are design tools, not design methodologies. Similarly, programs like logic simulators, design rule checkers, logic validation, and circuit analysis programs are design aids, not design methodologies. A design methodology defines procedures and establishes guidelines for utilizing design aids and design tools to design and lay out VLSI devices.

Tasks in VLSI design

Application-Specific Integrated Circuits (ASICs) are chips designed for dedicated, special-purpose, or user-defined applications. The devices are defined by system designers, with the resulting chip design often being categorized as top-down design. Following partitioning, the system designer will generally verify the functionality of the partitioned chip through a high-level simulator at the functional or Register Transfer Language (RTL) level. After verification is completed, the chip described at a functional or RTL level is transferred to the chip designer. This input contains the functional description that includes vectors used for system simulation. Starting at this point, virtually all design methodologies contain the following tasks:

- ☐ Partitioning and preliminary logic design
- ☐ Initial macrocell definition (RAM, ROM, register stack, multiplier, bit slice, ALU, etc.)
- ☐ Macrocell design and validation
- ☐ Logic design
- ☐ Logic simulation (including generation of test vectors)
- ☐ Layout of chip
- ☐ Fault simulation (where pressing schedules are required, it is not unusual to release a chip for fabrication before full designed fault coverage is attained).

Although these tasks are contained in virtually all design methodologies, this one reduces the overall design time because of the unique method for carrying out these tasks. First, the current conventional design methodology will be described.

Figure 1 illustrates how these tasks are performed in current conventional custom VLSI design methodology. For a device containing about 160,000 transistors, of which 35,000 are associated with random logic, typical elapsed times are shown in Fig. 1 for the logic design and layout.

Logic design

The first step in the overall logic design involves initial macrocell definition and logic design starting from functional or hierarchical description of the chip. During the logic design phase, test vectors are generated and logic simulation is implemented. For high-performance VLSI devices, prelayout dynamic calculations are generally accomplished using the delay and timing capabilities of the logic simulator design tool. Including the generation

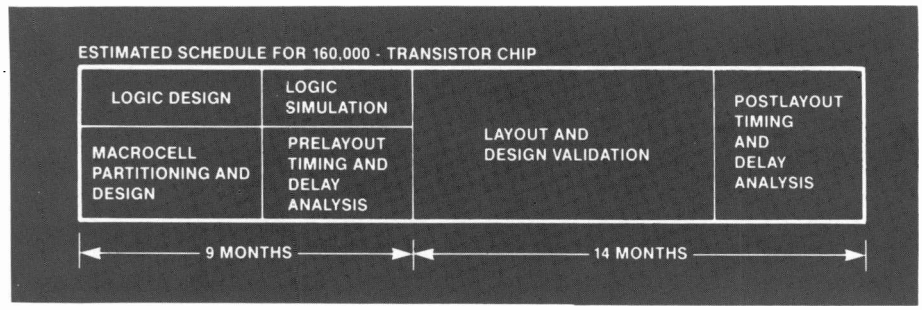

Fig. 1. *Current, conventional design methodology.*

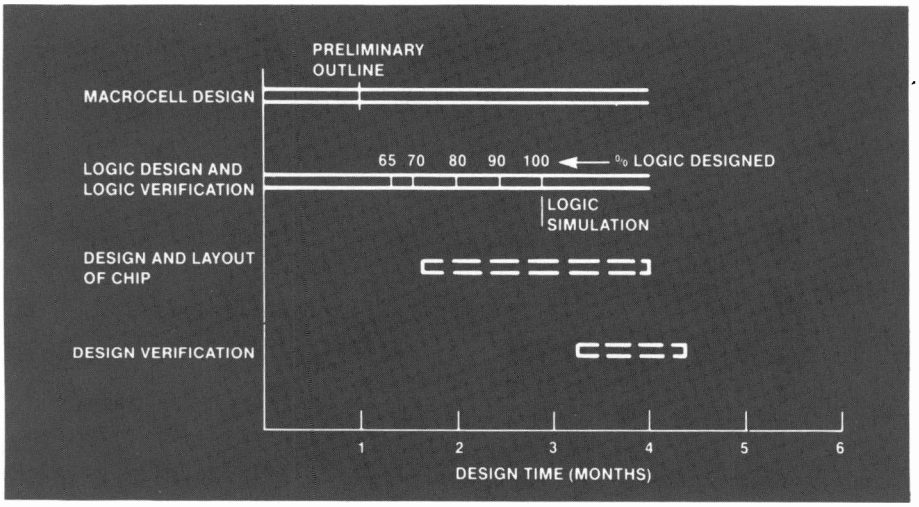

Fig. 2. *The new design methodology. Chip layout depends on macrocell or logic design.*

of additional functional macrocells, nine months of elapsed time is a typical design period that could be expected for chips of this size and complexity.

Layout and design validation

Conventional layout techniques generally start with the generation of a floorplan, which attempts to place the previously defined macrocells and the remaining logic in a relationship that will facilitate the next step, the interconnections (i.e., the routing). During the floorplan generation, additional macrocells will be defined. Because of the difficulty in interconnecting random logic functions on a workstation, or any system supporting interactive routing, it is not unusual to implement the random logic through functional blocks such as Programmable Logic Arrays (PLAs). To ease the routing problem, not only is an efficient and optimized floorplan necessary, but corresponding effort must also be addressed to the orientation of the I/O pins on all of the macrocell functions.

The only effective way to determine if an acceptable floorplan has been generated is to implement the interconnections. After sufficient routing has been implemented to evaluate the floorplan, it is often too late to generate a new floorplan, if that conclusion emerged as a result of the routing phase. As the requirement for higher density increases, this problem becomes proportionally more difficult.

Post layout timing and delay analysis

The next step in the design process is verification that the performance, propagation delay, and clock speed specifications have been met and that all race conditions, clock skew, and timing problems, if they exist, have been brought to the attention of the designer. Generally, a logic simulator containing timing analysis capability is used to measure the performance. However, for acceptable accuracy, the parasitic capacitance and resistance must be included to update the appropriate simulation models. Currently, tools are reported to be capable of doing this for custom, handcrafted designs.

After the model is updated for active load conditions, the timing analysis is repeated. If the dynamic performance fails to meet specification, corrective action must be taken. This corrective action could involve different drivers, different placement, editing of interconnections, and perhaps functional changes. In a conventional handcrafted environment, these changes and modifications are implemented in a much different manner.

Validation

Validating a layout involves a Design Rule Check (DRC) and verifying that the logic reflected in the output layout (artwork commands) is identical to the input logic. A DRC check of a conventional handcrafted layout involves the use of costly, computer-intensive programs such as the PDS or ECAD systems. Verifying the output logic is currently done by manual line-tracing to synthesize the logic from the traced logic a difficult, laborious, and error-prone system. Several systems such as ECAD and PDS have the ability to synthesize transistors and the logic gates from the layout and, by this means, provide a tool to verify the correctness of the output logic. These capabilities are now being evaluated for larger chips where major effort is now directed.

New design methodology using fully automatic placement and routing capability

This new design methodology is based on a fully automatic placement and routing capability, and other supporting tools that virtually perform all of the basic VLSI design tasks in parallel as opposed to the sequential mode that characterizes the conventional VLSI design methodology. This means, for example, that the design of a large macrocell (such as a RAM), logic design, and the layout of the chip are done in parallel. The full implementation of this design methodology and its radical departure from the standard design approaches is described and illustrated in the following sections.

