DMA AND CIG: A SHOTGUN WEDDING

Cary E. Wales and Michael A. Cosman
Evans and Sutherland Computer Corporation
Salt Lake City, Utah

ABSTRACT

The DMA data base and its future enhancements have been hailed as the general solution to creating visual data bases for CIG. DMA centered approaches have proven marginal or ineffective, however, in providing visual support of low altitude flight, since DMA stresses navigation rather than flying skills.

The imagery presented by the CIG system must tell the pilot where he is relative to his map, and where he is relative to the ground. These are two very different objectives, and emphasizing one may compromise the other. Heavy emphasis on 'capture criteria' (for map fidelity) has historically resulted in systems which do an inadequate job of supporting low level flight, a task which depends on high 'scene density'. This paper examines the capture criteria and scene densities required to support such missions, and quantifies critical aspects of a 'nap of earth' data base. We will examine various constraints inherent in CIG systems, and their influence on what can be achieved. We also will discuss data base definition strategies, including the use of DMA data, to see where traditional approaches have been deficient in providing the required visual cues. Then we will present an approach which combines DMA data, new mathematical methodologies and data base design strategies, and current production hardware, to meet both the capture criteria and scene density needs of 'nap of earth' missions.

INTRODUCTION

Whole mission simulation is a multifaceted technological challenge. The total integration of instrument, radar, sensor and visual reinforcement requires the mastery of a variety of technical disciplines. The task is made even more demanding because the processes which create and support various aspects of the simulation are becoming ever more complex and disparate, while the need to correlate these aspects gets steadily stronger.

Computer generated visual imagery is one of the most demanding of these facets. Visual imagery must correlate with and reinforce information being received from other sources, including radar, infrared and radio navigation systems. It must support those aspects of flight navigation which are based on looking out the window: general topography, major terrain characteristics and significant cultural features which serve as navigational waypoints must be recognizable and correlatable with other navigation processes. Finally, the visual imagery must support flight skills by conveying to the pilot an accurate sense of spatial and dynamic relationships. Users are requiring these capabilities to be achieved with increasing effectiveness in ever larger and more complex visual environment data bases. DMA data has been increasingly relied on to provide the real world and radar correlations required in such systems.

THE DMA CONTRIBUTION

The Defense Mapping Agency Digital Terrain Data Base proceeds from, and is designed to efficiently support, radar driven navigational needs. The DMA data source provides two basic categories of information in the form of terrain files and culture files.[1] The terrain files contain terrain elevation values, in meters, corresponding to sample points along a regular grid in latitude and longitude. Currently, terrain files are available in two levels of sample spacing, Level 1 and Level 2. The latitude spacing of the grid is 3 arc-seconds for Level 1 and 1 arc-second for Level 2, while the longitude spacing varies from 3 to 18 arc-seconds (Level 1) or from 1 to 6 arc-seconds (Level 2), in order to keep the grid approximately square. The absolute accuracy requirements for terrain data are specified as 130 meters horizontally and ±30 meters vertically for both Level 1 and Level 2.

Culture files are also available in two levels of detail: Level 1, which provides a generalized description and portrayal of planimetric features which meet relatively large minimum size requirements; and Level 2, which provides a more detailed description and portrayal of features which meet somewhat smaller minimum size requirements. The absolute accuracy requirement for culture file data is 130 meters for Level 1 data, and 26 meters, point to point, for Level 2 data, with the additional constraint that Level 1 and Level 2 feature detail must be compatible. Vertical accuracy specifications apply only to vertical obstructions of 46 meters or greater, and an accuracy of ±10 meters is specified for such features.

DMA AND NAVIGATIONAL CUES

The essential role of navigation cues in a visual data base is that of guiding the pilot along a course marked by a sequence of recognizable milestones, from a point of origin to a destination. The cues which are ultimately used will depend largely on the type of mission to be supported, the flight altitude domains, the strategic or tactical nature of the task and the characteristics of the terrain being overflown. Features which can be important navigational cues in an expansive desert landscape can be totally insignificant in a highly forested, hilly terrain. A segment of terrain designed to support high altitude navigation may not require the density of cues needed to support navigation at lower altitudes. Navigational needs will often be different in one portion of the data base than in others. The data base designer may capitalize on such differences in order to concentrate navigational cues where they are most significant.

