

NAVAL POSTGRADUATE SCHOOL

Monterey, California



THESIS

**AN AUTOMATED SPATIAL DECISION
SUPPORT SYSTEM FOR THE RELOCATION
OF ARMY RESERVE UNITS**

by

Mark A. Murphy

March, 1997

Thesis Advisors:

Daniel R. Dolk

George W. Thomas

Approved for public release; distribution is unlimited.

DTIC QUALITY INSPECTED 3

19971121 025

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instruction, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington DC 20503.				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE March 1997		3. REPORT TYPE AND DATES COVERED Master's Thesis
4. TITLE AND SUBTITLE An Automated Spatial Decision Support System for the Relocation of Army Reserve Units			5. FUNDING NUMBERS	
6. AUTHOR Mark A. Murphy				
7. PERFORMING ORGANIZATION NAME AND ADDRESS Naval Postgraduate School Monterey CA 93943-5000			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.			12b. DISTRIBUTION CODE	
13. ABSTRACT (maximum 200 words) This research analyzes the process used to evaluate potential relocation sites for Army Reserve units from the perspective of military readiness. A comparative decision model (based upon Multi-Attribute Utility Theory), augmented by a Geographic Information System (GIS), was designed and implemented in an automated Spatial Decision Support System (SDSS). This SDSS provides a flexible structure that can be generalized to serve as an executable conceptual model for a wide range of decisions containing significant geographic or location-related components. The Army Reserve Installation Evaluation System (ARIES) integrates several commercial software products in a seamless and synergistic manner. Data extracted from numerous large databases is spatially processed by a commercial mapping engine, and then overlaid onto a formal decision model. The decision maker can rely on a single, simplified interface that consistently applies the professional judgement of a panel of experts to produce standardized reports, or choose from a robust suite of methods for model management, sensitivity analysis, and the display of results. A process that previously required weeks can now be completed in minutes. More important, this approach improves the decision maker's effectiveness by conveniently providing insights into the nature of the source data and the decision process.				
14. SUBJECT TERMS Spatial Decision Support System, site location, Army Reserve, readiness, Multi-Attribute Utility Theory			15. NUMBER OF PAGES 140	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UL	

Approved for public release; distribution is unlimited.

**AN AUTOMATED SPATIAL DECISION SUPPORT SYSTEM FOR THE
RELOCATION OF ARMY RESERVE UNITS**

Mark A. Murphy

Lieutenant Commander, United States Navy

B.S., United States Naval Academy, 1982

Submitted in partial fulfillment
of the requirements for the degree of


**MASTER OF SCIENCE IN INFORMATION TECHNOLOGY
MANAGEMENT**

from the


NAVAL POSTGRADUATE SCHOOL

March 1997


Author:


Mark A. Murphy


Approved by:


Daniel R. Dolk, Thesis Advisor

Daniel R. Dolk, Thesis Advisor


George Thomas, Co-Advisor

George Thomas, Co-Advisor


Reuben T. Harris, Chair

Department of Systems Management

ABSTRACT

This research analyzes the process used to evaluate potential relocation sites for Army Reserve units from the perspective of military readiness. A comparative decision model (based upon Multi-Attribute Utility Theory), augmented by a Geographic Information System (GIS), was designed and implemented in an automated Spatial Decision Support System (SDSS). This SDSS provides a flexible structure that can be generalized to serve as an executable conceptual model for a wide range of decisions containing significant geographic or location-related components.

The Army Reserve Installation Evaluation System (ARIES) integrates several commercial software products in a seamless and synergistic manner. Data extracted from numerous large databases is spatially processed by a commercial mapping engine, and then overlaid onto a formal decision model. The decision maker can rely on a single, simplified interface that consistently applies the professional judgement of a panel of experts to produce standardized reports, or choose from a robust suite of methods for model management, sensitivity analysis, and the display of results. A process that previously required weeks can now be completed in minutes. More important, this approach improves the decision maker's effectiveness by conveniently providing insights into the nature of the source data and the decision process.

TABLE OF CONTENTS

I.	INTRODUCTION	1
	A. BACKGROUND	1
	B. RESEARCH OBJECTIVES	2
	C. RESEARCH QUESTIONS	2
	D. SCOPE	3
	E. PROPOSED SOLUTION	4
	F. THESIS ORGANIZATION	5
II.	THE RELATIONSHIP BETWEEN TPU LOCATION AND MILITARY READINESS	7
	A. MILITARY READINESS	7
	1. The Role of the Army Reserves in National Defense	7
	2. Defining Military Readiness	8
	3. Readiness in the Army Reserves	9
	B. RELATING UNIT LOCATION TO READINESS	11
	C. CHAPTER SUMMARY	12
III.	DEVELOPING AN SDSS FOR UNIT RELOCATION	15
	A. THE NATURE OF A DSS	15
	B. APPLYING THE DSS CONCEPT AS AN SDSS	17
	C. DEVELOPING A DECISION MODEL	19
	1. Identifying the Decision	19
	2. Classifying the Decision	20
	3. Identifying Decision Goals	23
	4. Constructing a Hierarchy of Goals	24
	5. Eliciting and Implementing Preferences	26
	D. MODEL ASSUMPTIONS	30
	1. General Assumptions and Simplifications	30
	2. Specific Assumptions Pertaining to Goals	31
	3. The Resulting Hierarchy of Goals	34
	4. Specific Assumptions Pertaining to Implemented Measures	37
	a. Measures of Facility Quality	37
	(1) Facility Age.	37
	(2) Facility Backlogged Maintenance.	38
	(3) Facility Condition.	38
	(4) Facility Operating Costs.	38
	(5) Facility Ownership.	39
	(6) Facility Weekend Usage.	39

b.	Measures of Fill Level	39
(1)	Area Average Manning	40
(2)	Area Drill Attendance	40
(3)	Area Loss Rate	41
(4)	Area Transfer Rate	41
(5)	Closing Unit Transfers	42
(6)	Competition	42
(7)	Distance to Area Maintenance Support Activity ..	43
(8)	Distance to Equipment Concentration Site	44
(9)	Distance to Recruiter	44
(10)	IRR Available	45
(11)	Reassignments	45
(12)	Recruit Market	45
c.	Measures of MOS Qualification Level	46
(1)	Available MOS - Closing Units	47
(2)	Available MOS - IRR	47
E.	CHAPTER SUMMARY	47
IV.	MODEL SUPPORT: USER INTERFACE AND DATA PROCESSING	49
A.	USER INTERFACE USED FOR SPECIFYING DECISION PARAMETERS	49
B.	SOURCE DATA	51
1.	Data Preparation	53
2.	Geocoding	54
C.	PREPROCESSING PHASE	57
1.	Filtering Based Upon Data Values	59
2.	Aggregating Data	60
D.	PROCESSING PHASE	61
1.	Filtering Based Upon Distance	61
2.	Queries	62
3.	Archiving	63
E.	EVALUATION PHASE	65
1.	Data Confirmation	65
2.	Outputs	67
a.	Goals Hierarchy	67
b.	Site Desirability Rating	67
c.	Ranking Results Matrix	70
d.	Sensitivity Analysis	71
e.	Comments	73
f.	Preference Set Summary	74
3.	Preference Sets: A Mechanism for Accommodating Different Decision Perspectives	75
F.	CHAPTER SUMMARY	77

V. STRUCTURE AND USE	79
A. SYSTEM ARCHITECTURE	79
1. Integrating Shell	80
2. Mapping Engine	81
3. Decision Model Solver	82
4. Data Preprocessor	84
B. ORGANIZATIONAL IMPLEMENTATION	85
1. SDSS User	87
2. SDSS Administrator	88
C. CHAPTER SUMMARY	89
VI. MODEL VALIDATION AND ENHANCEMENTS	91
A. MODEL VALIDATION	91
1. Historical Validation	92
2. Expert Validation	92
3. Parallel Validation	93
4. Analytical Validation	94
B. FURTHER RESEARCH AND ENHANCEMENTS TO THE PROTOTYPE	95
1. Internal Model and Data Improvements: Specific to USARC ..	95
2. External Improvements: General Applicability of the Methodology	98
C. CHAPTER SUMMARY	100
VII. SUMMARY AND CONTRIBUTIONS	101
A. SUMMARY	101
B. CONTRIBUTIONS	102
1. Specific Contributions to USARC	102
2. General Contributions	104
APPENDIX A. MULTI-ATTRIBUTE UTILITY THEORY	107
APPENDIX B. ADDITIONAL MODEL INPUTS	109
APPENDIX C. DECISION MODEL INPUTS (MEASURES)	113
APPENDIX D. QUALIFIED MILITARY AVAILABLE (QMA) FILE	117

APPENDIX E. SOURCE DATA TABLES	119
APPENDIX F. QUERY STATEMENTS	121
LIST OF REFERENCES	127
INITIAL DISTRIBUTION LIST	129

I. INTRODUCTION

This research analyzes the problem of evaluating potential relocation sites for Army Reserve Troop Program Units (TPU's) from the perspective of military readiness. A comparative decision model is designed and implemented in a prototype Spatial Decision Support System (SDSS). This SDSS not only accommodates the extensive refinements expected of a prototype, but also provides a flexible structure that can be generalized to serve as an executable conceptual model for a wide range of decisions containing significant geographic components.

A. BACKGROUND

The sponsor of this research is the Force Support Package (FSP) Readiness Office, a component of the U.S. Army Reserve Command (USARC). This group is tasked with assessing and improving the readiness of priority Troop Program Units (TPU's). A TPU is the basic building block of the Army Reserve force, typically consisting of about 150 reservists. The TPU's that are of most concern to the Readiness Office are in the FSP, which are the units designated for rapid deployment.

In this context, readiness primarily refers to personnel readiness, the ability to maintain a sufficient number of properly trained and qualified people. Numerous factors influence readiness and many of those factors are dependent upon a unit's location. One of the most significant location-related factors is recruiting market, for unlike the active services, reserve units must recruit exclusively from their local area. When a unit is struggling to maintain personnel readiness, sometimes the best solution is unit relocation. Relocation may also be necessary to support force consolidation or restructuring efforts.

The TPU relocation decision incorporates such a large number of factors that to address it in a comprehensive fashion quickly overloads the cognitive abilities of the unaided, human decision maker. In the past, these decisions were typically based upon a combination of intuition and narrowly focused studies. This ad hoc process produced results that often proved difficult to communicate, defend, and build consensus around. Frustration with the inadequacies of the current approach to such a complicated problem inspired the search for a systematic yet convenient, automated tool based upon a decision model.

B. RESEARCH OBJECTIVES

The objective of this study is to develop a formal decision model for the TPU relocation problem and implement that model as a prototype, computer based SDSS. This involves analyzing the nature of the problem and its environment, identifying the relevant decision factors, applying an appropriate decision model, developing the necessary assumptions and simplifications, designing and building an automated prototype, and, to a limited degree, designing an organizational implementation of the system.

C. RESEARCH QUESTIONS

This research will address the following questions:

Primary Research Questions

- How can the TPU relocation problem be structured using formal decision theory?
- What are the relevant decision factors and their relative importance?
- What assumptions and simplifications are needed to produce a tractable decision model?
- What software must be utilized and developed in order to build a prototype system?

Subsidiary Research Questions

- How do the chosen assumptions impact the validity of the model?
- How effective is the implemented approach to the problem?
- What are the limitations of the initial prototype?

D. SCOPE

Although this SDSS addresses but a single decision variable, unit location, the pertinent criteria and implications of this decision are difficult to define, even when bounded by the context of readiness. Despite the risks of suboptimization and oversimplification, it was necessary to impose some restrictive limits on the problem scope in order to produce a implementable system that meets USARC's needs and constraints.

The Army Reserve Installation Evaluation System (ARIES) does not explicitly address the pre-analysis phase of the decision. It is assumed that both the problem and the viable alternatives have already been accurately identified. Performing an ARIES evaluation on the current location of a struggling TPU can provide insight into location-related problems, but this SDSS does not help the decision maker weigh relocation against other action alternatives. It assumes that relocation has already been appropriately chosen, possibly in conjunction with other corrective measures.

In reality, most decision processes are iterative. Analysis of existing alternatives often leads to the discovery of new alternatives and, sometimes, a redefinition of the basic problem. ARIES is intended to be used as a means of structuring only a single cycle of this

process, although the flexibility of the system allows it to evolve with changing iterations as the model converges with reality.

Externally imposed restrictions on the relocation alternatives also limit the scope of the system. Only those facilities currently owned by the Army Reserve are considered as potential relocation sites (approximately 1,500 nationwide). This SDSS does not consider the benefits of obtaining and developing new locations since current legislation effectively prohibits USARC from taking such action.

To minimize cost and avoid expanded data maintenance responsibilities, USARC also specified that all model inputs would be drawn from existing data sources. As a result, only two-thirds of the decision model inputs could be automated. The decision maker is provided with the capability to manually input the data needed to support the other decision factors for incorporation into the evaluation process. USARC determined that a sufficient number of inputs were available to justify the development of ARIES, but it is important for the decision maker to be aware of the pertinent influences that are not automatically captured by the initial instantiation of the model. Both the supported and unsupported measures are discussed in Chapter III.

E. PROPOSED SOLUTION

The proposed solution to the TPU relocation problem is an integrated modeling environment, in prototype form, which augments a multi-criteria decision model with the spatial representation capabilities of a commercial mapping engine. The use of spatial processing and display in concert with a decision model classifies this as a Spatial Decision Support System (SDSS). A controlling shell, written in Visual Basic™ (VB), is used to

integrate the commercial decision model solver (Logical Decisions for Windows™ or LDW) with a mapping engine (MapInfo™), as well as provide a seamless and simplified user interface. The system, named ARIES (Army Reserve Installation Evaluation System), is designed to meet the sometimes conflicting needs of both the novice and experienced users. The chosen approach provides the flexibility needed for a developmental prototype and yields a scalable architecture which can be extended straightforwardly to incorporate many decisions involving multiple criteria, uncertainty, both spatial and non-spatial variables, and both objective and subjective inputs.

F. THESIS ORGANIZATION

The balance of this study is organized as described below. Chapter II discusses the relationship between unit location and military readiness. After describing the basic elements and characteristics of a DSS, Chapter III presents the theory used to structure the TPU relocation decision, and the practical details of mapping this theory to a formal decision model. Chapter IV details the interface used to specify the decision parameters as well as the data processing that must occur to produce the inputs to the decision model. Chapter V describes the overall architecture used to implement the decision model in an automated system. Also provided in that chapter are recommendations on the organizational implementation of the SDSS. Chapter VI provides various validation strategies and Chapter VII presents conclusions concerning the contributions of this project as well as recommendations for further study and enhancements to the prototype.

II. THE RELATIONSHIP BETWEEN TPU LOCATION AND MILITARY READINESS

A. MILITARY READINESS

1. The Role of the Army Reserves in National Defense

One of the key elements of strategic readiness is force structure (Betts, 1995). In the United States military, the overall structure has two major components: Active forces and Reserve forces. "Maintaining a high degree of peacetime readiness in terms of being able to go to war in a short period of time requires maintenance of a large Active force which is costly to maintain. On the other hand, relying largely upon Reserve and National Guard forces during peacetime, while less costly, extracts a penalty in terms of how quickly the United States can respond to a threat" (Dolk, 1995). To strike an appropriate balance between Active and Reserve forces, it is necessary to understand their expected contributions to national defense.

Under current scenarios, the U.S. Army Reserve (USAR) is considered a primary provider of combat *service* support for the Army, and a major provider of combat support. Given the decreasing size of the Active forces, the role of the reserves is increasingly important. The problems associated with a reduction in active strength are being exacerbated by an increased participation of the Active forces in operations other than war (e.g., humanitarian operations, peacekeeping operations), which often reduces the resources available for preparation of those forces in areas of conventional warfighting. The military readiness of early deploying Reserve units is of increasing importance to the U.S. military.

As the role of the reserves increases, there are growing recruiting difficulties brought on by a drop in the size of the prime military age group. In recent years, this problem is compounded by relatively strong and consistent economic growth, increasing the viability of the often competing civilian employment market. "Obtaining full reserve unit manning, a major requirement in maintaining desired levels of readiness, is becoming a more important goal at the same time that it is becoming more difficult to achieve" (Borack et. al., 1985).

2. Defining Military Readiness

In a broad sense, readiness is the ability to provide the needed resources within given time constraints. Betts (1995) suggests three categories of readiness: operational, structural, and mobilization.

Operational readiness is about *efficiency* and is measured in terms of how soon an existing unit can reach peak capability for combat. Operational readiness is assessed according to inward-looking standards: the absolute potential inherent in the unit and the difference between its actual capability and that potential. This standard has nothing to do with how many units at that level of efficiency might be needed to beat the adversary, or what larger number of units at a lower level of efficiency might still be able to fight successfully. It indicates how proficiently a unit may fight, but not whether it will win.

Structural readiness concerns *mass*; it is about how soon a force of the size necessary to deal with the enemy can be available. Structural readiness refers to the number of personnel under arms with at least basic training, the number of formations in which they are organized, the quantity and quality of their weapons, and the distribution of combat assets among land, sea, and air power. The standard for assessment is outward looking: the relative effectiveness needed to engage the enemy successfully.

Mobilization readiness refers to the "... small peacetime nucleus of military forces for structural expansion, and of the government administrative apparatus for coordinating the

changeover of the civilian economy to war production” (Betts, 1995). Of the three, operational readiness is the focus of this study.

3. Readiness in the Army Reserves

Given the unique challenges faced in reserve manpower supply, and the clear effects of manning issues on readiness, USARC relies upon fill level, Military Occupational Specialty (MOS) qualification level, and turnover rate as their primary indicators of unit readiness. These indicators provide basic metrics about whether a unit has a sufficient number of people and whether those people possess the needed skills to support the unit’s mission.

Three major issues distinguish the manpower supply to the Reserves from that of the Active forces: reliance on local labor markets, heavy dependence upon prior service recruits, and status as a secondary form of employment. The Active forces have the luxury of recruiting on a national basis and relocating their members as needed. The reserves, however, must draw their membership from the local population. They rely heavily on an environment over which they have little influence. A second issue which distinguishes the reserve manpower supply is a heavy dependence upon prior service personnel. Approximately half of Army Reserve accessions have prior military service, as compared with less than 10 percent for accessions to the active duty Army. A third aspect is the fact that the Reserves are generally not an individual’s primary job. Approximately 90 percent of reservists hold full time jobs (McNaught, June 1981; McNaught July 1981; and Burright, et. al., 1982). The high turnover rates experienced by the reserves are partially explained by its status as a secondary form of employment (Borack et. al., 1985).

One of the primary obstacles to operational readiness in the Army Reserves is a high turnover rate. The average annual turnover rate in Army Reserve units from 1992 to 1995 was 35 to 37 percent (Dolk, 1995). This turbulence typically has a deleterious effect on the quality and effectiveness of training, operational efficiency, and organizational cohesion, impeding a unit's ability to attain its "absolute potential". Because of the time needed for the initial training of new recruits and the retraining of transferred reservists, even those units able to quickly replace their losses and maintain their fill rates, can rarely maintain the levels of training and qualification deemed necessary to achieve operational readiness.

