A LOW-COST SOLUTION TO SIMULATED GROUND-BASED RADAR SYSTEMS

Robert J. Sawler
Hughes Training, Inc., Link Operations
Binghamton, NY

Douglas J. Mielke
Hughes Training, Inc., Link Operations
Binghamton NY

ABSTRACT

Successful simulation results in a true immersion into a synthetic environment. This is the goal for which all simulation engineers strive. Along with that goal comes bounds, limitations and design constraints. These considerations are the compromise between accuracy, realism, schedules, and cost. With the economic conditions and the trend toward cutbacks, producing quality, high-fidelity training devices at a low cost has become of paramount importance in the quest of winning contracts. Ascertaining that optimal approach can be a very difficult task for all parties involved in the process. All issues for each specific application must be addressed and a thorough understanding of problems facing the design engineers must be defined.

Today's challenge is to produce low-cost, computationally complex software systems for real-time radar simulation. Fortunately, there are now avenues for simulation designers to accomplish this, with the advent of inexpensive, mass-produced, high-powered processors that are currently available. This paper discusses a low-cost solution to a simulated ground-based radar system using PC-based technologies and off-the-shelf products. The paper starts with a review of classical approaches to radar simulation. It defines the problem facing design engineers who must choose the delicate balance between low-cost and high-fidelity simulation. It introduces the development methodologies that cover the up-front engineering design approaches. The paper then presents the design solutions for a particular application using Commercial Off-The-Shelf (COTS) and innovative graphical techniques. Finally, it makes recommendations concerning future directions of other applicable systems.

ABOUT THE AUTHORS

Douglas J. Mielke is a Systems Engineer with Hughes Training, Inc., Link Operations in Binghamton, New York. He is currently heading up Link's F-16 off-line development PC lab. His most recent assignment was the design, development and implementation of Link's Precision Approach Radar simulation. He has over seven years' experience in PC-based simulations of propulsion systems for the F-16. Mr. Mielke holds a Bachelor of Aerospace Engineering and Mechanics from the University of Minnesota.

Robert J. Sawler is a Senior Systems Engineer with Hughes Training, Inc., Link Operations in Binghamton, New York. He is currently working on Link's programmable Digital Radar LandMass Simulation (pDRLMS). He was the lead for Link's development of the Precision Approach Radar simulation. In the past, Mr. Sawler has been involved in the design, development, and implementation of motion and control systems for both fixed and rotary wings simulators. He has over ten years experience in motion and control systems and over four years experience in the networking of high and selected fidelity simulators. Mr. Sawler has a Masters of Electrical Engineering degree from Syracuse University and a BS in Electrical Engineering from Rutgers University. He has published several papers on the subjects of cue correlation, networking and computer-generated forces.

A LOW-COST SOLUTION TO SIMULATED GROUND-BASED RADAR SYSTEMS

Robert J. Sawler
Hughes Training, Inc., Link Division
Binghamton, NY

Douglas J. Mielke
Hughes Training, Inc., Link Division
Binghamton, NY

INTRODUCTION

The rising need for high-fidelity training in a time when customers are increasingly constrained by shrinking budgets has produced a requirement to create innovative training solutions that are both low cost and high fidelity. Traditionally, the level of fidelity of a training system has been directly proportional to the cost, with high-fidelity simulated systems requiring the support of large engineering teams, using expensive hardware solutions. In order to survive in today's marketplace, the fidelity must continue to increase while the cost decreases. These factors have created the opportunity in the simulation industry to focus attention on exciting new approaches to radar simulation design.

CLASSICAL APPROACHES TO RADAR SIMULATION

The classical approach of radar simulation can be broken into two parts. One approach is an effects level model, where the emphasis is on providing the correct display appearance. The second approach uses energy level modeling to track the emission of microwave energy, its interaction with the environment and the resulting signals. (D. Tucker and K. Collom) This paper focuses on the fixed ground-based radar systems employing the effects level modeling. This approach concentrates on the resulting screen display due to a radar operator's input, not on the energy as a function of the environment.