With the new design methodology, the schedule for generating a validated layout is generally independent of the chip layout. This approach operates in several modes, with the different design parameters being limiting factors on the schedule. The overall concept, together with several of its different operating modes, is explained and illustrated in the next section.

Figure 2 illustrates the basic concept of the new design methodology. Note that the design and layout of the macrocell, the design and verification of the logic, and the chip design and layout are accomplished in parallel. The design time to

complete the chip layout is dependent on the macrocell design or the logic design, whichever is longer. The chip layout is shown as a dashed line to emphasize that its length is a variable determined by the macrocell or logic design.

The basic design methodology operates as follows. Consider the design and layout of a chip described at a functional or Register Transfer Level (RTL), which contains a functional macrocell such as RAM, ROM Register Stack, Multiplier, Barrel Shifter, etc. Because of their functional nature, the macrocells can generally be defined earlier in the design cycle. At some earlier point, an estimate is made of its size, aspect ratio, and pin orientation (shown in Fig. 2 as the Preliminary Outline).

Logic design is initiated at the start of the program. Consider that point, as shown in the schedule, when about 70 percent of the logic (still unverified) has been defined. At this point, a netlist (interconnection list) is generated and entered into the automatic placement and routing program that proceeds to lay out the chip automatically. Obviously, since the logic is incomplete, the layout is treated as a preliminary layout. However, the layout can be reviewed and optimized using various support tools that are compatible with the automatic placement and routing capability.

The optimization of the layout continues until 100 percent of the logic is generated and verified. When 100 percent of the logic has been generated, logic simulation and fault grading can occur.

At this point, design verification can also be initiated. Assume that the design and layout of the macrocell is completed before the logic design and logic simulation. Therefore, the logic design becomes the limiting item. As soon as the logic is verified, the layout is validated using special-purpose design-rule checking tools and a program that automatically verifies the output logic.

The first layout that will be used to illustrate the new design methodology is a 38,500-transistor, signal-processing computation chip that contains two macrocells: an 18×18-bit multiplier and a 16×32-bit register stack. The chip was specified to the chip designer at a functional level; therefore, a complete detail logic design had to be implemented. In addition, because of the 50-ns requirement placed on the multiplier, a special optimized Booth algorithm was developed. It was estimated that the design and layout of the multiplier would exceed the logic design time. Therefore, the proposed schedule for the chip, shown in Fig. 3, indicates the multiplier design time to define the chip layout time. As given in the schedule, four months were allotted for the design of the multiplier. The logic design was completed and verified in about two months. In the first few weeks, the multiplier outline, aspect ratio, and pin assignment were generated on a preliminary basis. With the incomplete logic and the preliminary outline of the multiplier, a layout of the chip was generated using the automatic partitioning, placement, and routing capability of the RCA VLSI design system.

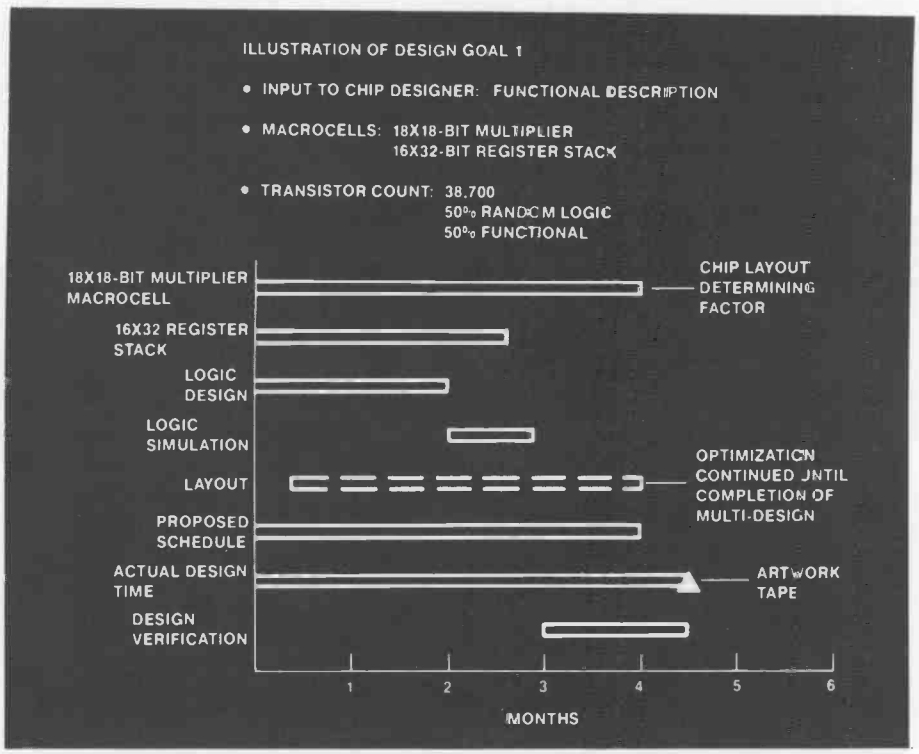

Fig. 3. *Design and layout of a signal processor computation chip.*

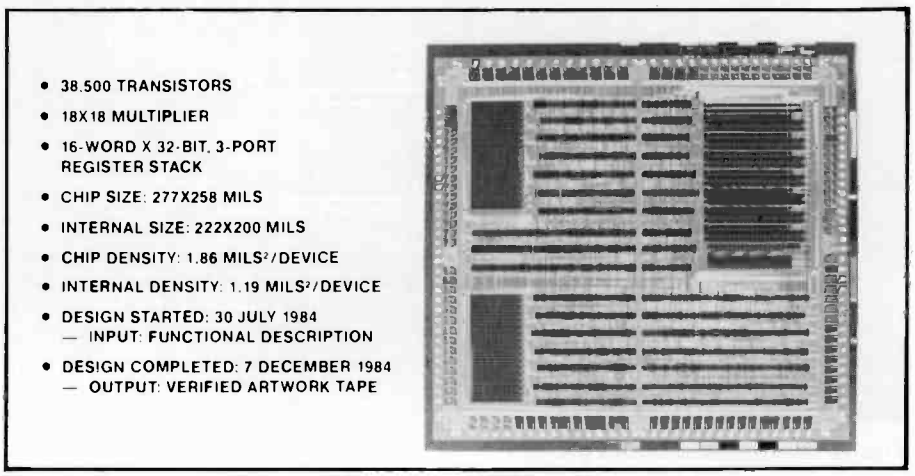

Fig. 4. *The Thetis-B signal processor computer chip.*

Coincident with the design of the logic and multiplier is the optimization of the dynamic performance and the topological design of the chip layout using a series of special-purpose design tools developed specifically for that purpose. For example, one of the tools used for dynamic performance optimization is the CRIPTIC program. This program identifies and selects the critical timing paths automatically. CRIPTIC automatically extracts the parasitics associated with the interconnections and performs delay calculations that include the effect of the layout parasitics.

In addition to the dynamic performance evaluation and optimization, the layout topology is also optimized while the logic

and the multiplier macrocell are being designed. This optimization is based on utilizing the automatic partitioning and placement capabilities to optimize the placement and therefore the layout. In this mode, an engineer might spend up to an hour providing inputs that will enable the placement to be optimized. The chip is then laid out using the automatic routing program. Again, the key to the new design methodology is the fully automatic routing program and the way it is used to allow design functions to be done coincidentally as opposed to sequentially.

