AD-A127 364 US COAST GUARD ALTERNATIVES FOR DISTRIBUTED DATA BASE MANAGEMENT SYSTEMS(U) NAVAL POSTGRADUATE SCHOOL MONTEREY CA E M FIEGEL ET AL. DEC 82 1/2

UNCLASSIFIED F/G 5/1 NL

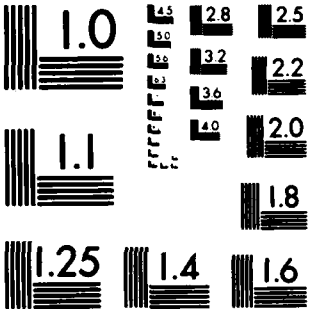

MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS 1963-A

NAVAL POSTGRADUATE SCHOOL
Monterey, California

AD A127444

DTIC
APR 28 1983
A

THESIS

U.S. COAST GUARD ALTERNATIVES FOR
DISTRIBUTED DATA BASE MANAGEMENT SYSTEMS

by

Edward Mark Fiegel
and
Stephen H. Goetchius

December 1982

Thesis Advisor: N. R. Lyons

Approved for public release; distribution unlimited.

DTIC FILE COPY

83 04 28 125

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE

READ INSTRUCTIONS BEFORE COMPLETING FORM

1. REPORT NUMBER

2. GOVT ACCESSION NO. A124 364

3. RECIPIENT'S CATALOG NUMBER

4. TITLE (and Subtitle)
U.S. Coast Guard Alternatives for Distributed Data Base Management Systems

5. TYPE OF REPORT & PERIOD COVERED
Master's Thesis
December 1982

6. PERFORMING ORG. REPORT NUMBER

7. AUTHOR(s)
Edward Mark Fiegel
Stephen H. Goetchius

8. CONTRACT OR GRANT NUMBER(s)

9. PERFORMING ORGANIZATION NAME AND ADDRESS
Naval Postgraduate School
Monterey, California 93940

10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS

11. CONTROLLING OFFICE NAME AND ADDRESS
Naval Postgraduate School
Monterey, California 93940

12. REPORT DATE
December 1982

13. NUMBER OF PAGES
120

14. MONITORING AGENCY NAME & ADDRESS(if different from Controlling Office)

15. SECURITY CLASS. (of this report)

15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)
Approved for public release; distribution unlimited.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)
Data Base Management Systems
Coast Guard
Information Resource Management

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

The United States Coast Guard is a relatively small federal agency tasked with a number of duties. Its multimission nature and low budget is driving the Coast Guard to realize better use of its resources. A pivotal factor in this goal is the investment in the information systems architecture of the future, today. Within the information architecture, data base technology plays an important role. It is to be employed in major operational and

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE
S/N 0102-014-6601

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

and administrative systems, as well as in the future Coast Guard District Minicomputer Procurement. The purpose of this thesis is to examine the alternatives available to the U.S. Coast Guard for implementing data base technology.

COPY INSPECTED

Approved for public release; distribution unlimited.

U.S. Coast Guard Alternatives for Distributed Data Base Management Systems

by

Edward Mark Fiegel
Lieutenant, U.S. Coast Guard
B.S., Lowell Technological Institute, 1970

and

Stephen H. Goetchius
Lieutenant, U.S. Coast Guard
B.S., U. S. Coast Guard Academy, 1977

Submitted in partial fulfillment of the requirements for the degree of

MASTER OF SCIENCE IN INFORMATION SYSTEMS

from the

NAVAL POSTGRADUATE SCHOOL
December 1982

Authors: E. M. Fiegel

Stephen H. Goetchius

Approved by: _____
Thesis Advisor

_____
Second Reader

_____
Chairman, Department of Administrative Sciences

D. M. Woods
_____
Dean of Information and Policy Sciences

ABSTRACT

The United States Coast Guard is a relatively small federal agency tasked with a number of duties. Its multi-mission nature and low budget is driving the Coast Guard to realize better use of its resources. A pivotal factor in this goal is the investment in the information systems architecture of the future, today. Within the information architecture, data base technology plays an important role. It is to be employed in major operational and administrative systems, as well as in the future Coast Guard District Minicomputer Procurement. The purpose of this thesis is to examine the alternatives available to the U.S. Coast Guard for implementing data base technology.

TABLE OF CONTENTS

I. INTRODUCTION 10

A. BACKGROUND DISCUSSION 10

B. RESEARCH OBJECTIVES 10

C. THESIS SCOPE, LIMITATIONS AND ASSUMPTIONS . . 11

D. METHODOLOGY 11

E. ORGANIZATION OF THE THESIS 12

II. RESEARCH FRAMEWORK 13

A. GENERAL 13

1. Coast Guard Missions 13

B. COAST GUARD ORGANIZATION 17

1. General 17

2. Operational Resources 23

3. Budget 24

4. Personnel 25

5. Assumptions and Constraints 26

6. ADP Equipment 26

C. TECHNOLOGY FRAMEWORK 36

III. ORGANIZATIONAL IMPLICATIONS 39

A. GENERAL 39

B. DIRECTION COAST GUARD IS HEADED 39

1. Text Handling 40

2. Messaging 41

3. Distributed Systems and Networks 42

4. Decision Support Systems (DSS) 43

5. Data Base Management Systems 44

C. INFLUENCING FACTORS 45

D. MINIMIZING NEGATIVE IMPLICATIONS 49

E. SUMMARY 52

IV. FUTURE TECHNOLOGY DEVELOPMENTS 53
A. GENERAL 53
B. ASSOCIATIVE COMPUTER ARCHITECTURE 54
C. MEMORY 56
D. DATABASE MACHINES 57
E. DISTRIBUTED DBMS 61
F. SUMMARY 63

V. CURRENTLY AVAILABLE DBMS ARCHITECTURES 65
A. GENERAL 65
B. RELATIONAL AND CODASYL NETWORKS 66
C. RELATIONAL - SYSTEM R 72
1. Capabilities/Characteristics 72
2. Critique 73
D. NETWORK - CODASYL TYPE - TOTAL 74
1. Capabilities/Characteristics 75
2. Critique 76
E. HIERARCHICAL - SYSTEM 2000/80 77
1. Capabilities/Characteristics 78
2. Critique 78
F. SUMMARY 79

VI. DATABASE MANAGEMENT SYSTEM STRATEGY 81
A. GENERAL 81
1. Present Situation 81
2. Planned Growth 83
B. DBMS OBJECTIVES 84
1. Multiple Views of Data 84
2. Performance 85
3. Minimization of Costs 86
4. Minimal Redundancy 87
5. Query and Report Generation 88
6. Integrity Controls 89
7. Security and Privacy 90

8. Three-Level Hierarchy 91
C. COAST GUARD INFORMATIONAL GOALS 92
1. Protection of Intellectual Development . . 93
2. Growth 97
D. HETERCGENEOUS VERSUS HOMOGENEOUS SYSTEMS . . .100
E. SUMMARY 102

VII. CONCLUSIONS 103

APPENDIX A: RESCURCES 106

APPENDIX B: SDD-1 108

LIST OF REFERENCES 117

INITIAL DISTRIBUTION LIST 120

LIST OF TABLES

I. Coast Guard Missions by Logical Department 16

II. Coast Guard Districts 18

III. Coast Guard Operational Resources 23

IV. Duties of DBA and Staff in a DBMS Environment . . 48

V. Data Definition, Manipulation, and Control Facilities 74

VI. Comparison of General Capabilities 75

VII. 1979 DataPro Survey of TOTAL Users 77

VIII. DataPro Survey of System 2000/80 Users 79

LIST OF FIGURES

2.1 Hierarchical Information 31
2.2 Transactional Information 32
2.3 U.S. Coast Guard Planned Information Architecture 33
2.4 Coast Guard NETWORKS 35
4.1 Associative System 55
4.2 Growth in Micoelectronics 57
4.3 Current DBMS Architectures 58
4.4 Future DBMS Architectures 59
5.1 Three-Level Hierarchy of a DBMS 67
5.2 Example of Two Relations 70
5.3 Owner-Coupled Set Occurrence 71
B.1 SDD-1 Architecture 108
B.2 Horizontal Fragment of a Relation 109
B.3 Vertical Fragment of a Relation 110
B.4 Conflict Graph 113

I. INTRODUCTION

A. BACKGROUND DISCUSSION

In 1981, The U.S. Coast Guard created a new Office of Command, Control and Communication. This Office has set forth the information architecture concept for the Coast Guard, which is committed to investing in the architecture of the future today so that the Coast Guard may become an "information corporation" of the 1990's. The Commandant of the Coast Guard has identified three critical success factors (CSF): intelligent terminals, data base management systems, and telecommunications networks. The intelligent terminals (standard terminals) are already being acquired and configured for applications. Data base management technology is to be employed in major operational and administrative systems under development, as well as in the future Coast Guard District minicomputer procurement. This thesis will address the Coast Guard's second CSF - data base management systems - looking at how the Coast Guard should implement database technology in a distributed mode, given the present heterogeneous hardware/software systems and planned system acquisitions.

B. RESEARCH OBJECTIVES

The primary objective of this thesis will be to provide the Coast Guard with a strategic plan for implementing data base technology given the present configuration of the Coast Guard and the status of systems available today. A secondary objective would be to present what future trends in data base technology are foreseen and the situation of present data base systems in general.

C. THESIS SCOPE, LIMITATIONS AND ASSUMPTIONS

The main thrust of this thesis will be an examination of the various Data Base Management System (DBMS) architectures available as well as an examination of the Coast Guard's requirements for data base technology. Future trends in DBMS technology will be examined, realizing that the systems acquired today will be affected by the technological advances which occur in the near future. The researchers did not confine the scope of the thesis to any one specific organizational element of the Coast Guard, but rather attempted to consider the requirements of the organization as a whole. The reader is assumed to have at least a cursory knowledge of computer systems nomenclature. (Note: The authors will not attempt to use a single spelling of either "data base" or "database" but will use both interchangeably since both spellings are used throughout current literature).

D. METHODOLOGY

The methodology employed in this research effort was primarily an observational type approach coupled with an extensive literature review of current books, periodicals, articles and journals, as well as Coast Guard directives, plans, and policy guidance. These literature reviews were augmented by interviews with appropriate personnel as necessary. These research techniques were appropriate because they furnished Coast Guard requirements while identifying the available technologies and future trends.

E. ORGANIZATION OF THE THESIS

The following is a breakdown of the various chapters included in this thesis:

Chapter II - This chapter provides the reader with the current framework or situation which the researchers faced in conducting their research.

Chapter III - This chapter discusses some organizational implications which may affect the Coast Guard's Command, Control and Communications program goals in general, and the data base goals in particular.

Chapter IV - This chapter discusses the future technology developments in data base management systems and how these developments will affect the currently available alternatives.

Chapter V - This chapter provides a detailed description of three currently available architectures and discusses the strengths and weaknesses of each.

Chapter VI - This chapter provides a strategy which the researchers believe the Coast Guard should employ in implementing database technology.

Chapter VII - This chapter provides a summary of the researchers' conclusions and recommendations.

II. RESEARCH FRAMEWORK

A. GENERAL

The purpose of this chapter is to give the reader an understanding of the framework which the researchers addressed in formulating this thesis. The authors will give a brief overview of the missions assigned to the U.S. Coast Guard and then look at how the Coast Guard is organized in order to accomplish these missions. The researchers will then look at the plans, policies and resources dedicated to the Command, Control and Communications program of the Coast Guard in supporting their mission objectives. A brief look at the current technology framework is included to describe the environment of the information systems industry facing the Coast Guard today.

1. Coast Guard Missions

The United States Coast Guard was created in 1790 by an Act of Congress which was initiated by the Secretary of the Treasury, Alexander Hamilton. The primary purpose of the organization was initially to collect tariffs for the country to offset the debt from the Revolutionary War. This organization was first called the Revenue Cutter Service and had an initial complement of 10 cutters with 40 total officers. In 1915 the Life Saving Service was combined with the Revenue Cutter Service to form the U.S. Coast Guard. The Lighthouse Service was amalgamated in 1939. At present the Coast Guard is tasked with performing the following missions:

SEARCH AND RESCUE (SAR): Locating and assisting persons and property in distress on the high seas and U.S.

jurisdictional waters. Charged by the National SAR Plan with total responsibilty for SAR in the North Pacific and most of the North Atlantic.

ENFORCEMENT OF LAWS AND TREATIES: Place territorial and adjacent waters under surveillance to deter illegal activities, acquire relevant data, and promote safety of life and property.

AIDS TO NAVIGATION: Operate and maintain a complex network of navigational signals (e.g. buoys, lights, radio beacons) covering all U.S. jurisdictional waters, and radionavigation signal systems (LORAN and OMEGA) meeting the requirements of the Department of Defense and the civil community.

ICE BREAKING: Clear passage through frozen waterways for ship movements and conduct scientific research in ice covered waters.

MARINE ENVIRONMENTAL PROTECTION: Perform surveillance of maritime areas for evidence of discharged oil, other hazardous substances and ocean dumping violations. Prevent unauthorized discharging and enforce applicable laws and treaties. Ensure necessary containment and removal of spills.

MILITARY READINESS: Maintain an effective force of personnel, facilities and equipment in a state of readiness for war or peacetime emergency.

RECREATIONAL BOATING SAFETY: Conduct safety patrols, especially in areas of high boating density, instruct state officials and the Coast Guard Auxiliary in methods to increase recreational boating safety, and promote boating safety through public contact.

COMMERCIAL VESSEL SAFETY: Develop safety standards and enforce them through vessel and equipment inspection, vessel documentation, investigation of accidents and violations, and licensing of seamen.

PORT SAFETY AND SECURITY: Enforce laws and safety regulations. Facilitate traffic movement. Investigate accidents and violations. Monitor loading operations and movement of vessels carrying hazardous cargo [Ref. 1].

The Coast Guard is the smallest of the U.S. Armed Services and is unique in that it performs these various peacetime missions as opposed to the other Armed Forces which really only perform wartime related functions. Having a multitude of missions has created some debate as to which department the Coast Guard should belong. Originally, the Coast Guard was part of the Department of Treasury, but in 1967 the Coast Guard was placed under the newly created Department of Transportation. Some observers have stated that the Coast Guard should be part of the Treasury Department while others feel it belongs with the Commerce Department. Another view points out that since the Coast Guard is ultimately an Armed Service, it should be part of the Department of Defense. Table I relates the various missions which the Coast Guard performs to the Federal Department to which that function would be connected [Ref. 2]. Reference 2 points out that there are various reasons for not putting the Coast Guard in either the Department of Defense, Commerce or Treasury and recommends that the Coast Guard remain a part of the Department of Transportation, with which the researchers concur. But it is observed that there will continue to be speculation of relocating the Coast Guard as long as it remains a multi-mission organization. In recent years this fact, along with being a relatively small organization, has contributed

TABLE I

Coast Guard Missions by Logical Department

MISSION	LOGICAL DEPARTMENT
Armed Force/Miltary Readiness	Defense
Maritime Law Enforcement / Navigation	Commerce or Transportation
-fisheries, sanctuaries, mammals	
-Navigation laws, aids to navigation	
-Navigation faciltation	
Maritime Law Enforcement	Justice
-Drug Interdiction	
-Immigration Laws	
Maritime Law Enforcement	Treasury
-Customs Laws	
Marine Safety Functions	Transportation
-Search and Rescue	
-Aids to Navigation	
-Commercial vessel, boating and port safety	
Marine Environmental Protection	EPA

toward giving the Coast Guard somewhat of a visibility problem. The end result of this visibility problem is that it has been difficult for the Coast Guard to get adequate funding and resources in order to properly carry out the missions with which they are assigned. In 1980, in combination with the Office of Management and Budget and the Office of the Secretary of Transportation the Coast Guard carried out a sophisticated zero-based personnel study. This study considered the missions currently assigned to the Coast

Guard and the program standards developed for current levels of operations. The study found that to continue the current level of operations without deterioration of plant and equipment, the Coast Guard would need between 9,000 and 15,000 more personnel than they currently have. The average age of Coast Guard ships is 27 years and there is approximately a $2-3 billion backlog in capital investment to its aging shore facilities [Ref. 3].