DESIGN OBJECTIVE - A LOW-COST, HIGH-FIDELITY PRECISION APPROACH RADAR SIMULATION

The design objective was to build a low-cost Precision Approach Radar (PAR) simulation system. A Precision Approach Radar is a fixed ground-based radar, which a trained controller can use to help guide both military and commercial pilots to a safe landing from as far away as forty nautical miles. The Precision Approach Radar had to be networked to an Air Traffic Control (ATC) tower simulation, as well as an Approach Radar simulation. The targets (aircraft) had to be correlated among all of the systems.

The Precision Approach Radar consists of two major pieces of equipment. The first piece is the transmitter/receiver antenna equipment and the second is the indicator equipment (the radar display box). The Precision Approach Radar has two parabolic reflector antenna assemblies (one Elevation, one Azimuth) mounted together on a large single antenna mount (ITT Gilfillan). Figure 1 illustrates the PAR antenna assembly.

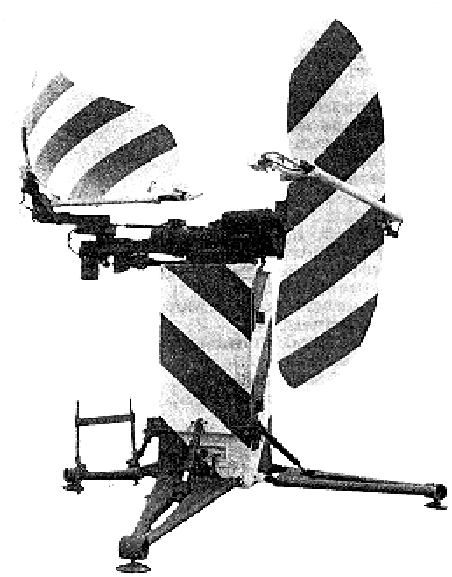

Figure 1
PAR Antenna Assembly

The Precision Approach Radar indicator equipment consists of a navigational computer, a marker generator, a radar set control panel, a power distribution panel and the indicator (radar display). Figure 2 is representative of the equipment contained as part of the PAR indicator group.

The specification required a Precision Approach Radar (PAR) simulation system that was a realistic representation of a real world PAR integrated with a simulated Air Traffic Control Tower and a Terminal Approach Radar. A high level list of the Precision Approach Radar simulation functional requirements are described in Table 1.

One task of the PAR simulation was to duplicate the rectangular (beta scan) radar sweep while also being required to present the targets (aircraft), reflectors, cursors and ground clutter with dynamic brightness. This meant that the levels of brightness had to vary as a function of the antenna sweep position while the sweep travels up and down the screen, and have the corresponding phosphor decay as displayed in the real world equipment.