Fill level is a related but separate issue from turnover rate. Even if a unit has a low turnover rate, if it is unable to replace its losses to maintain a desired level of equilibrium manning, it is not able to achieve operational readiness from the perspective of possessing sufficient human resources to accomplish its military missions. In undermanned units, vacancies may also force the available reservists to assume additional responsibilities, further reducing overall efficiency. For various reasons, units are activated as a group and so individual vacancies are not normally filled with supplements from non-deploying units. The fill levels of units in the FSP, those units designated for rapid deployment, are particularly important.

Military Occupational Specialty (MOS) qualification levels, although strongly influenced by turnover rates and fill levels, provide a slightly different indication of readiness. These qualifications are used to formally document attainment of the skills deemed necessary to support the unit's performance of its military missions. Two units with the same turnover rate and fill level can perform quite differently in this measure depending

upon whether they are able to recruit replacements that already hold these qualifications and how demanding the qualification process is for the MOS's of interest.

B. RELATING UNIT LOCATION TO READINESS

Intuition, experience, and various studies have indicated that many of the factors influencing unit readiness are in some way a function of location. Access to quality recruiting markets is often the most important factor in maintaining fill levels. Even the units with the highest retention rates can find it difficult to replace their losses if located in a meager recruiting market. Location also determines the distance to the nearest recruiting station, as well as the distance to the nearest support sites and major training facilities. Location even determines the compatibility between a unit's mission and the local civilian occupational structure, which for some types of units can be a significant readiness factor.

Experience and formal studies suggest that turnover rates are also heavily influenced by location-related factors. Turnover rate is primarily an issue of personnel retention. A study of separating reservists revealed that dissatisfaction over wasted training time was the most frequently cited reason for leaving the reserves (Boykin, et. al., 1980). Training efficiency can be related to location in a number of ways. The distances to weekend training (WET) sites, special training sites, and weapons qualification ranges determine the amount of available training time "wasted" in transit. When training involves equipment that is impractical to store at a unit's facility, the distance to the nearest Equipment Concentration Site (ECS) also becomes important. The distance to the Area Maintenance Support Activity (AMSA) is one indicator of the speed with which assistance can be provided when training

equipment breaks down. All of these locational issues can affect the efficiency of training, which in turn significantly influences morale, retention, and turnover.

Another group of location-related factors that may influence retention are the distances to various support sites. Reservists may become frustrated if they must travel excessively to receive pay and administrative services or to take advantage of commissary and exchange privileges. The relocation site selection also affects how crowded the facilities will be (primarily an issue for full-time support personnel) and the physical condition of the structures to be used. These factors influence morale and retention.

One unit's location can also influence the success of other units through economies of scale and the draw on regional resources. If units are widely dispersed, soldiers must either travel farther for their training or else training must be provided at a higher cost for a smaller number of individuals. If units are highly concentrated, the effective size of the recruit market can be significantly influenced by the competition from other units in the same area. Operational readiness is influenced not only through the characteristics of the relocation area but also by the distribution of other units.

C. CHAPTER SUMMARY

The forces of the Army Reserves play an increasingly important role in national defense, and yet it is becoming increasingly difficult to attract and retain a sufficient number of qualified individuals. Military readiness can be classified into two types, structural and operational. The focus of this project is on the operational readiness of those Army Reserve units scheduled for rapid deployment in time of conflict, but there are clearly implications for force-wide structural readiness as well.

The statistics chosen by USARC to serve as the primary indicators of unit operational readiness are fill level, MOS qualification level, and turnover rate. Performance in these areas can be related to numerous location dependent factors including access to preferred recruiting markets and distances to various training and support sites. This project is based on the premise that, holding all other readiness variables constant, it is possible to improve the operational readiness of some Army Reserve units by relocating them to preferred areas as indicated by a variety of location related attributes.

III. DEVELOPING AN SDSS FOR UNIT RELOCATION

A. THE NATURE OF A DSS

Just as there is no universally accepted definition for military readiness, there is also little consensus on the exact definition of a Decision Support System (DSS). Scott-Morton is often credited with being the first to articulate the concept of a DSS in the 1970's under the term "management decision system". Many definitions have been subsequently suggested, but most seem inappropriately restrictive based on issues such as usage patterns (Moore and Chang, 1980), system components (Bonczek et. al., 1980), or development processes (Keen, 1981). Turban (1990) suggests the following working definition which appropriately accommodates a wide range of systems:

A DSS is an interactive, flexible, and adaptable computer-based information system that utilizes decision rules, models, and model base coupled with a comprehensive database and the decision maker's own insights, leading to specific, implementable decisions in solving problems that would *not* be amenable to management science optimization models per se. Thus, a DSS supports complex decision making and increases its effectiveness.

Several characteristics of a DSS suggested by Turban are not explicit in this definition. A DSS:

- provides support for decision makers mainly in semistructured and unstructured situations by bringing together human judgement and computerized information.
- supports a variety of decision-making processes and styles; there is a fit between the DSS and the attributes of the individual decision makers.

- must be adaptive over time. The decision maker should be reactive, being able to confront changing conditions and adapt the DSS to meet these changes. A DSS must be flexible so users can add, delete, combine, change, or rearrange basic elements (providing fast response to unexpected situations). This capability makes possible timely, quick, ad hoc analyses.
- should be easy to use. User-friendliness, flexibility, and strong graphic capabilities can greatly increase its effectiveness. This ease of use implies an interactive mode.

In addition to these characteristics, Turban emphasizes some basic DSS concepts. The priority of a DSS is to improve the effectiveness (accuracy, timeliness, quality), rather than the efficiency (minimal use of resources), of a decision. Furthermore, even though an automated system may be able to improve decision quality, it can not replace the human decision maker. The DSS user should be provided with complete control over all steps of the decision making process, with the ability to override the computer's recommendation at any time. The interaction between human and computer leads to learning and should support a continuous process of developing and improving the DSS.

The complexity of an interactive, flexible, and adaptable system, capable of implementing these characteristics and concepts, can appear overwhelming at first. One approach to simplifying this system, at least conceptually, is to break it into constituent parts. A standard paradigm for decomposing the DSS architecture posits three major components: data, models, and a user interface (Sprague and Carlson, 1982). The data component includes a database, a database management system (DBMS), a data directory (dictionary), and a means of query. Similarly, the model component includes a model base, a model base management system, a directory, and a means of executing and integrating models. The user interface provides communications between the other two components and a user.

B. APPLYING THE DSS CONCEPT AS AN SDSS

Hoping to leverage previous USARC success with GIS applications, the original goal was to build an effective decision aid entirely within a GIS framework, using the geographic software to provide all three components: models, data, and user interface. USARC had previously developed a GIS application for tracking closing units which inspired the initial concept of a GIS-centric system able to display recruit market information, other units competing for the available recruits, and the support and training sites that influence readiness. By providing a geographic visualization, it was hoped that the impact of this spatial data on the TPU relocation decision could be better assessed. The system would also provide convenient access to the underlying data and a variety of query capabilities, but would still be used for little more than an automated map. Although a few unit locations could be easily identified as "well supported" or "poorly supported" by visually evaluating the local recruiting market and the clustering or sparsity of support and training organizations, the majority of the proposed relocation sites were not easily categorized from casual inspection. This problem was exacerbated as the number of decision criteria grew.

To improve the clarity and the utility of the displayed data, a suite of thematic maps was built to capture the multi-dimensionality of the problem. These maps employed a collection of symbols that varied in shape, size, color, intensity, and pattern to communicate multivariate data. The restrictiveness of this approach quickly became obvious. Capturing more than four or five distinct variables at a time strains the comfort, if not the limit, of human comprehension (see Tufte, 1983). This is somewhat analogous to the cross-tabulation situation for multiple variables where one table effectively shows only the

interaction between two of the variables with the others being held constant. Relying upon this strategy would have required the user to possess a prodigious cross-correlation ability in order to make sense of the data. Decision structuring was burdensome in this environment for it lacked a systematic means of comparing alternatives. That approach did not yield consistent and defensible support for the decision maker.

One of the most significant contributions of this research was the introduction of explicit decision models as a context for interpreting the spatial data. Rather than trying to capture the multivariate nature of the problem solely through spatial representations, the problem was mapped to a multi-attribute utility model, using a hierarchy of goals as a means of providing a decision structure. In this revised approach, the primary roles of the Mapping engine were to provide the initial user interface and facilitate the spatial processing of data (e.g., via geographic queries such as “find the number of facilities within a specified radius”). The outputs of the Mapping engine were transferred via an overarching program shell to a commercial multi-criteria software application (LDW), which ranked the alternatives. Chapter VI provides a detailed discussion of the system architecture.

This integrated structure finally provided all the key components of an SDSS which we called the Army Reserve Installation Evaluation System (ARIES). The model component was embodied almost entirely in LDW, which permits the use and seamless integration of multiple models and a variety of preference elicitation methods (e.g., Analytical Hierarchy Process, importance orderings, swing weights, tradeoffs, weight ratios). Responsibility for the data component was split between the Mapping engine software (MapInfo™) and Visual Basic™ (VB). Any data requiring spatial processing (i.e., used to

calculate distances or make proximity determinations) was managed by MapInfo™. Visual Basic™ provided management of the large source databases, data extracts, and interim tables. Although both MapInfo™ and LDW include modern Graphical User Interfaces (GUI's), the requirement for a single, simplified interface drove the development of the ARIES control screens using VB.

ARIES exhibits the previously mentioned characteristics of a DSS:

- It is a flexible modeling environment capable of improving the effectiveness of loosely structured, multi-criteria decisions, particularly those decisions with a significant geographic component.
- Because it easily integrates a variety of methods for eliciting and structuring preferences, it can accommodate a variety of decision making styles.
- The decision maker is provided significant control over the basic structure of the decision model and the overall system, permitting both to adapt over time.
- The seamless integration of decision software and a mapping engine provides an overall system that is user-friendly, flexible, and benefits from strong graphic capabilities.

C. DEVELOPING A DECISION MODEL

1. Identifying the Decision

At the heart of the ARIES architecture is a decision model. Before a model could be developed, it was necessary to clearly define the TPU relocation problem. Although USARC representatives suggested a variety of ways in which to understand the importance of unit location (e.g., distances, market supportability areas, overlap between units, etc.), it quickly became clear that the primary objective was to relate location to unit readiness. The

decision was made to develop a model that could isolate and evaluate the location sensitive portion of the unit readiness problem.

In order to evaluate a relocation site from the perspective of readiness, it was necessary not only to capture the appropriate characteristics of the new location, but also to incorporate characteristics of the relocating unit that indicated its needs or compatibility with a new area (e.g., size, Military Occupational Specialty structure, home addresses of members). For the prototype model, the expert panel made the assumption that the basis for evaluation would be a single reserve unit. One alternative to this, relocating a portion or derivative of a unit, is sometimes more appropriate, but that course of action was not directly addressed by the initial model. Another option was relocating multiple units to the same area. The interactions between these moved units could be constructive, destructive, or neutral from the readiness perspective. Chapter VII discusses the modifications that would be necessary to adapt ARIES to encompass these "micro" and "macro" options.

2. Classifying the Decision

The TPU relocation decision can be classified as semi-structured, because for the average decision maker, all phases of the decision are neither fully structured (i.e., routine, repetitive problems for which standard solutions exist), nor fully unstructured (i.e., vague problems which defy all standard solutions). Although most aspects of this decision are based upon easily defined calculations (e.g., distances, averages, sums), the subjective interpretation of the calculated values introduces considerable uncertainty.

This is an example of the type of problem that can be best supported by, "... bringing together human judgement and computerized information" (Turban, 1990). The

initial version of the model not only processes a large number of objective inputs, but also captures the judgements of an expert panel for the subjective aspects of this decision. ARIES provides the user with the flexibility to treat this either as a structured decision, relying solely on the stored, subjective inputs of the experts, or a semi-structured decision, allowing the user to specify new perspectives and preferences.

The next step was to select a model type (e.g., optimization, simulation, heuristics) to provide a framework for analysis, comparison and understanding. An optimization model was considered and rejected. Such a model could identify the ideal geographic coordinates for a relocated unit, but the constraint on USARC to only use sites currently owned by the government, diminished the usefulness and applicability of such an approach. In addition, producing a meaningful optimization model that could incorporate the large number of pertinent decision criteria would have been quite difficult. The additional cost and complexity of optimization was not warranted. Another alternative was to use multiple regressions to establish site desirability as a function of locational attributes, but in this case meaningful data-series for readiness are lacking as are time-series for the attributes. The *decision analysis* approach proved to be appropriate for this problem, primarily due to the finite, manageable number of alternatives which could be readily identified and directly assessed in terms of their value under each of the decision criteria. This approach provided a comparative model, thus satisfying the most basic needs of the USARC decision maker, a ranking of alternatives and insights into the evaluation process.

In order to apply the decision analysis approach, it was first necessary to identify the overall objective, or top-level goal, of the decision. The initial inclination was to simply

select military readiness. Further discussions with the expert panel revealed their desire to include other factors that had to be balanced with readiness, leading to a more general goal, site desirability. Lacking a direct measurement of desirability, the panel decomposed this goal into two subgoals, support of personnel readiness (the ability to maintain the desired number of qualified reservists at the proposed site) and site quality (a general assessment of the costs and benefits of a location that are only loosely related to readiness). Based upon the introduction of multiple, possibly conflicting objectives, the specific version of decision analysis chosen for this problem was a Multi-Criteria Decision Model (MCDM).

Using the MCDM approach, the two primary decision criteria were further decomposed to various location-related measures. The disparate units of the input measures (e.g., miles, number of people, dollars) raised at least two questions; what was the relationship between the input measures and the decision goals, and how should the inputs from such varied measures be combined? Utility theory was used to address both of these uncertainties. Utility functions were specified to convert the units of each measure to common utility units. This conversion reflects the relative desirability of all values of the measure, over an expected range, on a standardized scale. The way in which these common units are combined depend upon the degree to which a decision maker's preference for each measure is influenced by other measures. The underlying basis of this approach is a specialized version of MCDM known as Multi-Attribute Utility Theory (MAUT).

MAUT applies to problems that involve multiple decision dimensions and uncertainty. It is based primarily on the work of Keeney and Raiffa (1976) and uses either a linear or a multiplicative combination of utility ratings to provide an ordinal ranking of

alternatives. The ordinal scales are arbitrary, and consequently, it is not possible to say anything about the degree to which one alternative is preferred to another based on the final rankings. A MAUT function can be formulated if the input values are measurable and satisfy specific mathematical requirements (see Appendix A).

3. Identifying Decision Goals

A logical first step in applying MAUT is identification of decision objectives. Most decision literature uses the term *objective* when referring to a desired direction and *goal* when discussing quantifiable progress in that direction. The documentation for LDW uses the term *goal* when referring to what is more commonly known as an *objective*. It also refers to specific attributes of an alternative as *measures*. For purposes of this discussion, we will restrict our usage to the terms *goals* and *measures*.

MacCrimmon (1969) suggests three approaches for generating goals: (1) examination of relevant *literature*, (2) *analytical study*, and (3) *casual empiricism*. Examination of *literature* revealed myriad sources discussing the general decision of facility location, but robust conceptual models for military readiness of reserve units are largely lacking. Although no *analytical studies* were found that directly addressed the overall issue of military readiness for reserve units, a number of studies were available on various aspects of the problem (e.g., market supportability, reserve retention, commuter behavior) and these served as the basis for specifying a limited number of goals and relationships. As additional research is conducted, the structure of the ARIES decision model is such that it can easily accommodate the incorporation of new inputs and statistically estimated relationships.

Rather than rely on *casual empiricism*, professional judgement was chosen as the primary means of generating goals. This The relocation problem was discussed with a group of knowledgeable experts, who had been involved in similar decisions in the past. By analyzing the way in which these decisions were previously approached, a structure of pertinent goals and relationships was developed. Although this methodology risked formalizing the flaws of an informal approach, the primary alternative, initiating new analytical studies of readiness, was determined to have unacceptable risk, cost, and data availability. The combined experience of these experts promised to not only be more expeditious, but also more valid in many of the more subjective decision criteria.

4. Constructing a Hierarchy of Goals

Once the pertinent goals were identified, they had to be structured and prioritized. Some specific goals were actually components of broader, more comprehensive goals. Discussions during this phase of model development identified the previously mentioned desire to include decision goals that could not be directly related to readiness.

The overall goal was changed from military readiness to site desirability. If the decision maker could easily and consistently rank the alternatives based directly on their performance under the most comprehensive goal, there would be little need for automated support. Given the cognitive limits of the human mind, the number and complexity of decision factors involved in the TPU relocation problem would quickly overwhelm the unaided decision maker. An intuitive approach to this challenge was to divide the decision into more limited and manageable components (sub-goals) which are typically easier to evaluate.

Consistent with MAUT, goals can be decomposed multiple times to form a hierarchy of goals, where each branch terminates in some measurable attribute of an alternative (known in LDW as a *measure*). The appropriate degree of decomposition is normally influenced by the scope, explicitness, and relative importance of a goal, as well as the desired level of detail and objectivity. In general, the more a goal hierarchy is subdivided, the easier it is to identify inputs that can be objectively assessed. One of the priorities for USARC was to subdivide as necessary to accommodate objective inputs on each branch of the hierarchy. In addition to the motivations mentioned above, this approach was also driven by a desire for convenience and consistency. It allowed almost all input values to be extracted directly from existing databases thereby minimizing the inputs required from the user, who is asked to identify only the moving unit and the proposed relocation sites. The intent was not to impugn the validity or value of individual subjective inputs, but rather to standardize the analysis by relying on a panel of experts for the subjective interpretation of the available objective data.

Various principles were applied during the decomposition of goals. All pertinent components of higher goals were accounted for in one, and only one, of the subgoals or measures. This ensured that none of the key considerations were omitted, and avoided redundancies between measures that could result in an over- or under-statement of the effect of an individual measure on the overall decision. Another principle employed was the "test of importance" (Ellis, 1970). This test challenges each goal and measure on the basis of whether it can alter the preferred alternative. Challenging each factor helped to eliminate

those with little relevance to the decision which were included in the first cut of the decision model.

The hierarchy of goals resulting from the application of these principles is shown in Figure (1). This hierarchy represents the consensus of the experts and evolved over the course of this research. Constructing an explicit model forced the experts to recognize and express their underlying beliefs, priorities, and assumptions. As an example, the original intention was to model only those criteria that directly influenced readiness, but it eventually became evident that there were other considerations (e.g., the condition and cost of maintaining a facility) that could not be ignored. Every phase of the modeling process, from the definition of a hierarchy of goals, to the capturing of the decision maker's preferences, inspired introspection and discussion on the critical aspects of the problem. Even if this SDSS had never been implemented, the experts would have significantly benefitted from the insights gained during the modeling process.