The required spatial frequency of navigation cues tends to be relatively low at all altitudes. In real world missions the frequency of navigational waypoints is dictated by pilot workload constraints, and tends to be determined by flight time intervals of between 15 and 60 seconds. For tactical missions involving high speed low altitude flight, the spatial distribution of specific navigational features can be quite sparse while still providing an excess of waypoints which can be used during the mission.

The DMA terrain elevation data provides a natural starting point for the definition of large scale terrain topography for both the visual and radar simulations. A variety of approaches have been used to transform the data, correct errors or inconsistencies, and convert it to forms displayable by computer image generators. The nature and scale of the data provide sufficient correlation to fulfill navigational needs of the visual environment. However, a preliminary analysis of the cultural files suggests that much of the wealth of data contained in these manuscripts is not designed to support the visual data base design task, nor provide effective imagery for 'nap of earth' or low altitude flight. Positional information, dimensions, material descriptions and coverage factors are potentially useful in providing cultural navigational cues and generic terrain embellishment for optical flow, however.

OPTICAL FLOW AND LOW LEVEL FLIGHT

At high altitudes there is little danger of crashing into the ground, and ground topography is of interest primarily to support navigational goals. In flight profiles very close to the ground, however, a high percentage of visual cues must be designed to support terrain avoidance and flying skills such as those employed to stealthily approach targets without being detected by potential threats. The lower a pilot tries to fly, the more he depends on key characteristics of the visual scene in front of him. His attention gradually shifts from cockpit data sources to the outside world; as his work load increases he must extract from the visual scene better information at faster rates in order to avoid impact with the terrain. This ability depends on a complex set of visual factors, including the optical streaming of scene details away from a single fixed image point (the "aim point"), the rate of motion of image details, relative parallax shifts, and the emergence of successive levels of detail (including any absolute size cues).

Some work has been done in attempting to quantify these factors and transfer them into a computer generated visual scene. Early attempts at providing these cues with abstract repetitive surface texture (i.e. checkerboards) were comparatively unsuccessful. Once pilots had calibrated the texture to determine how many wavelengths above it they should fly, holding a constant height above terrain was very easy. Such training does not transfer well to the real world, since it is based on strong artificial cues which may be absent in the actual mission. More recent research[2] has suggested that a mixture of surface texture and three dimensional features is required, and that these cues should be distributed throughout the image plane so that a certain density of cues per solid angle is maintained. This last requirement ties the spatial density of cues to the minimum flight altitude, and suggests some simple rules for computing this density.

NATURE AND DENSITY OF NOE FLIGHT CUES

John Sinacori, as quoted from reference 2 above, suggests that "the dominant perceptual strategy used for NOE point to point flight is motion perspective (optical flow) augmented by other mechanisms, such as linear and aerial perspective, interposition, shading, texture and apparent/familiar size, if and when their corresponding stimuli are available." He establishes a suggested minimum distribution of textural elements in the visual scene which is quantitatively related to the mean height of the pilots eye above the ground surface. He suggests that appropriate stimuli to support the motion perspective perceptual strategy should consist of a random distribution of textural elements of varying size, shape and contrast, whose mean spacing is approximately one eye height. He cautions that the estimate is preliminary and emphasises that such an element spacing is deemed a minimum texture level needed to reveal surface shape using motion perspective. He further states that although regular patterns or highly coherent texture may reduce pilot workload they should be avoided. Such regular patterns as checker boards invite the prospect of learned responses and negative training.

Within the framework of the motion perspective (optical flow) mechanism, additional strategies suggest themselves which relate to other perceptual mechanisms in both a complementary and a supplementary role. Sinacori states that "what should reduce workload and/or improve performance is the addition of coherent objects that more easily facilitate a static perception of terrain surface shape and observer position relative to that surface." He further asserts that "The addition of vertical objects should also reduce workload and/or improve height holding performance", and suggests a mean separation of vertical objects of about three to five eye heights. The inclusion of 3D objects such as trees, bushes and rocks, in addition to contributing optical flow cues, provides parallax cues and information about relative sizes of known objects which can be particularly valuable at slower speeds or when performing hover operations.

Sinacori provides a specific example of an NOE data base designed around the preceeding notions. He uses a portion of the Fort Hunter - Liggett area to postulate a helicopter NOE mission near the Nacimiento valley. Adequate ground truth is achieved with about 1000 polygons per square kilometer (this density roughly corresponds with direct skinning of the DMA terrain elevation grid). A dense 2D polygonal enhancement is added to the terrain to further define surface shape and provide optical flow, and 3D enhancements (trees) are added at a density of between 6000 and 12000 trees (roughly 100000 polygons or 300000 edges) per square kilometer.