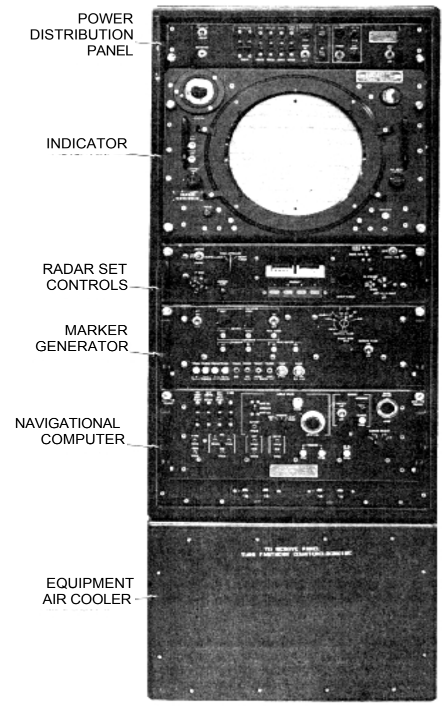

Figure 2

PAR Indicator Group

Table 1 Radar Simulation Functional Requirements

HIGH VOLTAGE	High voltage is switched on / off. This results in a display with no information from either azimuth or elevation antennas.
NARROW PULSE	The pulse length is reduced from 0.625microsec to 0.125microsec.
SAFETY CURSOR	A safety cursor is displayed below the glide path.
RUNWAY SELECT	This function provides the capability of switching between runways 1 and 2.
GLIDE PATH	The glide path display is adjusted for 3 degrees or 6 degrees.
RANGE	The range display is adjusted for 2, 10, or 20 nautical miles.
ANTENNA POLARIZATION	Selects the linear or circular antenna polarization.
FTC CONTROL	Selects between linear FTC, logarithmic FTC or linear video modes.
HEIGHT FINDER	Selects final approach or height finder operation. When height finder is selected it is possible to move the height finder cursor up and down. The altitude is continuously displayed in height finder counter.
ANTENNA ADJUST	Azimuth of the elevation antenna is adjusted to the right or left (azimuth of elevation cursor will move on the screen), elevation of azimuth antenna is adjusted up or down (elevation of the azimuth cursor will move on the screen).
LO TUNE	Adjust fine tuning of the local oscillator for elevation and azimuth antenna echo signals.
IF GAIN	Controls the gain of the IF amplifier for the azimuth and elevation displays.

The Design Challenge

The design objective was to build a Precision Approach Radar which, when reviewed as an individual task, seems to be a straightforward mundane engineering effort. The project became more complicated when the low-cost constraints were levied in an attempt to make the design fit into a tight budget. There are many radar model simulations and products available that can easily meet or exceed the technical requirements, but fall short of meeting the goal when factoring in the costs. The design challenge was established when the schedule of only six months was imposed on the project! In order to meet the challenges, effective development methodologies had to be implemented.

DEVELOPMENT METHODOLOGIES

In any successful program, four important methodologies will be found:

- Teamwork
- Communication
- Understanding the "Big Picture"
- Rapid prototyping

Teamwork / Communication

These tenets were the backbone for the development of the low-cost Precision Approach Radar simulation. The success of any program can be measured by meeting cost and schedule. A balanced team means having representatives of the appropriate disciplines with a sense of rapport, ownership, responsibility and a feeling that the project will "make a difference". (G.H. Boyle and B. Edwards) Many factors are involved in the process of forming the best team of people for a specific project. These factors can be dynamic. They have the potential to change from project to project, and even to change within a project. One point that does not change in effective teams, however, is communication. Good communication is directly proportional to the success of a team. It will be a momentous driving force on the schedule, cost and quality of the finished product. If the team does not have good communication, then it runs the risk of losing sight of where the project is going, and a project left to wander will surely go astray!

The "Big Picture"

Understanding the "Big Picture" and producing a quality product at a lower cost is the desired goal. Designing a training device requires knowledge of the mission in order to define the training environment in which the students will operate. Understanding the type of training that is important to the end user is a very significant factor. A training device can be based on the most rigorously derived mathematical models, and yet not be successful because it does not provide what the customer considers most important for training. Communication among the customer, training designers and engineering teams from the beginning will pay big dividends throughout the project. (G. H. Boyle and B. Edwards)

When specifications are created by a customer, they can be written at a very detailed level (as in military contracts) or at very high level (as in commercial contracts). In either case, a thorough understanding of customer *expectations* is of paramount importance.

The PAR project was truly a team effort, and the team's success was secured by the fact that its members worked and communicated well together. Throughout the development cycle, one simple philosophy was consistently applied: "every problem has a solution" (Tom Hanks). When problems arose, the team worked together to find a solution. This approach led to new avenues of innovative solutions.

Rapid Prototypes

A very important factor that must be considered when involving the customer in any design is that not all customers are engineers. Engineers cannot expect the customer to be able to visualize the end product by simply looking at code, or a complex math model defining everything about the radar. An important mode of operation during the development of the Low-Cost Precision Approach Radar was to have a working demo at all times. A philosophy of rapid prototyping new design approaches and always having a working demo was adopted. If it can be said that a picture is worth a thousand words, then a working demo is worth the customer's confidence. While design engineers have the ability to look into code and foresee the finished product, not all customers can. Having a working demo allows others not as intimately familiar with the simulation design to "see" into the code as well! A history of the project's development in the form of working demonstrations gives people the ability to perceive whether or not the product is heading in the desired direction, while at the same time produces tangible milestones