5. Eliciting and Implementing Preferences

From a practical perspective, the modeling discussed thus far merely provided a framework for the development of a multi-attribute utility (MAU) function. The MAU function is the tool used to calculate the overall utility, and thus the ranking, of each alternative. The initial utility function developed for the TPU relocation decision is of a multilinear form, containing both additive and multiplicative terms. LDW automatically constructs this function based on the preferences and interactions indicated by the user.

For situations where there were no interactions between measures, the corresponding utility functions assumed a relatively simple, additive form. The assessment of an n -measure




Figure 1. Complete Hierarchy of Goals

utility function was reduced to the assessment of n one-attribute utility functions and $n-1$ independent scaling constants. The necessary and sufficient conditions for additive utility functions were defined by the work of Pruzan and Jackson (1963), Fishburn (1967a, 1967b, 1968), and Pollak (1967). The conditions for additive independence, with n attributes (i.e., input measures) can be expressed as follows. "Attributes $X_1, X_2, \dots, X_n$ are additive independent if preferences over lotteries on $X_1, X_2, \dots, X_n$ depend only on their marginal probability distributions and not on their joint probability distribution" (Keeney and Raiffa, 1976). Using consequence and lottery tradeoffs with the expert panel it was determined that preferences for many measures could in fact be influenced by the level of other measures. An example of this was the preference for the number of potential new recruits living in an area being influenced by the number of available Individual Ready Reserve (IRR) members in that same area (IRR members are civilians with prior Army Reserve service). Having a large number of IRR members available to recruit from reduces the utility of alternative sources of recruits, such as non-prior service civilians. In addition, a third measure, the number of potential recruits from area units that are scheduled to close, also influences preferences. In these cases, the utility of the lottery on one measure depends upon the joint probability distribution of multiple measures. Other pairs of inter-dependent measures include:

- facility age and facility condition (condition based on visual inspection).
- the number of available people with the desired MOS from closing units and from the IRR.
- average loss rate and transfer rate associated with an area.

For measures exhibiting such codependent properties it was necessary to define multi-measure utility relationships, thereby introducing multiplicative terms and additional scaling constants into the overall utility function.

Using the interface tools provided by LDW, various levels of preference information were elicited from the expert panel. First, independent of the alternatives, measure levels were converted from their original units to the standardized units of utility by graphically defining Single-Measure Utility Functions (SMUF's). LDW offers four other methods for defining SMUF's but three of them involve pairwise comparisons of specific alternatives which would not have been appropriate for a reusable set of SMUF's (recall the desire to make this decision appear to be fully structured to some users). Utility is a measure of the worth or value that the decision maker attaches to an outcome or situation; in this case, the value is most often judged in terms of a perceived contribution to readiness. Utilities are subjective, context-sensitive, and may change over time, and so the utility functions defined by the expert panel serve as a default set that can be easily modified.

After defining the SMUF's, weights were assigned, indicating the relative importance of the measures and goals. LDW offers seven methods for specifying weights, but the expert panel felt most comfortable with directly assigning these values. Tradeoff methods were sometimes used to confirm or refine the chosen weights. These weights served as the basis for the initial scaling constants.

Determining the multiplicative terms and interaction scaling constants of the utility function required additional preference information from the decision makers. The values

were defined by responding to additional tradeoff or probabilistic questions. The general form of the multilinear utility function is discussed in greater detail in Appendix (A).

D. MODEL ASSUMPTIONS

1. General Assumptions and Simplifications

The creation of a manageable decision model for the TPU relocation problem required numerous simplifying assumptions to reduce the complexity of the overall problem. The major assumptions pertaining to the overall construction and use of the decision model include:

- Reserve unit location will have little or no influence on where people choose to live. People will not move just to be closer to the unit or relocate when the unit does.
- The “area of the proposed site” refers to the region within 50 miles of the facility. Measures such as the number of competitors and market for potential recruits assume that anything or anyone located outside of this area will have no effect on the relocated unit. Although the value of 50 miles can be easily changed to another constant, at least one study (Sohn and Thomas, 1991) has suggested that this distance should be varied based on the predicted commute behavior of the local population. The commute distance model developed by that study is not incorporated into the prototype SDSS.
- This model treats recruiting efforts as an uncontrollable attribute of the environment. Differences in the effectiveness of various recruiting commands are ignored. The only measure of recruiter support in this model is the distance to the nearest recruiting station. Based on the priority placed on active duty recruiting, it is assumed that the USAREC (U.S. Army Recruiting Command) will not relocate recruiting stations to accommodate the location of TPU’s. It is also assumed that the effectiveness of the recruiting effort varies with the distance to the recruiting station. The exact nature of this relationship is captured in the utility function associated with the Distance to Recruiter measure.

- Out of necessity, many measures are based only on numbers of people, ignoring the differences in contributions to readiness made by individuals. This simplification is easier to accept if one assumes that the *average* contribution to readiness is constant between the groups under consideration. If two TPU's have a 50 percent annual turnover rate, it is therefore assumed that they experience the same impact on readiness even though in reality, one unit may have lost a greater number of people who are considered more critical to the unit's mission.
- In some situations, only the distance to the *closest* support site (e.g., Area Maintenance Support Activity (AMSA), Equipment Concentration Site (ECS), recruiting station) is considered relevant. By using only the distance to the closest site, there is no consideration given to having additional sites available (within the 50 mile radius) even though they may be only slightly farther away.
- Distances are straight-line calculations and do not reflect the actual distance that would be traveled using available roads. This simplification was used because the data necessary to implement a more realistic approach was unavailable.
- Some measures were chosen as empirical indicators of the desirability of an area. For these measures (e.g., Area Drill Attendance, Area Loss Rate, Area Transfer Rate, and Average Area Manning), it is assumed that the values for the relocated unit will be better for those locations where the average value based on all of the units already located in the area is better. This approach ignores the significance of the type of unit being relocated (e.g., trucking, artillery, medical), assuming that all unit types are equally appealing to the local population, regardless of the local job structure. By awarding a high score to those locations with successful units in the area, the possible detrimental effects of competition is ignored in these measures. The effect of competition is captured in a separate measure.
- Although ARIES does not provide a rigorous, causal model of readiness, it is assumed that the hierarchy of screening factors provides a meaningful assessment of the propensity of a given location to support the achievement of high levels of readiness. This approach also assumes that other necessary contributors are present, particularly those that are not related to location. Even though different relocation sites will normally result in different reservists in key positions, it is assumed that readiness factors such as leadership are constant from site to site.

2. Specific Assumptions Pertaining to Goals

In addition to the general assumptions listed above, there were many other, more specific, assumptions that were made during the construction of the hierarchy of goals.

Model formulation was primarily conducted in an open forum with the USARC experts. This section presents the assumptions that surfaced during this process.

The overall desirability of a proposed site must weigh the site quality with the ability of the area to support personnel readiness. These two subgoals are intended to encompass all of the important facets of site desirability pertinent to the TPU relocation decision in a readiness context.

Pursuing increased objectivity, personnel readiness was decomposed based on two fundamental questions; is there a sufficient number of personnel and do they possess the necessary skills? As discussed in Chapter I, unit readiness is a complicated issue subject to numerous influences such as training, leadership and even individual talents. The ARIES model reflects the simplified approach used at USARC by treating fill level and MOS qualification level as necessary and sufficient measures of personnel readiness and the only factors that are sensitive to a change in unit location. This simplification requires the acceptance of a number of assumptions.

First, it must be assumed that all reservists make an equal contribution to readiness. Simply considering fill level does not differentiate between the contributions made by different individuals. Although location A may result in participation of superior individuals who clearly contribute to a higher level of readiness, if the predicted number of participating individuals at location B is the same, the two sites are considered equal in this criterion. This simplification was chosen to avoid the complexity of modeling individual talents and motivation.

The assumptions for MOS qualification level are similar. Individuals that hold an MOS are considered equal, and MOS's are treated the same except in units where membership in the top three MOS's account for more than half of the personnel assigned. For those units, only people holding the top three MOS's are considered when evaluating the available market. This reduces the problem of overemphasizing the value of a large number of available recruits who can only fill a small number of billets at the relocated unit.

MOS qualification level reflects the ability to fill the required number of billets in each required skill area. MOS qualification level cannot be directly measured, so various attributes of the area surrounding the proposed site were chosen as proxy measures. Although no effort is made to predict the exact number of reservists that the relocated unit will achieve in each MOS at the new location, preferred values for the proxy measures are assumed to indicate an improved probability for reaching the desired MOS levels.

One proxy measure used for this prediction is the total number of people currently holding the desired MOS's, who are assigned to closing units in the area of the proposed site. It is assumed that the fraction of these reservists who will transfer to fill the billets of the relocated unit will be consistent from site to site. Another proxy measure for MOS qualification level is the number of Individual Ready Reservists who reside in the area of the proposed site. The same type of assumption is made for the fraction of IRR members that will activate to serve at the relocated unit.

Measures of general market supportability, competition, training support, and facility quality are considered in the fill level, but not the MOS qualification level goal. In reality,

these factors influence both criteria, but the data necessary to model their influence on specific MOS's was not available.

3. The Resulting Hierarchy of Goals

Using the assumptions presented in the previous two sections, the expert panel constructed the hierarchy of goals shown in Figure (2). This represents their best assessment of all location-related factors that should be considered in the TPU relocation decision. Unfortunately, the data needed to support many of the input measures were not available. To minimize cost and avoid expanded data maintenance responsibilities, the only data utilized were those routinely stored on the USARC LAN. The expert panel decided that a sufficient number of inputs were available to justify the development of a prototype SDSS, but it is important for the decision maker to be aware which pertinent influences are not captured by the initial instantiation of the model.

Figure (3) shows the implemented version of the hierarchy of goals. Comparing Figures (2) and (3), it can be seen that one third of the thirty measures identified in the ideal model were not implemented for automatic processing in the prototype. Brief descriptions of the ten unimplemented measures and the reasons for their omission are provided in Appendix (B).

Some decision factors identified by the expert panel are not explicitly captured in either version of the hierarchy. One such influence is the complementary effects associated with other units, particularly those with similar MOS structures, operating in the area of the proposed relocation site. Although the detrimental effects of competition are captured, the potential salutary effects of nearby units are not. Area Reserve or National Guard units may




Figure 2. Complete Hierarchy of Goals




Figure 3. Hierarchy of Goals (showing only those measures with automated inputs)

benefit the relocating TPU through direct support, economies of scale, and recruiting spillover. The dearth of analytical study and the range of practical experiences on this topic made it difficult to reach a modeling consensus. It was not clear whether it should be combined with competition measures, implemented through preference sets (discussed in detail in the next chapter) or added as a new measure under the Fill Level and/or MOS Qualification Level goals. As further studies and data become available to support these relationships, an appropriate method for modeling them should become clearer.

4. Specific Assumptions Pertaining to Implemented Measures

For each of the measures that could be implemented there were specific assumptions associated with their choice and structure. Brief descriptions of each measure and its associated assumptions are provided below. Additional information on each of these measures can be found in Appendix (C).

a. Measures of Facility Quality

The measures in this category describe specific attributes of the facility (i.e., the building and the real estate). These values are extracted primarily from engineering databases and describe the age, condition, capacity, and costs associated with the major structures at a site. Each facility is uniquely identified by a facility identification code (FAC ID).

(1) Facility Age. The age of the primary structure located on the proposed site has several sources of ambiguity associated with it. Although most structures were acquired by the government when they were new, this is not always the case. In addition, these dates do not reflect major refurbishments that could be viewed as an

appropriate basis for resetting the construction date. Although there is a separate measure named Facility Condition that could be viewed as being redundant with Facility Age, the two are intended to capture different concerns. Facility Condition is a categorization (green, amber, or red) based upon a visual inspection of the structure and tends to concentrate on the short term improvements needed to make the structure serviceable to a TPU. Facility Age is intended to reflect an assumed long term structural degradation that may not be apparent through most visual inspections. Redundancy will occur when visual inspections are thorough enough to accurately reflect major structural problems. Measuring the age of a structure does not necessarily reflect the differing effects seen with age between different construction methods (e.g., wood frame, cinder block), environmental conditions, or maintenance habits.

(2) Facility Backlogged Maintenance. This measure indicates the total dollar value of backlogged maintenance. An underlying assumption of this number is that the importance of the maintenance listed in this category from the perspective of site desirability can be measured in terms of the dollar value needed to correct it.

(3) Facility Condition. The condition of the major structures located on the proposed site is rated as green, amber, or red. Use of this measure is based on the assumption that these ratings accurately reflect the current condition of the structures of interest to a relocating unit.

(4) Facility Operating Costs. This value indicates the average, normalized monthly operating cost, in dollars per square feet, for the facility. This measure

is intended to capture the undesirability of sites with high operating costs but does not reflect the advantage of assigning additional units to the site to help share the fixed costs.

(5) Facility Ownership. This measure indicates whether the proposed site is leased or owned. This measure is needed in order to model the preference for sites that are owned over those that are not. This information does not reflect any plans to purchase sites that are currently leased, or sell sites that are currently owned.

(6) Facility Weekend Usage. This measure reflects the number of weekends per month that the facility is currently being used by the TPU's assigned to the facility. Since most units require exclusive use of the facility one weekend every month, this value normally corresponds to the number of units assigned and is normally limited to four. Although some facilities may not be able to accommodate the full-time support staff for four units, and others may be able to accommodate more than one unit drilling on the same weekend, these are considered to be the exceptions and are not accounted for. An alternative, dichotomous approach could categorize facilities in two ways, those with space available (3 or less units assigned) and those without. The disadvantage of such an approach is that it does not capture other facets of facility quality, such as overcrowding which can impair unit readiness by forcing the storage of more training equipment at the ECS. Of course, there may also be efficiencies to be gained through the collocation of units. These issues should be captured in the utility function associated with this measure.

b. Measures of Fill Level

Measures under the Fill Level goal provide an indication of the ability to maintain a sufficient number of reservists. The fill level is dependent upon both the

accession rate and the loss rate. Measures relating to the recruit market and competition provide indicators of the expected accession rate for an area. Measures relating to training support are included as indicators of loss rate because training inefficiencies are a major source of dissatisfaction and a frequently cited reason for quitting the reserves. Also included under this goal are four empirical indicators, based upon the performance of units currently operating in the area of the proposed site. These measures provide indications for both the accession and the loss rates.

(1) Area Average Manning. The strength of each of the units in the area of the proposed site is determined by dividing the number of people assigned to that unit, by the number of personnel required at that unit. The average of these proportional strengths is calculated for the final measure.

This empirical measure of market quality is based on the assumption that a unit relocated to the area will experience a manning level similar to the average level for units that already exist in that area. This measure does not capture possible complementary or competitive effects between units, and does not address compatibility of the unit mission with the local workforce.

(2) Area Drill Attendance. A fractional measure of satisfactory drill attendance is calculated by dividing the total number of reservists who participated in 21 or more drill periods over the previous four complete quarters by the total number of people who were required to drill. The average rate is determined based on all units in the area of the proposed site. This dichotomous approach to attendance, makes no effort to capture the degree by which individuals exceed or fall short of the goal.

(3) Area Loss Rate. To determine a fractional loss rate for each TPU in the area of the proposed site, the number of people lost in the last year is divided by the current number of personnel assigned. The results for all of the units in the area are averaged to yield a single value.

Area Loss Rate is used as an indicator of the market quality in the area of the proposed site. This number is intended to indicate the influence of regional issues such as economy and disposition to military service. It is assumed that a relocated unit would experience lower loss rates for those areas where the current average loss rate is lower. This measure ignores the compatibility of the unit mission with the local workforce. For example, this model would conclude that the expected annual loss rate for a relocating medical unit would not vary between two areas that have the same average loss rates for non-medical units even though one proposed site is close to a hospital and the other is not.

This measure significantly simplifies the relationship between turnover and readiness by assuming that all losses have an equal effect. Typically, losing the members of some groups, due to their extensive training or importance of their skills to the unit mission, has a much more dramatic impact on readiness than the loss of people from the less critical groups. The implications for readiness vary even down to the individual level due to the differing talents, motivation, and abilities of each person. This measure does not capture these differences.

(4) Area Transfer Rate. For each TPU in the area of the proposed site, the number of people transferred to other units in the last year is divided by the current

number of personnel assigned to determine a fractional transfer rate. The results for all of the units in the area are averaged to yield a single value.

Similar to the Area Loss Rate, the Area Transfer Rate is used as an indicator of the market quality in the area of the proposed site. It is intended to measure the stability of the local market based on issues such as local economics and demographics, but since it does not distinguish between transfers inside and outside of the local area, it may be influenced by dissatisfaction with specific units. It is assumed that a relocated unit would experience lower transfer rates for those areas where the current average transfer rate is lower. Like Area Loss Rate, this measure ignores the compatibility of the unit mission with the local workforce.

(5) Closing Unit Transfers. This measure sums the total number of reservists assigned to all TPU's in the area of the proposed site that are scheduled to close. No differentiation is made based on the length of time until inactivation. Inactivations can be posted in the G17 database up to 18 months in advance. This measure assumes that the same fraction of the displaced reservists will transfer to the relocated unit regardless of the area or the compatibility between the moving and closing units.

(6) Competition. Measures of the available personnel market and the performance of other local units must be balanced by the competitive effects associated with introducing a new unit to the area. For this measure, the total number of positions that must be filled by all other USAR and ARNG (Army National Guard) units operating within the area of the proposed site is determined. For ARNG units, the number of competing positions

is based on the number of "required personnel" from each UIC. For TPU's, the position count is based on the "authorized strength" of each unit.

Experience of the expert panel indicated that competition was most intense from USAR and ARNG units. The decision to not include the competitive effects of other services was based both on experience and practicality, for the information needed on the other services was not readily available. Although an earlier study (Solnick and Thomas, 1990) found evidence to confirm the competitive effect of other USAR and ARNG units, it also found that these effects were most pronounced among units of the same mission type (e.g., trucking, artillery, medical) and that between dissimilar units there were actually complementary effects, possibly due to spillover from recruiting efforts. The prototype form of this model does not differentiate competitors based on their mission type.

(7) Distance to Area Maintenance Support Activity. The straight-line distance from the proposed site to the closest Area Maintenance Support Activity (AMSA) is calculated by MapInfo™ using a geocoded table of all AMSA sites. An AMSA is tasked with providing the maintenance expertise for all units in its region. Distance is used as a proxy measure for response time and support quality. A study conducted by the Naval Reserve (Boykin, Merritt, and Smith, 1980) supports the impression held by the expert panel that wasted drill time is a significant factor for those reservists who choose not to reenlist (in the cited study it was the most significant factor). The AMSA, through the maintenance and repair of training equipment, plays a key role in equipment availability which is directly related to the amount of "wasted" drill time. It is assumed that an AMSA that is farther away will require more time, on average, to complete repairs that are critical to the conduct

of training. When one piece of training equipment becomes unavailable one would expect that other training methods would be used but the assumption here is the best use of training time is always pursued (and sometimes frustrated by equipment condition) and so other options are less desired.