The bottom line with regard to this discussion is that the Coast Guard is a severely resource constrained organization (as all organizations would probably admit, but the researchers believe that the Coast Guard is especially so). Let us then proceed to look at how the Coast Guard is organized to carry out its missions and then look into how the Coast Guard handles its information resources.

B. COAST GUARD ORGANIZATION

1. General

The Coast Guard is an organization with approximately 39,000 military personel and 6,000 civilian employees. Coast Guard Headquarters is located in Washington, D.C. and is divided up into staff elements responsible for the various operational or support programs the Coast Guard carries out in support of its assigned missions. (The Search and Rescue program is an example of an operational program.) These staff elements are often referred to as program managers and they set the policies and objectives for their respective programs according to the guidance set forth by the Commandant of the Coast Guard, a 4-star flag rank.

The next organizational level in the Coast Guard is broken down into two Area Offices - the Atlantic and Pacific, located in New York and San Francisco respectively.

TABLE II

Coast Guard Districts

ATLANTIC AREA	PACIFIC AREA
1st: Boston, Ma.	11th: Long Beach, Ca.
2nd: St. Louis, Mo.	12th: San Francisco, Ca.
3rd: New York, Ny.	13th: Seattle, Wa.
5th: Portsmouth, Va.	14th: Honolulu, Ha.
7th: Miami, Fl.	17th: Juneau, Al.
8th: New Orleans, La.	
9th: Cleveland, Oh.	

Each Area is then subdivided into District regions. Table II shows the breakdown of Districts with the location of each of the District headquarters. Each District is then subdivided into Groups, Stations, and Captain of the Port zones depending on the needs of each particular District. The various afloat resources are distributed among the Districts with Coast Guard aircraft being operated out of 26 different air stations. The District Commanders are responsible for all operational activities within their geographic boundaries. Each Area office provides administrative and operational support and guidance when more than one District is involved in a common mission. Since almost every Coast Guard unit is considered multimissioned, management of operational, logistic and informational support can be a difficult task.

At this point, the researchers will briefly describe how the management of information resources fits into the Coast Guard organizational picture. Prior to 1981, management of information resources was essentiallly split

between three Headquarters' program offices - OTM: telecommunications division, EEE: electronics engineering division, and FIS: information systems division. OTM was responsible for the general operation and technical performance of the Coast Guard communications system. OTM also played a substantial role as technical advisor and 'marketing' agent for the Office of Operations in translating mission needs into system requirements for EEE to actually acquire. EEE was responsible for all electronics systems in use operationally with specific responsibility for computer/information systems when used in conjunction with these operational systems. FIS was responsible for almost all of the Coast Guard's computer systems and some special data communications systems. FIS was created in the 1950's as primarily a group of experts running and building applications on one central computer.

On 17 March 1981, the Commandant announced his decision to form a new Office of Command, Control and Communication (G-T) which was a direct result of the Decision Support Systems Study Group (DSSSG) report. The most significant point raised in this report was that full integration of the three divisions (OTM, EEE, FIS) was necessary in order to maximize the human and financial resources available to meet the unprecented demands for information technology.

The Office of G-T has been subdivided into several divisions to carry out the functions previously performed by OTM, EEE and FIS; in addition, it has added a new staff component, the Plans and Policy Division (G-TPP). This division's functions include: (1) managing and performing planning functions for the allocation of Coast Guard-wide resources to the G-T program, (2) Information Resources Management (IRM) and (3) Data Resources Management.

A broad definition for Information Resources Management (IRM) is the effective and efficient utilization and control of the resources (i.e. hardware, software and data) necessary to support the information processing requirements of an organization. Information is now seen as a primary resource of an organization in that it allows the members of the organization to answer questions and solve problems, thereby achieving organizational goals. As a primary resource, information needs to be effectively managed just as any other resource (i.e. manpower or capital) is managed by an organization.

As part of the 1981 formation of the Office of Command, Control and Communication, seven 5-year goals were established:

(1) Support Goal - Providing the support and maintenance of installed hardware/software systems occupies one-half of the Office of G-T. The goal is to continue and reduce the cost of this support program management with continual sensitivity to the difference between actions necessary to maintain system performance and those that improve it. The latter are also the responsibiltiy of the support forces, but improvements must compete with other worthwhile projects for resources. The SUPPORT and OPPORTUNITY goals are complementary; either one, if carried to excess, can render the other unachievable.

(2) Acquisition Goal - The Office of G-T is the designated support manager for acquiring major electronic/computational systems, and the goal is to develop an adaptive organization to carry out this mission with project management techniques. 'Acquisition' is the entire function, from concept through deployment and acceptance by the operating and support managers. 'Major' is a case-by-case decision reflecting funds, quantity, time, failure-risk, etc. A large-dollar replacement acquisition may well be a

support/maintenance action, while a small new-technology project could require a 'major acquisition' project.

(3) Opportunity Goal - The Office of G-T, in pursuing the above two goals, will have many opportunities (e.g., in allocation decisions, degree of budget emphasis, use of small discretionary resources, etc.) to tackle other goals of opportunity. For example, within the support/maintenance environment, a familiar target of opportunity is to seek improved performance in the course of a maintenance-justified action. This goal, coupled with the SUPPORT goal, is to aggressively pursue the use of information technology throughout the Coast Guard with these 'opportunity' resources. This service-wide use of information technology is to be in accord with the Commandant's long range goal-seeking efficiency through automation and effectiveness through decision support.

(4) Operation Goal - The Office of G-T is charged with the operation of four facilities to directly support Coast Guard Headquarters: Flag Plot, COMMCEN, National Response Center, and Terminal Center. The goal is to increase the efficiency of this support, particularly through use of information technology, without an increase in resources.

(5) Maritime Communication Policy Goal - The Office of G-T performs a unique function in government forums (domestic and international) in fields within this umbrella subject, such as spectrum management and radio distress policy/technology. The goal is to draw upon the increased technical and senior management resources of the new Office to increase Coast Guard effectiveness in this subject without new resources.

(6) Human Resources Goal - There are two dimensions to the goal for human resources - the general program management of the officer and enlisted specialty groups

under the cognizance of the Office of G-T, and of the personnel assigned to the Office staff at Coast Guard Headquarters. In the former, the goal is efficient use of these people in meeting most of the above goals. This requires considerable integration of previously separate specialties, development of new career patterns and retraining. In the case of Coast Guard Headquarters staff - particularly civilian staff - the goal is one of opportunity for personal growth and development along with the new Office. Such development would be through training and new positions in the new subjects being undertaken.

(7) Information Resources Management - An accesssory function in the support/maintenance area for information systems is 'data base administration'. This forms a core in the broader IRM subject which includes manual (paper) systems and other topics such as office automation, electronic mail, and teleconferencing. Future assignment of these responsibilities to the Office of G-T is intended; the goal is mature development of the initial data base concepts and inclusion within the next several years of the manual systems management into a true IRM role for the Office.

The researchers consider the IRM goal to be the most significant and revolutionary for the Coast Guard of the seven goals established by the Office of G-T. The other goals support already well established programs in the Coast Guard, whereas IRM is essentially a new one, and one which the researchers believe is essential if the 'information corporation of the 1990's' concept is to be achieved. Database technology is one aspect of IRM in that it supports the IRM goal. Therefore, recognizing the importance of IRM for the Coast Guard, the authors' will discuss IRM with regard to the Coast Guard's present situation, planned growth, and how database technology fits into the Coast Guard's IRM goal.

The following paragraphs show a breakdown of some of the factors which the researchers feel will be affecting the Coast Guard in the management of information resources and in the attainment of the above stated goals. Assumptions and constraints concerning these resources are also included.

2. Operational Resources

Table III shows the breakdown of the current level of major Coast Guard resources. Drastic changes in the level of these resources is not expected in the next five years especially in view of the budget situation presently

TABLE III
Coast Guard Operational Resources

CUTTERS: 250
AIRCRAFT: 173
BOATS: 2000
SHORE FACILITIES: 700

facing the Coast Guard. The Coast Guard is acquiring 11 new 270 foot Medium Endurance Cutters which will improve the average age of the fleet somewhat when they replace some of the most aged vessels. With regard to operating resources, the Coast Guard will continue to maintain equipment in the most effective manner and will seek to acquire the most technologically optimal systems available. Improved management of information resources is of high importance in allowing the Coast Guard to keep its operational resources

in the highest level of readiness possible. Discussion of information resources will be included below.

3. Budget

The current level of the budget for the Coast Guard is approximately $2 billion annually. Of this amount, a portion is spent in supporting the Coast Guard information resource requirements (either directly or indirectly). References 4 and 5, outline the budget and resource allocation process to be followed with regard to the Coast Guard's Command, Control and Communications Support program. These policies are consistent with the Coast Guard's Planning, Programming and Budgeting System (PPBS). The summation of these policies is as follows:

(a) Major information systems will be capitalized through Acquisition, Construction and Improvement (AC&I) budgeting. Acqusition will be done by Coast Guard Headquarters to program manager specifications. AC&I billets will be required, and the program manager will need a degree of experise in stating requirements.

(b) Small to medium-scale systems can be supported by Resource Change Proposals (RCP), by a program's internal funds, or by competition for the Headquarters Office of G-T's small, in-house resources.

(c) Small to medium-scale systems would normally be procured by the Office of G-T, although the program manager could do this himself following applicable standards.

(d) Installed "institutional" or Coast Guard-wide official systems will be under the Office of G-T's configuration and standards management.

(e) Operating expense follow-on resources for maintenance of informational management systems must be identified for the users for adequate provisioning.

At any rate, although the Coast Guard's expenditure for information resources will be increased in view of the previously discussed reorganization, the overall Coast Guard budget is not expected to increase in any significant amount, and may actually decrease when the effect of inflation is taken into account. The researchers therefore feel that the Coast Guard MUST make prudent decisions about the hardware and software it acquires today to improve its information resource management, in order that it does not alienate many of its members from future use of this equipment through bad experiences with the systems which are purchased now.

4. Personnel

As mentioned earlier, the Coast Guard is presently made up of about 39,000 military and 6,000 civilian employees. Appendix A shows the latest personnel resources presently assigned to the G-T program. The most important note to make concerning Coast Guard personnel in the information resource arena is that there is currently a critical shortage of qualified people. FY80 estimates identified a shortfall of approximately 90 professionals merely to maintain existing systems and complete development of already initiated systems. The 1981 reorganization only served to identify more shortfalls by opening up more billets in the information resource management program. Training and education are obvious parts of the solution to this problem, but meeting the Coast Guard's needs will require a strategic approach to the overall management and use of personnel.

5. Assumptions and Constraints

The users of the Command, Control and Communication (G-T) program are the entire Coast Guard plus parts of other goverment agencies and the maritime community. This thesis assumes no basic changes in the traditional organization or mission tasking.

There will be little or no expansion in the total resources available to the G-T program, although it will continue to compete vigorously in the budget process.

The core of the G-T program strategy will be one of centrally providing concept design, standards, acquisition and control of major systems or application of major resources throughout the entire Coast Guard.

There is a family of externally imposed requirements which impact on the manner of execution of the G-T program. These are, generally, the original Brooks Act which empowers GSA with sole ADP equipment procurement authority; the myriad of computer-related Federal procurement and Federal Property Management Regulations which implement the Brooks act; other Federal Regulations and various interagency agreements. These factors will be discussed in greater detail as appropriate [Ref. 4].

6. ADP Equipment

a. Existing Tools and Facilities

The following paragraphs contain a brief list of present facilities which in some manner support Coast Guard information processing. In the next section, the researchers will discuss the future information architecture developments being planned for the Coast Guard. The authors discuss the present resources in order to show that these resources consist of outdated equipment put together in disjoint systems, which appear to have evolved in response

to narrowly scoped requirements without consideration for consolidating the facilities which perform similar functions into more effective, general purpose systems. The researchers want to convey the image that Coast Guard information systems are currently behind the times, and can be better managed, but that positive steps are being taken in that Coast Guard is realizing its problems and is taking action to correct the situation.

Communications Stations and Radio Stations - These facilities provide personnel resources and equipment for data transfer shore-to-shore and ship-to-shore.

Electronics Engineering Center and Electronics Engineering Laboratory, Station Alexandria - these provide electronics engineering personnel expertise for field maintenance management, field testing and special systems development.

Command Centers (Communications Centers and Operations Centers) - these provide the nuclei for continuing command and control of all Coast Guard operations; provide a continuity of management of all activities during non-working hours; provide for the management of all electronic information flows into and out of District offices; and they will become increasingly dependent upon information technology.

Leased lines and GTE TELENET Packet Switching Network - Leased lines are procured from the telephone company and provide the equivalent of a pair of wires from point to point - these are now used extensively in landline communications and in the District data network; GTE TELENET is a shared, computerized communications system which the Coast Guard has selected as a critical augmenting resource for servicewide communications needs.

AUTODIN - currently AUTODIN is the backbone of DoD message and store-and-forward data communications, and provides the bulk of inter-district message communications for the Coast Guard; DoD does not charge for this service; DoD is currently designing a packet switching replacement, similar to GTE TELENET network; FAA is designing a similar one, called NADIN II, which the Office of the Secretary of Transporattion (CST) is considering for a Department - wide network.

Model 28 Teletype mechanical communications terminals - these terminals are physically and technologically very old and are becoming extremely difficult to maintain in operation; they have very low capacity (75 bits per second, about 100 words per minute) and possess no capability for intelligent message handling.

VHF/FM and HF Radio Communications Systems - installed radio networks provide the bulk of operational command and control for mobile units, and coordination with civilian marine traffic.

District Northern Telecom DATA 100 Remote Job Entry Terminals - these are high-speed, intelligent terminals which can be programmed in a short time to communicate over leased or dial-up lines with a variety of computer equipment in several protocols; they can also be programmed to execute simple processing locally in a language called RPG; these currently support the bulk of district data processing needs with the Transportation Computer Center (TCC) and local processing.

Word Processing Systems - in use at various locations in the Coast Guard serving primarily word processing and some list processing functions; some have higher processing capability; many can be used in a communications environment, and can be utilized as terminals; wherever appropriate and economical, these systems will remain in place.

Mini and microcomputers - various models exist throughout the Coast Guard supporting important functions for local needs; as long as these needs are being satisfied in an economic manner, these systems will remain in place.

Operations Computer Center (OCC) and Computer Centers at Aircraft Repair & Supply Center (ARSC), Academy, Coast Guard Yard, and Supply Center at Brooklyn, NY - These major processing centers will provide special processing needs to programs or facilities.

DOT and Other Governmental Computer Service Centers (i.e. TCC, Transportation System Center, FAA Aeronautical Center, National Institutes of Health, Walter Reed Army Hospital, Navy Fleet Numerical Oceanography Center) - these systems provide general-purpose computing resources to the Coast Guard, on a reimburseable basis, shared by other agencies.

Commercial Timesharing - Competitive single-user contracts, GSA Multiple-award schedule contracts, and DOT-wide multi-user contracts under negotiation; the advantages of this type of service are first, higher levels of analytical assistance and second, usually more reliable on-line, interactive service; this better service carries a higher cost.

For the interested reader, Reference 4 contains an inventory of the software systems cuurently in use throughout the Coast Guard.

b. Projected Coast Guard Information Resources Architecture

The purpose of this section is to give the reader an understanding of how the Coast Guard's existing resources and facilities will be expanded, reconfigured and utilized in the next 5-10 years. References 4 and 6 provide a more detailed description of this topic.