DESIGN SOLUTIONS

To meet the low-cost challenges, a PC platform was chosen. It was also desirable to remain independent of any special graphics boards. The decision was made to base the simulation on a 486 DX2 66Mhz PC with a standard SVGA graphics board, minimum hardware controls, a trackball and keyboard. The simulated radar display needed the look and feel of the real radar, complete with all screen functions, landmass, target displays, sweep rate and phosphor decay. The Precision Approach Radar simulation was part of an Air Traffic Control Tower simulation in which each student and/or operator station was run by a single PC. Unfortunately, the PC would not be dedicated to just the Precision Approach Radar, it would also be used for other training applications. Therefore, the design goal was to fit entire radar simulation on only one PC.

Some questions immediately came to mind:

- Which operating environment should be used -- Windows or DOS?
- What are the trade-offs between Graphical User Interface (GUI) and real-world controls?
- What simplifications could be used to achieve the single PC constraints?
- How to program dynamic sweep phosphor decay without eating up all of the processor time?
- How to display landmass data (use polygons or bitmaps)?
- How to get the landmass data (DMA data, topographical maps or photos of the radar screen)?

Windows vs. DOS

The Air Traffic Control simulation was operating under a windows 3.1 environment. The PAR system criteria dictated that the simulation was to fit within one PC and then communicated to the Tower simulation via the network software. This gave the opportunity to choose an operating system for the PAR simulation. DOS was selected because of the strengths and abilities to handle graphics and drawing applications without the overhead of the windows programming.

Design Simplifications

There were some significant simplifications that allowed the use of a single processor 486 PC. Most important was that a ground-based radar system has a fixed landmass with a 30-degree field of view. In addition, the landmass was more of an obstacle as opposed to something that was to be identified by the PAR operator. Finally, the radar sweep was in a relatively slow rectangular (b-scan) fashion instead of a fast circular sweep.

A single processor 486 PC with no special graphics accelerator chips was tasked with accomplishing a large amount of work. Early prototype attempts quickly exceeded the processor time with only a crude radar sweep (no landmass display, targets, or phosphor decay). The generation of a dynamic sweep was in the right direction, but the amount of CPU time it consumed was unacceptable. Many creative engineering ideas were employed to conquer the plethora of roadblocks along the way.

The tight budget constraint of the project continually forced the issue of choosing the balance between the fidelity level, meeting the specification and the resulting cost. On the other hand, a real program killer is unmet customer expectations. So the methodology adopted was that if you are not simulating it, do not try to make it look like you are. Functions modeled must either be of such low fidelity that the trainee does not try to equate them with the real world, or be equal to the fidelity of the actual system. (G. H. Boyle)

The following sections outline the most notable design solutions.

Graphical User Interface (GUI)

After an analysis of the radar man-machine interfaces was complete, the radar functions were grouped into two categories, based on accuracy, realism and customer inputs. The functions that required a high degree of realism would emphasis the touch and feel of the real-world equipment. The other functions would allow the design to be more flexible in terms of implementation. The simulated radar approach became a design with minimal hardware controls. The only real-world controls provided with the simulation were those that required constant use by the radar operator, where realistic touch and feel are most important. For example, when a Precision Approach Radar operator is guiding an aircraft in for a landing, their hands are almost continually on the antenna-steering joystick and the Azimuth and Elevation IF gain control knobs. These were critical man/machine interfaces. Other functions of the radar that are not as frequently used (such as High Voltage On/Off) are referred to as secondary functions. The simulation of the secondary functions allowed the development of a graphical user interface (GUI) to completely satisfy those training requirements. The GUI requirements imposed by the design were that it had to be a fast and efficient interface with a professional look and feel. The GUIs available under Windows 3.1 provided a standard look and feel, with a relative ease of programming. However, the Windows 3.1 software uses a large amount of overhead processor time, and an early decision ruled against it.

A third-party commercial software package solved the GUI interface problem by providing the capability to program buttons, gauges, and sliders with minimal effort. This commercial-off-the-shelf (COTS) package was extremely low cost, yet it still provided the desired professional look and feel. The Graphical User Interface developed for the low cost PAR is shown in Figure 3. A trackball was used to control input into the GUI. In conjunction with the trackball, the function keys were also programmed to duplicate the GUI functions to aid the more experienced users.