(8) Distance to Equipment Concentration Site. The straight-line distance from the proposed site to the closest Equipment Concentration Site (ECS) is calculated by MapInfo™ using a geocoded table of all ECS's. It is assumed that the all ECS's are equally desirable, the closest one will be able to accommodate all of the equipment used by the moving unit, and that sites other than the closest one offer no benefit. The rationale for including this distance in a readiness model is similar to that used for the Distance to AMSA measure in that it is based on the efficiency of training, and the negative effects of "wasted" training time.

(9) Distance to Recruiter. The straight line distance from the proposed relocation site to the nearest recruiting station is calculated by MapInfo™ using a geocoded table of all recruiting stations (RZA), and is used as an indicator of recruiter support. This approach does not account for the actual travel distance between the two, the quality of the recruiting effort in the area, or the contributions made by other than the closest recruiting station. It assumes that the Army Recruiting Command will assign an appropriate number of recruiters to existing stations and will achieve the same success rate in all markets. Although this approach entails a significant number of sweeping assumptions, it attempts to isolate an intuitive effect of TPU location without introducing the complexities associated with recruiting issues. Modeling interactions with the Army Recruiting Command

and the effectiveness of various recruiting strategies were considered outside the scope of the initial model.

(10) IRR Available. The number of IRR members living in the area of the proposed site is determined in this measure. There is no attempt to ascertain the applicability of their MOS to the moving unit or their time since serving in the active reserve. Because of their reserve experience they are considered to be good candidates for retraining should their MOS not be needed by the moving unit.

(11) Reassignments. The Reassignments measure reflects the total number of people currently assigned to the moving unit whose homes would be within 50 miles of each proposed relocation site. The exact location of each reservist's home is approximated by the centroid of the zip code in which they live. This approach assumes that all reservists currently assigned to the moving unit will continue their affiliation after the move provided that their new commute is 50 miles or less. It also assumes that reservists currently assigned to the unit will not travel more than 50 miles from their homes to serve at the relocation site.

(12) Recruit Market. The recruit market measure estimates the total number of people in the upper fifty percent of scores on the Armed Forces Qualification Test (AFQT) (a mental test given to enlisted Armed Forces recruits) who reside in the area of the proposed site. This number is intended to be an indicator of the ability of a relocated unit to recruit appropriately qualified members. There is no attempt to predict the actual number of individuals who can be recruited. The only qualifier placed on the market is a division into two groups: those who score above the median AFQT and those who do not. Markets

with the same total number of people predicted to be in the higher mental categories are considered equivalent.

This approach could be made more accurate by reducing the total number of available recruits by the number of people who are members of the IRR or assigned to closing units to avoid “double counting” of these individuals in different measures. Based on the inaccuracies that are inherent in the data sources and the size differences expected between these categories, this concern was considered insignificant.

Estimates of the number of individuals qualifying for each mental category by gender, race, and geographic unit, were obtained from work done by Kocher and Thomas, (1994). These equations correlated market screening factors such as the percent of the market in poverty, average parents’ education, and geographic region to performance on a carefully selected random testing program conducted across the country. Census results from 1990 were used to update the population data for each zip code and adjust the estimate of mental category membership based on the market screening factors. Additional details on this database can be found in Appendix (D).

c. Measures of MOS Qualification Level

Some measures are considered to be good indicators of the support that a location provides for maintaining a sufficient number of people with the needed Military Occupational Specialties (MOS’s). Maintaining a desired levels of manning involves both accessions and losses. These measures are based on the market for accessions with the needed MOS’s, available from the IRR and closing units. Although a number of location-

related factors were identified that influence retention, they are included under the Fill Level goal because they do not discriminate between specific MOS's.

(1) Available MOS - Closing Units. For all TPU's in the area of the proposed site that are scheduled to close, the total number of reservists who hold an MOS required by the moving unit is determined. No differentiation is made based on the length of time until inactivation. Inactivating units can be posted up to 18 months in advance. Underlying the use of this measure is the assumption that the fraction of reservists with the desired MOS's who are assigned to closing units and will transfer to a unit relocated into the area is constant from area to area.

(2) Available MOS - IRR. This measure counts the total number of IRR members who live in the area of the proposed site and hold an MOS that is needed by the moving unit. Like the Available MOS-Closing Unit measure, by using the total number of people in this category, the implicit assumption is made that the fraction of these people who will actually serve at the relocated unit is constant for all areas. This approach assumes that all MOS's are equally important regardless of how difficult they are to fill. Also not considered in this measure is the time since the IRR member last served in the active reserves.

E. CHAPTER SUMMARY

The TPU relocation decision, with its multiple, sometimes conflicting criteria and disparate inputs, was structured using Multi-Attribute Utility Theory. The hierarchy of goals and relationships between variables were defined by a panel of experts. The resultant form provides a multilinear equation for which the inputs are objective data drawn from existing

databases and the output is an ordinal ranking of relocation alternatives from the perspective of readiness.

Producing a tractable decision model involved numerous assumptions and simplifications. Although some assumptions were necessitated by a lack of pertinent data, most could be easily justified based upon the experiences of USARC experts. Some were supported by the results of previous studies.

The next chapter presents information on the data sources and data processing necessary to support an automated decision model. Also provided is a detailed description of the interface that is used by the decision maker to specify the problem and interpret the model outputs. Finally, the concept of preference sets, and the flexibility that they afford, is discussed in detail.

IV. MODEL SUPPORT: USER INTERFACE AND DATA PROCESSING

A. USER INTERFACE USED FOR SPECIFYING DECISION PARAMETERS

User interface requirements significantly influenced the overall design and implementation of ARIES. Each of the commercial software packages incorporated into this SDSS provides its own Graphical User Interface (GUI), but USARC desired a single interface that could conveniently accommodate the needs of both the novice and experienced users. This inspired the design of a simplified interface that consolidates many of the most important features offered by the software components. The interface presents the relocation decision as if it were fully-structured, requiring only basic objective inputs, but it does not inhibit an experienced user from performing the complex, semi-structured decision analysis that the constituent software packages are capable of supporting. Standardized outputs are available based upon the stored, professional judgement of experts.

The user interface is simplified in a number of ways (see Figure (4)). One way is through the use of a single input screen that requires only the minimum number of inputs necessary to define the relocation problem. The user is asked only to provide the unit identification code (UIC) of the unit being considered for relocation and the facility identification codes (FAC ID's) of the potential relocation sites, a maximum of five alphanumeric strings. In situations where there are no significant problems with source data, those minimal inputs are sufficient to complete an evaluation session and generate a package of standard reports. Another simplification is the ARIES toolbar, which provides a consolidated control mechanism for all the tasks involved in a basic evaluation session (e.g.,

ARIES v3.0

File LDW View Reports

Geoscript Geoquery Measures

PROPOSED FACILITY INFO

FAC ID: CA031 UIC: WS3MAA

UNITNAME:

WS3MAA	0419 MI DET	STRAT INT
WSGAH0	0024 MI BN	DET H INTE
WSGAK0	0024 MI BN	DET K

CITY: MOUNTAIN VIEW STATE: CA ZIP: 94043

ACCEPT

FACILITY SELECTION

Moving Unit	Facility One	Facility Two	Facility Three	Facility Four
UIC: WVE080	FacID: CA120	FacID: CA015	FacID: CA114	FacID:
FacID: CA039	UIC: W7VSAA	UIC: W50QAA	UIC: NONE	
0319 SC BN Co 8 SIGTEL COM	2373 SIGNAL DET US NSC	0729 TC CO MDW TRK CNTR	AMSA #13 (M)	
DUBLIN CA 94568	MOFFETT FED AF CA 94035	FRESNO CA 93706	STOCKTON CA 95203	

MEASURES COMPUTATION

USARC GeoREF

- CA031
 - WS3MAA
 - WSGAH0
 - WSGAK0
- CA120
 - W7VSAA
 - WQVQAA
 - WR08D2

ARIES Controls

Arise Clear Exit

Enter Facility Identifier for Proposed Facility 4

Zoom: 272 mi Custom Tool: FacID Info

Start ARIES v3.0 3:26 PM

Figure 4. Input Screen for Specifying Decision Parameters

process the inputs, display the evaluation outputs, generate the standard reports, shift to MapInfo™, return to ARIES, return to the input screen). All of the functions and programs launched from the ARIES toolbar can also be started using the Windows File Manager or the GUI's associated with MapInfo™ or LDW, but a consolidated toolbar helps the novice

user complete a standard evaluation with minimal training, while adding convenience for the experienced user.

The reason that UIC and FAC ID were chosen as the user inputs is because they are the most precise methods available for identifying specific units and facilities. However, the decision maker may not know these codes. It is more common for units to be referred to by their location (e.g., the unit in Oil City), or their type and location (e.g., the artillery unit just south of Pittsburgh). By including an interactive map as part of the input screen, the user can either type in the alphanumeric codes or graphically select an icon associated with the moving unit or relocation site. When an icon is selected, the data associated with that location (i.e., unit name, unit type, facility identification code, name of the closest town or city, state, and zip code) are displayed as a means of confirming the selection. Choosing an icon where more than one facility or unit are colocated yields a selectable list of all the facilities and units at that spot. To minimize confusion, a tree diagram is included which shows the name of each location, all the facilities at that location, and all units assigned to each facility.

B. SOURCE DATA

ARIES draws its inputs from a wide variety of large databases (see Table (500)). These sources are flat files independently designed to support a variety of separate, "stovepiped" systems. Consequently, there is little consistency between the data sources. For example, depending upon the database, a unit's identification code may be found under various synonyms, including fields named UIC, OWN_UIC, CURR_UIC, or UIC1. Not only do field names vary, but the formats of the data in those fields are also inconsistent. For

Source Table Name	Type of Information	Approximate Number of Records	Size (Mb)
SIDPERS-G19TRUE	Required positions	233,200	25
SIDPERS-G18CWE	Reservists	204,300	198
IRR (Individual Ready Reserve)	Prior service market	400,000	400
QMA (Qualified Military Available)	Non-prior service market	34,300	2.7
RPINFODT	Unit backlogged maintenance	8,000	25
FINANCE	Individual finance and drill records	17,300	3
COMMAND PLAN	Unit structural history	12,500	3.3
FYxxLOSS	Attrition records	8,900	151
GEOREF	Facility locations	1553	.2
INTEREST	Facility age	4,000	4.4
FPS	Facility condition and operating costs	1,600	5.8
COMPLEX	Facility use and ownership	1,600	1.3
RZA	Recruiter locations	1,800	.2
AMSA	Maintenance activities	190	.1
ECS	Equipment storage sites	30	.1
NGNON_CL	Competing positions in Army National Guard units	3,700	.5

Table 1: Source Data Tables

example, some tables contain nine digit zip codes while others use only five digits. The lack of common data standards complicated the development of an integrated data framework. Appendix (E) provides a more detailed summary of the source data tables and the data fields used by ARIES.

In addition to consistency problems, most sources also contain fields with incorrect and missing data. The expected format for a facility identification code is a two letter state abbreviation, followed by three numerals (e.g., PA 035, CA132). In addition to data in the proper format, facility identification codes sometimes appear as two letter state abbreviations with no numerals, "TBD", "NA", "N/A", or blanks. Even worse from the perspective of error detection, some fields have default values which cannot be easily distinguished from actual data. The error trapping philosophy adopted to handle these problems is to provide a data flag when any missing or obviously incorrect data is encountered.

1. Data Preparation

The ARIES prototype is currently implemented on a notebook computer, but eventually will be modified by USARC to operate on the local area network (LAN) at USARC Headquarters in Atlanta, Georgia. With this plan in mind, the only data sources used are those that either currently reside on, or can be conveniently migrated to, the USARC LAN. These databases are used for other purposes within USARC, so ARIES data requirements do not incur any new data administration responsibilities. The SDSS prototype operates on a static "snapshot" of the source databases and so does not have to contend with the issue of changing source data. On the operational LAN version however, the databases that supply inputs to ARIES may be updated as frequently as weekly. In this

environment, changes to the databases must be transparent to the decision maker. ARIES includes a module which extracts the needed data from the source databases and conditions it for use (e.g., filters out unneeded records, converts data to a common format). This data preprocessor is discussed in detail in the next chapter.

2. Geocoding

In addition to assuring that current data is located where ARIES can find it, some files must be geocoded before they can be used. Geocoding is a process that associates data with two dimensional spatial coordinates. Although a wide range of spatial systems can be accommodated (e.g., the floor plan of a building, the surface area of an image), the most commonly used convention, and the one used in ARIES, is a representation of the earth's surface. Geocoding is performed by a mapping engine which links geographical positions (defined by latitude and longitude) with attributes of objects located at those positions. The process appends latitude and longitude fields to the existing database structure.

Data is geocoded so that it can be used in spatial queries and displays. Examples of spatial queries are, "find the number of facilities within 50 miles of a given relocation site" or "return the distance to the nearest Equipment Concentration Site". Unlike non-spatial queries, which rely on data field values to relate objects, spatial queries relate objects based on distances or geographic regions. MapInfo™, augmented with MapBasic (a complementary programming language), can employ either technique. Once records are identified or categorized by their location, standard database operations can be performed on their associated fields. For example, ARIES not only can count the number of Army National Guard units within 50 miles of the proposed site, but also can sum the number of

authorized positions (using entries in the AUTH field of the geocoded NGNON_CL table) at each of those geographically selected units.

MapInfo™ offers two methods for geocoding individual database records, manual and automatic. Using the manual method, a user selects each record and associates it with a position by either choosing a point on a map or typing in coordinates. Reserve facilities, Area Maintenance Support Activities (AMSA), and Equipment Concentration Sites (ECS) were geocoded in this manner, providing a very accurate representation of their locations. Once a data table is manually geocoded, it can be used to automatically geocode other tables that share a location defining data field. In ARIES, the zip code field was most commonly chosen as the shared data field. Using a geocoded table of all zip codes that is provided with MapInfo™, other tables containing a zip code field (e.g., IRR, G18CWE) were automatically geocoded by matching zip code values and transferring the associated coordinates.

Using a field like zip code, which usually represents an area rather than a single position, as the basis for the automatic geocoding process introduces positional inaccuracies. When records are geocoded based on zip code, all data in the underlying database, regardless of which specific position within the zip code region they are actually associated with, are matched exclusively with the centroid point of the zip code region. The distance to every record in that zip code is calculated as the distance to the centroid of the geographical region. A radius query (defining a circular region of interest) that intersects a portion of a zip code region will not return any values associated with that zip code if the radius does not encompass the centroid of the zip code region. Similarly, a geographical query with a radius that circumscribes a centroid returns the data associated with the entire region despite the

possibility that much of the region may lie outside of the intersected area. Improved accuracy is possible by refining the granularity of the geocoding process (i.e., capturing data based upon very small geographical areas or manually geocoding the exact location), but in most cases, either the data needed to support such refinements was unavailable, or the effort of applying that data to so many records would have been excessive.

One alternative to using a centroid point to represent an entire region, is to divide the underlying data by the area to define a uniform density for each region, and then multiply that density by the area of intersection between the query region and the zip code region. This approach could provide some average improvement in accuracy, but that small improvement did not warrant the added complexity involved with implementing such a process. Examples of data that are geocoded based upon zip code are the estimated numbers of qualified recruits and the home addresses of both Individual Ready Reserve (IRR) members and the reservists currently assigned to the moving unit.

In most situations, the average distance to all points within a region, or even better, an appropriately weighted average, would provide a more appropriate calculation for distance. The use of centroids to represent areas introduces errors, which are more pronounced for shorter distances and larger areas. Although a variety of methods have been identified for calculating the distance between a point and an irregularly shaped area (Miller et. al., 1997), implementation of such calculations would require the use of very sophisticated software packages and, in this application, significantly increase the computational overhead.

The files that must be geocoded prior to running an ARIES evaluation session are AMSA, ECS, IRR, NGNON_CLOS, RZA, and GEOREF. Most of these files are relatively static and so the geocoding process is only required on an infrequent basis. The exception is the personnel file (G18CWE) which is updated weekly. Once all of the source data is prepared (i.e., stored in a location that is accessible by ARIES and geocoded if necessary), then an ARIES evaluation session can be initiated, starting with the data preprocessing phase.

C. PREPROCESSING PHASE

Even if all source databases were consistent and accurate, their number and sizes present considerable performance challenges for a PC-based, front-end processor. The first version of ARIES was a "proof of concept" that relied upon a monolithic, highly sequential approach. As the design was refined, a single evaluation session (which evaluates the current site and up to four relocation sites) was decomposed into three phases: preprocessing, processing, and evaluation (see Figure (5)). Functions were distributed between these phases so as to simplify the interface for the SDSS user and to improve processing efficiency (thus reducing the time required for a complete evaluation session).

In order to reduce the run time associated with an evaluation session, some of the functions associated with individual queries were grouped together in a preliminary phase of data preprocessing. In most cases, this functional redistribution eliminated redundancies during the processing phase.




Figure 5. Evaluation Session Phases

The first function performed during preprocessing is extraction of the needed data from source databases. Using Visual Basic™ to control JET SQL queries, temporary extraction tables are created containing only the fields that are used by ARIES. These extractions reduce the amount of data handled during the processing phase which helps to significantly reduce the run time for an overall evaluation. This approach conforms to a common programming axiom known as the "Principle of Least Privilege" which recommends granting only that access necessary to perform the task and no more. After extraction, two types of preprocessing are performed: filtering based upon data values and aggregation.

1. Filtering Based Upon Data Values

In some cases, data values contained in the extracted fields are used to filter out undesired records, further reducing the size of the tables. An example of such filtering is the screening of records from the FINANCE table, which is used to calculate the fraction of people assigned to area units with satisfactory drill performance for the previous year. The FINANCE database contains pay and drill attendance data on *all* Army reservists, including those who would not be expected to participate in drills for various reasons. After extracting the necessary fields, but before checking drill attendance numbers, records were eliminated from consideration for all *inactive* reservists and for non-prior service recruits who have only recently joined the unit (these people are typically unavailable for drills during their initial training period of nine months). Rather than continually applying these filters to the FINANCE table as each proposed relocation site is analyzed, they are applied only once,

during the preprocessing phase. This reduces the size of the table that is repeatedly manipulated during the processing phase.

