PUSHING THE LIMITS: INTEGRATION OF NEW TECHNOLOGIES IN GIS

JB Fong, DP Kennedy, TJ Brown, MJ Carlotto
The Analytic Sciences Corporation
55 Walkers Brook Drive
Reading, Massachusetts 01867

ABSTRACT

A GIS application that uses knowledge of terrain features and military doctrine to assist analysts in the performance of the labor intensive and time consuming activity of prioritizing likely areas for unit deployment is described. The system combines technologies which have only recently reached a level of maturity which permits their integration to solve problems previously only considered in concept. Object-oriented database and optical disk technologies have been integrated with a rule-based expert system to interpret the potential for unit deployment over a vast geographic area. The experience of constructing a prototype system has provided a wealth of information about the promise and limitations of the new technologies, and has advanced our ability to represent and reason with high resolution geographic terrain data over very large geographic areas.

INTRODUCTION

The Target Location Prediction Prototype (TLPP) is an expert system that uses knowledge of terrain features and military doctrine to assist military analysts in the performance of the labor intensive and time consuming activity of identifying likely areas for deployment of military units. This activity is currently performed manually using paper maps and a high degree of user visual processing and experience-based reasoning. The analysis is performed with crayons or felt tip markers on clear acetate overlays and paper maps. TLPP provides an interactive digital video medium to address the visual and annotation needs of the analyst, and provides a knowledge-based expert system that reasons about potential locations for unit deployment. The expert system uses a multilayer, digital terrain database that is among the most accurate and extensive of those constructed to date. The expert system determines the suitability of terrain to support the deployment of military vehicles, based on satisfaction of terrain constraints and military deployment doctrine.

TLPP represents an integration of technologies that have only recently reached a level of maturity that makes construction of such systems possible. The system integrates a rule-based expert system, object-oriented database, and write-once-read-many (WORM) optical

---

[1] The work is being sponsored by U.S. Engineer Topographic Laboratories and by DARPA under Contract Number DACA76-88-C-0002. The opinions expressed in this paper are those of the authors, and do not reflect any policies, either expressed or implied, of any granting agency.

disk drive technologies into an interactive environment for assisting military analysts. Development of the TLPP system has addressed problems of extraordinary database size and efficiency of data access. The sheer magnitude of data required (in excess of one million objects), the level of detail used to represent areas of potential interest, and the many diverse kinds of data stored in the system to support the reasoning process have pushed (and sometimes exceeded) the limits of available commercial technology.

Software development has been accomplished in CommonLisp to the maximum extent possible, in order to facilitate hosting the system on a variety of computing platforms. Strict adherence to CommonLisp has been relaxed when necessary to achieve processing efficiency or unique user interface capabilities. As shown in Figure1 the system is currently implemented on a Texas Instruments Explorer II+ Lisp machine workstation networked to a Texas Instruments microExplorer workstation (an Apple Macintosh II with a TI Lisp board installed).

Operationally, a typical session with the TLPP system involves defining a geographic area to be analyzed and indicating the type of unit of interest, defining the scenario (context) for processing, and performing predictions in the defined area. TLPP provides an analyst with the ability to quickly evaluate alternate scenarios; to accomplish "what if" analyses by allowing changes to rules in a terrain deployment rule base and re-running the location prediction process.

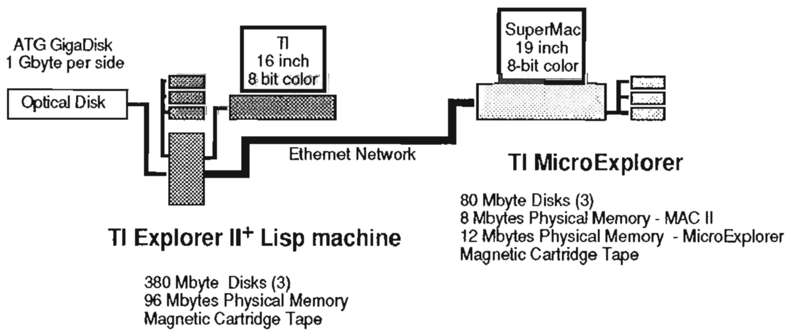

Figure 1 TLPP Hardware

The TI microExplorer provides a user interface which permits the analyst to a) define the scenario or context within which processing is to be accomplished, b) activate the rule-based location prediction process which assesses terrain suitability of areas within an analyst specified area of analysis, c) view and modify the knowledgebase which contains the rules used in deployment area terrain suitability analysis. The TI Explorer II+ Lisp machine workstation a) provides the computational resource for location prediction and analysis, b) displays map or image background and object oriented terrain feature overlays (roads, railroads, streams, classes of surface material, urban areas, etc.) from the GIS database, c) displays analysis results from the location prediction process as an overlay, d) provides for overlay object selection and description, and e) generates explanations for both positive and negative prediction results.