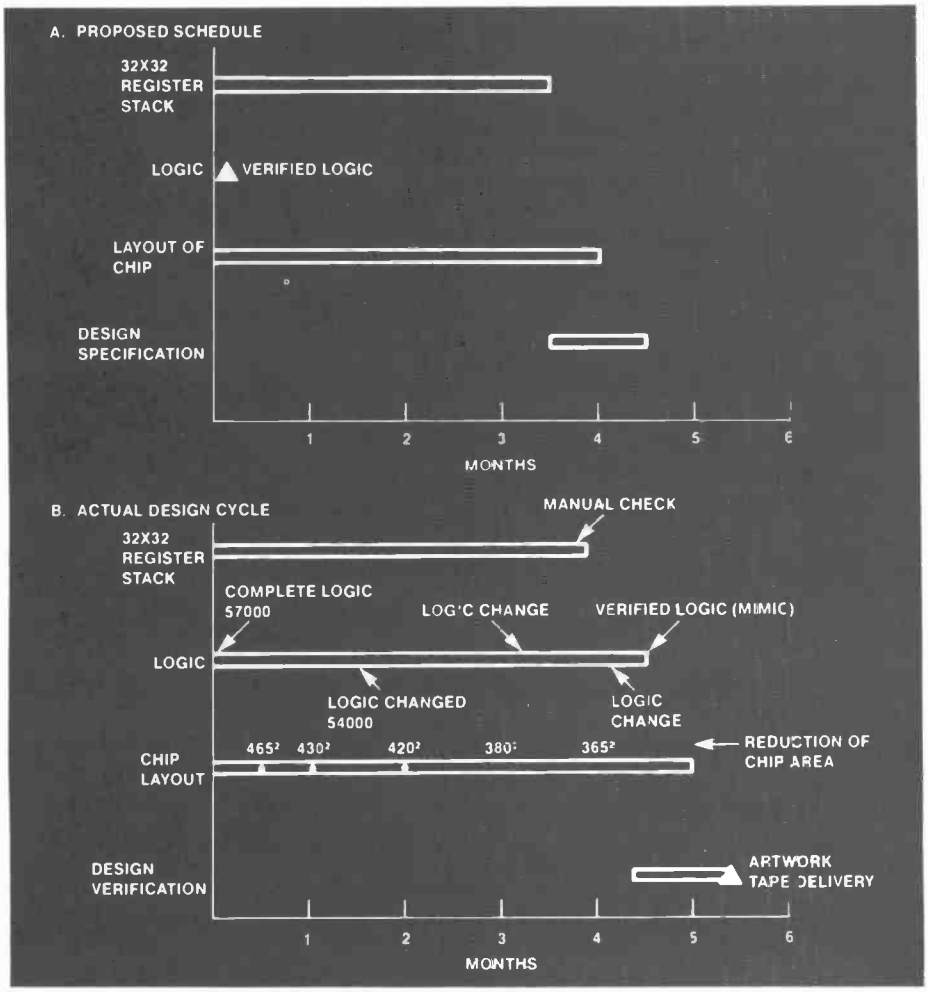

Fig. 5. *Microprocessor design.*

- 54,000 TRANSISTORS
- MACROCELL
  - 32X32 REGISTER STACK
  - STACK DENSITY = 55 MILS²/TRAN.
- CHIP SIZE (S&R) 365X368
  - S&R DENSITY: 2.48 MILS²/TRAN.
  - ACTIVE DENSITY: 1.79 MILS²/TRAN.
- DATE STARTED: NOVEMBER 1984
- ORIGINAL END DATE: 28 FEBRUARY 1985
- DATE DELIVERED: 7 APRIL 1985

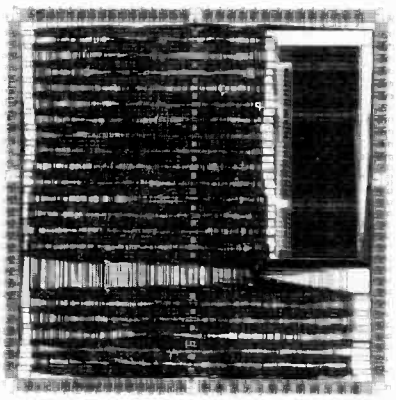

Fig. 6. *Multiple-instruction set microprocessor chip.*

As shown in Fig. 3, the multiplier, which was the schedule-limiting item, was done in four months. Within the next two weeks, the layout was completed, checked for logic correctness and design rules, and mask artwork generated.

Thus, a design that would normally take 12 to 16 months to lay out using the conventional design approaches was completed in 4.5 months. A layout of the actual chip is shown in Fig. 4. Note in the chip statistics that an overall chip density of less than two square mils per device was achieved for a chip containing almost 39,000 transistors, about one-half of which is random logic.

The second layout that will be used to demonstrate the new VLSI design methodology is a 54,000-transistor chip that is part of a sophisticated microprocessor chip set. This chip contains one macrocell, a 32×32-bit register stack. In this case, the logic was described as complete, but not

Al Feller recently transferred to the Microelectronics Center as Manager of the CAD and VLSI Design Laboratory. He is responsible for continuing development of a comprehensive system for automated design, layout, and testing of VLSI/VHSI circuits. He began this work at Advanced Technology Laboratories, where he was Program Manager for several CAD development contracts. In addition to developing RCA's Computer-Aided Design and Design Automation System (CADDAS), his group designed over 100 LSI circuits using five different technologies. Al received BSEE and MSEE degrees from the University of Pennsylvania. He is the recipient of a David Sarnoff Award for Outstanding Technical Achievement.
Contact him at:
Advanced Technology Laboratories
Moorestown, N.J.
Tacnet: 253-6544

verified (test vectors had yet to be generated). It was estimated that the design, layout, and complete verification and validation of the 32×32-bit register stack macrocell would take about four months. Thus, this was established as the original schedule for the entire chip design and layout (see Fig. 5a).

Although the logic was originally described as complete, it in fact was not. The logic experienced at least two major changes and was not verified as correct until more than four months after the start of the program (see Fig. 5b, which documents the actual design cycle). Thus, as given in Fig. 5b, the logic design, not the macrocell design, became from a practical viewpoint the limiting task for the design and layout of the chip.

The basic premise of this new design approach (that the layout itself is not the limiting task) is clearly illustrated by the chip layout task in Fig. 5b. Notice that after the third week of the program, a chip layout was produced whose size is 465×465 mils. This layout reflected the state of the logic design at the time and used a macrocell outline based on preliminary estimates. Within these constraints, the layout was complete and optimization techniques were initiated. First, while the logic and macrocell designs were proceeding, the area of the layout was being reduced by various techniques that are a part of the RCA VLSI design methodology. When the logic was completed and verified using the RCA logic simulator program (MIMIC), the chip size had been reduced to 365×365 mils, about a 30-percent decrease from the original layout. In addition, timing, delay, and critical path analysis and optimization were also done during the design period, as illustrated in Fig. 5b.

Within a month after the logic was finished and confirmed, the design, layout, and verification of the chip were completed and an artwork tape generated to be used by the mask-making facility to fabricate the masks. A layout of the chip is presented in Fig. 6.

The next two devices to be designed with this new design methodology contain 190,000 and 160,000 transistors, respectively. In each case, about 70 percent of the logic appears in the form of macrocells, while the remainder can be characterized as random logic. The schedule calls for a 5-month design cycle starting with just the outlines of the macrocells and at least 70 percent of the random logic. If a higher percentage of the completed logic is available, then either the schedule can be reduced or the chip area reduced for the 5-month schedule.