IMAGE GENERATOR ISSUES

The above example suggests strongly why historical attempts to simulate NOE imagery have proven inadequate. The typical design strategy has been to convert the DMA terrain elevation data set into a terrain skin by a straightforward linear process which results in a similarly sized output data set; the terrain is then further embellished with cultural and organic features to provide optical flow and navigational cues. With this strategy, the basic problem is how to produce, store, read, manage and display the resulting very large data base. The gulf between current capabilities and such a system is measured in orders of magnitude, not small incremental improvements. The following general discussion of these problems acknowledges that all existing CIG systems have analogous limitations.

Data Base Size

A rough metric of data base size would be to express the production and storage requirements in terms of bytes per polygon (or face). This metric would include the overhead of data base management structures, all raw geometric information, and topological information required to assemble the scene. A value of 200 bytes per face is close enough for our analysis. Referring back to Sinacori's NOE example, a modest 25 by 25 kilometer data base would require roughly 13 Gigabytes of information. A large area tactical data base of several hundred thousand square kilometers modeled to a density of 100000 polygons per square kilometer would be prohibitively expensive to produce, and would require hundreds of very large disks to store.

Disk Bandwidth and IG Memory

There is a direct relationship between data base density, aircraft velocity and the amount of data that must travel between the on-line (disk resident) data base and the IG memories each second. If the sample data base has to support speeds of 100 knots and a general visibility of two kilometers, then about four Megabytes of information would be moving from the disk to the IG each second. The amount of IG memory required to contain all scene details within this visibility limit would be nearly 250 Megabytes.

Data Base Management

The image generator must select from the entire data base two very important subsets of data. The first is that portion which must reside in IG memory because it is closer to the viewer than the system visibility limit. This subset is continually changing as the simulated aircraft moves within the overall gaming area, but will generally represent a small fraction of the total data base. The rate at which this subset changes may be comparatively low, so this task is typically performed as a background task over multiple field times in the general purpose computer which hosts the disk. The amount of work required to perform this cull is related to total data base content. Processes which are adequate for typical production data bases whose densities are 10 to 100 polygons per square kilometer will be inadequate for the sample NOE problem.

Culling and Display

From this data an even smaller subset is extracted which consists of scene details which may be within the display windows. This second subset is usually culled to include only the simplest permissible representations of scene elements (level of detail selection), and only those surfaces which are front faced and large enough to be useful in the displayed image. The second cull processes only the data which survives the first cull (and resides in the IG), but it must be performed once every field, and on a much richer and larger set of image elements. The IG must extract those scene elements which are maximally useful to the viewer, while rejecting many others, and while preventing the artifacts of this management from distracting the viewer.

SOME HELPFUL OBSERVATIONS

All of the above might suggest a very bleak outlook for those simulator users requiring an effective NOE capability. Incremental linear advances in CIG technology are not likely to catch up to these needs as fast as the needs are evolving. What is required is a different way of looking at the problem, and a totally different strategy for dealing with the requirements. Fortunately, some helpful observations spring to mind.

First, there are two fundamentally different types of scene detail to deal with. Optical flow to support flight skills is almost entirely supplied by dense generic details with a high spatial frequency content, while the navigational "specificness" of the data base is due to very low frequency characteristics of the terrain topography and related cultural cues. This suggests that we ought to try to discover ways to "re-use" optical flow related scene details, so that many visual representations can be drawn from a small number of actual elements. We must then find a way of combining these reusable optical flow details with a terrain underlayment which is specific enough to meet our navigational requirements.

Second, perhaps we can discover a data base management strategy that implements level of detail selection while also providing for the physical subdivision of geographic data base areas into smaller areas. This would give us a way to manage the data base so that rejection of unneeded scene detail occurs in pieces as large as possible, expediting the data base culling process.

CT5: EXPLORING PARTIAL SOLUTIONS

The Evans and Sutherland CT5 image generator marked a number of significant departures from historical practice in the CIG field. Primary among these was the processing of image details by area rather than scan line,[3] and the implementation of all data base management functions in the IG hardware.[4] Each of these notions brought fundamental changes to the character and capability of the image generator. The rigorous application of antialiasing technology in the Display Processor achieved new levels of image quality, while increased basic capacity throughout the IG was augmented by much more efficient data base management. These improvements were accompanied by a better than linear capacity response of the system in areas where NlogN or N-squared responses were the rule in previous systems.