2. Aggregating Data

In some cases, the level of source data granularity is finer than required by ARIES. In such situations, data aggregation in the preprocessing phase can significantly reduce the time required for calculations during the processing phase. One of the most dramatic examples of aggregation is performed on a file that contains a separate record for every Army reservist in the country (G18CWE). There is no reliable data field available in any of the source databases that indicates the total number of people assigned to each unit, and yet that number is needed for many of the calculations performed by ARIES (e.g., Area Loss Rate, Area Transfer Rate, Average Area Manning, Closing Unit Transfers, and Reassignments). The most accurate method available to obtain the number of assigned people is to count all personnel records in the G18CWE table associated with each unit. Rather than repeatedly counting these records, the counting operation is performed only once and the results are stored in a temporary table that contains only two fields, UIC and the number of people assigned. In addition to reducing the number of times that the large G18CWE table is queried, this approach also reduces the complexity of the queries used during the execution phase, which yields significant gains in efficiency.

A related thesis will analyze the resultant efficiency gains associated with the design decisions (e.g., coding structure, processing paths) in more detail. The time required for preprocessing (which supports the evaluation of up to four relocation sites) is typically 10 minutes or less on a 90 MHz Pentium processor. Evaluation of a single relocation site in a

metropolitan area that initially required two to four hours, can now be completed in 5 to 30 minutes as a result of preprocessing.

In addition to increasing the execution speed, preprocessing also converts source data tables from a variety of formats into Microsoft Access tables. This conversion permits consistent data handling in a format that is compatible with the programming language used for overall control, Visual Basic™. Once the data is screened, aggregated, and converted to a consistent format, it is stored in a standard location, and is ready for use in the processing phase, which further manipulates the extracted data to produce the inputs (i.e., a measures table) needed for the decision model.

D. PROCESSING PHASE

The processing phase begins with the previously discussed user inputs (UIC and FAC ID's). During the processing phase, the extracted data is further processed, using Structured Query Language (SQL) queries based upon the user's inputs. The output of this phase is a measures table, which contains the input values for the decision model measures. The independence of many of the queries and calculations makes this overall process a good candidate for the use of parallel processing or an object oriented approach, but the prototype is implemented using a predominantly sequential architecture based on practical programming concerns and experience.

1. Filtering Based Upon Distance

One of the first steps performed in the processing phase is additional data filtering. Unlike the filtering accomplished in the preprocessing phase, which was based upon data values, this filtering is based upon distances. In some situations, records associated with

distant locations can be immediately eliminated from consideration. The G18CWE file, which contains a record for every Army reservist (over 200,000 records nationwide), once again provides a good example. Using MapInfo™, all records for people residing more than 50 miles from a proposed site are eliminated. Following this screening, calculation of the value for the Reassignments measure (a count of the people assigned to the moving unit who live within 50 miles of the proposed relocation site) involves only a simple query that counts records associated with the moving UIC. Performing this calculation without the geographic screening would require a complex query comparing all 200,000 records against a list of area zip codes (often numbering in the hundreds), and then checking for matches with the moving UIC. The time associated with this particular query was reduced by a factor of 60 using geographic filtering.

2. Queries

Two forms of queries are implemented in this phase, spatial and non-spatial. The spatial queries are executed using MapBasic (a programming language associated with MapInfo™) and provide functions such as converting positions to distances, making proximity determinations (e.g., identifying the closest support or training sites), and classifying locations by regions (e.g., constructing a list of all competing units within the area of the proposed site). Non-spatial queries are implemented using JET SQL and provide the remainder of the model input data. Many of these queries involve multiple levels of complexity (i.e., using nested query statements). Appendix (F) provides a listing of all the query statements used in ARIES.

The use of two different query tools (i.e., MapBasic and JET SQL), and efforts to streamline coding through the elimination of redundancies, led to the creation of interim data tables. For example, a *spatial* query creates interim tables containing all zip codes within 50 miles of each proposed site. These tables are used in conjunction with a variety of *non-spatial* queries to populate the measures table. Other examples of interim tables passed from spatial queries to non-spatial queries are lists of all units and lists of all facilities in the area of a proposed site. Interim tables are also created for passing data from the preprocessing to the processing phase. An example of this is a temporary table containing the number of people assigned to each unit (previously mentioned in the aggregation discussion) that is created in the preprocessing phase and then used by five different queries in the processing phase. Interim tables are transparent to the user and are deleted after ARIES is closed.

3. Archiving

Archiving can further reduce the time needed for the preprocessing and processing phases by taking advantage of calculations performed during previous evaluation sessions. Two forms of archiving are implemented in ARIES: storage of complete evaluation sessions, which helps to avoid the processing phase completely, and storage of data on individual facilities.

The same measures table which is transferred to LDW is also archived in a "facility repository". The only difference is that the archived version includes an additional field for each alternative that contains the UIC of the moving unit. Eighteen of the twenty values for the measures table can be calculated based solely on the identity of the proposed relocation site. If any of those sites are ever proposed again, in a subsequent evaluation session with

a *different moving unit*, ARIES automatically extracts those 18 values from the facility repository and calculates the other two. If a new evaluation session is conducted for the *same moving unit*, the row in the measures table for any proposed facility that is repeated from the earlier session will be immediately filled with all 20 values, avoiding all calculations.

The other form of archiving saves an entire evaluation session. This is performed manually in the evaluation phase. By saving an entire session, the session can be retrieved and the decision maker can resume evaluation of that session without any preprocessing or processing of data. One hazard of this form of archiving is the possibility of using outdated data. Since archived information is perishable, when a new extraction is created with the preprocessor, all data stored in the facility repository is deleted. This is not true for entire sessions saved by the user. There is no automated control of these files and so the user is responsible for managing file names, storage locations, and file deletions.

When the various queries, archiving, and other manipulations associated with the processing phase are completed, the final data table is imported into LDW for use in the evaluation phase. The transition through the first two phases occurs automatically. After defining the problem (by identifying the moving unit and the relocation sites) the decision maker selects the "ARIES" push-button. A variety of status bars and progress lists are displayed as preprocessing and processing occur, but no further interaction with the decision maker is needed until the first screen of the evaluation phase.

E. EVALUATION PHASE

The evaluation phase is where processed input data (i.e., a measures table) is converted into evaluative information. This phase relies almost exclusively on a commercial decision software package (LDW). The ARIES toolbar offers icons that can be used to view and print out various LDW displays, bypassing the menu driven controls used by LDW. This toolbar simplifies user access to outputs that are expected to be used frequently in this context. Otherwise, LDW can be used just as it would as a standalone product. During the evaluation phase, the decision maker is provided access to a variety of comparative displays, preference sets, methods for modifying the model, and means of sensitivity analysis.

1. Data Confirmation

Following the automated preprocessing and processing of data, the user is placed in the LDW environment in the "matrix view", which displays an array of the decision model input data using a row for each alternative location, and a column for each measure (See Figure (6)). This screen provides the user with the opportunity to verify the model inputs before they are used to produce a ranking of alternatives. If all data sources are accurate, there should be no user action required at this point. Cells associated with source data that are missing or are obviously in error are filled with an error flag (-999). The user must replace all error flags with an appropriate value in order to produce valid results. Any flags that are not corrected will skew the results because LDW will interpret a "-999" as a valid input.

	Measure 1	Measure 2	Measure 3	Measure n
Alternative 1				
Alternative 2				
Alternative 3	<i>Decision model input values appear in the individual cells</i>			
Alternative 4				
Current Location				

Figure 6. Input Matrix

The idea of forcing the user to supply the values that are used to replace error flags was initially resisted because it seemed to be contrary to the objective of a simplified interface that can be effectively used by a novice. Consideration was given to automatically inserting values representing neutral utility (0.5 utility units) and providing the user with a list of measures where this was done. Discussions with the expert panel revealed the inappropriateness of that approach based upon the qualifications of the intended users. Although it should not be assumed that those who will use ARIES have extensive computer experience (hence the simplified interface), it is appropriate to assume that they are experts in the operation and evaluation of TPU's. Consequently, they should know where to find the missing data or be well qualified to provide satisfactory estimates. When source data is missing or obviously in error, the assumption is that forcing the SDSS user to locate or estimate that data will provide more accurate inputs than applying a default value. Further, preliminary experience has shown the "-999" convention to be very useful in identifying problems with "dirty" data in the source databases.

2. Outputs

One of the strengths of LDW is the variety of methods that it offers for the analysis and display of decision information. The expert panel felt that such diversity could be overwhelming to the novice user and so a limited subset of displays was chosen and made available through the ARIES toolbar. When accessed through the toolbar, these displays are also automatically printed and collectively form a standard reports package. This section focuses on the outputs provided in that standard package, explaining the significance of each display and the reason why it was considered important enough to be included. Information on the other graphical and tabular displays is available in the LDW documentation.

a. Goals Hierarchy

At the heart of the SDSS is a MAUT model that can be represented as a hierarchy of goals and measures. Figure (7) shows the "Goals Hierarchy View" produced by LDW for the prototype decision model. Measures (represented by ellipses) are the quantitative variables that describe the alternatives and goals are containers that hold measures and subgoals (represented by rectangles). Although the structure of the hierarchy should not change from one analysis to the next, this display is included as part of the standard reports package because it is a concise summary and an effective reminder of the model on which the decision process is based.

b. Site Desirability Rating

The primary output of ARIES is the overall Site Desirability rating which is calculated for each alternative and provides an ordinal ranking of the proposed relocation sites. The current location of the TPU is automatically considered and ranked along with the




Figure 7. Hierarchy of Goals

specified alternatives. The final scores do not reflect the degree to which one option is preferred to another, just the order of their desirability from the perspective of this model.

Figure (8) shows the Stacked Bar Ranking that is produced by LDW and is included as part of the standard reports package. Each alternative (the current site or one of four relocation sites) is represented by a bar whose length is proportional to the overall Site Desirability rating of the alternative. The contributions to the overall score made by individual measures are indicated by color coded segments of the bargraph. The length of a segment reflects both the importance of the associated measure (relative weight) and the desirability of the value that is provided as input to that measure.

Although comparing the lengths of the same segment from alternative to alternative is an accurate indication of the degree of preference, the same is not true when comparing the total lengths of all segments. If the bar for alternative A has a Recruit Market segment that is twice as long as the Recruit Market segment in the bar for alternative B, it is accurate to say for that specific measure, that A is preferred twice as much as B; it has




Figure 8. Stacked Bar Ranking Display

twice the utility to the decision maker. If, on the other hand, the *overall length* of the bar for A is twice as long as the *overall length* of the bar for alternative B, it cannot be concluded, in terms of overall desirability, that A is preferred twice as much as B. This distinction arises because the evaluation scales for each measure are determined independently and they are not directly compared with each other.

Although it is not possible to perform cardinal comparisons between overall utility scores, *changes* in those scores offer another valid method of comparison. Consider an example where alternative A has a utility of 0.4, alternative B has a utility of 0.5, and alternative C has a utility of 0.7. Not only is it proper to conclude that alternative C is preferred to alternatives A and B (i.e., ordinal ranking), but it is also valid to describe the increase in desirability from B to C as greater than the increase from A to B. Although not used in the prototype model, LDW offers an approach to preference elicitation that results in cardinal rankings. This method, Analytical Hierarchy Process (AHP), was not chosen for this model because it becomes unwieldy with a large number of measures for it requires entry of pairwise weight ratios for all possible measure pairs.

c. Ranking Results Matrix

This display provides a matrix of the utility results for all the alternatives under each measure and goal (see Figure (9)). The goals and measures are shown at the top, with their weights in the row below them. The alternatives are shown on the left edge and each cell represents the utility for the row alternative under the column goal or measure. This display is included among the output options because it provides the specific values upon which all graphical displays, like the Stacked Bar Ranking, are based.

	Goal 1	Goal 2	Goal n	Measure 1	Measure 2	Measure n
	.2	.3	.5	.4	.1	.5
Alternative 1						
Alternative 2						
Alternative 3	<i>Utility ratings appear in the individual cells</i>					
Alternative 4						
Current Location						

Figure 9. Ranking Results Matrix

d. Sensitivity Analysis

Sensitivity analysis is conducted in order to gain improved understanding of the model and the world that it purports to describe. It aids the decision maker when there is uncertainty over the accuracy or relative importance of information. It can show the impact of changes in either input information or model structure on the final results.

At the most basic level, sensitivity analysis shows the impact on the output variable of changes made to input variables. In ARIES, this can be accomplished by manually changing values in the input matrix and observing the effect on the output. The VB shell automatically populates the input matrix for the decision model, but these values can be changed through direct input in the LDW environment.

ARIES provides two types of sensitivity analyses when evaluating the weighting scheme, *automatic* and *dynamic*. *Automatic* sensitivity analysis, as implemented in the Sensitivity Graph display, shows the effect on any goal (normally shown for the

overall Site Desirability rating) that results from varying the weight assigned to any specific goal or measure over the range of zero to 100 percent (see example in Figure (10)).




Figure 10. Sensitivity Graph Display

Relative utility is shown on the vertical axis and the percent of the total weight assigned to the chosen goal or measure is represented on the horizontal axis. The graphed lines represent overall utilities for each of the alternatives at different weights. The highest line represents the most preferred alternative for a given weight. The slopes of the lines help the decision maker evaluate the sensitivity of each alternative to weighting changes. An intersection of the lines representing alternatives indicates a crossover point where a change in weighting results in a new preference in alternatives. A vertical line indicates the weight currently assigned to the chosen goal or measure (in this example, .12

weight for Measure 3). This display is not included in the standard reports package for practical reasons; the decision model implemented in the prototype contains over 30 goals and measures and so would require that many graphs to capture the sensitivity of all the associated weights. Although not represented on the ARIES toolbar, sensitivity graphs are one of the most powerful tools available for challenging the subjective weight assignments. They can be easily accessed through the LDW menubar (under "Results").

Dynamicsensitivity allows the user to quickly perform "what if" analysis on any weight. The display window is divided into two panes; the upper pane shows the current overall utilities for all alternatives and the lower pane shows the weights for the goals and measures (See Figure (11)). As the user temporarily changes the value of a weight, the lengths of the bars representing the utilities of the alternatives, vary in response. The weights of all other goals and measures vary proportionally to accommodate the specified weight change, ensuring that all global weights continue to sum to 100. Although the real value of this display is its ability to facilitate interactive experimentation with weights, when accessed through the ARIES toolbar, it is printed out as one of the standard reports to provide a convenient, graphical summary of weights.

e. Comments

This report provides a consolidated listing of the comments stored for each goal and measure. Like the Goals Hierarchy View, this report does not normally change from analysis to analysis but it is included in the standard reports package as a convenient reference. In the prototype, the comment space is used to document the significance and the method of calculation for each measure and goal.




Figure 11. Dynamic Sensitivity Display

f. Preference Set Summary

This matrix display provides a summary of the overall ranking results for all alternatives under all preference sets. Preference sets contain the utility functions and relative weights needed to rank the alternatives on the measures and goals. Different preference sets are used to apply different decision perspectives, or views. The Preference Set Summary display (Figure (12)) provides a convenient means of comparing the effect of different perspectives on the decision outcome.

	Baseline Preference Set	New Preference Set #1	New Preference Set #2
Alternative 1			
Alternative 2			
Alternative 3	<i>Overall utility ratings appear in the individual cells</i>		
Alternative 4			
Current Location			

Figure 12. Preference Set Summary Display

3. Preference Sets: A Mechanism for Accommodating Different Decision Perspectives

Multi-Attribute Utility Theory (MAUT) provides a framework for expressing a decision maker's preferences. Based on the model structure, weights, and interactions, a multi-measure utility function is designed to capture the approach and priorities of the human decision maker. This raises the critical question of whose perspective should be modeled?

For the TPU relocation problem, it is impossible to identify a single decision perspective that is applicable for all situations. The importance of each locational attribute may vary from group to group. The distance between a TPU and its headquarters may be more important to people at the Army Reserve Command than it is to current reservists, and it may be of no concern to potential recruits. Depending upon the decision perspective of the group, any one of these interpretations might be considered to be most important. Individuals within the same group are also likely to differ when assigning priorities to various measures. Furthermore, preferences are often situationally dependent. For example, as funding levels change, the relative importance of financial measures may also change.

Logical Decisions for Windows™ provides the user with a convenient means of organizing and changing these subjective judgements through the use of “preference sets”. Each preference set contains one version of each utility function, including single-measure utility functions (SMUF’s), multi-measure utility functions (MUF’s), and the associated weights. Preference sets can be created to reflect the priorities associated with different groups, individuals, and situations and stored for later use. The baseline preference set was defined by a panel of experts and is intended to be used as a common approach to the problem.

Some examples of issues that may inspire different perspectives and thus are candidates for separate preference sets are:

- Fiscal constraints. Based on the inevitable uncertainties involved in predicting budgeting trends it may be necessary to create multiple preference sets that adjust the weights of cost-related measures.
- Type of training required. Some MOS’s require formal instruction at consolidated training centers. The training for other MOS’s can be conducted locally, which is considerably less taxing on a unit’s limited resources. For units with a predominance of MOS’s that require formal training, a preference set could be created that elevates the importance of the prior service (pre-trained) market. This would favor relocation sites that are expected to minimize the difficulties and costs associated with formal training.
- Recruit/retention rates. For certain types of units, conditions, or areas, it may be particularly difficult to recruit or retain a sufficient number of qualified people. These situations could be used as the basis for preference sets that are more sensitive to the impact of market supportability and competition measures.

- Competitive/complementary effects based on type of unit. Although the prototype model does not explicitly address the differences in competitive effects that result from the degree of mission similarity between the relocating unit and the units that are already established in the proposed relocation area (see Solnick and Thomas, 1990 for a discussion of this issue), it may be appropriate to incorporate this information into a preference set. At a minimum, two preference sets could be created, one where the relocated unit is considered similar to the units in the new area, and another where it is not similar. The weights associated with measures that reflect competitive effects, closing unit transfers, and empirical market indicators (Area Drill Attendance, Area Loss rate, Area Transfer Rate, and Average Area Manning), could all be changed based upon the degree of similarity between units.

In general, any situation where assumptions or simplifications have been made is a likely candidate for a preference set, reflecting a different assumption or point of view.

F. CHAPTER SUMMARY

An ARIES evaluation session is performed in three phases: preprocessing, processing, and evaluation. The preprocessing phase is performed first, to extract and condition data from the source databases. After the decision maker supplies the identities of the moving unit and proposed relocation sites, the processing phase is performed automatically, without further user interaction. At the completion of the processing phase, a populated measures table is imported into the decision model, and the screen shifts to the LDW interface (in the "matrix" view). The measures table is also archived in a facility repository that is used to reduce the computations necessary during subsequent evaluation sessions that consider the same relocation sites.

The evaluation phase involves a confirmation of model input data and the use of decision analysis software. The most obvious output is an overall ranking of alternatives, but there are many tools available to help the decision maker gain insights into both the

problem being evaluated and the method of evaluation. Sensitivity analysis can be performed on both the input variable (TPU location) and the weighting scheme. Preference sets enable the user to quickly evaluate the impact of different decision perspectives. The explicit, formal model used in the evaluation phase provides an organized means for interpreting reality and communicating concerns and priorities. This approach facilitates improved analysis, provides a baseline for debate, and supplies a tool for building consensus.