Summary and conclusion

A new design methodology for ASICs exists and is available. Where conventional design methodologies implement the various design steps in sequence, this new design methodology implements many of them in parallel. In addition, when using this new design methodology, chip layouts are initiated with incomplete and unverified logic. This clearly accelerates the chip design schedule because generation of the complete logic of the test vectors and verification of the logic is frequently one of the most schedule-limiting design steps. The key to this design methodology is the fully automatic placement and routing capability that permits the interactive layouts to be implemented until the logic is complete and verified. If density optimization is required, it can be achieved by additional application of the iterative layout procedure until the desired density is achieved.

It is reasonably clear that as Application Specific Integrated Circuit devices start to exceed the 50,000-transistor level, and if excessive and unacceptable design times are to be avoided, present conventional design methodologies must be replaced by those that have properties and characteristics of the new design methodology just described.

Looking back . . .

At the time of publication of our previous issue (Nov./Dec. 1985), **George Whitley**, RCA Laboratories, author of "Cut and clinch innovations in printed circuit board assembly," was assigned to a special project at Consumer Electronics Operations, Indianapolis, Indiana. George has since returned to RCA Laboratories, Princeton, N.J., and anyone wishing to contact him in connection with his article may do so at Tacnet: 226-2709.

Patents

Aerospace and Defense

Hayman, J.
Multilevel thresholding for target tracking apparatus—4550435

Astro-Electronics Division

Altman, W. P.
High-energy, single longitudinal mode hybrid laser—4554666

Ganssle, E. R. | Miller, C. P.
Reflector antenna mounted in thermal distortion isolation—4550319

Broadcast Systems Division

Bendell, S. L.
Television camera with solid-state imagers cooled by a thermal servo—4551760

Hamalainen, K. J.
Adaptive automatic scan tracking system—4550351

Consumer Electronics Operations

Blatter, H. | Amaral, J. E.
Microprocessor self-turn-off arrangement for a consumer instrument—4544923

Carlson, D. J.
RF Diplexing and multiplexing means—4555809

Chen, K. J.
Television receiver with high voltage protection circuit—4544954

Fling, R. T.
Binary divider as for a digital auto flesh circuit—4550339

French, M. P.
On-off arrangement in a microprocessor controlled remote transmitter for a consumer instrument—4544924

Hettiger, J.
Trilevel sandcastle pulse encoding/decoding system—4549202

Norley, R. R. | Sendelweck, G. K. | Chen, K. J.
On-screen display—4549218

Sendelweck, G. K.
Automatic contrast reduction circuit for a teletext or monitor operation—4549217

Shanley, R. L., 2nd
DC stabilization system—4549203

Shanley, R. L., 2nd | Yost, T. D. | Lagoni, W. A.
Keyed DC stabilization system with protection from error introduction during vertical sync interval—4554577

Shanley, R. L., 2nd
Control system for luminance/chrominance signal processing circuits—4554588

Vanbreemen, B.
Vertical color shift correction in a rear projection television screen—4544946

Willis, D. H.
Error compensated control system in a video signal processor—4554578

RCA Laboratories

Acampora, A.
Circuitry for correcting motion induced errors in frame comb filtered video signals—4553158

Altman, T. N. | Fedele, N. J.
Low cost monotonic digital-to-analog converter—4544911

Batterman, E. P.
Signal transition enhancement circuit—4553042

Blackstone, S. C. | Jastrzebski, L. L. | Corboy, J. F., Jr.
Vertical igfet with internal gate and method for making same—4546375

Botez, D. | Connolly, J. C.
Method of making a laser array—4547396

Carlson, C. R.
Spatial prefilter for variable-resolution sampled imaging systems—4554585

Chin, D.
Auto-tint circuit for a TV receiver—4544944

Corboy, J. F., Jr. | Jastrzebski, L. L. | Blackstone, S. C. | Pagliaro, R. H., Jr.
Method for growing monocrystalline silicon on a mask layer—4549926

Derendinger, M.
Apparatus for applying a layer of powder to a surface—4550680

Dieterich, C. B.
Video disc encoding and decoding system providing intra-field tract error correction—4549228

Dischert, R. A. | Moles, W. H. | Rhodes, R. N. | Walter, J. M.
Digital video transmission system—4550337

Ettenberg, M.
Optical recording medium and information record and method of making same—4547876

Gange, R. A.
Line cathode heater and support structure for a flat panel display device—4551648

Gibson, J. J.
Demodulator of sampled data FM signals from sets of four successive samples—4547737

Gibson, W. G. | Plotnick, M. A. | Chen, T. Y.
Video disc encoding and decoding system providing intra-field track error correction—4546389

Haferl, P. E.
Switched vertical deflection circuit with bidirectional power supply—4544864

Haferl, P. E.
Variable picture size circuit for a television receiver—4547708

Hinn, W.
Video signal DC restoration circuit—4549214

Hurst, R. N., Jr. | Dischert, R. A.
Apparatus for correcting errors in color signal transitions—4553157

Jastrzebski, L. L. | Ipri, A. C.
Vertically integrated igfet device—4554570

Kao, Y.
Auto flesh circuitry as for a digital TV receiver—4554576

Kern, W.
Deposition of borophosphosilicate glass—4546016

Pen and Podium

Recent RCA technical papers and presentations

To obtain copies of papers, check your library or contact the author or divisional Technical Publications Administrator (listed on back cover) for a reprint.

Advanced Technology Laboratories

G.O. Abrams | J.K. Peters
A VAX-Based Personal Computer Ethernet Network—Presented at the 1985 Fall DECUS Symposium, Anaheim, Cal. (12/10/85) and published in the *Proceedings*

G.J. Ammon | J.A. Calabria | D.T. Thomas
A High-Speed, Large Capacity "Jukebox" Optical Disk System—Presented at the 7th IEEE Symposium on Mass Storage Systems, Tucson, Ariz. (11/4-7/85) and published in the *Proceedings*

J.W. Dempsey
Relevant Help, User Modeling, and Uncertainty—Presented at the U.S. Army Research Institute Symposium on Robotics and Artificial Intelligence, Austin Tex. (11/6-7/1985) and published in the *Proceedings*

M.L. Levene
A High Data Rate, High Capacity Optical Disk Buffer—Presented at the 7th IEEE Symposium on Mass Storage Systems, Tucson, Ariz. (11/4-7/85) and published in the *Proceedings*

J.I. Pridgen | D.P. O'Rourke | R.O. Yeager | A.I. Bramble
Agile, Low Power CMOS/SOS Frequency Synthesizer—Presented at the 1985 Government Microcircuit Applications Conference, Orlando, Fla. (11/5-7/85) and published in the *Proceedings*

D.C. Smith | R.N. Putatunda | S.A. McNeary | J.C. Crabbe
HAPPI: Fully Automatic Design for Submicron, Subnanosecond IC Designs with 50k to 400k Transistors—Presented at the Government Applications Conference, Orlando, Fla. (11/5-7/85) and published in the *Proceedings*

R.T. Strong | E.G. Watson | J.R. Tower
Thick Film Hybrid Application to Cryogenic Infrared Imaging Arrays—Presented at the ISHM Conference, Anaheim, Cal. (11/11/85) and published in the *Proceedings*

Astro-Electronics Divison

T.A. Morris
Omnistar On-Orbit Servicing and Module Exchange—Presented at the NASA, JSC, Satellite Services Workshop (11/6/85)