GEOGRAPHIC INFORMATION SYSTEM

The GIS component of TLPP is a geographic information system which provides access to objects which represent terrain features within the area of interest, and which computes spatial attributes of and between related terrain feature objects during the location prediction process. The TLPP system utilizes G-Base™ (a registered trademark of GRAPHAEL, Inc.), a commercially available object oriented database for the TI Explorer II+ LISP machine, to provide access to digitized maps, satellite (LANDSAT) images, and terrain features represented in the form of objects and attributes derived from maps, imagery, and Digital Terrain Elevation Data (DTED).

Digital Terrain Database

The total geographic area covered by TLPP's database is now somewhat larger than 30,000 km^2, and will soon exceed 100,000 km^2. The number of objects in the database currently exceeds one million. Providing management and efficient access to this large database has severely taxed available commercial object oriented database technology.

Organization of data in the database reflects the most common ways in which the data is used for display and user interaction capabilities of the TLPP system. A typical TLPP display consists of a background and several user selectible overlays. Digitized maps and satellite imagery are used as background images for display. Other layers are comprised of objects such as roads, railroads, surface material classes, streams, and builtup areas.

Within the object oriented database there is a class definition for each kind of data object from which instances can be created. A hierarchy of classes is defined in which subclasses inherit properties from their parent super-class parents. In addition, at any class, new properties may be defined which in turn are inherited by sub-classes. Instance objects are instantiated from certain classes in the class-hierarchy. These instances have the properties of the class of which they are a member, whether those properties were defined first at that class or inherited from some super-class. Actual values assigned to the properties defined for instances of a class are what distinguish one instance of a class from another instance of the same class. Two kinds of properties are defined: terminal and structural (non-terminal). Terminal properties are used to express descriptive information about an object, generally through simple values like numbers, character strings, symbols, arrays, or lists of simple values. For example, the *area* of one forest object might be 20 km^2 while the *area* of another might be 35 km^2. Structural (non-terminal) properties are used to represent relationships between objects in the database and are essentially pointers through which objects are connected. For example, one urban area object might be *adjacent-to* a particular forest object.

The topology of the database model of TLPP is a connected graph, shown in Figure2. Classes are depicted as named bubbles in the diagram. Unlabeled arcs between classes represent class-to-subclass links. Labeled arcs in the diagram indicate relations between model classes (structural properties in G-Base jargon). Terminal properties associated with a class, such as *area* or *slope*, are not shown on the model diagram, but are inherited through the class-to-subclass link.

To make it easier to generate data for the database as the system was being built and to permit identification of appropriate data to load into memory for a given area of analysis, a "tile" grid was defined over the region covered by the database. Within each tile, a map, an image, and other layers or overlays of objects are created and stored in the database. Though the database is organized in this way, TLPP presents the user of the system a "seamless" view of the geographical area. Areal features which cross tile boundaries are created at runtime by merging areal sub-features from separate databases. Values of attributes of the composite object are computed in the GIS by appropriate combination of corresponding property values of its components.