A Hierarchical Data Base Structure

The decision to implement the data base management function as a hierarchical tree was largely driven by the need to sort scene details in a front to back order. A powerful extension of list priority techniques, Cellular Priority, was developed to address the CT5 sorting problem. Its key property was the hierarchical subdivision of the overall data base priority problem into a tree of many nested simple problems. Within such a process visibility decisions about major data base elements propagate automatically to smaller elements. A similar property attends the data base management and culling process. Some key ingredients of this structure are particularly important to the NOE data base management problem.

A cell is a volume in model space represented by a decision node in the data base tree. The primary function of the cell is to provide a choice between either a simple or a complex representation of some portion of the data base, based on the distance between the cell and the viewer. Either choice can consist of a null (the item is so far away that nothing need be displayed), an object (a collection of polygons and light strings), or a mesh. A mesh is a collection of cells with a rule for their visual ordering. The mesh is thus the primary structure in the cellular priority process. The mesh can also be used to subdivide a region of model space into several smaller regions; this is the primary mechanism used to control the degree of detail that will be unfolded before display or rejection of data base elements.

The data base tree is the structure of meshes and objects organized by their association into cells; it is processed by testing cells and ordering meshes. Processing ceases along any subpath in the tree when the cell choice indicates either an object or a null. Processing will continue when the cell choice indicates a mesh. The order of processing in the tree is determined by the ordering rules contained in each mesh. At every stage of processing the IG attempts to prove that the item being processed is unneeded because it does not appear in any display window. The rejection of an item means the termination of processing along that subpath in the tree; unneeded portions of the data base tree are thus pruned off as early in the process as possible.

Capacity of the Object Manager

Processing of the data base tree occurs in the object manager (OM). The major task of the OM is to reject as much of the data base as possible, as soon as possible, and in pieces as large as possible. As the overall content of the data base increases, so does the content of the pieces which will be tested, and in many cases subsequently rejected. Each additional level of data base tree structure exacts a fixed amount of processing time, while increasing the total data base content by a multiplicative factor. Tree trace processing time thus grows linearly with data base depth, and logarithmically with total data base content.

Instancing in the IG

The CT5 image generator supports the reusing of data base items in a general and extensive fashion. A mesh or an object can be referenced from multiple cells in a data base, and each reference can include a positioning vector to plant the mesh or object in a particular place. All data base structure and scene detail which ultimately results from an instance will also be moved to the new location. By cascading levels of instanced meshes, very extensive and complex data base structures can be quickly created. For example, an entire forest can be built from a single IG resident tree and a very compact nested and instanced mesh structure. This allows a nearly total decoupling of memory requirements from apparent displayed density, and significantly reduces the amount of data which must be transferred from the disk.

Paging from Disk to IG

The object manager traces the data base tree to identify data blocks which may be needed to draw the image, but which are not currently in IG memory. This is done as a background task over several field times. The pager algorithm performs memory block allocation and garbage collection as well, and communicates data block requests to the general purpose computer. Data transfers from the disk to the IG occur by direct memory access. The hierarchical data structures which support data base management within the IG also suffice to control disk to IG processes; no secondary software management structures or tables are needed. This process can support a very high rate of continuous data flow from the disk to the IG.

CT5A: THE NEXT STEP

The CT5A system was the logical outgrowth of a characterization study of CT5. This study resulted in a better understanding of the data base management problem and how CT5 solves it, identified major new data base design and development strategies, and proposed a set of system modifications to significantly increase basic capacities while effectively supporting the new modeling approaches.[5] These new features significantly improve the ability of the system to address the NOE simulation problem by allowing visual data bases of over 500000 polygons per square kilometer to be processed, as the following discussion will illustrate.

Fade Level of Detail

Fade level of detail helps mask the switch between two levels of detail by fading the retiring scene element out as the replacement element is faded in. The transparency capability of the system is used to effect the switch, which is based on range and can occur at several selectable rates. This greatly reduces the visual distractions associated with level of detail switching, and allows these transitions to occur much closer to the viewer. This helps the system concentrate on processing details which are large enough to be valuable to the viewer, and significantly increases the apparent richness and complexity of the image.