The words information and data have various meanings and sometimes the words are used interchangeably. The Coast Guard has decided to refer to data as static entities with assigned definitions. Information is the enhanced meaning of value associated with these data when they are somehow processed or utilized by human or machine intelligence. As an introduction to the Coast Guard's future information resource (IR) architecture, it is useful to think about information in three global ways in which the organization uses it - namely hierarchical, transactional and local information.

Figure 2.1 shows hierarchical information flowing from the smallest unit, to the top of the Coast Guard, and ultimately beyond to both Congressional and Executive recipients. Hierarchical information supports the management control and strategic control functions of the organization.

Figure 2.2 shows transactional information, based upon data groupings that flow intact from point-to-point in the organization and usually cause or support rapid activity.

Local information is anything and everything that any individual or organizational element chooses to utilize when it is not institutionally required to do so. Strategic design of the information architecture recognizes the use and need of local information for maximum user benefit. The architecture must provide some reasonable facility for each level of the organization to meet its local information needs.

Figure 2.3 shows the merger of traditional record communications and future data communications. The various networks are simply treated as resources to move data, and it becomes a network management task to insure that operational transactions are handled with adequate

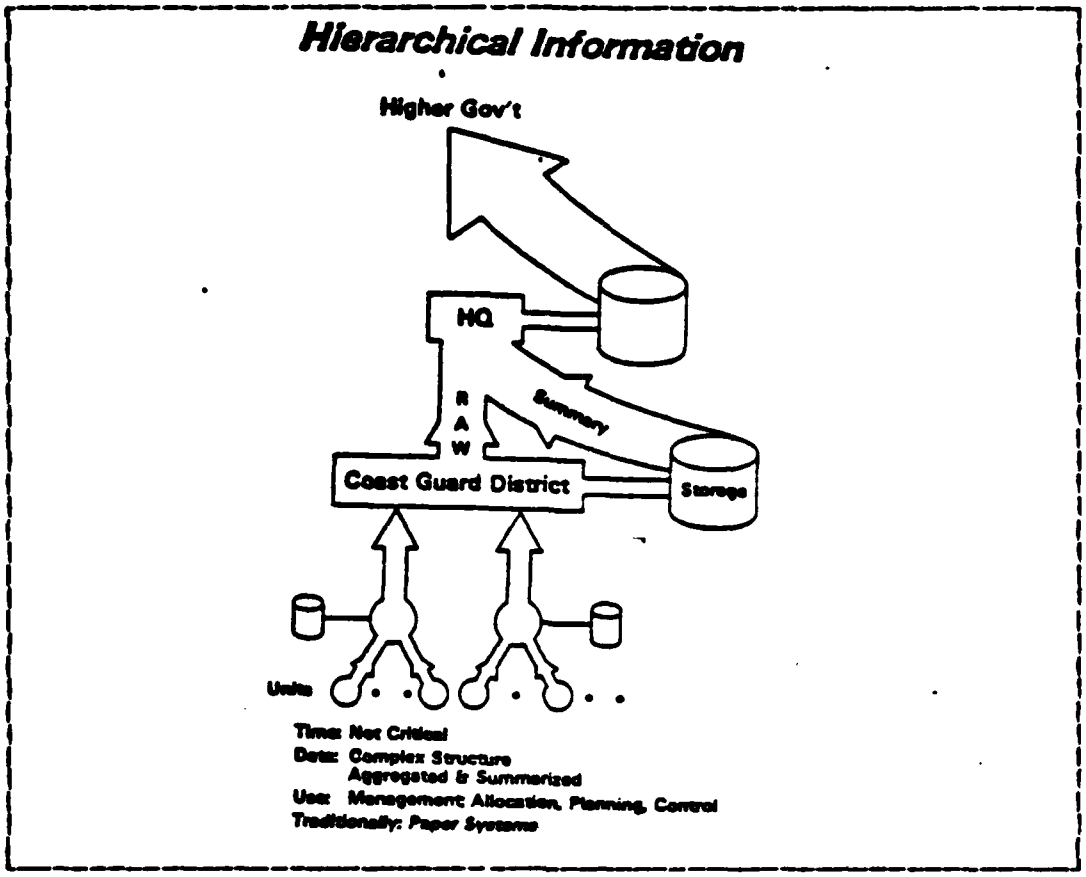

Figure 2.1 Hierarchical Information.

priority. By implementing most data communications as an enhanced version of record communications, upward-compatible and graduated growth is achieved, as well as the loosest possible dependence between physically separate information resource elements.

An important feature of the 'information resources architecture is the new Coast Guard Standard Terminal. This piece of equipment has been chosen to be the primary 'standardized' computer hardware device in the Coast Guard for the 1980's. It can be viewed as a relatively inexpensive device with three essential features:

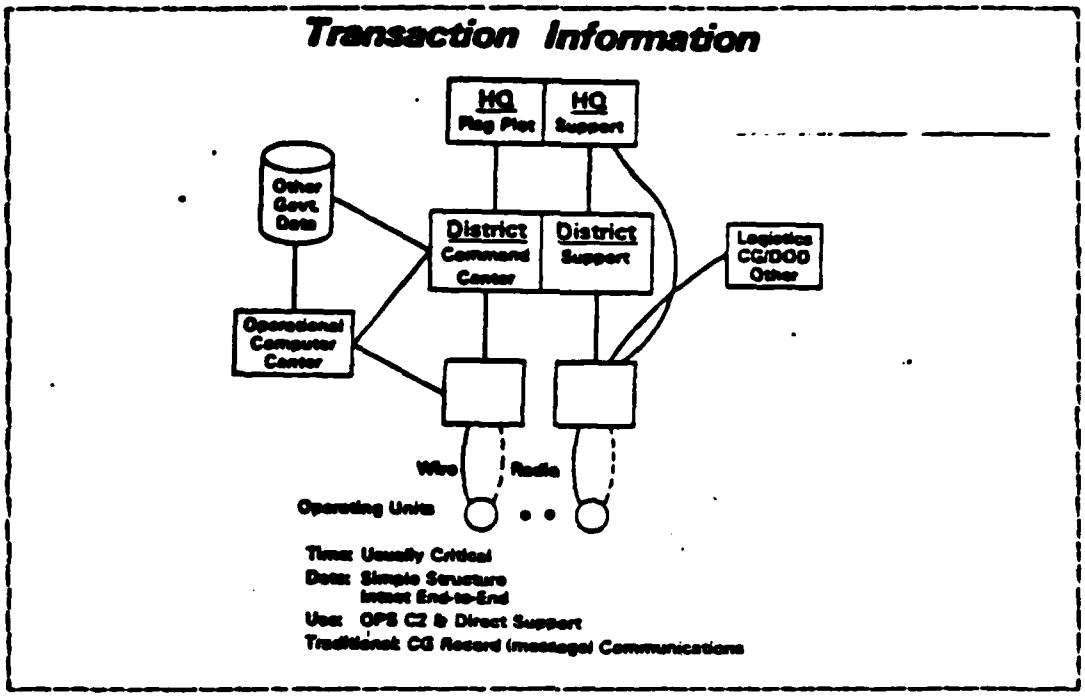

Figure 2.2 Transactional Information.

(a) Communication - the terminals can communicate over all of the channels planned for Coast Guard use, with software and hardware supplied by the vendor. They can also be clustered among themselves for sharing of processors and data, with the wire interconnection being the communication channel.

(b) Human Interface - With its CRT screen, software and keyboard it can easily be used by all levels of Coast Guard personnel. The special function keys and electronic generation of fill-in-the-blank forms make it capable of being viewed as only marginally different from paper and pencil, functioning as the source of data for all three types of information.

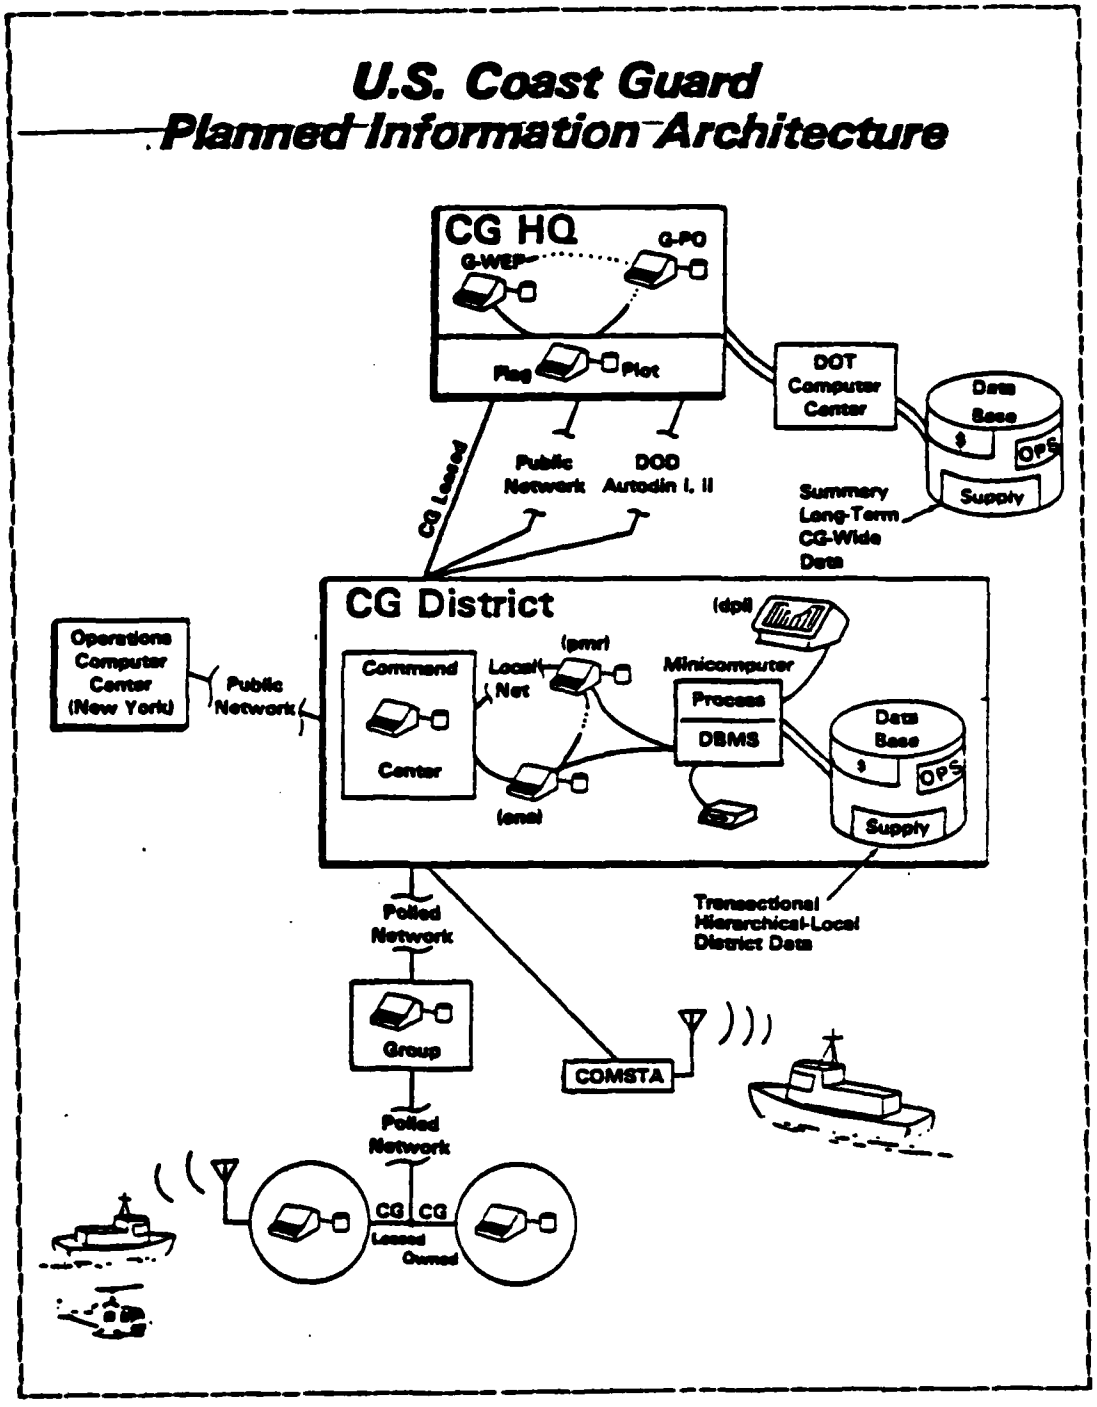

Figure 2.3 U.S. Coast Guard Planned Information Architecture.

(c) Process/Storage - The medium sized microprocessor (Intel 8086) present in each terminal gives them the ability to meet nearly all local information needs for the 1980's. In addition, standing alone it can perform many of the formal transactional and hierarchical processing functions [Ref. 4].

Three categories of communications networks will be used by the Coast Guard in the 1980-1990 time frame. These are an outgrowth of the Coast Guard's command hierarchy and the record message store and forward communication mode. These three categories of networks are national, regional and local networks as shown schematically in Figure 2.4.

NATIONAL NETWORK - The National Network is a generic term which encompasses those means of interconnecting major Coast Guard national switching nodes. (The major switching nodes are presently Districts, communications and radio stations. Any possible combination of unconnected networks may be chosen in the future but today's Coast Guard 'National Network' consists of SARLANT/SARPAC (Coast Guard search and rescue systems), AUTODIN, Secure Command and Control Network, (all three interfaced by manual torn-tape operation), TELENET, and the leased Information Network.

REGIONAL NETWORK - the regional network will provide connection and information exchange between the regional nodes and local nodes (end user terminals or local area networks). This is analagous to the District or Group Teletype net. (Both National and Regional Networks may ultimately end up using the same network technolgy as ARPANET.)