N. LaPrade | H. Zelen
Solid State Power Amplifiers for Communication Satellites—Presented at WESCON '85, San Francisco, Cal. (11/19-22/85)

K.V. Raman | A.J. Calise
On Modal Decoupling Insensitivity—Presented at the AIAA, ASME, IEEE American Control Conference, Boston, Mass. (6/19-22/85)

S.S. Seehra | G.J. Brucker | C.K. Bowman
Study to Establish Data Sheets for CMOS Devices in Space and Nuclear Applications—*IEEE Transactions on Nuclear Science* (12/85)

Automated Systems Division

R.L. Cahoon | N. Meliones | B.M. McDermott
Computer-aided design of a lightweight electronic rack—*RCA Engineer*, Vol. 30, No. 6 (Nov./Dec. 1985)

R.C. Guyer
Mini Laser Rangefinder—*RCA Engineer*, Vol. 30, No. 6 (Nov./Dec. 1985)

S.C. Hadden
Test and Monitoring Developments for US Army Gas Turbines—The Technical Cooperative Program (TTCP) Meeting, NAVAIR, Arlington, Va. (9/85)

J.F. Martin
STE-ICE....A Soldier's Friend—*The Ordnance Magazine*, U.S. Army (8/85)

RCA Laboratories

R. Amantea
An Introduction to Modeling and Simulation—*RCA Review*, Vol. 46 (9/85)

R. Amantea | B. Hwong
Integrated Simulation of CMOS Transistors—*RCA Review*, Vol. 46 (9/85)

J.K. Butler | D.E. Ackley | M. Ettenberg
Coupled-Mode Analysis of Gain and Wavelength Oscillation Characteristics of Diode Laser Phased Arrays—*IEEE Journal of Quantum Electronics*, Vol. QE-21, No. 5 (5/85)

C.R. Carlson | R.W. Klopfenstein
Spatial-frequency model for hyperacuity—*Journal of the Optical Society of America A*, Vol. 2, page 1747 (10/85)

W.R. Curtice | M. Ettenberg
N-FET, a New Software Tool for Large-Signal GaAs FET Circuit Design—*RCA Review*, Vol. 46 (9/85)

G.M. Dolny | H.R. Ronan, Jr. | C.F. Wheatley, Jr.
A SPICE II Subcircuit Representation for Power MOSFETs Using Empirical Methods—*RCA Review*, Vol. 46, (9/85), and presented at the Power Electronics Design Conf., October 15-17, 1985, Anaheim, Cal.

B.R. Epstein
A Program to Test Satellite Transponders for Spurious Signals—*RCA Review*, Vol. 46 (9/85)

T.J. Faith | R.S. Irven | E.P. Bertin
High-current and thermal-shock testing of $TaSi_2$-polycide/Al-alloy composites—*J. Vac. Sci. Technol. B.*, Vol. 3, No. 5 (Sept./Oct. 1985)

M.A. Lampert | R.U. Martinelli
Buffering of charge by the coulomb condensate in non-linear poisson-boltzmann theory—*Chemical Physics Letters*, Vol. 121, No. 1,2 (11/1/85)

P.A. Lyon
Desorption Mass Spectromety-Are SIMS and FAB the Same?—*ACS Symposium Series 291*, American Chemical Society, Washington, D.C. (1985)

C.W. Magee | E.M. Botnick
On the use of secondary ion mass spectrometry in semiconductor device materials and process development—*Mat. Res. Soc. Symp. Proc.* Vol. 48

D. Meyerhofer
Simulation of Microlithographic Resist

Processing Using the SAMPLE Program—*RCA Review*, Vol. 46 (9/85)

S.M. Perlow
New Algorithms for the Automated Microwave Tuner Test System—*RCA Review*, Vol. 46 (9/85)

H. Schade | Z.E. Smith
Mie scattering and rough surfaces—*Applied Optics*, Vol. 24, page 3221 (10/1/85)

H. Schade
Enhancement of oxygen adsorption on silicon following electron irradiation—*Applied Surface Science 24*, North-Holland Amsterdam (1985)

C. H. Steinbruchel
On the sputtering yield of molecular ions—*J. Vac. Sci. Technol. A*, Vol. 3, No. 5 (Sept./Oct. 1985)

L.K. White
A Modelling Study of Superficial Topography for Improving Lithography—*J. Electrochem. Soc.*, Vol. 132, No. 12 (1985)

Missile and Surface Radar Division

A.K. Agrawal | W.E. Powell
Monopulse Printed Circuit Dipole Array—*IEEE Transactions on Antennas and Propagation*, Vol. AP-33, No. 11 (11/85)

W.M. Atwood
Helping the FAA Modernize our Nation's Air Traffic Control System—RCA's Role—*Trend*, Vol. 25, No. 8 (12/85)

Solid State Divison

C.F. Wheatley, Jr. | G.M. Dolny
COMFET—The Ultimate Power Device; A General Study of Power Devices—*Solid State Technology* (11/85)

W.D. Williams | R. Duclos | N.J. Magda
A new approach to surge suppression—*Electronic Component News* (9/85)

Engineering News and Highlights

Dusio named Division VP, Engineering at ASD

The appointment of **Emilio W. Dusio** as Division Vice President, Engineering, has been announced by Eugene M. Stockton, Division Vice President and General Manager, RCA Automated Systems Division.

A 24-year RCA employee, Mr. Dusio is responsible for Automated Systems Division's wide range of engineering services. The position was previously held by Eugene M. Stockton, who was promoted to Division Vice President and General Manager last fall.

Mr. Dusio comes to Automated Systems Division from RCA Astro-Electronics Division, where he served as Manager, Astro Design Engineering. Previously, he held a number of engineering management positions in electronic systems, software engineering, the TIROS-N and DMSP meteorological satellite programs, and various special projects.

Mr. Dusio received Bachelor's and Master's degrees in Electrical Engineering from New York University. He also holds a Master's degree in Computer Science from Monmouth College. Before coming to RCA in 1961, he worked for two years as an engineer at ITT Federal Labs. Mr. Dusio has written a number of technical papers on software and spacecraft testing, and is a member of the IEEE, AIAA, and ACM.

Charles A. Schmidt is Group Vice President, RCA Government Communications Systems

Charles A. Schmidt has been elected Group Vice President, Government Communications Systems. A 35-year RCA employee, Mr. Schmidt previously was Division Vice President and General Manager, RCA Astro-Electronics Division, a position he had held since 1981.

In his new position, Mr. Schmidt is responsible for the RCA Communication and Information Systems Division and RCA Government Volume Production, both located in Camden, N.J. The two organizations provide a wide range of electronic products and systems to government and commercial customers. He reports to John D. Rittenhouse, Executive Vice President, Aerospace and Defense.

After Joining RCA in 1950, Mr. Schmidt held a number of engineering and management positions at Government Communications Systems. Before moving to RCA Astro-Electronics, he was Manager of Integrated Communications Systems, with responsibility for the Navy's Trident submarine Integrated Radio Room and other advanced communications programs. Most recently, Mr. Schmidt was largely responsible for the formation of EOSAT, a joint venture of RCA and Hughes Aircraft that has assumed responsibility for commercial operation of the U.S. Government's Landsat program.

Mr. Schmidt received a Bachelor of Science degree in Electronic Physics from LaSalle College in 1965. He also is a grad-

uate of the Harvard Business School's Advanced Management Program.