Using Multiple Levels of Detail

In a visual scene, the most complex representation of an image detail will only need to be presented when the viewer is very close to it. At larger distances a much simpler representation is adequate and at still larger distances only the proper suggestion of bulk, shape and color is needed. In an extensive evenly distributed array of such details, most of them can be displayed using the simplest version. For a constant number of on screen polygons, the geographic density of such an array can be much higher if multiple representations are employed. Additional leverage is obtained by using fade level of detail to further shorten the transition distances.

Using Multiple Levels of Scale

In order to provide usable cues over a range of flight altitudes the data base must contain a selection of cue details of varying sizes and spacings. At close ranges the smaller cues will provide the bulk of near scene detail; at longer ranges the small cues will have been managed out of the scene, and the larger cues will be needed. One result of this strategy is that the angular distribution of cues in the observers field of view is nearly homogenous throughout the image, instead of being severely compressed against the horizon.

A Mathematical Design Methodology

A rigorous mathematical process can be used to design such a data base. The mission flight profile is used to establish the number and types of levels of scale, and the density requirements of each level. Components of the data base are then identified and characterized as to number and complexity of levels of detail. Level of detail transition strategies are defined, taking advantage of factors which can help minimize IG load. A global instancing strategy is developed to separate reusable generic scene features from specific details, and minimize the total modeling effort. The interaction of each data base portion is analyzed, and a composite IG load level is computed. The data base development can then proceed, confident that mission visual requirements will be met without overloading the system.

THE BASIS SET NOTION

The pursuit of structured and modular data base development strategies grew out of a desire to decouple the data base development cost from the geographic extent and total content of the data base. A first step in this direction was to separate the problem into two distinct portions: the specific terrain underlayment, and a reusable set of organic and cultural 2D and 3D components which could be used to embellish this base terrain. The actual ornamentation of the terrain would add those specific features which are navigationally important, and additional generic enhancements to increase realism and optical flow.

Lineal features such as roads and rivers presented an interesting problem. They had to be visually continuous across terrain boundaries, while adhering to the topological properties of the terrain underlayment. A process for providing these features was developed which represented roads by instancing predefined road segments. These segments were modeled in a set of road types, segment lengths, and segment orientations. From a comparatively small set of these elements specific, if somewhat stylized, roads can be built.

Three dimensional specific terrain presented a still more difficult problem. The fundamental elements of a strategy were identified during the course of extensive experiments aimed at the NOE problem. What ultimately emerged was the concept of a basis set: a set of scene components which exhaustively satisfy a class of local boundary conditions. We developed methods for identifying and defining a terrain basis set from a given set of constraints and fidelity requirements. Algorithmic ways of choosing terrain basis set elements to provide required specificness were explored, and basis sets with comparatively few elements were found to provide adequate ground truth and topographic complexity.

USING TERRAIN BASIS SETS

A terrain basis set consists of a number of patches of terrain which can be made to fit together in various ways by translating each element with respect to its neighbors. In the general case, basis set elements need not be either planar or similarly shaped, but the selection process must ensure that adjacent patches have compatible properties along their shared boundary. Because each patch can be offset in Z as well as X and Y, terrain features much larger than any one patch are easily created; the degree of fidelity preserved in this process depends on how large the basis patches are relative to the topographic features of the terrain. Many algorithms for skinning real terrain with a basis set are possible; DMA terrain elevation data provides one particularly good starting point.

Each terrain patch can be thought of as a place holder in the terrain skin. Its visual representation can be a simple polygon, or a complex scene component with embedded level of detail and dense 3D embellishment. The patch can provide additional subdivision of the terrain surface to increase scene complexity, although this additional complexity will be generic in nature. The basis set can include terrain patches with various lineal features, such as roads and rivers modeled at various orientations. Specific 2D and 3D features can be associated with the instance of a patch to provide navigationally specific embellishment.

Implementation of a terrain basis set brings other significant changes to the data base development process. The total amount of geometry that needs to be modeled is dramatically reduced, since only one version of each basis set element must be created. Thus the modeler can spend much more time optimizing the artistic aspects of each patch while dealing thoroughly with how the patch loads the IG. This approach substantially separates the geometric modeling problem from the organizational and hierarchical problems, and in effect puts a lid on the total amount of scene elements which must be created. The data base designer is thus freed of a great deal of modeling busywork, while being strongly induced to concentrate on and optimize the hierarchical aspects of the data base.