LOCAL NETWORK - The local network will provide connection and information distribution (exchange) between the

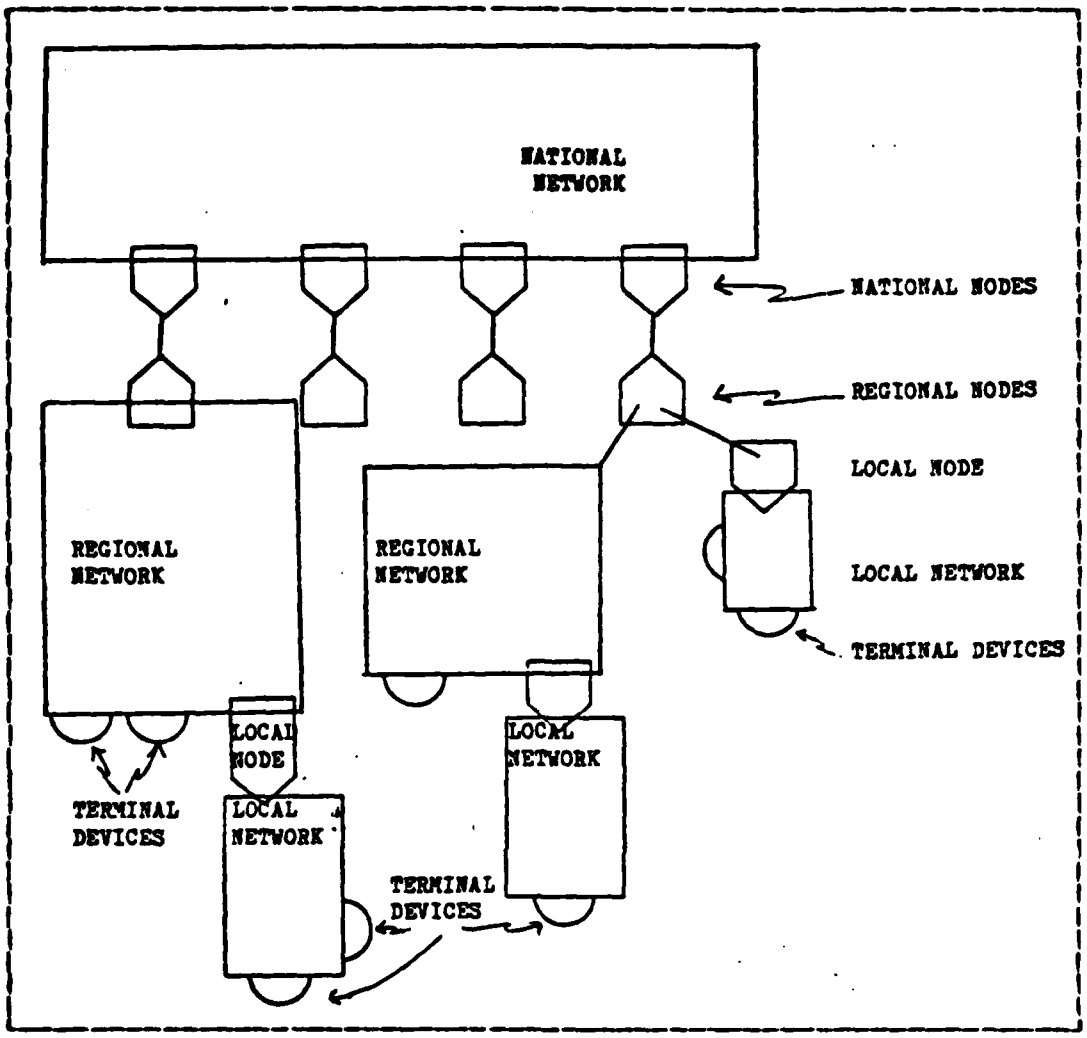

Figure 2.4 Coast Guard NETWORKS.

local network, or local area network (LAN) as they are popularly called and connect users intra-building, intra-facility or intra-industrial park. Examples of present LAN's are a shared logic word processing system connecting all offices in a District Office, a loop of Model 28 Teletypes controlled by an HP9825 or on-base administrative teletype systems. In the future these

LAN's will be much more capable of handling large volumes of information and a heterogeneous mix of users (such as Wang Net) [Ref. 6].

The Coast Guard's last main feature of the information resource architecture is that of data resource management. This concept deals with the fact that data is an important resource of an organization and that the management (the collection, storing and processing) of the data is a crucial determinant for the success of any planned information resource architecture of that organization.

The focus of this thesis is concentrated in looking at how the Coast Guard addresses the problem of information resource management via DBMS technology. More will be discussed concerning planned growth for Coast Guard information systems in Chapter VI.

C. TECHNOLOGY FRAMEWORK

The researchers will now discuss how the current 'state-of-the-art' has evolved in computer or information systems environment specifically with regard to data resource management.

When computers were first introduced in the 1950's hardware cost was the most significant factor. The computers were relatively slow (especially when compared with today) and designed primarily for the solution of sophisticated mathematical number crunching problems. As the speed of computers improved, the ability of computers to process a large amount of repetitive data was recognized and computers began to be utilized increasingly for 'management-related' processing (i.e. payrolls, personnel, inventories, etc.). But still throughout the 1960's, the computer was viewed as a large, expensive machine kept in a central location and operated by specially trained, dedicated personnel.

Efficiency of the machine was stressed due to its high cost, and if it just so happened that the reports a manager was getting were not responsive to his needs, then getting a new report created was not an easy task. This was because one had to get the analysts and programmers together from the data processing department, tell them what was necessary, and have them write a new program to do it. If it didn't turn out the correct way, a person might not go back to change it because of all the trouble to do so. Data in this type of situation was probably centrally located, but as discussed, getting to the data was often not very easy.

With the revolution of integrated circuits and silicon chips, the speed of computers increased greatly while the price of computers decreased even greater. Mini and then microcomputers were introduced with recognition that not all situations require a large mainframe and often it is advantageous to have several smaller computers instead of one great big one. 'Power to the people' has been used to describe this philosophy of providing a small, low-priced computer directly to a user in order to be more responsive to his particular needs. As data sharing needs grow, these independent smaller computers can be interconnected to form networks or distributed systems. 'EFFECTIVENESS' (utilizing the computer to achieve optimal assistance to users in meeting their needs) rather than 'EFFICIENCY' (achieving optimal utilization of the computer) becomes the important factor with regard to this type of computer architecture.

But, as is the case in a changing world, when you solve old problems, new ones often appear. As the use of computers grew, and more people used them, more programmers appeared and often different sections or departments in the same organization developed their own programs. This was in part because the organization wanted to make the computer more responsive and give the end users better control, but

problems with data redundancy and consistency appeared. Each department began maintaining their own set of data, and often this data was repeated in several departments.

This led to the development of data base management systems in order to allow sharing of data, to control data integrity, to allow greater data independence, and to improve data recoverability [Ref. 7]. These data base management systems are sophisticated software products which provide for effective management of data resources, and when combined with a special query language can be utilized by a non-specialist user to quickly access the data base in a responsive manner. Most of these DBMSs are designed for a single computer. The Coast Guard has decided on a policy of distribution of computer resources in which the information resources (hardware, software and data) are placed in the hands of the commander who exercises responsibity and control over the forces they support. Therefore the Coast Guard faces a problem of managing data resources given a distribution of those resources.

III. ORGANIZATIONAL IMPLICATIONS

A. GENERAL

In this chapter, the researchers would like to take a look at the overall direction the Coast Guard is headed with respect to information systems and database technology, and then relate this direction with what is happening in the information systems industry at this time. The authors will discuss some organizational impacts which may be expected as a result of the actions taken by the Command, Control, and Communication program of the Coast Guard. Finally, the researchers will provide some recommendations for planning for and dealing with these impacts in order to minimize any negative effects.

B. DIRECTION COAST GUARD IS HEADED

The direction the Coast Guard appears to be taking in the information resource management (IRM) area seems to be one of greatly increasing the importance of and improving IRM throughout the organization. This is evidenced by the reorganization and consolidation of the organizational elements involved in this program, and the commitment of spending more dollars on upgrading IRM capabilites. The major goals of the Command, Control and Communications program were outlined in Chapter II of this thesis, and the Coast Guard appears to be pursuing these goals in a systematic manner with controlled growth, coordinated acquisition policies, and expansion of education and training opportunities for personnel in this area. The researchers feel very strongly that the overall direction of improving the effectiveness of IRM by the Coast Guard is highly valuable,

somewhat overdue, and will put the organization in a better position to overall manage its total resources. The use of data base technology is just one avenue for increasing the effectiveness of an organization's IRM. What are other related technologies, and how does the Coast Guard stand with respect to these concepts?

Reference 9 discusses some of the new technologies in the IRM field which are available at this time. These include:

1. Text Handling

Most organizations have realized the importance of having most printed text in machine readable form in order to speed up the production, update and correction of printed documents. Often referred to as word processing systems, these type systems can be obtained as a microprocessor dedicated to this task or as a feature of a general purpose computer system. The Coast Guard has implemented word processing systems in varying degrees throughout the organization. These range from the simple, stand-alone dedicated systems to sophisticated, interconnected, multiuser systems. The Office of G-T has decided that where cost-beneficial, these systems can be continued. The standard terminal is an example of a general purpose system which can be used for word processing.

One common extension of word processors is the ability to electronically send documents between processors in different locations in lieu of having to manually transport the hardcopy. Therefore, it is necessary for communication standards to be established in order that the machines can communicate with one another over the established protocols. Standards for interfacing dedicated word processors with larger general purpose computer systems should also be required. The Coast Guard has recognized the

importance of such standards and has promulgated the protocols to be utilized by the affected equipment in Reference 4.

2. Messaging

This concept is related to the feature discussed in word processors on electronically sending documents, but should be considered as a more general feature than just sending documents. Messaging conveys the idea of using the information system as a means to send communication between individuals without having person-to-person or telephone contact. This concept is often referred to as electronic mail. An example of messaging would be a system which allows text messages to be sent from terminal to terminal enabling a person to use the terminal directly for reviewing incoming messages thereby cutting down on actual paperwork and speeding up the communication in the process (i.e. no busy signals). Another example might be a variation which would allow recorded voice messages to be dispersed by the computer, similar to what a telephone answering machine does for incoming calls, allowing the message to be sent to more than one destination. This would relieve the originator of the time consuming task of calling a lot of other people and getting many busy signals.

The idea of messaging can be extended to having a computer conference where the messages are sent back and forth in an online mode, thus alleviating the need to have actual personnel travel to one geographic site in order to conduct the conference. Since some people find conferencing by computer too impersonal and hard to follow, the next extension is to transmit video signals, or teleconferencing. The Coast Guard presently has some systems capable of electronic mail type messaging, although long distance systems are not in place at this time. Continued growth in this

area can be expected given the rising costs of transporting personnel and the decreasing costs of computer and teleconferencing systems.

3. Distributed Systems and Networks

The concept of distributing computing resources (hardware, software, and data) is currently a major topic throughout the computer industry. A tremendous amount of research is being conducted in looking at how computer networks and distributed systems can improve the effectiveness and user satisfaction of today's information systems. The researchers will use the following definition of a distributed system: first, the system should possess two or more geographically separated processors, second, the processors should be linked, and third, the processors should serve a single organizational entity. Computer networks can be generally defined as a connection of two or more computers which are able to communicate between one another. Examples of computer networks include the Advanced Research Projects Agency (ARPA) Network and IBM's System Network Architecture (SNA). Note that a distributed system is a computer network, but not every computer network is a distributed system [Ref. 8].

As previously mentioned, the Coast Guard has decided on a principle of distribution of computing resources to some extent. The Coast Guard will build upon message store-and-forward communication as its principal mode of interconnection. True computer processor interconnection in the style of IBM's SNA will be discouraged for the near term. The reason for this policy is based on the principle of prudent evolution. There is no forseen great payback to the Coast Guard for implementing computer networking of the SNA-type systems at this time, especially considering the high technical-skill and supports costs of SNA [Ref. 4]. An

important point the Coast Guard should consider when implementing its present systems is to design for future change and recognize that computer networking and distributed systems are the direction information systems are headed. Therefore, the Coast Guard should strive for compatibility and standardization as much as possible in order to have systems capable of becoming networks when such a configuration is considered economically feasible and desirable.

4. Decision Support Systems (DSS)

Decision Support System is also currently a highly discussed topic in the computer field today. Some consider it to be the latest 'buzz-word' in the industry, picking up where Management Information Systems (MIS) left off. Many critics feel that DSSs are no different than MISs and that the new term has simply evolved so that salesmen can sell more systems. But many others, including the researchers, feel that DSS is an independent concept, and that it refers to building an information system to be readily utilized primarily to support the decisions which an individual in the organization has to make. Some components of a DSS would probably include: an interactive query or dialogue capability, a database management system, a modeling capability, and a graphics capability. Some observers consider the modeling capability as the feature which distinguishes a DSS from an MIS, while others believe its distinction to be that it is primarily designed for the particular indivdual or job position for whose decisions it is meant to support.

At any rate, the researchers feel that the Coast Guard will most likely be investing in decision support systems in the future, as these systems seem to fall nicely into the goals and objectives which the Coast Guard has set for its Command, Control and Communications program. The researchers are aware of at least one DSS in use in the

Coast Guard which is installed in the First Coast Guard District in support of the operations center controller for search and rescue decision making.

5. Data Base Management Systems

This thesis discusses in some detail the status of, goals and objectives of database management systems. The researchers believe that the Coast Guard is realizing the importance of DBMS, and specific recommendations regarding a strategy for implementing database technology are included in Chapter VI.

The Coast Guard is in varying stages of installing database technology. The Marine Safety Information System (MSIS) has been in development for over 10 years and will contain a sophisticated DBMS. With regard to the District minicomputer procurement, the Wilson-Hill study, [Ref. 10], has been viewed as a 'jumping off' point for developing the logical data base design to be utilized by the Coast Guard Districts. A pilot project for designing the logical database is underway in the First Coast Guard District and it has been reported that they are attempting to integrate the Districts' database requirements for this design with those of other Coast Guard organizational elements. This will hopefully ensure a standard Coast Guard data dictionary, which the researchers believe is essential for the success of this project. A target date of April 1983 has been projected for the initial draft of this design.

The authors will next discuss some of the factors which may influence the Coast Guard in pursuit of the goals of the Command, Control and Communication program.

C. INFLUENCING FACTORS

The researchers believe that the following factors - organizational change, resistance to change, technological change, government procurement regulations, and support requirements - will most likely have some impact on the Coast Guard as a result of installing information in database management systems. The authors will proceed to discuss these implications in more detail and describe what effects they may cause in an organization.

The first factor which may have an impact on the Coast Guard in implementing database technology and other information resource management goals, in fact, already has occurred - organizational change. As was mentioned, one of the initial events which triggered the Coast Guard's emphasis on information and data resource management was the reorganization of the Headquarter's divisions of OTM, EEE and FIS into the Office of G-T. But there are other less evident ways in which the introduction of data base technology may change the organization. These changes will be in the way in which particular individuals perform their job and also in the informal organizational structure of the Coast Guard. Hopefully, the new systems will improve or make easier the tasks which Coast Guard personnel carry out. This is one of the reasons the new systems were introduced in the first place. The informal structure of the organization refers to the fact that in order to get things accomplished, the formal organizational structure is not always strictly followed, and informal working relationships often crop up. By introducing new technology, changes in this informal structure can be expected as new users of the systems interact with other users who either work in a similar manner with the system, or are more knowledgeable. The effects of the changes can be beneficial or harmful depending upon the productivity of these relationships.

The next factor which might influence the Coast Guard in this area is resistance to change. By this, the researchers are referring to the idea that no matter how much one prepares for changes in the organization, there should be some expectation that not every one is going to accept the new system (i.e. there will be resistance). These may appear such as a senior enlisted radioman who becomes disgruntled at the thought of having a computer take over the job he used to do; or providing a computerized system, such as computerized inventory system, to an employee who developed the old system and is highly familiar with it. These types of personnel can be expected to 'buck' the new system and therefore their reactions should be considered in order to make the system work, because without their interest, involvement and support, the new system will most likely fail.

Other factors which should be considered are the impact of new technology, and the effect of Government procurement policies and regulations. Technology in the computer industry seems to continue to improve upon itself at an amazing rate. Yet forecasters predict that technology in the computer industry will contiue to get better, and therefore the Coast Guard must be aware of this improving environment and be ready to take advantage of it. The authors believe that the Coast Guard should strive to acquire homogeneous systems whenever possible. This will be discussed in Chapter VI. It is noted that the Brooks Act seems directly contrary to this goal in the interest of stimulating competition in the industry. Other regulations such as GSA having sole procurement authority over ADP equipment must be taken into account when planning for and acquiring new systems. One might think that the bureaucratic procurement regulations of the Federal government would be something that the Coast Guard should be well aware

of; but, as discussed, for the most part the Coast Guard is really only now purchasing its own computer systems having leased timesharing service primarily in the past. Therefore the Coast Guard should take heed of the lessons learned by other Federal agencies, particularly the Department of Defense, with regard to the long lead times involved in procuring ADP equipment and plan early to spend a considerable amount of effort in order to keep up-to-date systems. Another major factor the Coast Guard must contend with is the support requirements necessary for these systems to properly function. These requirements include: personnel, facilties, training, maintenance (of both hardware and software), security, and perhaps most importantly, dollars. The first couple of requirements are somewhat obvious in that everyone should realize that computer systems need people to run them, buildings to put them in, cables to run between them, etc. The aspect of training can often be overlooked, especially when one considers the importance of follow-on training. Table IV [Ref. 11] is an example of the training needs of the database administrator and staff in a DBMS environment. This list indicates the extensive training required to ensure proper operation of the system, and also the need to have competent professionals working with the system in order to educate the non-sophisticated user (of which there is likely to be a majority of in the Coast Guard, at least initially).