Figure 3
Graphical User Interface Layout

Creative Engineering

With the pressures of a short schedule and a small budget looming overhead, the use of rapid prototyping was employed extensively. Brainstorming was used to create ideas and then eliminate unfeasible candidates. The surviving approaches were researched, coded and their results were quickly evaluated for their applicability to achieving the overall goal. This type of engineering development environment led to many creative solutions such as palette animation, bitmap management, and bitmap generation. These techniques proved to be the winning solutions for the low-cost, high-fidelity radar simulation!

Palette Animation

The biggest question with the low-cost radar was how to program a dynamic phosphor decay. Phosphor decay can be described as the fading of color intensities between radar sweep updates. Drawing every pixel for every color intensity change would require an enormous amount of graphics power. Using a technique called palette animation allowed the appropriate dynamic phosphor decay to be achieved with the radar sweep using minimal processor time.

Palette animation is a technique used to cycle colors through the graphics card hardware palette using very little processor time. To understand palette animation, an amount of background in PC graphics is required. Every pixel on the graphics display is represented by an (X, Y) coordinate and a hardware palette index. A system with 256 colors means that there are 256 indices associated with a hardware palette. Each of these indices points to the red/green/blue combination of a color. To achieve animation, pixel locations on the PC screen are assigned with specific color palette numbers in such a way that, as the indices change, one gets the sense of movement without ever writing to the screen.

For the radar simulation, the PC screen was divided into two areas: the left side of the screen was reserved for the radar display and the right side was reserved for the GUI. Figure 4 shows the PC screen layout.

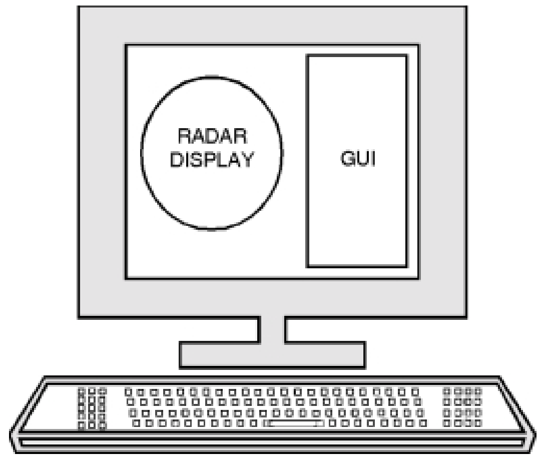

Figure 4

PC Screen Layout

The radar display section of the screen was then subdivided into two areas to represent the Elevation and Azimuth radar displays on the PAR. Figure 5 illustrates the Elevation and Azimuth areas of the radar display.

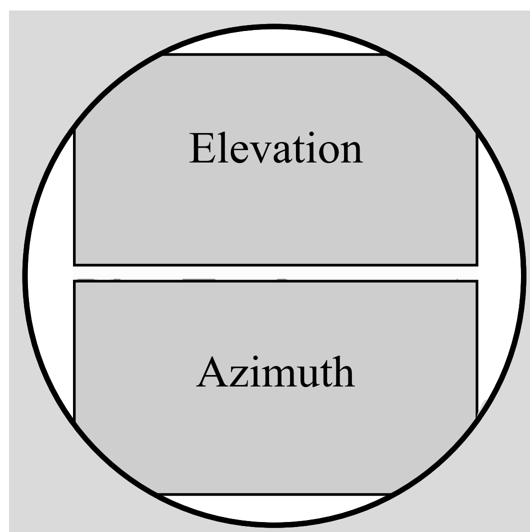

Figure 5

PC Radar Display Layout

An example of an Elevation and Azimuth radar screen displaying the ground clutter, glide slope, safety cursor, PAR minimum marker, runway centerline for the 2 nautical mile scale is depicted in Figure 6.