The basis set approach also alleviates several of the major IG bottlenecks which heretofore prevented effective NOE imagery. The terrain basis set and the embellishment components constitute a comparatively small amount of data, and most of this data will probably reside in the IG memories all the time. Traffic at the disk interface is greatly reduced, and most of the disk to IG transfers involve blocks of management and structure data, not geometry. Because the basis set elements enjoy a very high reusage factor, the apparent visual density of the scene is much higher.

A SPECIFIC NOE EXAMPLE

Let's see how all of these strategies can be brought together to solve a specific problem. We will reconsider the original Hunter - Liggett NOE problem, and derive a data base design which provides the high level of visual cueing required without exceeding the capabilities of the CT5A. DMA data will be used to define the offset vectors and assign associated basis set patches.

Our goal is to display 3000 polygons in the forward hemisphere of the pilots field of view. This will provide an adequate reserve of system capacity for complex target areas and detailed threat aircraft. We will need to have about 10000 active polygons in the vicinity of the viewer in order to display 3000 after FOV and backface culling. The data base will have four levels of scale, designed to support flight over an altitude range of three to 100 meters. One of these is the terrain skin itself; the other 3 are various sizes of trees, shrubs and rocks. We will arbitrarily divide the total polygon allocation to give the terrain 2500 polygons and each level of scale 2500. The data base design will anticipate an overall size of 25 by 25 kilometers and a general visibility range of two kilometers.

If all 2500 of our terrain polygons are used to define a single level of detail terrain skin within the two kilometer visibility range, the average area of each polygon will be about 71 meters square. We will instead provide two levels of detail. Within 350 meters of the viewer the average terrain polygon area will be 34 meters square; between 350 and 2000 meters an average of five of these polygons will be gathered together and replaced by a single polygon. The low level of detail polygons will be specifically defined while the high level of detail will be instances chosen from a terrain basis set. This will require the compilation and storage of about 108000 polygons for the terrain low level of detail; about 2200 of these will reside in the IG memories at any one time.

The smallest level of scale will be provided by randomly distributed shrubs and rocks up to three meters in size. We will provide three visual representations for each item: three, seven and 15 polygons. These might correspond to a simple tetrahedron tree, a tree with a trunk and shadow, and a tree with several tiers of leaves. We will use some of this capacity to provide 2D ground textural cues by depicting collections of rocks and pebbles; this will help provide the proper spatial distribution of cues when the pilot is closest to the ground. Since we are using fade level of detail, the level of detail changes can take place at 20 and 40 meters, and the simplest representation can be entirely removed from the scene at 80 meters. With these parameters the spacing of ground texture cues (rocks) is about two meters. Vertical cues (trees) are available every nine meters, and the total number of active polygons is the desired 2500.

Each terrain patch has an average of 29 of these small items on it, with each item requiring 15+7+3 polygons, for a total of about 725 polygons. Scene elements for the middle level of scale will be about twice as large, with corresponding increases in spacing and transition ranges, and the large level of scale will be two times bigger still, depicting trees up to 12 meters in height. These other two levels of scale will add another 181 and 45 polygons respectively, to each terrain basis patch. Thus the total number of polygons required to create each terrain patch is about 950. If we decide that a basis set of 31 patches will provide adequate terrain fidelity, the total number of polygons in the entire basis set will be about 29000. Nearly all of these will probably reside in the IG memories; adding in the 2200 low level of detail terrain polygons gives a total memory requirement for this data base of about 31000 polygons. Data rate between the disk and IG memories would be about 30 polygons per second. Note also that the overall size of the data base can be expanded without changing the amount of IG memory required or increasing traffic at the disk.

The density of the resulting data base is around 520000 polygons per square kilometer, counting only the most complex version of each scene element. The optical flow requirements suggested by Sinacori have been met, and this data base can be expected to properly support flight down to three meter eyeheights. The total amount of real modeling effort is very small, since most of the detail is generic and can be produced by compile time software tools. None of the IG limitations have been reached, and a large cushion of performance has been left to allow substantial enhancement of local target areas. The terrain fidelity is as good as DMA level 2 data will provide.