Maintenance of hardware is not considered to be an overwhelming constraint, as hardware systems have become more and more reliable, but backup procedures must be planned for in the case of hardware failures. The biggest factor in the maintenance area is that of software maintenance. Here again it must be pointed out that the Coast Guard is relatively new in the computer business and therefore there is not a lot of experience in dealing with computer systems.

TABLE IV

Duties of DBA and Staff in a DBMS Environment

DESIGN AND ADMINSISTRATION

Define schemas and subschemas

Select and maintain data model

Select and maintain DBMS software

ADMINISTRATION

Liaison with users

Training and assistance of users

OPERATION

Formulate and enforce procedures for security & privacy

Initiate and enforce procedures for recovery & integrity

Define, create, update, and retire data

MONITORING

Measure and monitor performance of DBMS resources

Log and monitor usage of DBMS resources

Schedule usage

Monitor security threats

As these systems become available, more and more programs (software) will inevitably be written. Therefore software engineering principals need to be stressed early on so that the software written today can be maintained tomorrow. Experience from DoD has shown that software maintenance is a tremendous concern and is probably the largest expenditure on DoD computer systems today. DBMS should aid in this area by introducing the concept of data independence into programs written for Coast Guard computers equipped with such systems.

Security refers to the concept that the physical components must be protected from threats such as theft, vandalism, and natural disasters, as well as protecting the privacy and confidentially of the data stored within the system. Security should be a primary consideration of any computer installation and DBMS should assist by providing access control, integrity checks, security and auditing features.

The last item mentioned is probably a result of taking care of some of the other factors which were discussed, and this implies that it takes money to properly install and operate an information system. The Coast Guard must realize that it is going to take a considerable amount of additional funding in order to properly operate, maintain and keep current the computer information systems it is installing today.

D. MINIMIZING NEGATIVE IMPLICATIONS

The folllowing are some general recommendations and comments which the researchers feel would assist the Coast Guard in minimizing negative implications associated with introducing information systems.

The first comment is that the Coast Guard must design for change. The Coast Guard must take into account the effect which charging technology will have on its information systems, as well as predict that user needs will most likely expand and therefore change as they become more familiar with the systems as time goes on. Therefore, the Coast Guard needs to design systems which are flexible and changeable. This concept has been repeatedly stressed in software engineering, of which modularization and changeability are key principles.

Along with this 'design for change' concept goes the principle of requiring good documentation for all the systems which are acquired. This concept is especially crucial in a military organization where personnel rotation policies guarantee that almost all the people who are present when a new system is installed will probably be gone after 2-4 years. This factor indicates that the documentation of the system had better be sufficient and up-to-date in order for new personnel to understand, operate and maintain the system. The Coast Guard can learn another lesson from DoD, which has reported many instances of poor documentation of systems which after a short while, led to a situation where no one really knew how the system worked. Therefore, if the Coast Guard expects to be able to understand and maintain its systems, good documentation is vital.

Another factor which will help minimize negative implications is for the Coast Guard to maintain top level support of its information management goals. If top level support exists, then the chances of success are much greater; however, if this support is not evident, then failure is almost assured. This statement seems most obvious, but it appears to the researchers that it is often overlooked. One problem involved here may be that the technology associated with the information systems which are being introduced today is probably much different than what most top level managers are experienced with. This could lead to intimidation of the senior level supervisors to the point that they will avoid using computers in order that they do not appear ignorant in front of their subordinates. If the senior level managers are not using the systems, then their subordinates may not feel inclined to use them and eventually the funds to support the systems might be cut or even shut off completely. Therefore top level support can be seen to be an important ingredient to the success of a system, but how can this senior level support be achieved?

This question leads into the next point which is the importance of human needs. The researchers feel that one of the best ways in which to achieve top level support for a system is to have them involved with and using the system. It has already been mentioned that senior level personnel may be intimidated by a computer, and it should be noted that they probably don't have a lot of extra time to spend learning how to use the system. Therefore one way to get them to use it is to make the system easy to learn and easy to use. Reference 12 points out some excellent general design principles which should characterize a user oriented interactive system. These principles include:

1. Self explanatory - the system should display sufficient information to the user without having to go to another source for explanation.
2. Self helping - the system should provide a checking of user inputs and reminders or further instructions on user request.
3. Simple interfacing - the system only requires actions which are short, simple and obvious.
4. Interaction by anticipation - the system anticipates all possible desires of the user in advance and provides prompts or menus so that the user can select the desired option in lieu of explicitly specifying it.
5. Optional verbosity - the system provides another level(s) of detail so that the user can obtain more detailed information on a particular item, if so desired [Ref. 12].

By utilizing these design principles, most systems would be easier to learn by most users. But even with these features, some personnel are going to require a lot of 'hand-holding' and training to be a satisfied user. Another important item to remember when designing a system is to always include representatives of the end users in the design committee. This will often turn up obvious problems and peculiar needs by the end user representatives which would not have turned up from the other design committee members and thereby saving a lot of expense altering the

system just after installation. This membership on the design committee also helps promote a feeling of early involvement with the new system which can help cut down on the resistance to change.

The bottom line is that system designers must take into account the needs, capabilities, and shortcomings of the human user in order that the system be responsive to user requirements, easy to learn and easy to use.

E. SUMMARY

In summary, this chapter pointed out several organizational implications such as organizational change, resistance to change, and support requirements which the Coast Guard should consider when planning for information systems. The researchers believe that the Coast Guard is headed in the right direction with its information and data resource management policies, and that good implementation practices are being followed by the Coast Guard. Coast Guard system designers must take into account human needs and shortcomings so that the information systems it employs are usable as well as efficiency and productivity improvements.

The following chapters of this thesis will proceed to discuss data/database management techniques and trends, and then investigate the alternatives open to the Coast Guard. The researchers will match the Coast Guard's particular needs against the various DBMS alternatives available and make recommendations as to the direction the Coast Guard should take in this area.

IV. FUTURE TECHNOLOGY DEVELOPMENTS

A. GENERAL

Before proceeding further, it is advantageous to explore the new technologies that are evolving today. In the field of computing, it is not surprising to see advances in technology that make the equipment, policies, and practices obsolete in just five or ten years. For example, the computer industry's architecture has progressed through four generations of machines in thirty years - from electrical accounting machines (EAM) to vacuum tubes to transistors to integrated circuits. Similar growth has been experienced in the logical view and management of data; flat files with record-at-a-time serial access, to indexed sequential files capable of random access to data, to sophisicated data bases. The physical level is no exception either as 1K bit chips were used in the late 1960's while today 64K bit chips are standard in most machines with 256K bit chips available on some.

Any strategy must take into consideration the impact of technology. Therefore, this chapter will discuss those items that today promise to have the most impact on database management systems (DBMS) in the next five to ten year time frame. The most promising areas are in mass storage hardware and software, memory, database machines, and distributed database management systems.

Although the conceptual views of DBMSs have gone through technological changes, these views have more or less settled down into three basic models - hierarchical, network or plex, and relational. These models will be discussed in a later chapter.

B. ASSOCIATIVE COMPUTER ARCHITECTURE

Traditionally, access to data in main memory has been accomplished by hardware address. This is a necessary consequence of the von Neumann computer architecture. However, the most desirable way to access non-numeric data is by value. Programmers and users of computer systems, especially those employing DBMSs, think of problems and solutions in terms of value, not hardware address. As the computer industry matured, a number of methods were used to convert an external value into an internal computer or peripheral device address. Most of these methods are still in use today, e.g. sequential, indexed, hashed, and set access methods.

As data bases became larger and larger, and access was required to more and more data, computer performance was kept to an acceptable level by brute-force improvements in hardware speed; i.e. integrated circuits, more bits per chip, etc. The one basic element that never changed was the von Neumann architecture. Today, computer scientists are beginning to rethink the foundation on which computer science has been built and are investigating contents addressable memory, or associative storage.

Associative processing systems generally have the following five basic components - data register, mask register, data array, word select register, and search results register. As shown in Figure 4.1 [Ref. 13] the Data Register contains the value to be searched, while the Mask Register indicates what part of the Data Register is to be used for the search. The Data Array contains the data to be searched, the Word Select Register indicates which words in the Data Array are to be checked, and the Search Results Register contains a '1' bit for those values in which there was a match. A Multiple Match Resolver indicates which

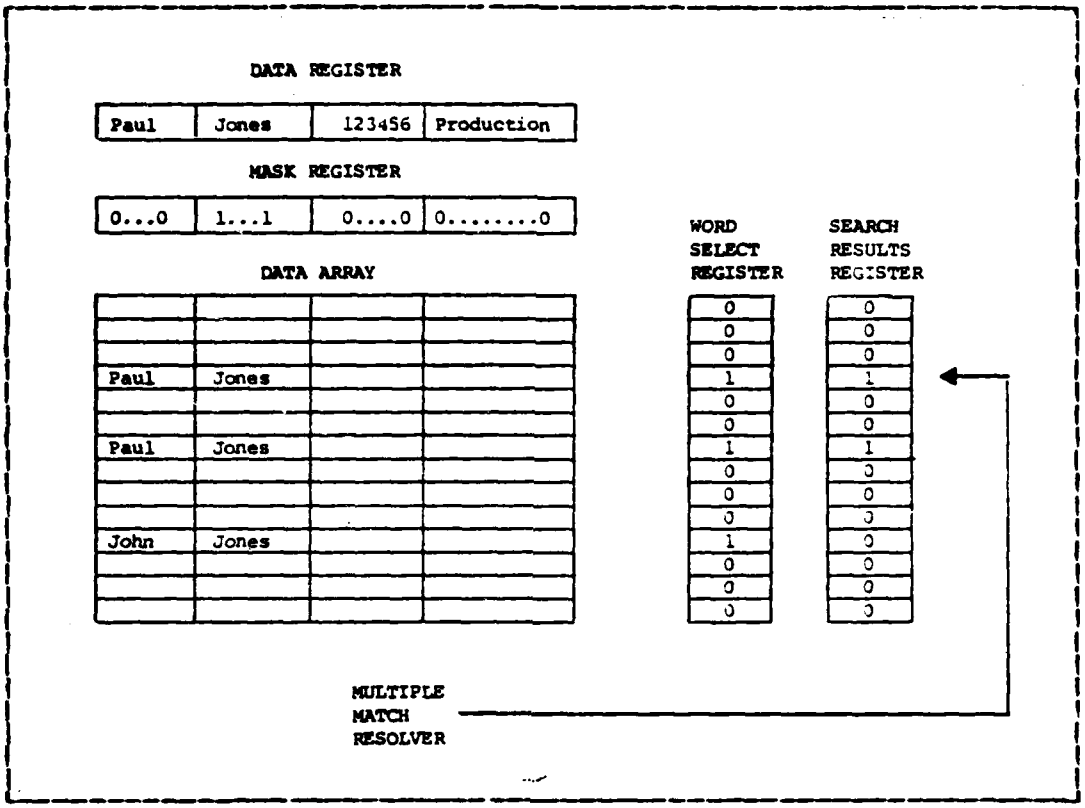

Figure 4.1 Associative System.

value was found first. From a casual inspection of Figure 4.1 it is apparent that the size of the Data Array contains an upper bound which limits the size of a data base that can be searched using parallel associative store in a single storage cycle. What is typically being done is a sequential associative search on large portions of sequentially stored data. This gives the logical effect of parallel associative store if not the speed.

This type of architecture is being experimented with today in a number of ways. One idea is to build an intelligent controller based upon associative processing, and place it between a general purpose computer (which contains the

DBMS software) and the mass secondary storage device of the data base. The intelligent controller would accept a parsed request from the computer, and process the transaction against the data base. Another approach is to incorporate this architecture into the computer itself. More will be said of this later.

C. MEMORY

The driving force behind the rapid advances in the computer industry has been microelectronics. The growth depicted in Figure 4.2 [Ref. 14] was predicted in 1964 by Gordon Moore, the director of research at Fairchild Semicondutor. To date the curve has been pretty accurate. Whether or not this trend continues is speculative, but current technology has not reached its limit with regards to physical laws. Further, as the number of components per circuit increases, the cost per bit drops; one cent per bit using 1K chips in the late 1960's to one tenth of a cent per bit using 64K chips in 1979. These developments, especially with respect to main memory storage, have great consequence for DBMSs, because main memory serves as a buffer area to operate secondary mass storage. The larger the main memory, the fewer I/O interrupts required to accomplish a process. Hence, performance requirements for large data bases can be kept to within acceptable limits as sizable portions of a data base can be kept in main memory without having the processor page data in and out of secondary mass storage.

Secondary storage devices too have been developing at a similar rate. Charged coupled devices (CCD) are available in 64K chips with 256K chips soon becoming available. Bubble memory is available, but it is not as popular as CCD, perhaps because it has slower access times for large amounts of data. Research is still on-going in this area.

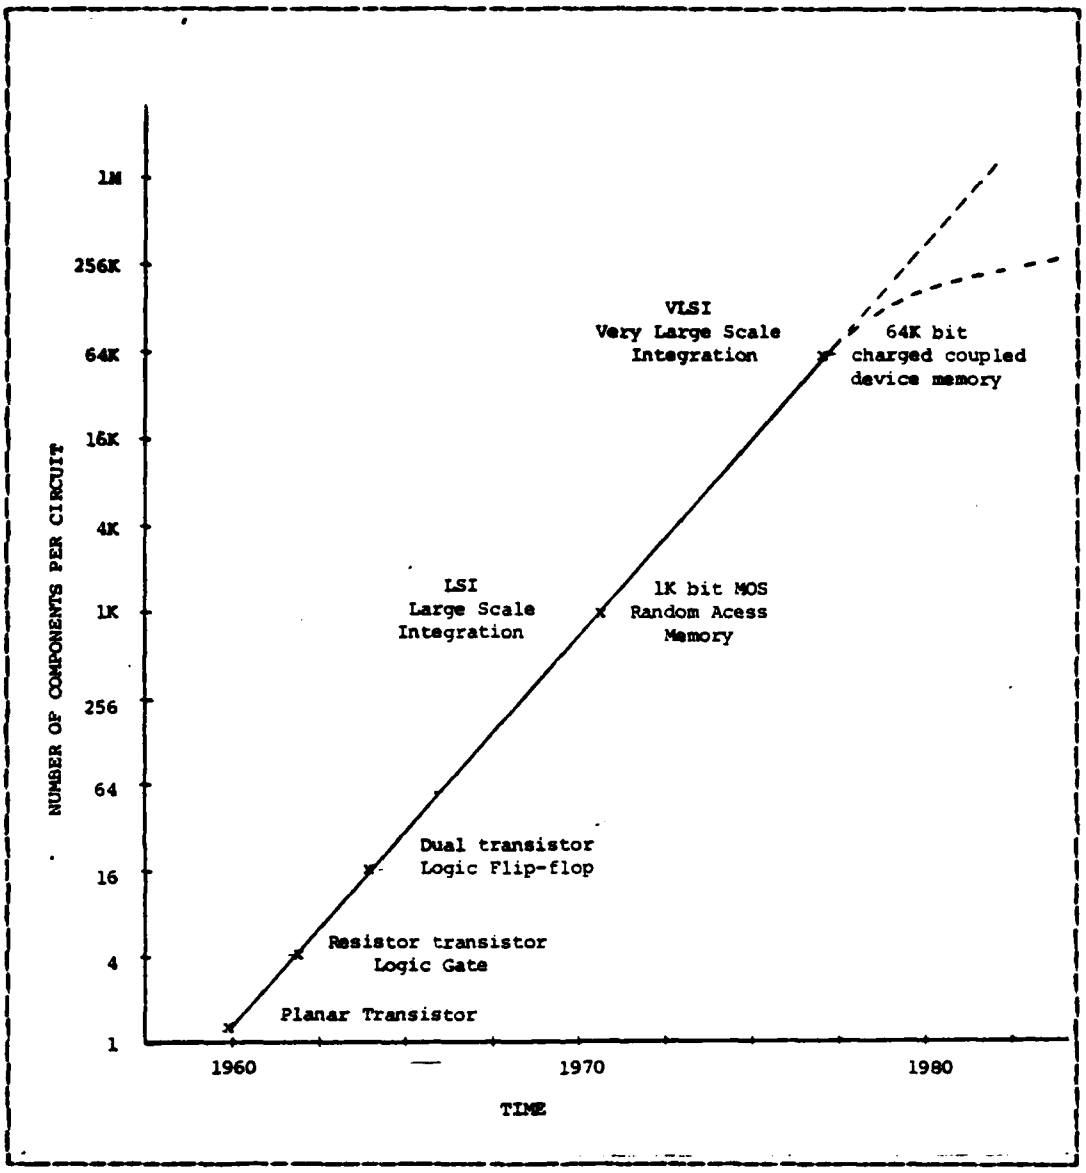

Figure 4.2 Growth in Micoelectronics.