V. STRUCTURE AND USE

A. SYSTEM ARCHITECTURE

The ARIES architecture is composed of four components, a mapping engine, a decision model solver, a data preprocessor, and an integrating shell. The mapping engine and model solver are commercial software programs, and the preprocessor and shell are original code written in Visual Basic™. Figure (13) shows a simplified representation of the system architecture.




Figure 13. System Architecture

The overall system architecture is designed to provide useful, standardized results, through a simplified user interface, without impairing the ability to perform sophisticated analysis. Once the problem is specified by the decision maker, the many steps necessary to

extract, cleanse, and process the data used to populate the measures table (i.e., preprocessing and processing) are well defined, and can be accomplished without additional input from the user. The predictable, structured nature of these tasks makes automation effective. The system shell relieves a user of the burden of understanding the individual software systems and associated data transfer protocols. In the evaluation phase, ARIES offers a wide variety of analysis and display techniques. For the novice user, a subset of these options are conveniently accessed through a consolidated toolbar. Achieving an appropriate balance between simplicity and functionality was one of the most significant challenges faced during the design of the system architecture.

1. Integrating Shell

An overarching program shell was created to control the entire evaluation process, which includes coordination of independent, commercial software programs. The shell is also used to produce a consolidated user interface, in some cases by repackaging interface elements from the other software components. It is written in Visual Basic™ and employs a variety of communication protocols to pass control information. The shell performs, with the aid of the mapping engine, the functions previously described as the processing phase.

Although coding of the integrating shell was started using Delphi, it was soon shifted to Visual Basic™. Not only did the shift accommodate the expertise of the USARC group that will be maintaining the system, but it also permitted the interface to be designed in a format similar to other USARC applications.

Access and Excel were selected to provide database and spreadsheet functions. These programs are compatible with Visual Basic™ (all three are Microsoft products) and

were already owned by USARC as part of an office software suite. Because these products provided the necessary functionality, and the decision maker does not directly interact with them in the ARIES context, little effort was made to identify or evaluate alternatives.

2. Mapping Engine

MapInfo™ is a commercial mapping package that is used as a graphical input tool and provides for the spatial definition and processing of data. It converts positions to distances, makes proximity determinations, and classifies objects by geographical region. Using the OLE communications protocol, Visual Basic™ is able to pass data to MapInfo™, launch a MapBasic program that executes spatial processing, and retrieve the processed results.

The documentation for MapInfo™ describes the program as a Geographic Information System (GIS). Like the definition for a DSS, there is little consensus on the exact definition of a GIS, but many experts would not place MapInfo™ in that category. Although it provides a wide variety of methods to spatially process and display data, it provides little support for the interpretation of data. ARIES, as an entire system that enhances the features of MapInfo™, conforms to a broader definition of a GIS.

The choice of MapInfo™ as the mapping engine was relatively easy. MapInfo™ satisfied all of the known and anticipated functional requirements, was already owned by USARC, had proven to be well supported and documented, and would minimize the need for additional training. The primary alternative, ArcInfo, is a much more powerful and sophisticated spatial processing tool than MapInfo™, but it also incurs higher financial costs,

additional computational overhead, increased complexity, and a steeper learning curve for the user.

3. Decision Model Solver

Logical Decisions for Windows™ (LDW) is used as the decision model solver. There is very little interaction between the shell and LDW. The only data passed between the two are the 20 measures and a city name for each alternative. The only other interaction is message passing via the ARIES toolbar, which permits the display and printing of selected LDW outputs. Visual Basic™ uses WIN32 API routines as the primary mechanism for communicating with LDW. This provides access to most, but not all, functions that are available to a typical user employing keyboard and mouse selections. A few of the steps that could not be controlled by WIN API routines, were executed by simulating manual menu selections (i.e., character streaming).

Selecting a commercial software package to serve as the decision model solver was challenging. LDW was chosen primarily for its superior implementation of the underlying decision framework, Multi-Attribute Utility Theory (MAUT). Although other products, such as Which & Why, also apply MAUT, they do not support the construction of utility functions that automatically convert quantitative input values to common utility units¹. Another important feature, particularly for a prototype product such as ARIES, is flexibility. LDW supports a wide variety of complementary techniques to elicit user preferences (e.g., ordinal

¹ Comparison of DSS software products was based primarily on literary review, particularly the "Decision Analysis Survey" in the August 1996 edition of *ORMS Today*. In some cases, the conclusions of these references were confirmed through direct experimentation with the software packages.

criteria ranking, tradeoffs, swing weights, direct graphical and tabular inputs) and also supports a significant alternative to MAUT known as the Analytical Hierarchy Process (AHP). AHP accommodates problems where the inputs can only be measured on a subjective scale. In ARIES, subjective inputs are stored as part of the model, so that only objective inputs are required during execution. As the model evolves, there may be situations where objective inputs are either unavailable or inappropriate. ARIES supports the ability to seamlessly incorporate methods using subjective inputs (e.g., AHP) providing a flexibility that may prove to be important to the evolution of this system. LDW was also judged to be superior (for this application) to other products in terms of overall ease of use and the available options for displaying the results graphically.

Given these fundamental strengths, the selection of LDW was still challenging. The primary reason is that LDW does not contain the necessary Data Link Libraries (DLL's) to support standard communication methods such as Dynamic Data Exchange (DDE) or OLE. A windows based communication protocol (WIN32 API) is used to achieve some coordination, but the 16-bit architecture of LDW limited the available control methods. Although well suited to the construction of an appropriate decision model, because of its meager support of these protocols, the current version of LDW is a poor candidate for integration with other software packages. Other than a basic ability to import and export data tables (in an LDW specific format), it offers little coordination with outside entities.²

²As an example, when manually selecting a display, the user is often presented with a variety of display options. For those options that are selected from a menu, character streaming can be used to simulate the user's input. For those options that must be selected with a mouse (using an option button), there is no convenient method to automate the desired selection.

Other decision software programs that fully implement standard communication methods were considered, but the features and flexibility offered by LDW outweighed its problems with integration.

4. Data Preprocessor

The data preprocessor, like the shell, is written in Visual Basic™. ARIES and the preprocessor are separate programs because their purposes are fundamentally different. The preprocessor is primarily used to extract the needed fields from the source databases, convert them to a common format, perform some basic manipulations, and then store them in extract tables. These steps are *independent of the specific units or sites being evaluated*. The shell, on the other hand, *uses the decision maker's inputs* (i.e., moving unit and proposed facilities) to select specific values from the extract tables. These values are either directly transferred to, or used to calculate inputs for, the LDW measures table.

The preprocessor also provides various functions for the administration of data. If the location of a source database changes, the preprocessor provides a convenient means for updating the extraction process. It can also be used to change the queries that actually accomplish the extraction. For informational purposes, it lists all fields, table names, and table indices that must be present to support the processing performed by the ARIES shell.

There is very little coordination required between the shell and the preprocessor. They are separate, executable programs that are not meant to be run concurrently. When the preprocessor is executed, it always stores its output tables in the same file location. When an evaluation session is run, the shell merely retrieves those extract tables from the standard location. The same extract tables can be used to run multiple evaluation sessions, but the

preprocessor can be executed whenever the source databases change to ensure that the extract tables are based upon current data.

B. ORGANIZATIONAL IMPLEMENTATION

ARIES is intended to help the objective decision maker understand the problem environment, surface underlying beliefs, and develop conviction for one of the alternatives. In reality, there often will be times when decision makers embark upon this formal analysis with their minds already made up. Keeney and Raiffa suggest four legitimate uses for a DSS in such situations. First, the formal analysis can provide *psychological comfort* for the decision maker; it can corroborate intuition. Second, it can aid in the *communication* of the decision, the basis for the decision, and the underlying structure or logic. The third use, *advocacy*, takes communication one step further, convincing others of the reasonableness of a proposed location. Fourth, formal analysis facilitates *reconciliation* between conflicting arguments. The model structure exposes the key elements of the decision and can help to sort out the merits of the conflicted approaches. It is also possible, even with a seemingly obvious outcome, that formal analysis will uncover unique insights that can motivate new alternatives or a revised model structure. (Keeney and Raiffa, 1976)

The flexibility of ARIES allows the decision to be conveniently tailored to meet the objectives of the user. The analysis performed for psychological comfort, for example, might be quite different from that used as a tool for advocacy. A personal analysis may rely on sensitive information and opinions that would be improper to openly share when seeking advocacy. The user who is unsure of his opinions can either rely on the baseline preference

set, apply alternative preference sets representing other perspectives, or simply experiment with weights and utilities as a method of learning and exploration.

To achieve successful implementation of this system, USARC should develop an organizational change strategy that capitalizes on the current discontent with manual methods, establishes credibility and builds consensus for the use of the automated system, and integrates this SDSS into the organizational culture. This is a complicated issue that is not addressed in depth in this study, but some basic recommendations are provided below.

One implementation concern is security, primarily from the perspective of data integrity. In the standalone prototype, ARIES provides all users with the same access rights. As the prototype is migrated to a multi-user environment on the USARC LAN, access control can be implemented using policies, training, and/or external programming methods (e.g., controlling the access rights on specific files through workstation or network operating systems). Without access control, it will be very difficult to maintain the integrity and version control of the baseline model and preference set.

Another implementation concern is the assignment of system responsibilities to members of the USARC staff. To maintain this SDSS, there are some functions that, if performed by every user, would unnecessarily complicate the evaluation session and may introduce inconsistencies. For these reasons, it has been recommended to USARC that they designate a SDSS Administrator, and assign him or her responsibility for these common functions.

1. SDSS User

Although not prevented nor discouraged from accessing the full flexibility of this SDSS, it is expected that the novice user will rely heavily upon the simplified user interface created with Visual Basic™. This allows a user with limited computer or application training to produce standardized results. As discussed previously, the TPU location decision is semi-structured, but ARIES permits the decision maker to treat the problem as if it were fully structured, eliminating the need for additional subjective inputs.

As users build experience and confidence with the system, they can take advantage of many evaluation tools provided by LDW. For example, they can see the effects of alternative inputs, perform sensitivity analysis on the relative weights, and experiment with different utility functions. Alternative decision perspectives can be stored as personalized preference sets and shared with others. A user can even create personalized versions of the decision model (hierarchy of goals), modifying the number and nature of decision measures and goals, as well as redefining the relationships between them. Although it may be difficult to add new measures to the *automated* input, for this would involve modifying the shell coding, they can be easily inserted using *manual* techniques. If a measure is added to the model (with the associated adjustments to the preference set), then at the beginning of the evaluation phase when the user is presented with a matrix of inputs, a new column will appear for the added measure, and the cells of that column will be populated with zeros. The user need only type in values for that measure to replace the zeros.

Another LDW tool available to the user that enhances model flexibility is the "measure category". A measure category can be used to combine related pieces of data

(called sub-measures) into a single measure. This feature can create measure inputs that are the weighted average or sum of other numbers. For the preprogrammed model inputs, these types of calculations are executed by the VB shell. The category feature allows some degree of additional flexibility in the form of mathematical manipulations that can be implemented without modifying the shell. The category multipliers (i.e., values multiplied with each sub-measure) are included as part of a preference set.

What a typical user should not be permitted to modify are the *baseline* model and the *baseline* preference set. These items will evolve with time, but the evolution should be centrally managed to ensure that logic and consistency are maintained. The typical user also does not need to be involved with the preparation of data sources. As the underlying databases are refreshed, renamed, and moved, forcing each user to accommodate these changes introduces unnecessary complexity and opportunity for errors. Some data preparation, like the geocoding of large files, can also be very time consuming. These responsibilities should be assigned to a SDSS Administrator.

2. SDSS Administrator

A SDSS Administrator should be assigned responsibility for maintaining centralized control over source data, the baseline decision model, and the baseline preference set. Before any data processing can occur, several steps must be taken to ensure that source data is ready for use by ARIES. The data preprocessor must have the correct file names and directory paths for the source data in order to perform automatic extraction. File names and locations are changed periodically, and so the preprocessor must be updated to reflect those changes. Another data preparation issue is geocoding. Some files (i.e., AMSA, ECS, IRR,

GEOREF, NGNON_CLOS, RZA, and G18CWE) must be stored in geocoded form before running the preprocessor. The SDSS Administrator should track the updates to the source files and create new geocoded versions when appropriate.

For the baseline decision model and preference set, the SDSS Administrator should provide centralized control of the evolutionary process. As the SDSS is used and validated against reality, there should be adjustments made to both the baseline model and preference set, but these adjustments must reflect the consensus of the key decision experts or policy makers. An elicitation process similar to the one conducted for this study should accompany any significant modifications to the baseline approach.

One of the most significant limitations on model flexibility involves the automation of inputs for additional measures. This process should be managed by the SDSS Administrator or other system expert. Eliminating the effect of existing automated inputs can be easily accomplished with weights, but automating new inputs involves rewriting the Visual Basic™ (VB) coding. In some cases, the code revision would be relatively easy. For example, a simple transfer of another data value to a new column in the measures table could involve the addition of a single SQL query. It is possible in other situations that, what might initially appear to be an easy revision, involves a change in the underlying structure of data tables and data passing. A related thesis will provide detailed documentation of the programming structure.

C. CHAPTER SUMMARY

The overall architecture provides simplified access to powerful tools for decision support. The four major components of ARIES are: a data preprocessor, a mapping engine,

a decision model solver, and an integrating shell. These components were either written or purchased with integration in mind. The data preprocessor is a separate, executable program, written in VB, that extracts and preprocesses data from the source databases. The mapping engine is a commercial software product, MapInfo™, that provides for the spatial definition and processing of data. The integrating shell, another VB program, performs all other data processing, provides a simplified user interface, and coordinates the efforts of the other components. The final component, the decision model solver, is implemented through Logical Decisions for Windows™, which provides the evaluation of processed data.

Although not a focus of this project, some organizational implementation issues are also addressed. Reasons for, and uses of, an evaluation are discussed. The issues of data integrity, access control, and assignment of system responsibilities must all be addressed by USARC. A recommendation is made to establish the role of SDSS Administrator, a person who can maintain centralized control over the baseline decision model and preference set, and assume responsibilities that would unnecessarily burden the SDSS User.

VI. MODEL VALIDATION AND ENHANCEMENTS

A. MODEL VALIDATION

The specific structure and content of the decision model are based primarily upon the professional judgements of experts. Consequently, validation of the model involves validation of these judgements, a process that will be conducted by USARC. Although an important part of the modeling process, the scope of this research does not include a formal validation of the decision model. Rudimentary checks were conducted to ensure the reasonableness of the outputs and four strategies for more extensive validation are provided below.

Before a meaningful validation of the decision model can occur, there are numerous problems with the source data that must be corrected. During USARC acceptance testing, the frequent receipt of error flags accentuated the fact that many of the underlying data tables contain missing or obviously erroneous data. Some of these problems were traced to files that were truncated as they were transferred to the development team, but many instances indicate a general level of inaccuracy. Even more disconcerting than these obvious errors are the subtle problems, like missing records, that can skew the results with no obvious warning.

Validation of this model is expected to be a slow, evolutionary process. USARC has been given the tools and knowledge necessary to gradually improve the entire system. The methods suggested to aid this process are: comparison with historical results, further expert validation, direct comparison with the current process, and use of additional analytical study.

1. Historical Validation

It is difficult to link historical unit performance (based on readiness metrics) solely to unit location. Because the ARIES decision model includes only those readiness factors which are location-related, any real world relocation results used for validation of the model results must ignore all factors that are not directly reflected in the model (e.g., leadership, training program). The difficulty of isolating objective locational results implies that validation cases will also require subjective evaluation before being used as a validation data point. As an example, for a unit that was relocated and subsequently performed well, a person will have to assess whether the new location was a significant factor in the improvement, or whether the negative impact of a poor location choice was masked by improvements in areas such as leadership or training. Experts must choose validation cases where experience and judgement suggest that location was a significant factor in the success or failure of a unit and use those cases, either assuming that location was the only factor or somehow negating the influence of other factors. There should also be clear measurements of unit performance so that a consensus can be reached on what constitutes a strong or weak performance.

2. Expert Validation

Another validation strategy is to expand the expertise upon which the model is based. Although the original panel was composed of experts representing a variety of interests at USARC (e.g., overall readiness, training, leadership, personnel management, facilities engineering), it was still a small group (varying from three to six people on different aspects of the model). The chosen experts reached consensus on most issues with relative ease.

Incorporation of the judgement of additional experts may improve the model.

Suggestions for methods of eliciting additional expertise are provided below. One way to validate the model is to guide different experts through the process of constructing a decision model and comparing the new version with the current model. First, a consensus should be reached on the overall objective of the relocation decision. Then the experts could be asked to list the important factors in that decision. Influence diagrams proved to be helpful in the original sessions. Those factors judged to be location related could be compared against the current decision model and discussions of the differences should provide meaningful feedback.

Another validation option is to supply the experts with scenarios to analyze. These scenarios can be either based on reality or contrived with carefully selected differences. The experts could score each alternative in the relocation scenarios on a numerical scale similar to that used by ARIES. In addition to a score, they could supply the reasoning behind their assessments. The judgements of multiple experts could be used to define a distribution of responses providing a more comprehensive basis for comparison with the results of the SDSS. The display and analysis tools provided by ARIES offer a powerful tool for developing an understanding of the nature and significance of the differences.

3. Parallel Validation

Another validation strategy is to perform parallel analyses using both ARIES and the current evaluation process. There are a number of obstacles to the effectiveness of this strategy. Because most of the significant decision makers involved with the current process provided inputs to the decision model it is unlikely that this strategy will identify many

significant differences. Furthermore, having participated in the modeling process, it is likely that the informal, mental models that these people apply to the decision process will now be very similar to the formal model used by ARIES. Also, as mentioned before, because of the large amounts of data and tedious calculations, it is impractical to try to calculate some of the same decision inputs without an automated tool like ARIES.

Although it is difficult to manually simulate the ARIES process, where this strategy may provide the most valuable inputs is a comparison between the outputs of ARIES and the results obtained from less formal, more intuitive evaluations. This approach may identify significant factors that are not considered in the SDSS or suggest a significantly different set of preferences.

4. Analytical Validation

The ARIES decision model can also be validated through additional study and analysis. Most of the utility functions in ARIES were produced using professional judgement and are only loosely based on rigorous, statistical analysis. The ideal validation tool would be a comprehensive study on the relationship between location and unit readiness, the findings of which could be directly translated into a relocation decision model. Until a comprehensive study is conducted, the results from more selective analyses can help refine pieces of the ARIES model. Empirical studies could result in more precisely defined yield functions, or relative weights. If enough evaluation case studies were available, neural network techniques might prove useful in estimating the implicitly applied relative weights. Regression analyses might offer more accurate estimates of the individual utility functions. As formal studies provide additional insights on the various location related readiness

factors, they can be used as the basis for repeating the preference elicitation process. Some studies may even suggest new ways of modeling some aspects of this problem. If desired, those models can be incorporated as inputs to the established hierarchy of goals.