A SAMPLE BASIS SET

A preliminary NOE terrain basis set is shown in figures 1-4. The amount of geometry included in the basis set is designed to provide a nominal amount of cueing to support NOE flight at altitudes of around three meters everywhere in the data base. This sample was designed to closely accomodate Sinacori's recommended design

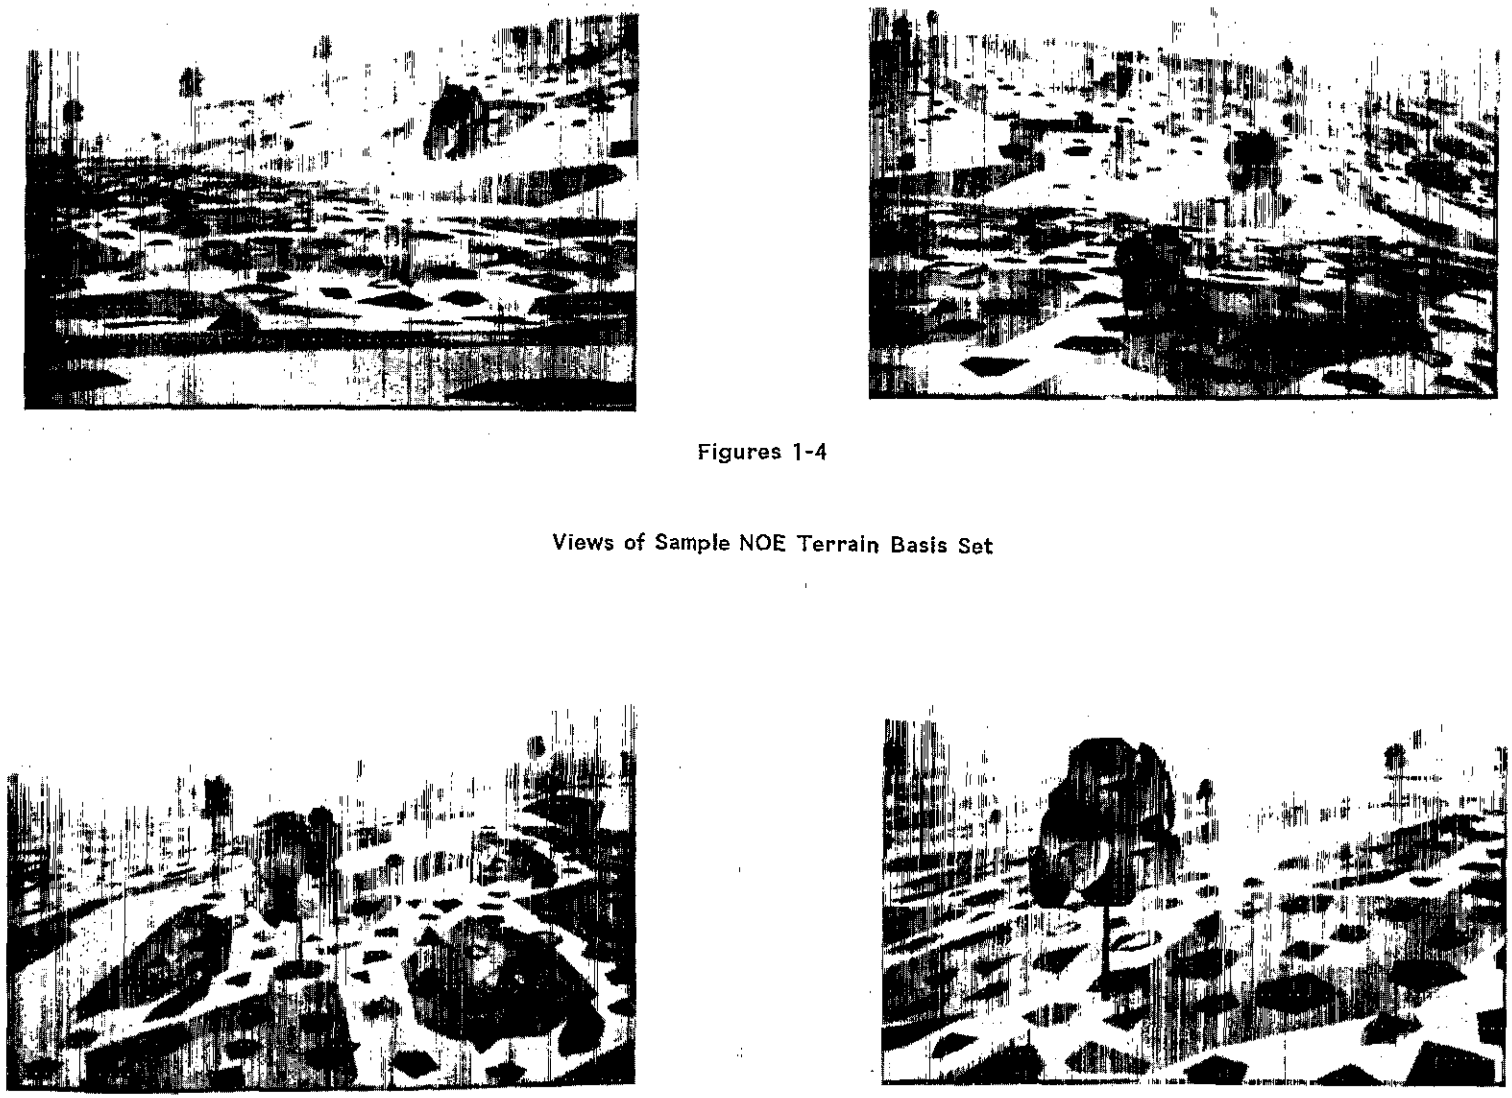

Figures 1-4

Views of Sample NOE Terrain Basis Set

parameters. Cultural and natural detail can be added, as well as specific navigational cues and target deployment areas. Many variations of terrain basis set geometry can be designed and used in the data base interchangeably. The sample photographs were taken with a horizontal channel field of view angle of 70 degrees.

SOME PROBLEMS REMAIN

These techniques add a very useful new tool to the data base designers kit; like other tools, it may not be the best solution for every problem. While the above strategy gets us a long way down the road, some significant problems remain. We have separated much of the geometric data from the management information, and have found ways to reuse the geometric data with fairly high leverage. The data base now derives its specificity from the structural data which organizes the basis set elements; while this is a much more efficient and compact type of data, there are still limits to the amount that can be created, stored and processed. The NOE example above will exceed these limits if the total geographic area is greatly increased. One solution is to decrease the geographic density of the structure information by using larger basis set terrain patches. Trading off terrain complexity for larger gaming areas may well be a viable solution, since a larger data base implies a higher flight velocity and altitude.

The inclusion of specific lineal features such as roads continues to be difficult and inefficient, and adding specific 3D cultural features requires additional effort. The process described above does allow the merging of specific scene details onto the instanced patches, but much of the advantages are reduced. This is also an area where the DMA data source is not especially helpful; the Level V and subsequent forms promise to alleviate some of this problem. As we continue to define better and more automated data base development strategies, and hardware which processes such data bases more efficiently, better cooperation between DMA and those who specify, develop or use these data bases can be very beneficial.