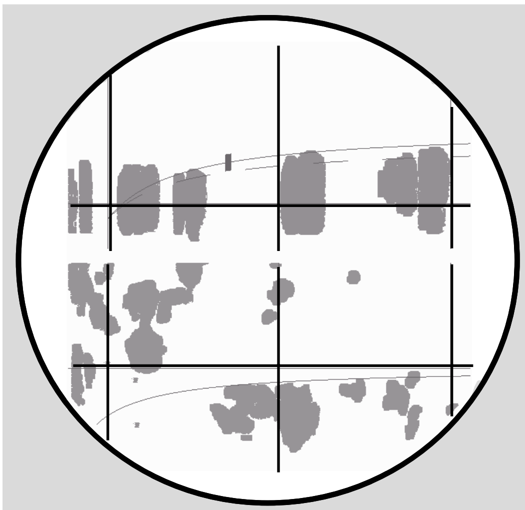

Figure 6

PC Radar Screen Display

For the radar simulation, the Azimuth and Elevation sections of the PC screen were assigned multiple rectangular segments. Figure 7 is and illustration of rectangular segments within the Elevation area.

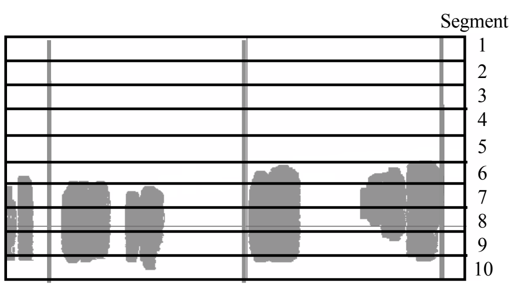

Figure 7

Elevation Segments Example

To take advantage of the palette animation and apply it to radar simulation, the radar screen was broken down into multiple segments. Each segment received palette index numbers belonging to one of three sets. One set of numbers was used for screen functions. Another set of numbers was used for landmass and weather. The final set of numbers was for targets. A color palette index of zero was assigned for non-colors and was set to the background color (black). Figure 8 shows an example of how color palette index assignments were done for a section of the Elevation radar display.

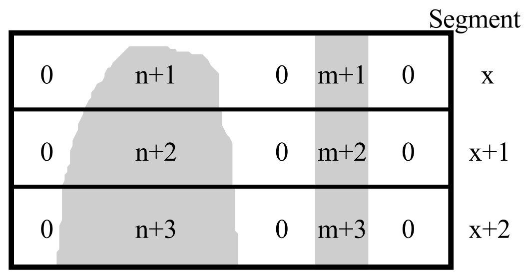

Figure 8

Palette Index Assignments

To ascertain the illusion of a radar sweep was the next task. By scheduling the rate of the sweep, the brightness of the sweeps leading edge segment could be set to the intensity determined by the current radar function settings. All other segments would reduce the intensity according to the following first order lag equation:

Palette[index] = Palette[index] * gain

where: gain ~ .9 to.99

The resulting display was a radar sweep with a realistic phosphor decay effect!

Screen functions, landmass (clutter), and targets kept different ranges of color palette indices so that groups of indices could increase or decrease intensities independently of the other display features, thus giving the effect of multiple layering. The screen functions always received or shared the highest intensity, thus appearing to be on top. Landmass and target color palette intensities could be at different levels, thus giving the illusion of target on top of clutter or target behind clutter.

Another feature to this approach is the ability to adjust many of the important parameters such as phosphor decay rate, sweep rate, illumination intensities and color just to name a few.

Bitmap Management

Even though the landmass was fixed on the display (except when the antenna assembly was rotating) the effects of the radar function settings could change the shape, size, and intensity of the landmass. For example, the IF Gain can make the ground clutter grow or shrink. Multiple bitmaps of each possible radar function combination were needed. However, this turned out to be a prohibitive number of possibilities. Color palette manipulation proved to eliminate many combinations and significantly reduced the number of bitmap combinations, but still left a large number to deal with. The reduced total number of bitmaps required for the simulation is shown in Table 2.