D. DATABASE MACHINES

Today when people talk of database machines, there is no common agreement on the architecture. In fact anything from a master-slave configuration (a general purpose processor

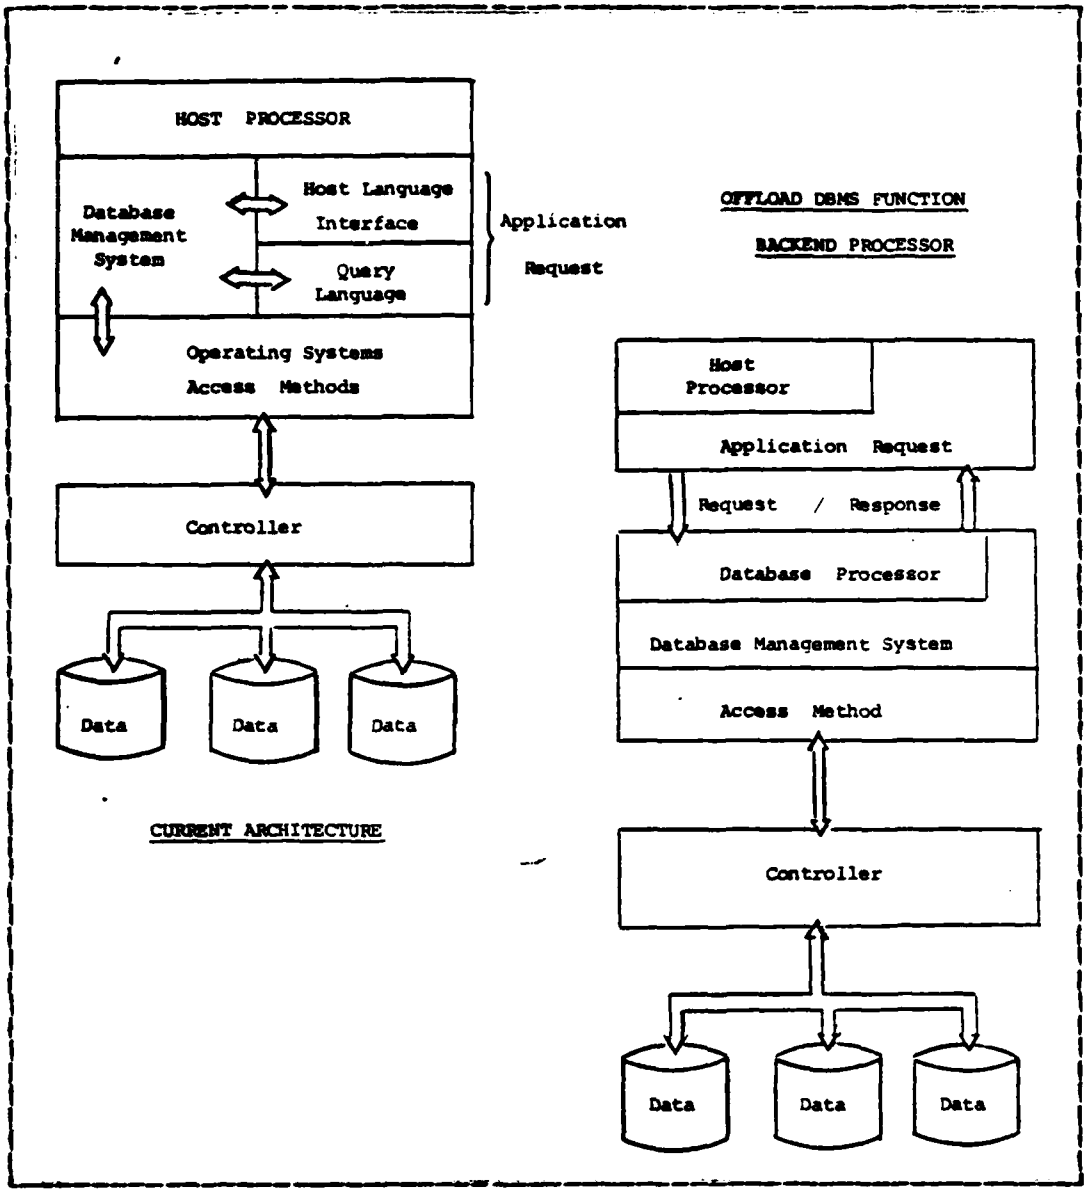

Figure 4.3 Current DBMS Architectures.

which performs all of the data base activities) to a special purpose back-end with intelligent controller, is considered

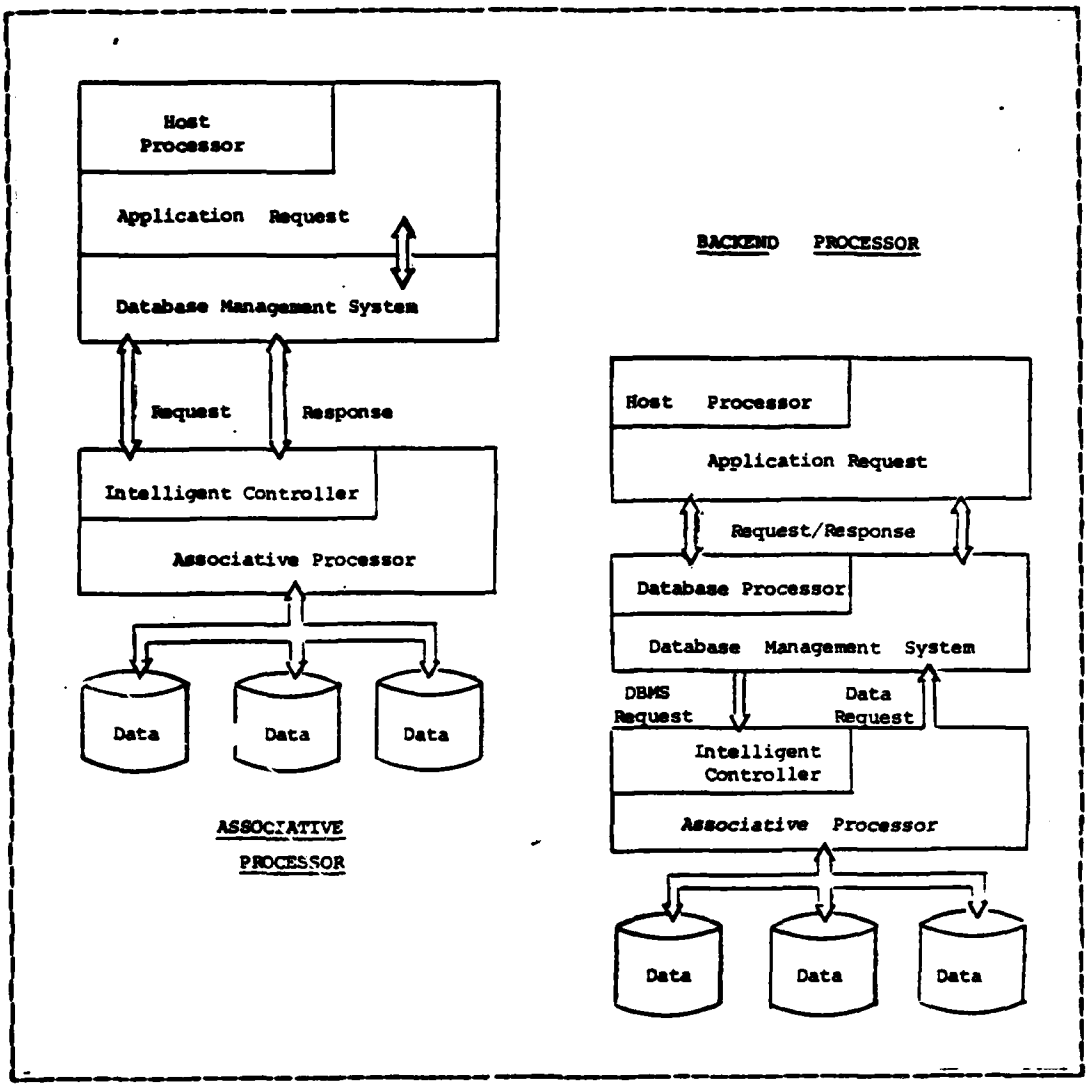

Figure 4.4 Future DBMS Architectures.

to have a database machine architecture. Figures 4.3 and 4.4 are some of the ways DBMSs are integrated in computer architecture, both present and future. The underlying concept of a database computer, regardless of the given configuration, is to off-load a portion or all of the DBMS system functions from the general purpose CPU. This

approach probably had its impetus in the front-end processors which isolated complex and sophisticated communications and message processing capabilities, and placed those functions in a specially designed separate processor.

There are a number of factors today which make the offloading of DBMS functions very desirable. First, user application requirements are requiring larger and larger data bases. The very volume of data is beginning to create performance problems with the length of time necessary to locate the data. In some cases the overhead for indexes and tables consumes more than the data itself. Typical issues any DBMS must address are (1) the organization of an integrated data base, (2) the storage locations for data in the system, (3) the location of data in the system, (4) the control of concurrent addresses and (5) the mechanisms and structures to provide security and integrity. In addition, more and more functions are being required in the DBMS. As the variety of these functions of a DBMS increase, so does its overhead with an accompanying reduction in response time and throughput.

The first attempts to ease the overhead and performance problems of DBMS's consisted of using a general purpose back-end computer to replace database management software, on I/O routines, and on-line secondary storage. Although this has provided a partial solution to the bottleneck problems, the following factors have mitigated against its popularity. The marginal value of a general purpose back-end computer's performance are in most cases much less than the marginal costs of the machine, it is not uncommon to have multivendor coordination difficulties, and reliability of the total system went down as a result of using both the host and back-end (if one goes down the whole system crashes.)

By specialization, i.e. associative addressing in an intelligent I/O controller (discussed earlier), in conjunction with special purpose hardware, significant improvements result which ease the performance bottlenecks of present-day software-laden DBMS. The reason, in a special designed processor the need for floating point hardware, long-word-size registers, and fast multiply and divide logic are not required. Instead, byte manipulation and high speed I/O capability to a wide variety of storage devices can be accommodated, providing the potential for significant improvements. There is much research needed in this area, but it shows some promise, and is attracting some of the best minds in industry and the academic environment. For example, Sperry Univac is working with Dr. David K. Hsiao in exploring the commerical potential for his DataBase Computer [Ref. 13: p. 1].

E. DISTRIBUTED DBMS

Another approach being explored to solve the bottleneck performance problem of DBMS software, distributed data-base technology, attacks the problem from an altogether different perspective. The major benefits claimed are increased data availability to the end user and reduced exposure to total system failure due to hardware or software failure [Ref. 15]. Because data processing and storage capabilities are distributed to the location of data origin or end use, faster and easier access to the data can be provided than under a traditional centralized DBMS. However the facilities for data organization access and control must be extended to include the following CODASYL recommendations in a Network DataBase Management System (NDBMS):

1. Intercept a user request and determine which nodes to send it to for processing. The majority of user requests should use local data and not require the NDBMS. These may go to the local DBMS directly or be passed to it by the NDBMS.

2. Access the network directory (which may possibly be remote) for the above purpose.

3. If the target data are on multiple nodes, coordinate the use of these nodes.

4. Manage the communication between its node and DBMS's in other nodes.

5. Provide data and process translation support in a heterogeneous distributed data-base environment. Heterogeneousness in this context implies differences between hardware (and software) elements in each node in the system.

Thus in distributed data base technology, faster access to data introduces more complexity to the overall system.

Maintaining data base accuracy with concurrent updates to the same data item, maintaining the internal consistency of the data base, and maintaining the consistency among the various copies of the data base, collectively comprise what is known as concurrency control. This is the primary research area in distributed database technology. One may think that having heterogeneous hardware and software components might create the most difficulty in a distributed system, but this is not so. Heterogeneousness does however, add to the complexity of the system. It is believed that before a distributed DBMS can achieve wide acceptability, the update synchronization or concurrency problem must be solved.

A number of methods or techniques are being explored to solve the update synchronization problem [Ref. 16]. The first is global locking where all the copies of the data are locked, the update is applied to every copy, and all the locks are then released. This approach effectively solves the concurrency problem, but at an unacceptable price in high system overhead. The second method is dominant copy synchronization where one copy in the distributed system is dominant and updated first. This dominant copy then controls the rest of the updates at the different nodes. This method is acceptable for systems having low update

frequency, but becomes the source of a bottleneck if updates are high, or if updates are time critical. The third method utilizes timestamps instead of locks. The idea presented is that timestamps are appended to every data item. The data at the node can be updated if the timestamp is greater (more recent) than the timestamp on the data items at the node. The concern with this approach is the storage burden created by the timestamp being appended to every data item. A timestamp for every record versus data item effectively cuts storage requirements, but some needless conflicts are incurred and some updates are delayed with this approach.

The Department of Defense has contracted with Computer Corporation of America (CCA) to investigate the feasibility of utilizing distributed database technology for command and control situations. CCA's result has been the development of the System for Distributed Database (SDD-1) which appears to have elegantly solved the concurrency control problem. As this development has generated much interest in the research field, a brief description is provided in Appendix B which provides insight to not only the concurrency problem in distributed database technology, but also the technical problems of query processing and reliable writing.

Research is on-going in this area and much work has to be done before it can be effectively used by an organization that is not technologically sophisticated.

F. SUMMARY

The requirements into the foreseeable future will require DBMSs to deal with higher and higher transaction rates, larger data bases, and more complex queries to minimize the number of reports received by the end user. Figure 4.3 depicts those architectures that are commonly available today. The most common architecture is having the

DBMS, host language interface, and query language capabilities collocated. Where the CPU is utilized more than 45% to perform data base functions, off-loading the DBMS functions into a back-end processor as in Figure 4.3 may be cost effective [Ref. 17]. But this is only a partial solution to the growing requirements.

The technological developments in associative processing, very large scale integration (VLSI) chips, and database machines, will make architectures like Figure 4.4 possible with much better cost/performance ratios. These configurations are not commercially available today. Further Coast Guard requirements do not call for them at the present time. But a DBMS strategy must recognize that future requirements will outstrip what is available today. Solutions to expanded requirements will more than likely find their source in the emerging technology, and any corporate strategy must take cognizance of this.