CONCLUSIONS

The 'nap of earth' mission places extraordinary demands on current CIG technology. Historical methods of producing and displaying such data bases have been inadequate. A fundamentally different way of dealing with the problem has been examined. This new approach uses the concept of a basis set to create specific terrain from reusable terrain patches and scene components. This methodology is supported by new data base design and production strategies, and by powerful features within the CT5A hardware. Together these provide the capability to produce and display visual data bases which effectively support 'nap of earth' training missions. DMA data can be used in this process to provide terrain fidelity and correlation with other simulation processes, and has the potential to substantially automate the data base production process.

REFERENCES

1. Defense Mapping Agency, "Product Specifications for Digital Landmass System (DLMS) Data Base", First Edition, July 1977.

2. Sinacori, John B., Draft Final Report, "Research and Analysis of Head Directed Area of Interest Visual System Concepts." NASA Ames, November 1982.

3. Schumacker, Robert A., "A New Visual System Architecture." Proceedings from the 2nd. IITEC Conference, November 1980.

4. Cosman, M. and R. Schumacker, "System Strategies to Optimize CIG Image Content." Proceedings of the IMAGE II Conference, June 1981.

5. Mayer, N. and M. Cosman, "Enhancing the Computer Generated Illusion." Proceedings from the 4th. IITEC Conference, November 1982.

ABOUT THE AUTHORS

Mr. Cary E. Wales has been active in the development of CIG visual systems at Evans & Sutherland for over 8 years, and has contributed extensively to data base design and production technology. Mr. Wales holds a Bachelors degree from Beloit College.

Mr. Michael A. Cosman has provided technical support in the development of CIG visual systems at Evans & Sutherland for over 9 years, working primarily on image quality and system architecture issues. Mr. Cosman holds a Bachelors degree from Brigman Young University.