B. FURTHER RESEARCH AND ENHANCEMENTS TO THE PROTOTYPE

Research and enhancements to the prototype system can be divided into two categories: those that improve the usefulness of the system in the context of the TPU relocation decision, and those that improve the ease with which this framework and methodology are applied to other situations. The improvements that are applicable to this specific decision context are suggestions for further work to be performed or coordinated by USARC. Modifications intended to improve the ease with which this system is applied to other complex, spatial decisions are being studied and implemented in two related theses.

1. Internal Model and Data Improvements: Specific to USARC

Although a variety of validation methods have been suggested, responsibility for the validation and evolution of the decision model primarily rests with the members of USARC. Some system enhancements, such as improving the validity of source databases, are relatively obvious, and are already being pursued. Through discussions with the expert panel and objective study of the relocation problem, some less obvious improvements were also identified.

One of the most significant opportunities for improvement is extension of the model to fully address the relocation of unit derivatives. All calculations in the prototype are based upon the assumption that an entire unit is being considered for relocation. In some situations, it may be more appropriate to move a derivative (i.e., a subset of the positions

assigned to a unit) to a new location. Although the decision model could be easily adjusted to accommodate this variation, an addition to the user interface would also be needed, for the user would have to specify the exact composition of the derivative (i.e., the numbers and types of billets to be moved). One way to implement this change would be additional screening on both the extract of the personnel file (produced by the preprocessor) and the geocoded version of this file used by the mapping engine. The processing performed by the shell would then remain the same for MOS related calculations (i.e., Reassignments, Available MOS-Closing Unit, and Available MOS-IRR).

Another decision variation that could be added is the concurrent relocation of multiple units to the same location. Although this situation can be handled with the prototype by performing multiple, individual analyses, such an approach does not properly reflect the cumulative draw on available resources. Analogous to the approach described above, which suggested treating derivatives as small units, the data for multiple units could be aggregated and treated as a single unit. It would also be necessary to adjust some of the Single-Measure Utility Functions, for many were based on an assumed, average unit size.

When evaluating issues that are dependent upon Military Occupational Specialty (MOS), the prototype either considers all of the MOS's assigned to the moving unit or the three largest MOS groups (if they comprise more than 50% of the total number of reservists assigned to that unit). Previous research (Dolk, 1995) suggests that critical MOS profiles can be identified for most units. Adequate manning and performance in these MOS's are crucial to the unit's ability to complete its mission. The concept of critical MOS's refines the relationship between military readiness and objective measures of the recruit market

(particularly Available MOS-Closing Units and Available MOS-IRR) and retention. As further research explores the critical MOS issue, it should be considered for incorporation into the logic of ARIES.

Another enhancement that promises both significant effort and significant gain is a measure of compatibility between the unit's mission and the local civilian occupational structure. Experience shows, for example, that medical units perform much better in measures of personnel readiness when located near a civilian hospital. For many other types of units, this occupational compatibility does not appear to be a concern. In fact, some reservists highly value the differences between their full time employment and reserve responsibilities. Research that further defines this relationship and identifies supporting data sources can contribute to the validity of ARIES.

One of the underlying assumptions of this model is that all reservists make an equal contribution to readiness. To implement a system without this assumption would require an assessment of the abilities of individual reservists. Such an assessment could be used to weight the inputs to measures that are currently based upon numbers of people. One assessment tool that is available is the current system of formal personnel evaluations (Officer Evaluation Reports and Non-Commissioned Officer Evaluation Reports). Research that assesses the correlation between personnel evaluations and direct contributions to readiness may be appropriate before incorporating this factor into the ARIES model.

Further differentiation of Individual Ready Reserve (IRR) members based upon rank has also been suggested. The current approach assumes that all IRR members offer equal benefit to the moving unit as potential recruits. In actuality, most IRR members would

reenter the reserves in the mid- to high-level ranks because of their prior service, but reserve units, structured in a typical military hierarchy, have the greatest numerical need for personnel in the lower ranks. The data needed to more accurately match the available IRR market to the needs of the moving unit are available.

Similar to the previous discussion on preference sets, assumptions used to simplify reality for model construction (e.g., consideration of only the closest support sites, treatment of all recruiting stations as equal) are likely candidates for further research and improvements. In some situations, that research may simply prove the validity of the assumption or simplification. In other cases, it may establish new relationships, and possibly even new models, that can be directly integrated into the ARIES decision model.

2. External Improvements: General Applicability of the Methodology

The prototype implements considerable flexibility, but a number of improvements that would significantly enhance the generalizability of this framework appear to be technically feasible. During the construction of ARIES, model development and the automated implementation of that model were two distinctly separate steps. The general structure of the model was established using manual elicitation techniques (e.g., influence diagrams, tradeoff questions) and facilitated discussions. Some model refinements were then accomplished through independent use of a mapping engine and a decision model solver. Once the decision inputs were well defined, the structure necessary to automate those inputs was designed and hand coded. What is envisioned to improve the convenience and general applicability of this structure is a single, automated modeling environment that includes tools to assist in all stages of this process.

The enhanced decision modeling environment would begin with problem definition tools. A cooperative system is envisioned, one that allows multiple decision makers to participate in the definition of the problem and help them reach a consensus. It would include mapping and data visualization tools to assist this effort. The system would then help them select an appropriate decision model, and assist in the application of that model to the specific decision of concern, including the cooperative specification of preferences. Unlike ARIES, this enhanced system would also allow the decision maker to directly create the code necessary to automate the inputs. Using fourth generation computer languages (4GL's), the enhanced system would automatically generate coding based upon judiciously selected user inputs. The outputs would be similar to those currently offered but would include a better display of sensitivity analysis to decision inputs and additional geographical displays of the decision outputs.

As the overall design of this enhanced system is being refined, limited improvements are being implemented to ARIES as a means of migration. Currently, although the data preprocessor permits the user to view the queries used for data extraction and conditioning, it does not permit them to be directly modified through the standard interface. To implement a query change, the Visual Basic™ source code must be modified, compiled and used to replace the extraction program. A related thesis is improving this process by allowing the user to view and modify queries from the standard interface, and subsequently automating the coding, compilation, and replacement steps.

System access and portability issues are also being considered. ARIES is currently implemented on a personal computer, but to enhance its general applicability it should be

migrated to a multi-user system such as a LAN, intranet or the Internet. This system would be a valuable addition to *ReadiNet*, a proposed Internet-based system for recording, sharing, executing, and integrating readiness data and model resources (Dolk, 1995).

C. CHAPTER SUMMARY

Although the scope of the current research does not include formal validation of the decision model, it does suggest a number of validation strategies for use by the system owners. The four suggested strategies are historical, expert, parallel, and analytical validation. Using these techniques, USARC could gradually calibrate the decision model until it accurately reflects an appropriate set of factors, priorities, and preferences for relocation decisions.

In addition to the improvements suggested through validation, many other potential improvements have also been identified during the course of this research. Enhancements that could directly benefit this system as it applies to the TPU relocation decision are the modeling of unit derivatives, critical MOS's, specific contributions of individuals, and the match between a unit's mission and the local civilian job market. From a more general perspective, there are many enhancements, such as an automated code generator and a cooperative tool for eliciting preferences, that would facilitate the application of this methodology to other decision contexts.

VII. SUMMARY AND CONTRIBUTIONS

A. SUMMARY

This research analyzes the Troop Program Unit (TPU) relocation decision and suggests application of a formal decision model based upon Multi-Attribute Utility Theory. Drawing on the experience of an expert panel, the overall decision objective (site desirability) was decomposed into a hierarchy of goals. Underlying the structure of the model are many assumptions used to simplify the real-world complexities of this decision. The branches of the hierarchy terminate with thirty decision factors, which are objective, measurable attributes of the proposed relocation sites. Because some of the data needed to supply the inputs for these decision factors are not currently available to USARC, only two-thirds of the inputs are supplied in an automated fashion.

Construction of a prototype Spatial Decision Support System (SDSS) to address the relocation decision, involved the use of two commercial software programs, Logical Decisions for Windows™ as a decision model solver and MapInfo™ as a mapping engine. Coordination of these products was achieved with an overarching program shell written in Visual Basic™. The shell also provides a composite interface for the specification of the decision parameters and assists the inexperienced user to perform standardized evaluations, which rely upon the stored judgements of experts.

The unifying theme of this research is integration. At the forefront is the symbiosis between decision models and mapping software resulting in both a GIS and DSS of unusually broad scope and general applicability. The integration of large, disparate,

“stovepiped” data files is another significant benefit of this approach. And finally, the use and coordination of Commercial off-the-Shelf (COTS) products to implement a modeling environment rich with flexibility, reinforces the overall theme. The confluence of these components as described in this work constitutes an innovative contribution to the study and practice of decision support systems.

B. CONTRIBUTIONS

This research suggests a general methodology for supporting complex site location decisions, and describes a specific application of that methodology. At a theoretical level, this research provides a template for the integration of data, models, and commercial software components. At the applied level, a working prototype Spatial Decision Support System (SDSS) was delivered to USARC that offers significant improvements in the TPU relocation decision process.

1. Specific Contributions to USARC

The primary benefit of this system is that it helps the decision maker make better decisions. The real value of ARIES is not solely its ability to rank alternatives, but rather, the ways that it enhances the effectiveness of the human decision maker. It offers a variety of complementary methods for eliciting preferences, analyzing decisions, and displaying results. Decision makers can choose those methods that best fit the situation or support their decision style. Compared to the status quo, this SDSS significantly improves the convenience, consistency, clarity, timeliness, and comprehensiveness of the site evaluation process.

USARC has conducted similar location evaluations numerous times in the past and spent weeks accumulating the necessary data. The informal and inconsistent manner in which that data was evaluated, made it difficult to apply consistent standards as well as defend and build consensus for the site selections. Because of politics and a large number of potential stakeholders, relocation decisions often receive intense and emotional opposition based on very specific concerns. In some cases, these specific concerns are either not explicitly considered in an informal process or are difficult to communicate in the overall decision context. External challenges are often hard to address because relocation decisions are neither fully documented nor extensively challenged as part of USARC's internal approval process. As reports recommending relocation are routed through key decision makers, it is often difficult to dispute the findings. Raw data is relatively inaccessible and the analysis tools necessary to thoroughly challenge the conclusions are not conveniently available.

ARIES is a powerful tool for improving the way in which the TPU relocation decision is made. Data extraction and presentation, a task that used to require weeks of effort, can now be completed in minutes. By applying an explicit model, the decision process is more consistent and understandable. Even if an automated system had never been built, the modeling process provided significant benefits, for it forced the experts to surface their underlying beliefs, assumptions, and priorities. This effort not only improved understanding of the problem, but also enhanced the ability to communicate the decision process to others, making it easier to justify, defend, and build consensus for site selections.

Finally, and perhaps most importantly, ARIES supports a *comprehensive* evaluation, including what the experts consider to be *all* of the pertinent decision factors. Using manual methods, it was impossible to incorporate all of these inputs. Many of the decision inputs involve tedious calculations based upon large amounts of data. Even when these calculations are done by a computer, a large number of individual results must be mentally integrated without the assistance of a formal decision framework. By reducing the time and effort necessary to evaluate the alternatives, ARIES encourages the decision maker to engage in a deeper, systematic questioning and experimentation before reaching a final conclusion. This SDSS helps the decision maker gain insights into both the specific relocation alternatives under consideration, and the process used to reach a decision.

The decision maker is provided with a structure, but that structure is not rigid. Although a novice user may treat this as an Expert System (ES), relying primarily on the expert judgement elicited during model construction, it is hoped that most users will actively challenge the baseline model and preference set. Simplification of reality into a meaningful model is such a complicated process that the baseline model, as implemented in the prototype, is only a first approximation. One of the strengths of the ARIES system is the ease with which alternative viewpoints and approaches can be accommodated, supporting an evolutionary model development.

2. General Contributions

In addition to the specific benefits afforded to USARC, this project offers some general contributions to the field of Decision Support. Most importantly, it suggests a flexible modeling environment that relies on the synergistic integration of spatial processing

with decision modeling for complex site location decisions. This structure can be generalized in a straightforward manner to accommodate a wide variety of decision problems, particularly those with significant spatial components.

This research also describes a methodology for integration of numerous, large, dissimilar databases, a variety of decision models, and incompatible, commercial software packages. Until a software product is marketed that offers this range of functionality, or until software communication standards are more universally implemented, ARIES provides a practical approach to the technical challenges of integration.

The flexibility of this architecture supports its application to a wide variety of site location problems. Potential military applications include determining recruiter locations, the homeporting of naval vessels, hazardous waste storage, approval of berthing sites for nuclear vessels, and Base Realignment and Closure (BRAC). The ARIES framework accommodates decisions with multiple criteria, uncertainty, both spatial and non-spatial factors, and both objective and subjective inputs.

APPENDIX A. MULTI-ATTRIBUTE UTILITY THEORY

Multi-Attribute Utility Theory (MAUT) applies to situations with uncertainty and multiple, often conflicting, objectives. For situations where the decision variables are independent, a weighted addition of utilities (i.e., linear) is used to produce an ordinal ranking of alternatives. When interactions between the variables exist, multiplicative terms are introduced, resulting in a multi-linear overall utility function. A general form is this equation is:

$$u(x) = \sum_{i=1}^n k_i u_i(x_i) + \sum_{i=1}^n \sum_{j>i} k_{ij} u_i(x_i) u_j(x_j) + \sum_{i=1}^n \sum_{j>i} \sum_{k>j} k_{ijk} u_i(x_i) u_j(x_j) u_k(x_k) \\ + \dots + k_{123\dots n} u_1(x_1) u_2(x_2) \dots u_n(x_n)$$

where:

1. u is normalized by $u(x_1^0, x_2^0, \dots, x_n^0) = 0$ (the least preferred level of all measures) and $u(x_1^*, x_2^*, \dots, x_n^*) = 1$ (the most preferred level of all measures).
2. $u_i(x_i)$ is a conditional utility function of X_i normalized by $u_i(x_i^0) = 0$ and $u_i(x_i^*) = 1$.
3. The scaling constants can be evaluated by:

$$k_{123\dots n} = 1 - \sum_i u(x_i^0, \bar{x}_i^*) + \dots + (-1)^{n-2} \sum_{i,j>1} u(x_i^*, x_j^*, \bar{x}_{ij}^0) \\ + (-1)^{n-1} \sum_i u(x_i^*, \bar{x}_i^0)$$

APPENDIX B. ADDITIONAL MODEL INPUTS

One third of the inputs needed for the 30 measures in the original hierarchy of goals could not be supplied by the prototype in an automated fashion due to source data problems. The table below provides a brief description of the ten unautomated measures and a listing of the source data needed to provide the necessary inputs. In some cases, the source data exists but not in a database format that can be transferred to the USARC LAN. In other cases, the needed data has not yet been collected, documented, or created.

Model Input Measure	Needed Data
Percent Administrative Space (Full-Time): This measure indicates what percentage of the relocating unit's need for full-time administration space can be accommodated at the relocation site.	-A list of the full-time administrative space requirements of each unit (or a method of calculating this value). -A list of the full-time administrative space available in each facility (adjusted by the requirements of the units already assigned to the facility)
Percent Administrative Space (Part-Time): This measure indicates what percentage of the relocating unit's need for part-time administration space can be accommodated at the relocation site.	-A list of the part-time administrative space requirements of each unit (or a method of calculating this value). -A list of the part-time administrative space available in each facility (adjusted by the requirements of the units already assigned to the facility)
Percent Motorpool Space: This measure indicates what percentage of the relocating unit's need for motorpool storage space can be accommodated at the relocation site. Overflow is normally taken to an Equipment Concentration Site. (ECS)	-A list of motorpool space required for each unit. This might be produced through a list of vehicles and a method of converting to required storage space (vehicle footprint). -A list of the motorpool space available at each facility (adjusted by the requirements of units already assigned to the facility).

Model Input Measure	Needed Data
Percent Storage: This measure indicates what percentage of the relocating unit's need for equipment (other than motorpool items) storage space can be accommodated at the relocation site.	-A list of the equipment storage space required for each unit. This might be produced through a list of equipment and a method of converting to required storage space (equipment footprint). -A list of the equipment space available at each facility (adjusted by the requirements of units already assigned to the facility).
Distance to headquarters: This measure provides one indicator of the responsiveness of a headquarters to the needs of a unit.	-A list of which headquarters to which each unit reports. This list must be geocodable, including either a geographic field (e.g., zip code), or a facility identification code for the headquarters (FAC ID's can be cross referenced with a geocoded facilities file)
Civilian Labor Market: This measure provides an indication of the similarity between a unit's military mission and the local, civilian occupational structure.	-Data on the civilian occupational structure of the entire U.S. - A table that converts Military Occupational Specialties to an appropriate civilian labor code.
Percent Facility Usage: This measure indicates what percentage of a moving unit's need for personnel space can be accommodated at the relocation site. This is primarily an issue for full time staff.	-A list of the number of full time personnel assigned to each unit. -A list of the personnel space available at each facility (adjusted by the requirements of units already assigned to the facility).
Distance to the Weekend Training (WET) Site: This measure indicates the distance that must be traveled to reach the site used for routine, weekend training that cannot be performed at the unit's facility.	-A list of all weekend training (WET) sites. This list must be geocodable. -If the closest WET site is not always acceptable, a table is needed that can match units to appropriate WET sites.
Distance to Special Training: This measure indicates the distance that must be traveled to reach the sites used for special, mission or MOS related, training.	-A list of all special training sites. This list must be geocodable. -A table that can match unit missions, or MOS's with the appropriate special training site.

Model Input Measure	Needed Data
Distance to Weapons Qualification Range: This measure indicates the distance that must be traveled to reach the appropriate weapons qualification range.	-A list of all weapons qualification ranges. This list must be geocodable. -If the closest weapons qualification range is not always acceptable, a table is needed that can match units to appropriate qualification sites.