V. CURRENTLY AVAILABLE DBMS ARCHITECTURES

A. GENERAL

One can categorize database systems by the form of the data model that is used, e.g. relational, hierarchical, and network. The primary characteristic of a relational model is that all information in the database is represented in a single uniform manner, namely in the form of tables. This data model supports many-to-many relationships of data. The characteristic of the hierarchical model is that information is represented by a simple tree structure. Each node in the tree can be likened to a single file; however, these files are a more complex object than the tables of the relational model. First, the file may contain more than one type of record, and second there are links connecting occurences (instances) of these records. Hence, the data in a hierarchical model is represented by records and links in a one-to-many relationship. The characteristic of the network model is that information is represented by records and links as in the hierarchical model, but the network model has the capability of modelling many-to-many relationships like the relational model.

Since the difference between the network (sometimes referred to as CODASYL (Conference on Data System Languages)) and relational approaches to DBMSs are such a controversial topic at the present time, this chapter begins with a discussion about the distinctions between the CODASYL and relational approaches to database management. The only goal here is to present the views of the authors and what they understand the difference to be in the approaches. After a somewhat lengthy discussion of the relational versus

CODASYL approaches, a discussion of a particular data model from each of the three current available DBMS architectures is presented. The objective is to provide the reader with a feel for what each data model offers - its strengths and weaknesses - not to evaluate all DBMSs available on the market today. At the end of the chapter the consequences of choosing one DBMS architecture over another are discussed as a prelude to the next chapter on a database management strategy.

B. RELATIONAL AND CODASYL NETWORKS

To understand the difference between the relational and CODASYL (network) approach to DBMSs, it is necessary to understand the three-level hierarchial view of a data base. The Standards, Planning and Requirements Committee (SPARC) and the American National Standards Institute (ANSI) published an interim report in 1975 defining this three-level structure. The programmers view of data is referred to as the <u>external schema</u>. The overall logical view or what one may consider the corporate view of data is called the <u>conceptual schema</u>. Finally the physical organization of the data itself is known as the <u>internal schema</u>. Figure 5.1 depicts the three different levels.

The internal schema is one whose records are manipulated by READ, WRITE, and DELETE commands of COBOL and PL/I. At this level one notices that a file has not only data, but flag and pointer fields to handle overflow in a hash file. (There exist other methods to organize the data physically; however, for purposes of illustration the file under discussion is assumed to be a hash file. The idea of a hash file organization is to divide the records of a file among buckets, which each consist of one or more blocks of storage. If v is a key value, h(v) indicates the number of

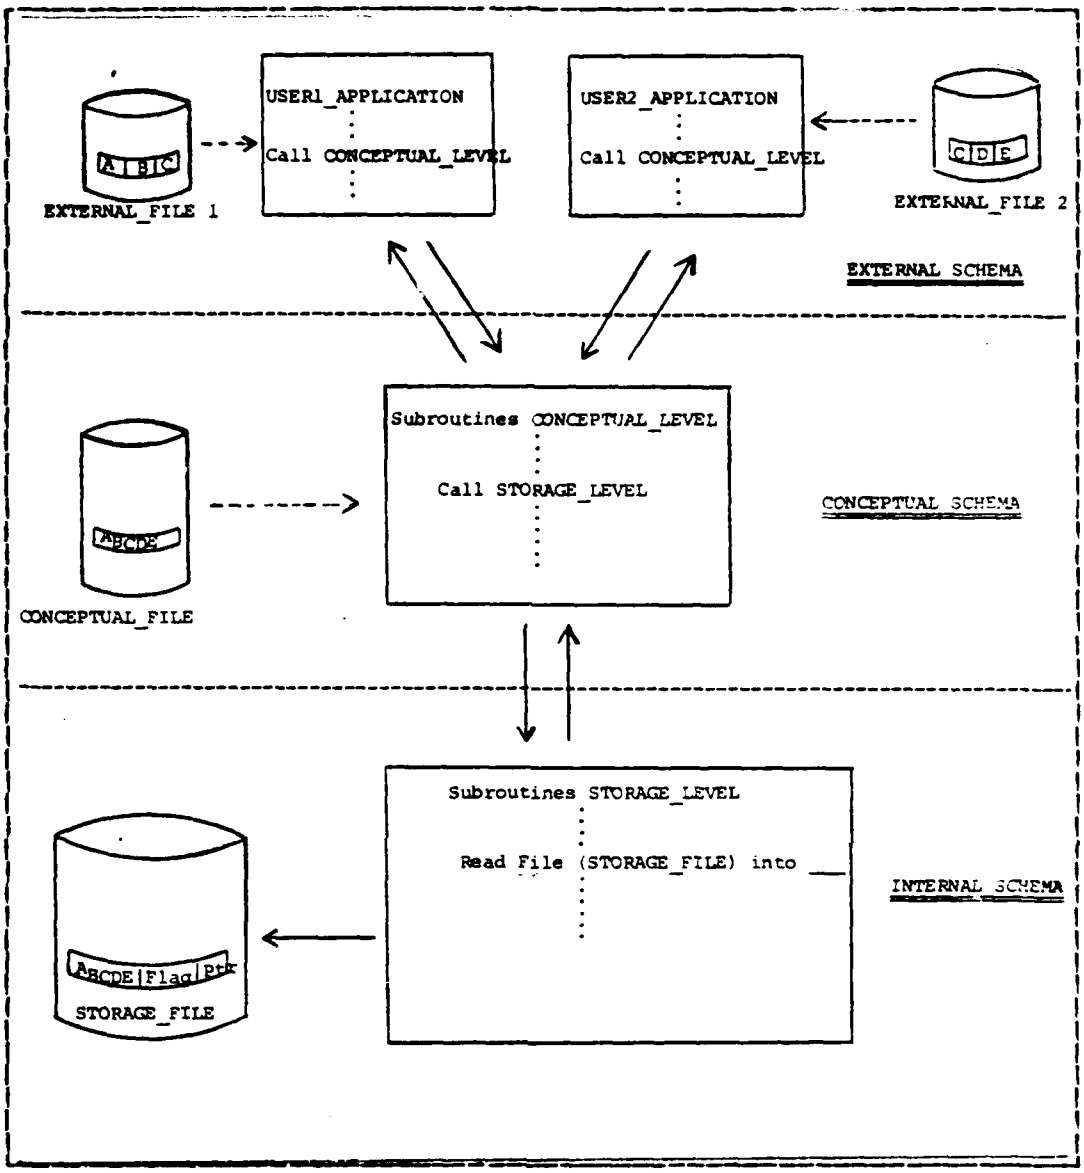

Figure 5.1 Three-Level Hierarchy of a DBMS.

the bucket in which the record with key value v is to be found.) At the conceptual schema level the flag and pointer fields are of no interest; in fact they would just clutter the file with meaningless data for the data base

administrator or the corporate user. At the highest level, the external schema, each programmer or user is only concerned with a portion of the conceptual schema. Thus external files and the fields of an external file are a subset of the fields of the underlying conceptual schema. The external schema contains no flags or pointers from the internal schema, only data of interest to the data base administrator or corporation. It should be pointed out that the fields in the external file may be derivable from the conceptual schema. As an example the field 'age' in the external schema is derivable from the field 'date-of-birth' in the conceptual schema. (Note these are two different data elements; one is probably represented as '480605' (conceptual schema) and the other as '34' (external schema)).

There are some important consequences flowing from this type of design. If for some reason the database became used heavily for sequential access, reads and writes, it should be possible to change the internal schema without affecting either the data base administrator's view or use of the data. Also the programmer or user (the external schema) should not become embroiled or tangled with changes made at the internal schema level. This is commonly referred to as <u>physical data independence</u>. Another closely associated concept is <u>logical data independence</u>. Here programmer B can add another field to his external schema. If the field is in or can be derivable from the conceptual schema then not much effort is required to provide Programmer B with the desired field. However, if a field is requested that is not in the conceptual schema then the conceptual and internal schemas will have some work to do. Logical data independence is achieved when programmer A's application is not affected by the changes required to make a new field available to Programmer B. The benefits that follow from

physical and logical data independence are lower maintenance cost and less programmer time. Both the relational and CODASYL data models have embraced the three-level concept.

The core of understanding the difference between these data models lies in the conceptual schema. One can view the conceptual schema three ways:

1. The same as the internal schema without its implementation of flags and pointers.
2. A common deromminator logical file for all users.
3. A collection of conceptual records, each of which describes an entity of the real world.

In the relational model, a relation (table) is the same as a conceptual schema in which no two records (sometimes referred to as tuples) are the same (i.e. no duplicate records). Further, a relation can be displayed as a two-dimensional matrix, where a row of the matrix is unique and corresponds to a record. (It is important to realize that a relation cannot be compared with a conventional file as the file is a storage file that physically exists. This is not so with the relation which can only be understood in terms of conceptual schemas.) Just because a relation is equivalent to a conceptual schema, where no two records are the same, does not mean that there exists conceptual schemas that are not relations. For instance, any number of repeating groups or fields can exist and are permitted in a CODASYL conceptual schema. If there are repeating groups or duplicate records in the conceptual schema, then the conceptual schema is not a relation. Hence the underlying difference between the two approaches.

The conceptual schema of the relational model imparts simplicity as external schemas derived from conceptual schemas are also relations. Further, the simplicity of the relational's external schemas make them easy and convenient to manipulate via application programs. The relationship

between the conceptual matrices are not explicitly specified (i.e. named) in the conceptual schema for a relation, the specification of the connection fields supporting the relationship being necessary and sufficient. For example it is the connection fields OPS_CNTL in AREA_COMMANDER and OP_CNTL

OPS_CNTL	HQ
LANTAREA	New York
PACAREA	San Francisco

AREA_COMMANDERS

SHIP_NO	OPS_CNTL	SHIP_TYPE
WAGB-10	PACAREA	ICEBREAKER
WAGB-11	PACAREA	ICEBREAKER
WHEC-716	LANTAREA	HEC
WHEC-721	LANTAREA	HEC
WHEC-723	PACAREA	HEC
WMEC-618	LANTAREA	MEC
WMEC-620	PACAREA	MEC

VESSELS

Figure 5.2 Example of Two Relations.

in VESSEL from Figure 5.2 that determines the relationship between the two conceptual matrices. This does create some difficulty when designing a retrieval language which is suitable for the nonmathematical user. Consider:

```
SELECT FROM AREA_COMMANDER WHERE
   OPS_CNTL IS IN
      (SELECT OPS_CNTL FROM VESSEL
         WHERE SHIP_TYPE = 'ICEBREAKER')
```

Here one is asking for the AREA_COMMANDER record whose OPS_CNTL field values from VESSEL records for which the SHIP_TYPE is ICEBREAKER. Very convoluted to say the least.

In contrast, in CODASYL, the files of the external and internal schemas can become very difficult to manipulate via external programs if the records of the schemas contain varying numbers of repeating groups of fields. The underlying reason is that the data structures are allowed to vary. In a CODASYL conceptual schema the records are typically grouped in two distinct ways. First, records of the same type are grouped together. This grouping is similar to the grouping of tuples (rows) in a relation; however, the CODASYL files that make up the conceptual schemas don't have to be relations, i.e. may contain repeating groups. The second method of grouping is into owner-coupled sets. In a set based on two files like in Figure 5.2 each owner record is grouped with its member record giving the grouping

OPS_CNTL	HQ	<->SHIP_NO	OPS_CNTL	SHIP_TYPE
LANTAREA	New York	<->WHEC-716	LANTAREA	HEC
	New York	<->WHEC-721	LANTAREA	MEC
	New York	<->WMEC-618	LANTAREA	MEC
PACAREA	San Francisco	<->WAGB-10	PACAREA	ICEBREAKER
	San Francisco	<->WAGB-11	PACAREA	ICEBREAKER
	San Francisco	<->WHEC-723	PACAREA	HEC
	San Francisco	<->WMEC-620	PACAREA	MEC
AREA_COMMANDERS			VESSELS	

Figure 5.3 Owner-Coupled Set Occurrence.

depicted in Figure 5.3. The relationship between the records of the two files is clear. CODASYL requires that this grouping in Figure 5.3 be given a name in the conceptual schema. The CODASYL query could be expressed in terms like:

Retreive the owner record of each owner-coupled set occurance in which all but one of the member records have a SHIP_TYPE value equal to ICEBREAKER.

Hence, as the relational approach asserts that groupings of files is not necessary in a conceptual database, CODASYL requires it. For the non-mathematican certainly CODAYSL's grouping is easier to comprehend.

In summary then, relational databases require a unique key for every record type, a means of representing set type relationships, and propagation of prime keys down the structure [Ref. 18]. The propagation of prime keys down the structure makes it very difficult for the user to form a correct query via the retrieval language of the database system. In contrast CODASYL does not force a data base designer to propagate prime keys down the structure; thus allowing English-like queries on the database. However, there is nothing in CODASYL to prevent conceptual files from being designed with repeating groups, duplicate records, etc., thus creating data structure problems in the future when structures have to be changed.

C. RELATIONAL - SYSTEM R

On 30 January 1981 IBM announced a product called SQL/Data System. This system is a relational DBMS. To the users of SQL/DS the data in the data base appears to be structured in the form of tables. The SQL/DS DBMS is the result of IBM's work on the prototype System R.

1. Capabilities/Characteristics

Some of the claimed capabilities of SQL/DS are automatic navigation through the data base and automatic optimization. The user of the DBMS specifies what they want, not how to get to the data they want. This

description of 'how' is performed by the very high-level language, SQL. This language has the ability to manipulate collections of records not just single records. SQL also acts as a host language interface to COBOL, PL/I, and Assembler, i.e. SQL statements can be freely mixed with source statements from COBOL and PL/I. Besides being a query language SQL encompasses data definition, bulk loading, updating, granting authorization, and database recovery. The goal of automatic optimization is achieved by allowing the user to dynamically change an index or key.

SQL/DS appears to support the American National Standards Institute ANSI X3 SPARC study group's interim report on the three-level hierarchy view of a DBMS. This recommendation is that a DBMS should support an external view of the data base which insulates the details of the physical organization of the data base from the logical view of the data. This should in turn, insulate the overall logical structure of the data base from the programmer or user.

As in most DBMSs, SQL/DS provides for data security via access control/authorization, concurrency control for both batch and interactive processing, recovery, and directory management. System resources consumed depend on the complexity of the data base and the nature of the transactions; however, it is claimed by IBM that performance can be kept to within acceptable limits.

2. Critique

During the past decade there has been much research and discussion about relational databases as the result of the papers and theoretical work done by Dr. Edgar F. Codd. But, the number of commercially available relational DBMSs is rather limited. Table V and Table VI provide a comparison between three (3) commercially available relational

TABLE V

Data Definition, Manipulation, and Control Facilities

Feature	SQL/DS	ORACLE	INGRES
Data Definition Support			
Variable Length Rows	Yes	Yes	Yes
Null Values Supported	Yes	Yes	No
Dynamically Add New Columns	Yes	Yes	No
Support Stored Commands	Yes	No	No
Dynamic View Definition	Yes	Yes	Yes
Automatic Elimination Of Duplicate Rows	Optional	Optional	Yes
Dynamic Data Base Expansion	Yes-Utility	Yes-Utility	No-O/S Function
Data Manipulation Support			
Storage of View Text	Yes	Yes	Yes
Append Table to Itself	No	No	Yes
Support for User Files	No	No	No
Specify Disk the Table is To Be Placed On	Yes	No	No
Maximum Size Char. Field	32767	255	255
Define Synonyms	Yes	No	No
Sort Query Results	Yes	Yes	No
Data Conversion Functions	No	No	Yes
Help Facility	Yes	No	Yes
Integrity Constraints	No	No	Yes-Single Table
Three-Walking Syntax	No	Yes	No
Outer-Join Syntax	No	Yes	No
Data Control Facilities			
Access Permission for Tables and Views	Yes	Yes	Yes
Permission Structure	Decentralized	Decentralized	Centralized
Resource Authorization	Yes	No	No
Password for Data Base Access	Yes	Yes	No

databases [Ref. 19]. It is apparent that these three (3) systems provide a level of capability to support the Coast Guard's requirements; however, because they are relatively new products, items such as vendor support, documentation, and user satisfaction cannot be determined.