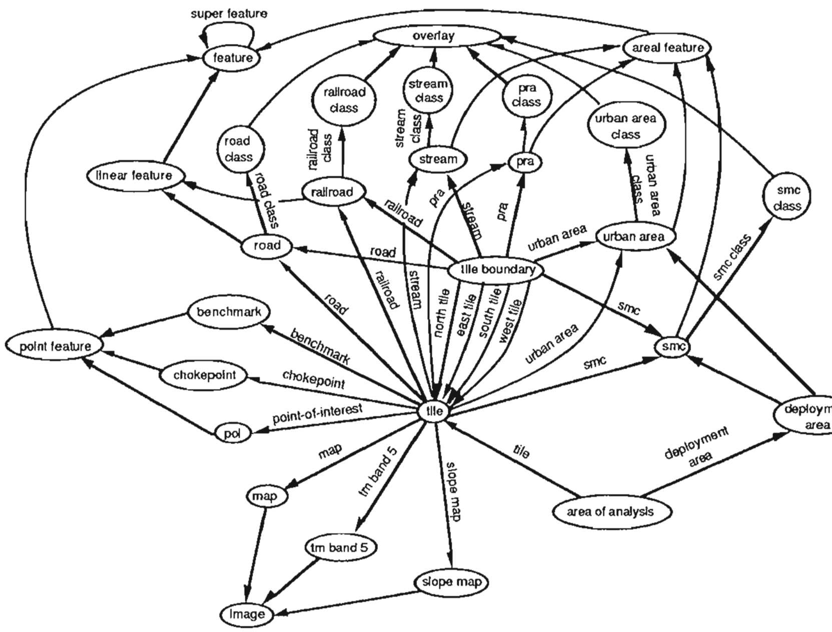

Figure 2 TLPP GIS Database Model

The database for a tile is built by creating instances of specific classes from the database model. Once this is done, attributes or properties of the instances are assigned values which may be accessed/modified during processing. In true object oriented programming sense, class instances also inherit methods from the class from which they are made. This feature of object oriented database technology enables objects in the database to respond to messages, usually by performing some computation and returning a value to the sender of the message. For instance, one could send a message to an instance of the surface material class asking it to return the value of its *area* property . The method responding to that message would look to see if the instance already has a value assigned to its *area* terminal property and, if so, return that value; otherwise the method would compute the area (from a raster representation (Carlotto and Fong, 1988) of the spatially distributed object in TLPP), store the value in the *area* property of the object and return the value to the sender of the message.

Optical Disk

There is a tradeoff in efficiency of processing and speed of retrieval of information from the database. Processing is much faster if all the data needed is resident in memory. But, for a limited amount of physical memory, it is necessary to tolerate database accesses to retrieve what is needed, when it is needed. Each object in the database is represented by a set of properties or attributes. Some of these property values are simple enough to store with the object in the database. Other data structures such as arrays for maps and images, and raster representations of spatially distributed objects proved to be too numerous and large to store directly in the object oriented database. An optical disk is used as the storage medium for this type of data. Layout of data on the optical disk is shown in Figure3. Addressing information for data structures stored on the optical disk is maintained with the corresponding object in the object-oriented database.

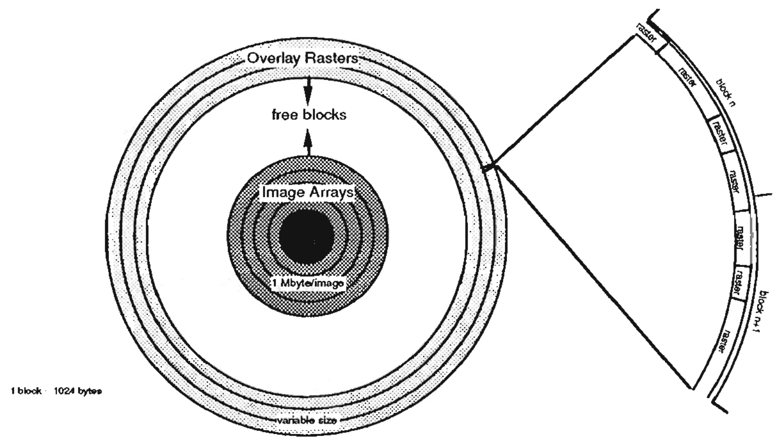

Figure 3 Optical Disk Storage Layout

Raster data, which is a form of run length coding of spatial features, is stored in variable sized chunks in the outer tracks of the 1 Gbyte optical disk platter. Image arrays, for maps and satellite imagery are stored uncompressed in 1 Mbyte chunks on the inner tracks of the platter. The two storage areas expand toward each other as data is added to the platter.

Maps and images are somewhat easier to maintain since each data structure is an uncompressed (for the time being) 1 Mbyte image, requiring 1024 contiguous blocks of optical disk storage transform the centroids of the objects to a global coordinate and compute the distance from global positions.