APPENDIX C. DECISION MODEL INPUTS (MEASURES)

#	Measure Name LDW name (ARIES name)	Units	Source Database	Fields Used (field name in the source database)	Type	Length
1	Fac Backlogd Maint (FAC_BACKLOGD_ MAINT)	dollars	RPINFODT	FAC_ID CWE_TOTAL	float (Single)	6.2
2	Fac Operating Costs (OPERATING_COST)	dollars/ sq. feet	FPS	COST_PR_SF FAC_ID	float (Single)	6.2
3	Facility Age (FAC_AGE)	years	INTEREST	DATE_ACQ FA_ID	integer	2
4	Facility Condition (FAC_COND)	no units	FPS	FAC_COND FAC_ID	character	5
5	Facility Ownership (FAC_OWNED)	no units	COMPLEX	GOVT_OWN FAC_ID	character (Boolean)	1
6	Competition (COMPETITION)	number of competitors	CMD_PLAN	FACID, UIC	integer	3
			G17	UIC		
			G19TRUE	OWN_UIC		
			NGNON_CL	ZIP, AUTH		
7	Area Drill Attendance (AREA_DRILL_ ATTEND)	Fraction of reservists with sat. drill atten- dance	CMD_PLAN	FACID, UIC	float	3.1
			FINANCE	PAY_STAT, NPS_IND, DOG, CURR_UIC, UTAxQCFY, UTAxQ1PF		
8	Area Loss Rate (AREA_LOSS_RATE)	Fraction of reservists lost in previous year	G17	ZIP, UIC	float	3.1
			G18CWE	UIC		
			FYxLOSS	UIC1, TRMN		

#	Measure Name LDW name (ARIES name)	Units	Source Database	Fields Used (field name in the source database)	Type	Length
9	Area Transfer Rate (AREA_XFER_RATE)	Fraction of reservists transferred in previous year	G17	TCCZIP, UIC, RECSTAT, TYPEORG	float	3.1
			CMD_PLAN	UIC, FACID		
			G18CWE	UIC		
			FyxxLOSS	UIC1, TRMN		
10	Avg Area Manning (AVG_AREA_MAN)	Manning level	CMD_PLAN	FACID, UIC	float	3.1
			G18CWE	UIC		
			G19TRUE	OWN_UIC		
11	Dist to Recruiter (DIST_TO_RECRTR)	miles	RZA	-	float	3.1
12	Closing Unit Xfers (TOTAL_AVAIL_ CLOS)	people	CMD_PLAN	FACID, UIC, TIER	integer	5
			G17	UIC, TIER, RECSTAT, TYPEORG		
			G18CWE	UIC		
13	IRR Available (IRR)	people	IRR	ZIP	integer	5
14	Recruit Market (RECRUIT_ MARKET)	people	QMA	ZIP, MWCAT12, MWCAT3A, MBCAT12, MBCAT3A, MHCAT12, MHCAT3A	integer	6
15	Reassignments (REASSIGNMENTS)	people	G18CWE	UIC, ZIP	integer	4
16	Dist to AMSA (DIST_TO_AMSA)	miles	AMSA	-	float	3.1
17	Dist to ECS (DIST_TO_ECS)	miles	ECS	-	float	3.1

#	Measure Name LDW name (ARIES name)	Units	Source Database	Fields Used (field name in the source database)	Type	Length
18	Fac Weekend Use (FAC_WKND_USED)	number of weekends used per month	COMPLEX	RS_WKND_PM FAC_ID	integer	1
19	Avail MOS Clos-Unit (MOS_AVAIL_CLOS)	people	CMD_PLAN	FACID, UIC	integer	5
			G17	UIC, TIER, RECSTAT, TYPEORG		
			G18CWE	UIC, PRI, ZIP		
20	Avail MOS IRR (IRR_MOS)	people	IRR	ZIP	integer	5
			G18CWE	UIC, PRI		

APPENDIX D. QUALIFIED MILITARY AVAILABLE (QMA) FILE

The Qualified Military Available (QMA) file is the only data source used by ARIES that was not supplied by USARC, but rather was supplied by the Naval Postgraduate School (NPS). Each record contains estimates of the number of people residing in each zip code who fall into four mental groupings and three race-ethnic groups (Black, Hispanic, and White Plus Other). The mental groupings are based upon the mental categories used by the U.S. Army Recruiting Command. The groups are categories one and two, category three "a", category three "b", and category four. The extract of QMA used in ARIES contains estimates for males from the ages of

Probability distribution are estimated for mental groups within race-ethnic group categories. (Probabilities add to one for each race-ethnic group in each file). These distributions are based on logistic regression equations estimated with NLSY data. Sociodemographic information from recent adjustments to the 1990 census is utilized in constructing these probability distributions.

Zipcodes, unlike counties, are frequently redefined. Some zipcodes which appear in this file were not included in data files containing sociodemographic information or in the set of zipcodes used by population forecasters. When sociodemographic information was not available for a zipcodes, the appropriate FIPS-level (county) input values were substituted. In some cases, zipcode-level estimates for 1990 based on 1988 sociodemographic information were substituted. The source of sociodemographic information for each zipcode is indicated by a code in the file.

The probability distributions in these files may be multiplied by population estimates for the appropriate gender/race-ethnic group segment of 17 to 21 year old high school graduates to yield specific-year estimates of qualified military available population by zipcode.

The file contains 34,265 records with the following data elements:

COLUMNS	DATA
1 - 5	Zipcode
6 - 10	FIPS
11 - 15	White plus other: Mental category 1 and 2
16 - 20	White plus other: Mental category 3A
21 - 25	White plus other: Mental category 3B
26 - 30	White plus other: Mental category 4 and below
31 - 35	Black: Mental category 1 and 2
36 - 40	Black: Mental category 3A
41 - 45	Black: Mental category 3B
46 - 50	Black: Mental category 4 and below
51 - 55	Hispanic: Mental category 1 and 2
56 - 60	Hispanic: Mental category 3A
61 - 65	Hispanic: Mental category 3B
66 - 70	Hispanic: Mental category 4 and below

APPENDIX E. SOURCE DATA TABLES

Table Name	Fields Used by ARIES	Size (Mb)	No. of records	Key Field	Purpose
AMSA	N/A	.1	190	N/A	Used in geocoded form to provide the location of all Area Maintenance Support Activities.
CMD_PLN	FACID,UIC, EDATE	3.3	12,476	UIC	Primarily used to convert FAC_ID's to UIC's.
COMPLEX	GOVT_OWN, FAC_ID, RS_WKND_PM	1.3	1554	FAC_ID	Provides data on each facility (whether it is owned by the government and the number or weekends that it is currently in use each month).
ECS	N/A	.1	30	N/A	Used in geocoded form to provide the location of all Equipment Concentration Sites.
FINANCE	PAY_STAT, NPS_IND, DOG, CURR_UIC, UTAXQCFY, UTAXQ1PF	3	17,293	SSAN	Used to determine the fraction of reservists who have participated in 21 or more drill periods in the four previous complete quarters.
FPS	COST_PR_SF, FAC_ID, FAC_COND	5.8	1,561	FAC_ID	Provides the condition and the operating costs associated with all facilities.
FYxLOSS	UIC1, TRMN	151	8,828	UIC	Used to determine the fraction of reservists who were lost by either attrition or transfer by all of the area units.
GEOREF	N/A	.2	1553	FAC_ID	Provides the location of all facilities.
G17	UIC, RECSTAT, TYPEORG, TCCZIP, TIER, UNITNAME, TCCCITY, TCCSTAT	3.3	5,319	UIC	Used to determine a valid list of UIC's and a list of closing UIC's. This table also supplies descriptive data on each unit.
G18CWE	UIC, ZIP, PRI	198	204,299	SSN	Used to determine who is assigned to each unit.

Table Name	Fields Used by ARIES	Size (Mb)	No. of records	Key Field	Purpose
G19TRUE	OWN_UIC	25	233,211	OWN_UIC PARA_NBR LINE_NBR POSN_NBR	Indicates the number of positions assigned to each unit.
IRR	ZIP, PMOS	7.3	140,000	SSAN	Used in geocoded form to indicate where all Individual Ready Reserve members live.
INTEREST	DATE_ACQ, FAC_IDSTR	4.4	3,963		Provides the date of initial acquisition of each facility (used to calculate facility age).
NGNON_CL	ZIP, AUTH	.5	3,673	UPC	Used to indicate the number of competing Army National Guard positions.
RPINFODT	FAC_ID, CWE_TOTAL	.1	7,982	WO_ID	Provides the cost of correcting backlogged maintenance.
RZA		.2	1793	RSID	Used in geocoded form to indicate the location of all recruiting stations.
QMA	ZIP, xxCAT12, xxCAT3A	2.7	34,265	ZIP	Provides a measure of the market for new recruits by zip code.

APPENDIX F. QUERY STATEMENTS

Name: COMMAND PLAN

Purpose: Obtain a valid list of UIC's. Only those records with an action date in the past or within the next 13 months are considered.

Select: DISTINCT UIC, FACID AS FAC_ID, EDATE INTO CMDPLAN

From: COMMANDPLAN

Where: (FACID <> "N/A") AND (FACID <> "TBD")
AND (FACID <> "") AND (LEN(FACID) > 2)
AND ((LEFT(EDATE,4) = '1998' AND
MID(EDATE,5,2) <= '02') OR
(LEFT(EDATE,4) <= '1997'))

Order By: UIC, EDATE DESC

Group By:

Index On: UIC As UIC Primary: No Unique: No

Name: COMPLEX

Purpose: Obtain a list of facilities indicating which are owned by the Army Reserves (i.e., not leased) and how many weekends per month each is currently in use.

Select: FAC_ID, GOVT_OWN AS FAC_OWNED,
RS_WKND_PM AS FAC_WKND_USED
INTO COMPLEX_

From: COMPLEX

Where: LEN(FAC_ID) = 5

Order By:

Group By:

Index On: FAC_ID As FACID Primary: Yes Unique: Yes

Name: FINANCE

Purpose: Obtain a count of all UIC's in this database.

Select: "W" & LEFT(CURR_UIC,5) AS UIC, COUNT(CURR_UIC) AS
UIC_TOTAL INTO FINANCE_

From: FINANCE

Where: CURR_UIC <> "" AND NPS_IND = NULL AND PAY_STAT='A'

Order By: CURR_UIC

Group By: CURR_UIC

Index On: UIC As UIC Primary: No Unique: No

Name: FINANCE_QTR

Purpose: Obtain drill attendance statistics for the previous 8 quarters.

Select: "W" & LEFT(CURR_UIC,5) AS UIC,
UTA1QCFY, UTA2QCFY, UTA3QCFY,
UTA4QCFY, UTA1Q1PF, UTA2Q1PF,
UTA3Q1PF, UTA4Q1PF INTO FINANCE_QTR

From: FINANCE

Where: CURR_UIC <> "" AND NPS_IND = NULL AND PAY_STAT='A'

Order By: CURR_UIC

Group By:

Index On: UIC As UIC Primary: No Unique: No

Name: FPS

Purpose: Obtain the Facility Condition and Cost per Square Foot for each facility

Select: FAC_ID, FAC_COND, COST_PR_SF INTO FPS_

From: FPS

Where: FAC_ID <> ""

Order By: FAC_ID

Group By:

Index On: FAC_ID As FACID Primary: Yes Unique: Yes

Name: FYxxLOSS

Purpose: Obtain the number of losses in the previous fiscal year at each UIC.

Select: UIC1 AS UIC, COUNT(UIC1) AS UIC_TOTAL INTO FYxxLOSS

From: FY_LOSS

Where: TRMN = 'LOSS'

Order By: UIC1

Group By: UIC1

Index On: UIC As UIC Primary: Yes Unique: Yes

Name: FYxxXFER

Purpose: Obtain the number of transfers in the previous fiscal year at each UIC.

Select: UIC1 AS UIC, COUNT(UIC1) AS UIC_TOTAL INTO FYxxXFER

From: FY_LOSS

Where: TRMN = 'TRFD'

Order By: UIC1

Group By: UIC1

Index On: UIC As UIC Primary: Yes Unique: Yes

Name: G17

Purpose: Obtain a valid list of all the UIC's, including unit descriptive data.

Select: UIC, UNITNAME, TCCCITY AS CITY, TCCSTAT AS STATE,
LEFT(TCCZIP,5) AS ZIP, TIER INTO G17Natl

From: G17

Where: (RECSTAT <> "1") AND (TYPEORG <> "2") AND UIC <> ""

Order By: UIC

Group By:

Index On: UIC As UIC Primary: Yes Unique: Yes

Name: G18

Purpose: Obtain a list of the Zip Code, MOS, and UIC of every Reservist.

Select: UIC, LEFT(ZIP,5) AS ZIPCODE, LEFT(PRI,3) AS MOS INTO G18Natl

From: G18_CWE

Where: PRI <> "" AND UIC <> ""

Order By: UIC

Group By:

Index On: UIC As UIC Primary: Unique:

Name: G18_UIC

Purpose: To obtain a list of the number of reservists assigned to each UIC.

Select: UIC, COUNT(UIC) AS UIC_TOTAL INTO G18Natl_UIC

From: G18Natl

Where:

Order By: UIC

Group By: UIC

Index On: UIC As UIC Primary: Yes Unique: Yes

Name: G19

Purpose: Obtain a list of all UIC's.

Select: OWN_UIC AS UIC, COUNT(OWN_UIC) AS
UIC_TOTAL INTO G19Natl

From: G19

Where: OWN_UIC <> ""

Order By: OWN_UIC

Group By: OWN_UIC

Index On: UIC As UIC Primary: No Unique: No

Name: INTEREST

Purpose: Obtain a list of facility aquisition dates.

Select: FAC_IDSTR AS FAC_ID, DATE_ACQ INTO INTEREST_

From: INTEREST

Where: FAC_IDSTR <> "" AND ABB_TYPE = "USARC (MB)" AND NOT
IS NULL (DATE_ACQ)

Order By: FAC_IDSTR

Group By:

Index On: FAC_ID As FACID Primary: Yes Unique: Yes

Name: RPINFODT

Purpose: To obtain the backlogged maintenance costs for each Facility.

Select: FAC_ID, SUM(CWE_TOTAL) AS MAINT_COST INTO RPINFODT_

From: RPINFODT

Where: FAC_ID <> ""

Order By: FAC_ID

Group By: FAC_ID

Index On: FAC_ID As FACID Primary: Yes Unique: Yes

Name: VALID UNIT

Purpose: Obtain a list of Valid Units from the GEOREF file.

Select: FAC_ID, FAC_TITLE AS UNITNAME,
FAC_CITY AS CITY, FAC_STATE AS STATE,
LEFT(FAC_ZIP,5) AS ZIP INTO VALID UNIT

From: GEOREF

Where: FAC_ID <> ""

Order By: FAC_ID

Group By:

Index On: FAC_ID As FACID Primary: No Unique: No

LIST OF REFERENCES

- Betts, R. K., *Military Readiness: Concepts, Choices, Consequences*. Brookings Institution, 1995.
- Bonczek, R. H., C. W. Holsapple, and A. B. Whinston, "The Evolving Roles of Models in Decision Support Systems," *Decision Sciences*, Vol. 11, No. 2, 1980.
- Borack, J. I., Mehay, S. L., and G. W. Thomas, *Estimating Local Area Manpower Supply for the Reserves*, Naval Postgraduate School Technical Report 54-85-006, Systems Management Department, Monterey, CA 93943-5103, September 1985.
- Boykin, M. L., Merritt, H. L., and R. L. Smith, *An Empirical Analysis of Retention Within the United States Naval Reserve*, Naval Reserve Retention Study, New Orleans, LA, 1980.
- Burright, B., Grissmer, D., and Z. Doering, *A Model of Reenlistment Decisions of Army National Guardsmen*, Rand Corporation, 1982.
- Dolk, D. R., *Readiness and Resources in the U.S. Army Reserve: A Research Study Plan*. Naval Postgraduate School Technical Report NAVAL POST-GRADUATE SCHOOL-SM-95-005, Systems Management Department, Monterey, CA 93943-5103, August 1995.
- Ellis, H. M., *The Application of Decision Analysis to the Problem of Choosing an Air Pollution Control Program for New York City*, Unpublished doctoral dissertation, Graduate School of Business Administration, Harvard University, Cambridge, MA, 1970.
- Fishburn, P. C., "Interdependence and Additivity in Multivariate, Unidimensional Expected Utility Theory," *International Economic Review*, 1967a.
- Fishburn, P. C., "Additive Utilities with Incomplete Product Sets: Application to Priorities and Assignments," *Operations Research*, 1967b.
- Fishburn, P. C., "Utility Theory," *Management Science*, vol. 14, 1968.
- Keen, P. G. W., *Value Analysis: Justifying Decision Support Systems*, MIS Quarterly, March 1981.
- Keeney, R. L., and H. Raiffa, *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*, John Wiley and Sons, New York, NY, 1976.
- MacCrimmon, K. R., *Improving the System Design and evaluation Process by the use of Trade-off Information: An Application to Northeast Corridor Transportation Planning*,

RM-5877-DOT, Rand Corporation, Santa Monica, CA, 1969.

McNaught, W., *Projecting Future Accessions to the Selected Reserve Components*, N-1563-MRAL, Rand Corporation, Santa Monica, CA, June 1981.

McNaught, W., and C. Francisco, *The Supply of Enlistees to the Selected Reserve Forces*, N-1562-MRAL, Rand Corporation, Santa Monica, CA, July 1981.

Miller, H. J., and A. Okabe, *Exact Computational Methods for Calculating Distances Between Objects in a Cartographic Database*, Proceedings of the Western Regional Sciences Association meeting, February 1997.

Moore, J. H., and M. G. Chang, *Design of Decision Support Systems*, Data Base, Vol. 12, Nos. 1 and 2, Fall 1980.

Pollak, R. A., "Additive von Neumann-Morgenstern Utility Functions," *Econometrica*, 1967.

Pruzan, P. M. and J. T. R. Jackson, "On the Development of Utility Spaces for Multi-Goal Systems," *Saettryk af Erhvervsøkonomisk Tidsskrift*, 1963.

Sohn, S. Y., and G. W. Thomas, *U.S. Army Reserve Membership Commute Distance Model*, Naval Postgraduate School Technical Report NPSOR-91-14, Operations Research Department, Monterey, CA 93943-5103, 1991.

Solnick, L. M., and G. W. Thomas, *Competition and USAR Unit Success: A Preliminary Examination*, Proceedings of the ORSA/TIMS National Meeting, May 1990.

Sprague, R., and E. Carlson, *Building Effective Decision Support Systems*, Prentice-Hall, 1982.

Tufte, E. R., *The Visual Display of Quantitative Information*, Graphics Press, Cheshire, Connecticut, 1983.

Turban, E., "Decision Support and Expert Systems," *Management Support Systems*, Macmillan Publishing Company, New York, NY, 1990.

INITIAL DISTRIBUTION LIST

	No. Copies
1. Defense Technical Information Center	2
8725 John J. Kingman Rd., STE 0944	
Ft. Belvoir, VA 22060-6218	
2. Dudley Knox Library	2
Naval Postgraduate School	
411 Dyer Rd.	
Monterey, CA 93943-5101	
3. Headquarters, U.S. Army Reserve Command.	2
ATTN: LTC Cleven	
3800 North Camp Creek Parkway SW	
Atlanta, GA 30331	
4. Professor Daniel Dolk, Code SM/Dk	3
Department of Systems Management	
Naval Postgraduate School	
555 Dyer Rd Bldg 330	
Monterey, CA 93943-5000	
5. Professor George Thomas, Code SM/Te	3
Department of Systems Management	
Naval Postgraduate School	
555 Dyer Rd Bldg 330	
Monterey, CA 93943-5000	
6. LCDR Mark A. Murphy.	3
153 Aaron Way	
Marina, CA 93933	