Mr. Schmidt is an Associate Fellow of the American Institute of Aeronautics and Astronautics (AIAA), a member of the Board of Governors of the National Space Club, and a member of the American Defense Preparedness Association (ADPA), the Air Force Association (AFA), the Association of the U.S. Army (AUSA), the Navy League, and the Armed Forces Communications and Electronics Association (AFCEA).

Staff announcements

Aerospace and Defense

Preston N. Shamer, Manager, Advanced Program Development, announces the appointment of **Richard H. Smith** as Manager, Advanced Intelligence Programs.

Donald L. Gilles, Staff Vice President, Employee Relations, announces that the following individuals are assigned as full-time Trainer/Facilitators for the Aerospace and Defense Quality and Management Improvement training programs: **Frank G. Adams,** Missile and Surface Radar Division; **Robert E. Park,** Price Systems; and **Joseph Wylen,** Communication and Information Systems Division.

Automated Systems Division

Eugene M. Stockton, Division Vice President and General Manager, Automated Systems Division, announces the appointment of **Emilio W. Dusio** as Division Vice President, Engineering.

Communication and Information Systems Division

Donald J. Parker, Director, Digital Communications and Recording Systems, announces the appointment of **William Blackman** as Manager, COMSEC Programs.

Consumer Electronics

Elliott N. Fuldauer, Manager, Materials, announces the appointment of **Gary L. Dyer** as Manager, Production Control.

Harry Anderson, Division Vice President, Program Management, announces his organization as follows: **Grant A. Adkins,** Manager, Development Projects; **David G. Campbell,** Administrator, Development Projects; **Dudley W. Jaggar,** Administrator, Development Projects; and **Grant A. Adkins,** Acting Administrator, Development Projects.

John M. Weisel, Manager, Strategic Sourcing, announces the appointment of **James G. Bishop** as Manager, Monitor Sourcing.

Willard M. Workman, Director, Product Engineering, announces the appointment of **Tom W. Branton** as Product Manager, Color Television.

Electronic Products & Technology

George D. Prestwich, Staff Vice President and Corporate Quality Executive, announces the appointment of **Paul W. DeBaylo** as Director, Quality Processes and Measurement. In this capacity, Mr. DeBaylo will work closely with all Major Operating Units in developing, implementing, and improving processes and measurement techniques designed to achieve RCA's "Total Quality" objective.

Microelectronics Center

Dennis R. Rickmon, Manager, Computer and Design Services, announces his organization as follows: **Vicki J. Fowler,** Manager, Integrated Circuit Artwork Design; and **Louis J. Judice,** Manager, Engineering Computer Services.

NBC

Michael Sherlock, Executive Vice President, Operations and Technical Services, announces the appointment of **Edgar Fereira** as Manager, System Implementation.

RCA Laboratories

James C. Miller, Acting, CAD System Services, announces his organization as follows: **Rodney L. Angle,** Manager, VLSI Design Systems, and **William M. Cowhig,** Manager, VLSI Design Services.

Dr. Karl Knop, Director, Research, RCA Laboratories Ltd. (Zurich), announces his organization as follows: **Michael T. Gale,** Group Leader, Optics, **Dr. Gunther Harbeke,** Group leader, Solid State Physics; and **Dr. Hans W. Lehmann,** Group Leader, Materials Synthesis and Evaluation.

Louis S. Napoli, Director, Integrated Circuit Research Laboratory, announces the appointment of **Norman Goldsmith** as Staff Scientist.

Solid State Division

Herbert V. Criscito, Division Vice President, Marketing, announces the organization of a newly created Strategic Partnering function as follows: **Lucien P. DeBacker,** Director, Strategic Partnering-Rockwell/RCA A&D/SCI; **Alfred Marmann,** Director, Strategic Partnering—Europe; **Thomas C. McNulty,** Director, Strategic Partnering—Chrysler/Delco Remy; and **Melvin L. Petersen,** Director, Strategic Partnering—Ford.

Professional activities

ATL paper wins award

The International Society for Hybrid Microelectronics named the posterboard presentation "Thick Film Hybrid Application to Cryogenic Infrared Imaging Arrays" the outstanding paper of the November 1985 conference. The paper, which focused on system requirements for mounting and packaging silicon Schottky barrier diode CCD detector chips and related successful design solutions, was a joint effort of ATL's **R. Thomas Strong, Edwin G. Watson,** and **John Tower,** and consultants **Robert R. Bigler** and **Samuel Goldfarb.**

Schloss awarded PhD

Rober Schloss, ATL, successfully defended his dissertation on the semiconductor laser, and received his PhD in Electrical Engineering from MIT in February.

Zarodnansky receives BEE

David Zarodnansky, Associate Member, Technical Staff, RCA Laboratories, received the BEE (with honors) from Villanova University. He is presently working on the MEE.

Two receive P.E. licenses

Nancy Graves, Consumer Electronics Operations, has received Indiana Professional Engineer's license number 21432.

John P. Larocco, Missile and Surface Radar Division, has been awarded New Jersey Professional Engineer's license number G30681.

New SSD MOSFET handbook available

A new handbook from RCA Solid State Division, *Circuit Ideas for Linear ICs,* contains 102 practical circuits using one or more RCA linear ICs and MOSFETS. Each circuit is accompanied by a full schematic and a brief functional description.

The handbook is 34 pages long, and is divided into nine sections. Included are sections on timing, measurement, modulation, power control, data conversion, communications, and alarm/monitoring. For more information, contact **Lon Cantor,** Tacnet: 325-6423.

Farrey is Engineer of the Month

William Farrey, Software Engineering, Communication and Information Systems Division, was named Engineer of the Month for August in recognition of his contribution to the development of sophisticated software that is presently being used on the GPCP, Watchmate, and Data Concentrator systems. Mr. Farrey developed IBM PC programs that allow the user to enter wire lists directly into the computer. The Cable Marker program automatically formats the "to" and "from" information.

Keith recognized by Eta Kappa Nu

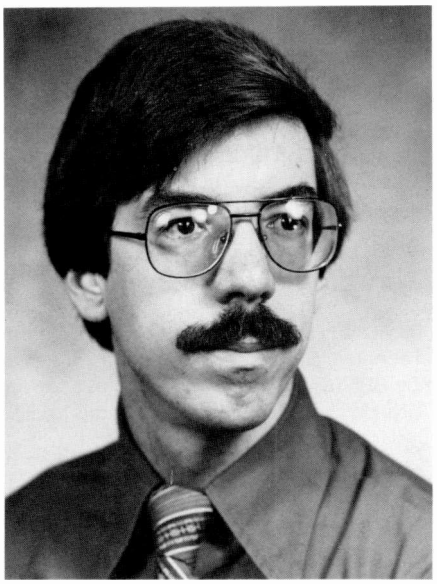

The 1985 Eta Kappa Nu (Honorary Electrical Engineering Society) Jury of Award has selected **Michael J. Keith,** Member, Technical Staff, RCA Laboratories, for Honorable Mention in the award program recognizing outstanding young electrical engineers. Mr. Keith is being recognized for "his contributions to the fields of computer-composed music and teletext systems, and for his involvement in church and cultural activities." **Mark G. Adamiak,** a Senior Engineer at American Electric Power Service Corp., Columbus, Ohio, has been named the winner. Two other young electrical engineers have been selected for Honorable Mention:

Harold A Hoeschen, AT&T Bell Laboratories, Holmdel, N.J.