Table 2 Bitmap Requirements

Total Bitmaps Required	
Range Settings	3
Runway Direction	2
FTC Effects	2
Narrow Pulse Effects	2
IF Gain Levels (**)	16
Number of Segments	52
Total Bitmaps	**19,968**

** Using color palette manipulation, 64 levels of intensity could be achieved, along with the required 16 size and shape changes

Calculating landmass changes in shape and size would prove to be processor intensive and could not be achieved in real time. A solution was then developed to create all the bitmaps off-line and store them on the hard drive. Calculations showed that, for this application, it would still require up to 60 Mbytes of disk storage.

One negative side effect of storing 60 Mbytes of bitmaps on the hard disk was the 486 PC slow access time for retrieval that could cause a noticeable slowdown in the radar sweep. Most of this problem could be solved by disk caching, but did not resolve the slowdown when reading in a new set of bitmaps. The goal was to have zero slowdown rate for any combination of radar function settings that the user required.

A successful design solution was to compress all of the bitmaps such that they could fit in the available 8 megabytes of random access memory (RAM), and then to uncompress them in real time. With this approach, the first concern was to find a method to compress the data by a factor of 10:1. The data looked relatively random and that kind of compression ratio looked ambitious.

It was realized that each bitmap consisted of one segment, and each segment consisted of 2 color numbers: one for background, and the other for the Landmass (refer to Figure 8). Therefore, a run length compression routine could be implemented on each bitmap (segment) which counted the number of like colors and stored the count. This compression was extremely effective and allowed a 30:1 average compression ratio!

Once the bitmaps were compressed, a retrieval method had to be developed. One approach might be to create a random access file structure. With this approach, each bitmap's location would be calculated by a formula and a file pointer located the proper address index.

For our application, however that method had a major problem. The random access files required equal length storage addresses, which meant that the storage space had to be the size of the largest compressed bitmap. This would cut the efficiency down to a 3:1 compression ratio, which would not work for our approach.

Since each bitmap had a different size, ranging from 6 bytes to 2000 bytes, a more efficient data storage method was required. The sequential file storage method was optimal for size storage, but retrieval speed would be very slow. To solve this problem, two files were created. Figure 9 shows the bitmap file retrieval method that was ultimately implemented. One file stored the starting address and length of each bitmap section in a random access format. The second file stored the bitmap data sequentially and its address index was calculated by the following formula:

$$index = (((((Rwy\_direction * 2) + Range) * 3 + Narrow\_pulse) * 2 + FTC\_mode) * 2 + IF\_Gain) * 16 + Segments$$

Where:

Rwy_direction = 0 (Left Rwy) or 1 (Right Rwy)

Range = 0 (2nm), 1 (10nm), 2 (20nm)

Narrow_Pulse = 0 (Off), 1 (On)

FTC_mode = 0 (Off), 1 (On)

IF_Gain = 0 to 15 (16 levels)

Segments = 0 to 51 (26 Elevation, 26 Azimuth)

The radar has many filter combinations which produced the corresponding 20,000 bitmaps for the simulation. The uncompression routine developed allowed any of these bitmaps to be uncompressed in real time and then displayed on the PC screen with no delays to the simulated radar sweep or phosphor decay.

Bitmap Generation

The options of creating bitmaps using computer-generated polygons or actual photographs of the radar displayed were evaluated. To produce better realism the photograph option was selected.

A major problem was how to generate the bitmaps from all of the photographs of the real radar display. Each photograph could be scanned into the computer using a high-resolution scanner and overlaid onto the PC simulation of the radar display, but several problems made this very costly. First, each photograph was distorted from the original display. Each photograph had reflections of background lights blended in with the landmass display. Finally, each photograph had all the radar's screen functions on display. Each of these photographs would have to be manually enhanced so that only the landmass information would remain, and be of the correct size and placement when overlaid on the simulated screen display. Furthermore, all of the combinations of radar functions were not available in photographs and these missing sets would somehow have to be generated.

The key to solving this problem was in realizing the fact that this particular type of radar landmass data was an obstacle in the radar operator's way, and that the operator was always trying to reduce or eliminate the landmass data, not trying to identify it. With this in mind, it was determined that the landmass display need not be an exact representation of the real world, but a close approximation. This greatly simplified the problem from enhancing hundreds of photos to enhancing just one photo and writing an off-line routine that simulated all radar function settings based on that photo. The off-line routine used a single scanned photograph and generated all 20,000 scaled and compressed bitmaps in less than 15 minutes, but that is a whole different paper!