LOCATION PREDICTION

The second major component in TLPP is a rule-based expert system which applies rules from a target deployment rule base to data obtained from the geographic information system database, within the context of a description of the current situation and processing requirements (scenario) provided by the analyst. (Figure 4) Assessment of terrain suitability is accomplished by sifting through a set of candidate areas to see which ones have acceptable terrain features to support deployment of a particular type of military unit, eliminating from consideration those areas which fail to meet specified constraints on deployment.

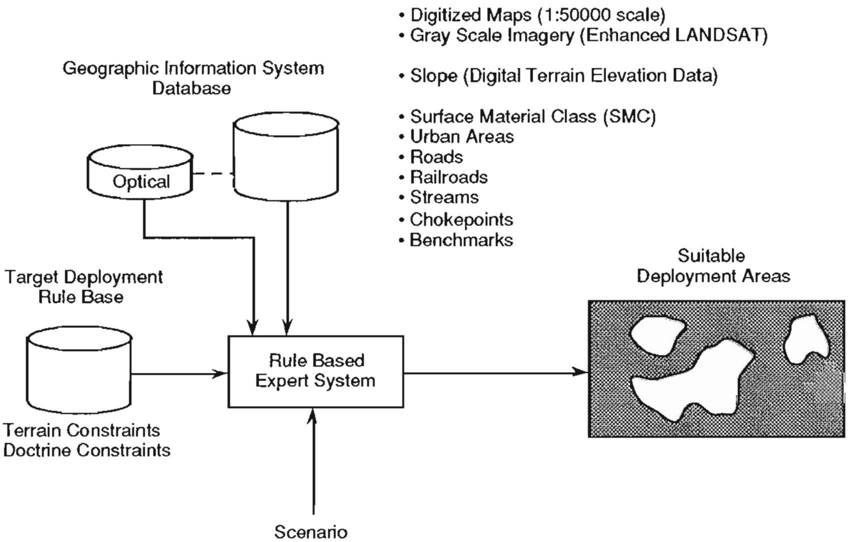

Figure 4 TLPP Operational Data Flow

Suitable deployment areas are displayed to the analyst as an overlay on a high resolution color screen . The background of the display may be either a digitized map or a gray-scale image derived from LANDSAT data. A number of terrain feature overlays (see Figure 4) may also be selected for display over the background. Those overlays are constructed from object descriptions stored in TLPP's GIS database. The multi-layer display is object oriented. An interactive interface supports object selection for attribute description and for explanation of location prediction analysis results. TLPP can generate an explanation of the evidence in favor of or against an area being considered suitable for deployment.

Terrain Constraints

TLPP's knowledge base contains rules which are composed of terrain feature or military doctrine constraints which affect area. Whole blocks are used and no fractional parts are left over as in the case of the raster data structures. Address information stored in the database to effect retrieval of image arrays therefore consists of only a platter-id, and a starting block number for the image.

Platter-id information is used to distinguish one optical disk platter from another with respect to the same database. The same object may be available from more than one optical disk platter, and will have address information for each platter on which it is represented stored with the object in the database. As the database system grows or is updated, or copies of information contained on one platter is copied to another, different platters may receive the same data, but it may not reside in exactly the same place on each platter. The addressing scheme we have used in the database accounts for this real-world practicality.

Implementation Considerations

This project has pushed its object oriented database system to the limit and in some cases further than its design originally permitted. As a result some restrictive design decisions were forced upon the prototype. For example, a single database could not be spread across multiple physical disk units. Our database is of such size it cannot be accommodated on one disk -- therefore it had to be artificially split into a set of smaller databases, each covering a portion of the geographic area of interest, with redundant storage of data at the boundaries between the sub-areas. This is inefficient in terms of stored data and increases database maintenance and integrity problems. Also, the application code, because of the artificial partitioning, had to be written so as to ensure the appropriate database was open for access (the database does not permit more than one application /database to be open at a time) and to switch databases cleanly if the necessary database is not available for immediate access. Finally, user requests which require data from two database subareas had to be trapped and politely denied, even though no logical barrier existed to prevent such a request from being honored.