Harvard S. Hinton, AT&T Bell Laboratories, Naperville, Ill.

Two others are being recognized as finalists.

Mr. Adamiak, Mr. Keith, and the others receiving Honorable Mention, including those being recognized as finalists, will be honored at the 50th Anniversary Award Banquet on Monday evening, April 21, 1986, at the Union league, 140 South Broad Street, Philadelphia. Each winner will receive an appropriately inscribed certificate, presented by the President of Eta Kappa Nu. Mr. Adamiak's name will be engraved on a bowl that is kept at IEEE Headquarters.

The Award Banquet will include a special 50th Anniversary Celebration. All past winners and those who had received Honorable Mention are being encouraged to attend the banquet. As part of the event, three past winners, one from each of three eras of the award (1936-52, 1953-69, and 1970-85) have been invited to address the audience briefly.

Since 1936, Eta Kappa Nu has annually recognized outstanding young electrical engineers. The purpose of this recognition is to "emphasize among electrical engineers that their service to mankind is manifested not only by achievement in purely technical affairs, but in a variety of other ways. It holds that an education based upon the acquisition of technical knowledge and the development of logical methods of thinking should fit the engineer to achieve substantial success in many lines of endeavor."

In the past 50 years, fifty young engineers (including the 1985 winner) who were less than 35 years of age, and who had received their Baccalaureate degree less than 10 years before, have received the award, and 108 others of similar characteristics have received Honorable Mention. The most recent RCA employee to be named the winner is John G.N. Henderson (RCA Laboratories), who was selected in 1977. The most recent employee to receive Honorable Mention is Robert P. Parker (RCA Consumer Electronics Operations), who was selected in 1984.

If you would like further information concerning this award program, please refer to page 20 of the July/August 1984 issue of the *RCA Engineer,* or contact **Jim D'Arcy,** RCA Astro-Electronics Division (Tacnet 229-2359). If you have ever been a member of Eta Kappa Nu, the RCA Technical Excellence Center (Princeton) would like to know. Please inform **Robin Deal** (Tacnet 226-2410) and include your chapter.

RCA contributes to TV handbook

The most recent edition of the *Television Engineering Handbook,* edited by K. Blair Benson (McGraw-Hill, New York, 1985, 1478 pages) has several RCA authors and contributors. It has a Foreword by Donald G. Fink, editor of the first edition (1957), and a Preface by K. Blair Benson. The RCA authors and contributors are:

Oded Ben-Dov, Broadcast Systems Division (coauthor, Chapter 8)

A.D. Cope, RCA Laboratories, retired (contributor, Chapter 11)

Robert N. Hurst, Broadcast Systems Division (contributor, Chapter 17)

Anthony H. Lind, RCA Camden, retired (coauthor, Chapter 17)

Robert G. Neuhauser, Broadcast Systems Division (author, Chapter 11)

Krishna Praba, Broadcast Systems Division (coauthor, Chapter 8)

Dalton H. Pritchard, RCA Laboratories (author, Chapter 21).

New AIAA Associate Fellowships

The American Institute of Aeronautics and Astronautics has announced that the following engineers at Astro-Electronics Division have been upgraded to Associate Fellows: **L. Berko, R. Buntschuh, R.J. Cenker, R. deBastos, E.L. Elizondo, S.M. Fox, J.A. Frohbieter, E.R. Ganssle, S.J. Gaston, P.G. Goodwin, J.J. Horan, A.J. Manna, R. Miller, S. Palocz, G.E. Schmidt, Jr., R.D. Scott, J.F. Seliga, J.F. Swale, P.C. Wise, L.P. Yermack,** and **H. Zelen.**

Third-quarter MSRD winners

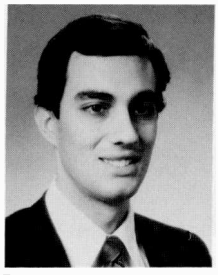
Bannar

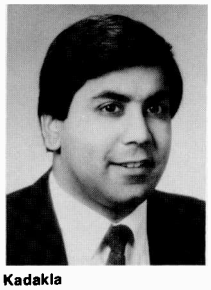
Kadakia

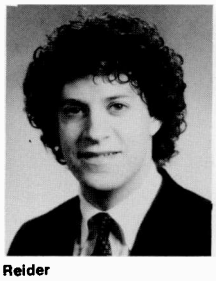
Reider

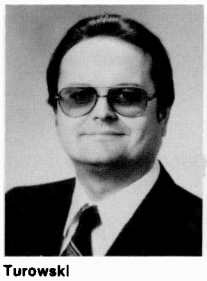
Turowski

The four winners of the third-quarter Missile and Surface Radar Division Technical Excellence Awards are as follows:

Carl P. Bannar—for special dedication and contributions to the integration of the Element Test Function (ETF) maintenance system into the AN/SPY-1B radar. Mr. Bannar personally resolved a number of complex anomalies, enabling the ETF system to function smoothly with the AN/SPY-1B signal processor. His effort was a major factor in the timely and successful testing of the ETF in the formal Navy AN/SPY-1B Operation Evaluation.

Sarit Kadakia—for major contributions to the development and integration of the AN/SPY-1B radar tactical programs, leading to the successful Operational Evaluation. Mr. Kadakia's perserverance was particularly effective in troubleshooting and problem isolation in hardware/software interfaces as well as those between computer programs. Another key contribution was his modification of software to enable operability testing to begin before the availability of the Operational Readiness Test System.

Robert L. Reider—for exceptional skill and initiative in developing an advanced detection algorithm that affords significantly increased detection performance against very low cross-section targets in the presence of point clutter. By applying a highly selective approach to multiple-pulse doppler sequence correlation, Mr. Reider's approach overcomes severe ghosting limitations of conventional detection algorithms. Initial tests of the new algorithm indicate a significant increase in target-handling capability.

Bernard M. Turowski—for conceptual configuration, detailed implementation and test verification, and team leadership in the successful design of the Multiple Althorithm Programmable Processor (MAPP). This powerful new radar signal processing assembly provides programmable processing functions at processing speeds of 960 million operations per second. Mr. Turowski's leadership of and dedication to this project over an extended period firmly establishes RCA's leadership position in radar and signal processing.

Third-quarter CE TEC Awards

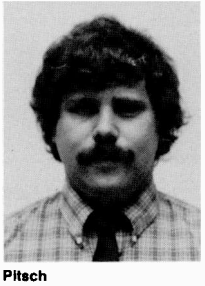
Pitsch


Benson

There were two winners of the third-quarter 1985 Consumer Electronics Operations Technical Excellence Awards. The winners, announced on December 11, are:

Robert Pitsch—for making a CAV (Dimensia) compatible satellite receiver in less than three weeks to support a corporate announcement concerning future DBS business. Although the satellite receiver and descrambler that were used were off-the-shelf items, they were not compatible. The units also had to be interfaced to a microcomputer controller that communicated with the front panel controls of the new package and the Dimensia control bus. Bob designed the necessary hardware and software, procured the equipment, wrote the software, and delivered the project on schedule.

Walter Benson—for coordinating the funding, design, and development of a 14-station in-line hot stamp system for plastics finishing operations on TV cabinet fronts. This system is the first in the industry and the project required extensive interaction between vendor and plant personnel. Mr. Benson organized appropriate training and maintenance programs for both manufacturing and maintenance personnel. He has demonstrated an exceptional degree of thoroughness and professionalism in every detail of the project and has been responsible for the successful completion of the program.