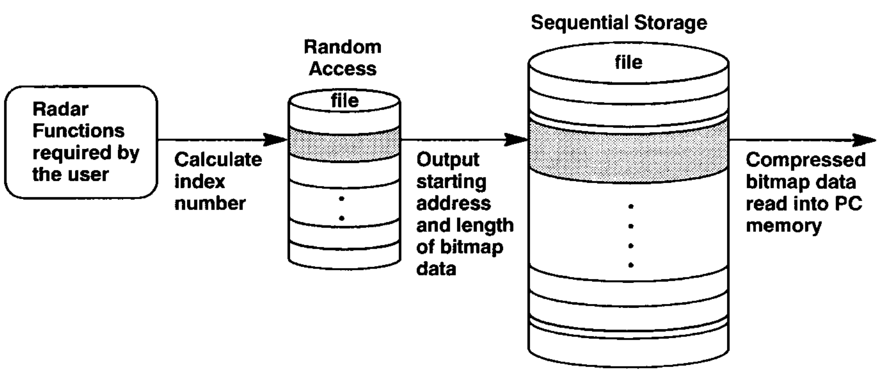

Figure 9
Bitmap File Retrieval

DESIGN IMPLEMENTATION

Applying current PC technologies, such as palette animation and bitmaps, to a radar project produced many favorable results. The outcome was a Low-Cost Precision Approach Radar simulation designed, built and implemented on a single 486 PC with a standard SVGA graphics board. The user interface to all simulated PAR functions was a combination of a Graphical User Interface and real-world controls. The Az/El antenna controls were realized using a four-position momentary toggle switch, while the Az/El IF Gain controls were simulated using stacked analog potentiometers.

The software was written using the 'C' language and operated in both a standalone and an integrated mode. The radar simulation produced a radar display with a phosphor decaying beta scan sweep while implementing all the required radar functions. Because of some creative engineering and the use of palette animation and real-time uncompression of bitmaps, the software was able to run the entire radar simulation with approximately fifty percent spare CPU time! As an additional feature, using a reduced set of landmass data, a representative implementation was able to fit on a single 3.5" 1.4 Mb floppy. The project concluded in approximately six months and was accepted with no formal customer discrepancies.

FUTURE DIRECTIONS

The future directions of this approach are focused on the applicability in the following areas:

- Circular sweep at a limited rate.
- Landmass that moves but at a limited rate -- ground vehicle- or ship-mounted radar.
- Any type of radar that has a phosphor decay (sweep or no sweep).
- Larger landmass databases with added PC memory.

CONCLUSIONS

There are many radar model simulations and products available that have the ability to meet and or exceed the vast majority of the technical requirements of any radar simulation project until factoring in the cost and schedules. The rising need for high-fidelity training in a time when customers are increasingly constrained by shrinking budgets has produced a requirement to create innovative training solutions that are both low cost and high fidelity. Fortunately, there are exciting avenues for creative simulation designers to accomplish this, using the low-cost high-powered processors, graphical techniques and off-the-shelf hardware and software products.

ACKNOWLEDGMENTS

Mr. John Christ, Hughes Training Inc.

Mr. Janne Enarvi, Finland CAA

Mr. Paul Kurash, Binghamton FAA

Mr. Seppo Matikainen, Finland CAA

REFERENCES

Boyle, G. H., Edwards, B. Low Cost Trainers: Lessons For The Future. 14th I/ITSEC Conference, November 2-4, 1992.

Hanks, Tom. The movie Philadelphia 1993.

ITT, Gilfillan. Quadradar Mark V Model A Trainee's Guide. September 1974.

Miller, D. M. Foundations for Tactical Training: A Challenge to Industry. AIAA Flight Simulation Technologies Conference, Boston, MA, September, 1989.

Tucker, D., Collom, K. Energy Level Modeling: A New Approach To Real-Time ECM Radar Threat Simulation. 15th I/ITSEC Conference, November 29 -- December 2, 1993.