Spatial Operators

The spatial operators compute relationships of and between various features in the geographic database. There are unary operators that compute the geometric properties of area, length, width, centroid, perimeter, etc. The n-ary operators compute relationships between features, such as distance between, containment, adjacency, etc. Within the GIS the basic operators compute and represent information in a pixel space. At the interface level to the GIS these values are transformed into units of measure, such as meters or kilometers. The details on implementation of the spatial operators that work with the raster representations and images are described elsewhere in these proceedings (Carlotto and Fong, 1990). The implementation of these operators has been extended to work in the environment of the tile organization of the database (Digital Terrain Database section). For example, one forest object might extend across two tile databases and the area for the total object would be the sum of the areas computed for the individual pieces. The GIS operators make this operation tranparent to the rest of the system. An n-ary operator that would compute the distance between two object centers would deployment of a given military unit. The following rule is an example:

A candidate deployment area

never has an area less than 5 km^2
almost always has an area between 40 and 100 km^2
never has a slope of more than 15 degrees
almost always has a slope of less than 10 degrees
usually has a slope of less than 6 degrees

The rule above has five phrases, which describe two constraints: one dealing with the *geometric area* and the other with the *slope* of the candidate deployment area. Phrases in a rule which mention the same attribute represent a single constraint whose conditions are combined with a fuzzy logic (Zadeh, 1979) during inference. This kind of logic permits interpretation of the degree of success in satisfaction of a constraint. A *composite membership function* is defined for each constraint and is applied during reasoning to the actual value for the attribute of the candidate area being tested. Thus, candidate areas can be more or less successful in satisfying constraints, and still be carried along to the next test, provided their *success scores* stay above an annihilation threshold.

Relations between objects in the TLPP database can also be expressed in rule form:

The distance between the candidate deployment area and any class 1 road

is always less than 20 km
is usually less than 15 km

<u>Control Strategy</u>

The location prediction process starts with a set of candidate deployment areas in the geographic region being analyzed. Each member of the candidate deployment area set is considered suitable for deployment of a specified type of military unit until it fails to meet some deployment constraint. The degree to which a candidate deployment area meets a given constraint is determined by the *composite membership function* which returns a score. As constraints are applied to the set of candidate deployment areas, scores are accumulated as a measure of successful evidence. Candidate deployment areas which develop low accumulated scores (thresholded) are dropped from the set of candidate deployment areas. Those areas which continue to satisfy constraints with a degree of success remain in the set of candidate deployment areas and are passed on as a group to the next constraint filter. After all constraints applicable to deployment of the specified unit have been applied, the remaining successful set of candidate areas are denoted deployment areas. Accumulated score is an indication of the relative goodness of one area with respect to another as a deployment area.

SUMMARY AND FUTURE WORK

The TLPP prototype has integrated significant new technologies in the realm of GIS systems. The state-of-the-art in optical disk and object-oriented database technologies integrated with AI technology is such that stable, commercial quality software to support them on a variety of platforms is in short supply. The problem is only compounded by a lack of reliable products for machines such as the TI Explorer Lisp machine. Significant compromises in design and implementation had to be made to use available commercial software. This experience has shed some light on requirements for next generation software and hardware to support GIS applications on the order of the scale of TLPP.

Optical disks need to be faster and provide for greater storage capacity, preferably with the ability to re-write areas of the disk as needed. Concepts such as juke-boxes provide added capacity but need to be faster to support significant GIS activities.

Object oriented databases must provide access over networks, and allow a database to be spread across multiple disk units. Much more efficient access methods and data representations need to be developed. Data backup and restoration facilities need a lot of work to be able to handle database applications of the size represented by the database in TLPP. The ability to permit parts of a database to be designated read-only and stored on optical disk at the same time other parts of the database are manipulated in read/write fashion from multiple conventional Winchester technology disk drives would be a welcome improvement. In general, commercial software in this area must become much more reliable and continue to improve the efficiency with which large data sets can be maintained and accessed.

Processing efficiency is a concern. The kinds of operations typically performed by GIS systems often lend themselves to parallelism of the sort provided by advanced computing architectures such as the Connection Machine. Implementation on this type of machine may provide the needed increase in run-time computational power for GIS operations. The computing environment must still provide the necessary facilities for implementation of an expert system and support for large database storage, management, and access.

REFERENCES

Carlotto M.J., and Fong J.B., "OMEGA: An object-oriented image/symbol processing environment," SPIE, Vol. 1003, November 1988, Cambridge MA.

Zadeh, L.A., "A Theory of Approximate Reasoning," Machine Intelligence 9, Eds. Hayes, et al, 1979