Best is second-quarter winner at MSRD

William Best has received a Technical Excellence Award for outstanding contributions to the integration and testing of the EDM-4 Signal Processor. Mr. Best came to the assignment as a junior team member with no previous integration experience, and has become one of the team's most valued assets after only 18 months. Of particular note are his expertise in microprocessor firmware analysis and his ability to troubleshoot all modes of the system.

PBG fourth-quarter TEC Awards

Palm Beach Gardens presented two fourth-quarter Technical Excellence Team Awards:

C. Hardy and **A. Ladick**—for the conception, design, and development of a high-temperature test apparatus for packaged integrated circuits. This resulted in accurate testing of packaged devices at elevated temperatures, reducing characterization time and manufacturing costs.

P. Chiovarou, H. Foxman, and **E. Jordan**—for establishing a process and procedure to prepare ion scan glass wafers to monitor ion implanters. This resulted in a reduced cycle time for sample preparation at a substantial cost reduction.

Eight at Somerville receive 4Q TEC Awards

The Somerville Technical Excellence Committee has announced eight winners of its fourth-quarter Technical Excellence Awards:

Tom Pampalone and **Frank Kuyan**—for their efforts in developing a photo-resist

system that minimizes resist dimensional variations over steps in 1.25-micron design rule circuits. This current generation of resist has allowed variations of less than ±1 micron at gate level.

R. Isham—for developing a "Sparkle Test System" addressing the need to quantatively define the upper frequency limits where any high-speed A/D converter ceases to function properly. Besides the ability of the system to detect full-scale errors, it can also be set for errors as low as one or two bits.

James J. O'Keefe—for developing a very flexible, operator-friendly control language program that automatically writes an RCAP file containing the device model parameters, the topology, and command files for single or multiple 3-micron standard cells. Although developed in conjunction with the small ICBM project, it can be applied to all families of standard cells and gate arrays, and may also be expanded to include multiple logic families. Autogate virtually eliminates the chances of error in preparing RCAP files for simulations of standard cells.

I.E. Martin—for establishing a library of RCAP model parameters for the 3-micron Rad Hard CMOS/SOS process. This has enabled the Hi-Rel Design Group to evaluate the Standard Cell Library for functionality and propagation delay as a function of temperature and radiation level and a number of process variations. Although this work was done in conjunction with the small ICBM contract, it can be universally applied to the study of any application of the Rad Hard CMOS/SOS technology.

R. Giordano, M. Low, and **M. Hagge**—for developing the AIU circuit on the first cut, and in a record time of five months. Due to their efforts, Consumer Electronics was able to proceed with an aggressive development schedule of their next-generation chassis. Due to the fact that this circuit worked on the first cut, it can represent sales of $7.5 million over the period 1987 to 1991, and it opens the doors to new business with Consumer Electronics and other consumer manufacturers.

SSD third-quarter TEC Awards

Left to right: Carl Turner, Joe Siwinski, Tuan Bui, Carmine Salerno.

Left to right: Carl Turner, Chester Leoszewski, Carmine Salerno.

There were three winners of the Solid State Division third-quarter Technical Excellence Award:

Tuan Bui and **Joseph Siwinski**—for designing and developing an 8×8-bit parallel multiplier as the first building block type of a DSP line of devices. The device will be built in both QMOS and SOS technologies.

Chester Leoszewski—for developing and implementing a full ac circuit probe test capability at wafer probe. This method of circuit test has led to yield improvements of from 50 to 90 percent on complex circuits.

Editorial Representatives and Technical Publications Administrators

Contact your Editorial Representative at the Tacnet numbers listed below to plan your *RCA Engineer* article and to announce your professional activities.

***Technical Publications Administrators, responsible for review and approval of papers and presentations, are indicated here with asterisks before their names.**

Advanced Technology Laboratories

		Tacnet
*Merle Pietz	Moorestown, New Jersey	253-6429
Ed Master	Moorestown, New Jersey	253-6436

Aerospace & Defense Staff

*Theresa Law (Approvals only)	Cherry Hill, New Jersey	222-5833

American Communications

*Anita Labib (Approvals only)	Princeton, New Jersey	258-4346
Carolyn Powell	Princeton, New Jersey	272-4194

Astro-Electronics Division

*Frank Yannotti	Princeton, New Jersey	229-2544
Carol Coleman	Princeton, New Jersey	229-2919

Automated Systems Division

*Gene Galton	Burlington, Massachusetts	326-3775
Linda Kriesman	Burlington, Massachusetts	326-3001

Communication and Information Systems Division

*Dan Tannenbaum	Camden, New Jersey	222-3081
Harry Green	Camden, New Jersey	222-2423
Bruce White	Camden, New Jersey	222-3908

Consumer Electronics Operations

*Eugene Janson	Indianapolis, Indiana	422-5208
John Hopkins	Indianapolis, Indiana	422-5217
Larry Olson	Indianapolis, Indiana	422-5117

Corporate Information Systems & Services

*Sue Handman	Cherry Hill, New Jersey	222-6242

Global Communications

*Dorothy Unger	Piscataway, New Jersey	335-4358

Microelectronics Center

*Susan Suchy	Somerville, New Jersey	325-7492

Missile and Surface Radar Division

*Don Higgs	Moorestown, New Jersey	224-2836
Graham Boose	Moorestown, New Jersey	253-6062
Jack Friedman	Moorestown, New Jersey	224-2112

National Broadcasting Company

*Bob Mausler	New York, New York	324-4869

Network Services

		Tacnet
*Bill Brown	Princeton, New Jersey	272-7601

New Products Division

*P.G. Bedrosian	Lancaster, Pennsylvania	227-6880
Bob McIntyre	Ste Anne de Bellevue	514-457-9000

Patent Operations

George Haas	Princeton, New Jersey	226-2888

RCA Cylix

Paul Dixon	Memphis, Tenn.	739-0214

RCA Laboratories

*Julie Dann (Approvals only)	Princeton, New Jersey	226-2061
Eva Dukes	Princeton, New Jersey	226-2882

RCA Records Division

*Greg Bogantz	Indianapolis, Indiana	424-6141

RCA Service Company

*Murray Kaminsky	Cherry Hill, New Jersey	222-6247
Dick Dombrosky	Cherry Hill, New Jersey	222-4414
Ray MacWilliams	Cherry Hill, New Jersey	222-5986

RCA Technical Excellence Center

*Tony Bianculli (For corporate approvals)	Princeton, New Jersey	226-2111

Solid State Division

*John Schoen	Somerville, New Jersey	325-6467
Sy Silverstein	Somerville, New Jersey	325-6168
Virginia Pate	Somerville, New Jersey	325-6704
Harold Ronan	Mountaintop, Pennsylvania	327-1473 or 327-1264
Dick Morey	Palm Beach Gardens, Florida	722-1262
John Young	Findlay, Ohio	425-1307

Video Component and Display Division

*Ed Madenford	Lancaster, Pennsylvania	227-6444
Lou DiMattio	Scranton, Pennsylvania	329-1435
Nick Meena	Circleville, Ohio	432-1228
J.R. Reece	Marion, Indiana	427-5566

RCA Engineer

A technical journal published by
the RCA Technical Excellence Center
"by and for the RCA engineer"

13 Roszel Road, P.O. Box 432, Princeton, N.J. 08540-0432
Address correction requested

BULK RATE
US Postage
PAID
Phila., Pa.
Permit No. 2906

Printed in the U.S.A.
Form No. RE-31-1