D. NETWORK - CODASYL TYPE - TOTAL

One of the most popular DBMSs on the market today is Cincom's Systems, Inc. - TOTAL. They claim over 3,000 users of their DBMS. TOTAL supports the network data model by allowing a large number of files to be linked together.

TABLE VI

Comparison of General Capabilities

Feature	SQL/DS	ORACLE	INGRES
Index Support			
Primary Index Required	No	Yes	No
Support Dynamic Index Def.	Yes	No	Yes
Support Multicolumn Indexes	Yes	No	Yes
Support Variable Length Keys	Yes	Yes	Yes
Physical Ordering of Data (Clustering)	Yes	Only on Initial Load	Only on Initial Load
Change Clustering Index	No	No	Yes
Locking			
Levels	Row,Page,Table, Data Space	Row,Table	Table
Automatic Lock Escalation	Yes	No	No
Storage Statistics Available	Yes	No	Yes
Optimization			
Access Method	Compiled	Interpretive	Interpretive
Path Selection Methods	Evaluate Indexes, Scans and Two Join Methods	Evaluate Indexes	Evaluate Indexes,Scans and One Join Method
Transaction Management			
Mult. Statement Transactions	Yes	Yes	No
User-Controlled Backout	Yes	No	Yes
Support Nested Transactions	No	Yes	N/A
Disk Recovery Techniques	Roll Forward & Backout	Roll Forward	None
Host Language Interface			
Type	Precompiler	Call	Call
Languages Supported	PL/I, Cobol and Assembler	Fortran,Cobol,PL/I 'C' and Assembler	Pascal & 'C'
Support for Multiple Cursors	Yes	Yes	No
Computer Environment			
CPUs	43XX,370,303X	PDP 11/XX, VAX	VAX
Operating Systems	VSE	RSTS,RSX-11M,IAS, Unix, VMS	VMS

1. Capabilities/Characteristics

Multiple phases are required to produce TOTAL'S DBMS. First, data base generation is necessary. This is done by describing files, file relationships, data elements, physical file characteristics, and buffer usage via the high-level Data Base Definition Language (DBDL). After data and files have been formatted the Data Manipulation Language

(DML) statements can be embedded in a host language to begin actual processing of data. Access to the database is possible via FORTRAN, COBOL, PL/I, RPG II, and Assembler. A simple query language, T-ASK, and report generator, SOCRATES, is provided, but SOCRATES can only be executed in batch mode.

Common features found in other DBMSs are also found in TOTAL. One example is that new applications and files can be added to an existing data base or a new relationship can be defined for existing files without impacting previously implemented applications and programs. TOTAL provides for integrity through its automatic recovery system which restores the data base to its original condition prior to the computer system or program failure. Password protection is provided at the file level. TOTAL supports both interactive and batch processing. Inconsistent updating is prevented via the use of record locking, i.e. when updating (writing) a record, other processes are denied access to the record.

2. Critique

TOTAL is popular, and can be run on many machines. It does not fully support the three-level hierarchy view of ANSI X3 SPARC, specifically the external schema. It is capable of handling large amounts of data. There are a number of versions of TOTAL which are supported, i.e. Basic, Extended, Central, and Multithread. The Extended and Central versions do not have fully re-entrant code, thus making the systems a bit slow in a multiprogramming environment where multiple I/O's are handled.

Table VII represents a 1979 Datapro survey of proprietary software users and their reaction to TOTAL [Ref. 20]. From the table, TOTAL is in general a popular DBMS, and could support the Coast Guards' requirements. It

TABLE VII

1979 DataPro Survey of TOTAL Users

TOTAL Large Scale Systems:

	Excellent	Good	Fair	Poor	W.A.
Overall satisfaction	30	51	5	1	3.3
Throughput/efficiency	20	48	17	2	3.0
Ease of installation	26	41	17	3	3.0
Ease of use	25	50	12	0	3.2
Documentation	14	41	27	5	2.7
Vendor technical support	18	36	26	7	2.8
Training	10	38	22	10	2.6

TOTAL Minicomputer Systems:

	Excellent	Good	Fair	Poor	W.A.
Overall satisfaction	12	19	4	0	3.2
Throughput/efficiency	1	27	7	0	2.8
Ease of installation	14	17	4	0	3.3
Ease of use	10	16	9	0	3.0
Documentation	5	17	10	3	2.7
Vendor technical support	10	14	8	3	2.9
Training	6	16	8	3	2.8

is worthy to note that TOTAL runs on many minicomputers, e.g. Prime 350-750, PDP-11, VAX-11, and others.

E. HIERARCHICAL - SYSTEM 2000/80

System 2000/80 supports a number of computer systems and operating environments, and handles hierarchical tree data base structures. The system was first introduced by MRI Systems Corporation in 1970. In 1979, MRI was acquired by Intel Corporation. An interesting feature is that Intel provides a Data Base Assist Processor (DBAP) which incorporates the Intel semi-conductor disk unit with the System

2000/80 database management system. This is a variation of the database machine mentioned in Chapter IV.

1. Capabilities/Characteristics

An Integrated Data Dictionary (IDD) supports System 2000/80 vis-a-vis maintaining information about the processing environment. An English-like query language, QUEST, is provided to generate reports, and a report writer is available for sophisticated batch report generation requirements. A Programming Language Extension (PLEX) allows commands to be embedded in host programming languages like FORTRAN, COBOL, PL/I, and Assembly Language; this allows further processing of data from the data base.

System 2000/80 supports both batch and interactive processing. Also journaling, rollback, recovery, and security are standard features of the DBMS. Security is provided for at the data base, record, or item level by restricting accesses to update, retrieval, or search. Finally, sixty-three command streams can be handled concurrently and interleaved at the I/O level.

2. Critique

System 2000/80 runs on many machines and environments including IBM 360, 370, 303x, and 4300 series; Univac 1100 series; and CDC 6000/Cyber Series. Boolean logic capabilities make it amenable to ad-hoc queries. The biggest drawback is that it is not available on minicomputers and it requires 180K bytes or more of main memory. The results of a DataPro survey [Ref. 20: p. 70E-526-01b] are provided in Table VIII. In general, System 2000/80 can support Coast Guard's requirements, except for the minicomputer requirement.

TABLE VIII

DataPro Survey of System 2000/80 Users

	Excellent	Good	Fair	Poor	W.A.
Overall satisfaction	12	15	3	0	3.3
Throughput/efficiency	3	16	11	0	2.7
Ease of installation	5	11	7	0	2.9
Ease of use	11	13	6	0	3.2
Documentation	6	15	8	1	2.9
Vendor technical support	9	15	3	0	3.2
Training	11	14	0	0	3.4

F. SUMMARY

In general any DBMS from one of the three data models can support the Coast Guard's requirements. These requirements are detailed in Chapter VI. The drawback of the hierarchical data model implementation is their memory requirements. IBM's IMS (IBM's hierarchical DBMS), and Intel's System 2000/80, mentioned previously, cannot be implemented on minicomputers. Ease of use is somewhat limited because, once the conceptual schema is defined, it becomes very difficult to modify. Also, the conceptual schema must be many-to-one and there are many applications which do not fit this mold. On the positive side, the hierarchical data model effectively supports very large centralized data bases, where efficient implementation is the constraining factor. This data model also, has been used for a decade in both industry and government with some success; hence there is a lot of knowledge about its use and implementation.

The network data model, unlike the hierarchical data model, can be implemented on minicomputers. The network data model does support large data bases efficiently. This model too, has been in existence for better than a decade. The problem for an unsophisticated user is the implementation details that must be known to use the DBMS. Access paths and the physical data organization must be defined prior to use by the implementor. Although the data model supports many-to-many relationships of data, it is done only with difficulty.

The relational data model is relatively new and based upon the theoretical papers of Dr. Codd. There are not many commercially available relational DBMSs in existence today. Because of this, what is known about the relational model is restricted to the research laboratories and the academic environment. The relational models that do exist support small and medium sized data bases. The exceptions are ORACLE, INGRES, and SQL/DS. Implementation of large data bases are usually plagued with performance problems. On the positive side, the relational models are easy to use, i.e. the conceptual and external schemas can be modified easily to fit any situation. This makes it ideal for supporting Decision Support Systems (DSS). Undoubtedly, the relational model will play a more and more important role in the future. First, more will become known about them as they are implemented commercially. Secondly, most of the research at universities and R&D centers are concentrating almost exclusively on the relational model.

VI. DATABASE MANAGEMENT SYSTEM STRATEGY

A. GENERAL

This chapter will present a broad strategic plan for the Coast Guard's use and implementation of a database management system (DBMS). To help motivate this scenario, the Coast Guard's present situation is outlined followed by a short discussion of their planned growth in the database technology arena. Items such as goals, database management system requirements, data dictionaries, migration paths, and approaches are discussed. It is the authors' belief that for the implementation of any strategy the Coast Guard employs, the technology must be commercially available and have proven itself over the past few years. The Coast Guard can ill afford to tackle some of the state-of-the-art technology with its meager resources. However, the actions taken today will have an impact on the Coast Guard's ability to use new technology later when it becomes available on the market. A wrong decision today could effectively lock the Coast Guard out of very promising alternatives in the future years. What follows are the authors opinions about how to achieve a database management system technology today in the Coast Guard without locking the Coast Guard out of the changing technology.

1. Present Situation

All strategies must begin with an inventory and evaluation of the present situation. Clearly the Coast Guard is behind the proverbial 'eight ball' in the ADP environment. Even with the Commandant's attention and reorganization of Coast Guard Headquarters to provide a

flag-rank office to oversee and direct Coast Guard-wide information requirements, there is still a long way to go.

To date the Coast Guard's use of DBMSs has been rather haphazard. A few years ago the Coast Guard bought Cullinane's database management system, IDMS, for Coast Guard-wide use. Because there were no Coast Guard users, IDMS was turned over to the Department of Transportation's Computer Center (TCC). Today, if the Coast Guard would like to use it, they have to pay a surcharge for both the use of the DBMS and the TCC's computer. Most applications in use today are using file structures that were popular in the late 1960's and early 1970's, and are not taking advantage of current database management technology. Probably the most representative example is the Standardized Aids to Navigation Data System (SANDS) [Ref. 5]. The system was developed in the late 1960's using the file structures of that date, i.e. flat, fixed, sequential access record structures. As a result, today the system is of little use or benefit as the data cannot be changed to meet the requirements of the users. As an aside, SANDS is being revised to see if the data, collected over the past 15 or so years, can be of any benefit. The one bright spot in this picture is the development of the Marine Safety Information System (MSIS). Although it has taken over ten years to get to its present state, and it is not fully implemented; database management technology plays a pivotal role in its development and implementation. MSIS uses Cincom's DBMS, TOTAL. The use of a data dictionary allows control of over 2500 data elements throughout the system. Also, Civil Engineering (G-ECV) is currently developing a data base, data dictionary, and data directory using database technology [Ref. 5: p. 5-32].

2. Planned Growth

Headquarter's Offices are planning new systems at a phenomenal rate. A glance through Commandant's Instruction M5230.81, Automated Data Processing (ADP) Plan, provides the interested reader with the size and scope of the new efforts.

The most ambitious new undertaking is the 'District Minicomputer Project.' It is the intention of the Commandant to procure minicomputers and place them in each of the twelve District Offices to provide an expanded ADP capability at major field commands. This project is to be incremently developed with initial attention being focused on identifying and responding to the critical information and management needs of the District Offices. Phase one of the project is to perform the required planning, identify information needs, conduct a feasibility study and provide for resource decisions. Phase two provides for the development of specifications and standards. Phase three initiates procurement action, performs system development, conducts database formulation and test/integration. Finally, phase four commits the system to operation and provides for refinement and expansion. This project is receiving top management control from the Office of Command Control and Communication (G-T) with the attendant bureaucractic control and standardization.

Once it is fully operational, the Marine Safety Information System (MSIS) will provide the Coast Guard with its first fully matured Management Information System. The Operational Computer Center (OCC) located at Governors Island, New York, using PRIME 750s, besides running the Automated Mutual-Assistance Vessel Rescue System (AMVER), the Computer-Assisted Search Planning System (CASP), Ice Plot (ICEPLOT), and the Enforcement of Laws and Treaties

(ELT) database (limited), will support the Hazard Assessment Computer System (HACS), the Spill Cleanup Equipment Inventory System (SKIM), and the ELT database (expanded). Thus Coast Guard's near term future requirements will more than double today's present capacity, while quadrupling the number of programs and files necessary to support the information needs for the rest of the decade.

B. DBMS OBJECTIVES

A database management system must address the basic issues of (1) the organization of an integrated database, (2) the storage locations for data in the system, (3) the location of data in the system, (4) the control of concurrent addresses, and (5) the mechanisms and structures to provide security and integrity. Aside from the DBMS software, other appendages are considered necessary - a data dictionary, and good software for interrogating, searching, maintaining, and generating reports from the database. These two items will be discussed later in the chapter. What follows is a discussion of the main objectives/requirements that a DBMS for the Coast Guard should support.

1. Multiple Views of Data

The Coast Guard will be, in a majority of cases, developing databases by subject area as proposed by [Ref. 10]. This is in contrast to developing data bases by application areas. Different programmers will have and require different conceptual views of the same data. As an example, if 'UNDERWAY_HOURS' were in a data base called OPERATIONS, then a programmer for Naval Engineering may want to use this data element with the 'VESSEL_AGE', and 'LAST_REPAIR_DATE' to determine a maintenance repair schedule. A programmer for Operations may want to use

'UNDERWAY_HOURS' from the database OPERATIONS along with 'STANDARD_UNDERWAY_HOURS', and 'AVAILABLE_VESSELS' to determine deployment schedules. The DBMS must be able to support both views of this single data element. Therefore the DBMS must be able to derive from both the data and logical relationship, the logical files that are necessary. Further, the programmer should not have to be concerned with how the data is physically shared. Transparency for the programmer is the main concern.

2. Performance

The Coast Guard will have requirements for both interactive and batch processing via the DBMS. The more complex queries and reports will, in all likelihood, be submitted for batch processing, while the relatively simple ad-hoc query will be submitted interactively. Wilson-Hill's Report, [Ref. 10: p. 4-15], indicates that approximately 90% of Coast Guard's use will be interactive. It is not unreasonable that other systems will have a similar mix; however, they will probably not be as high for interactive processing. Hence, database applications designed for use by a terminal operator, and queries of the data base, must have a response time appropriate for man-machine dialogue, i.e. response times of 3 to 5 seconds 85% to 90% of the time. [Ref. 21]

In addition, the DBMS must be capable of handling the throughput. Initially, in any DBMS whether for the District Minicomputer Project, for the Operational Computer Center (OCC), or for one of the 'super' minicomputers at Coast Guard Headquarters, the throughput will be relatively low. As a minimum each 'super' minicomputer should be capable of supporting 150 users with 20 concurrent users. Since these are the requirements determined and supported in Wilson-Hill's Report, [Ref. 10: p. 4-17], it seems

appropriate that this be a basic requirement for any DBMS used by Coast Guard. Needless to say, if an application or subject area data base requires higher performance then that should be the level. The Marine Safety Information System (MSIS) is one such exception. This system provides the upper end of performance throughput, as their requirements are for over 250 users with up to 50 concurrent users.

3. Minimization of Costs

There are many cost features which interact when considering a DBMS. The first is the cost of the DBMS software itself. Depending on the features in the DBMS, the costs may be rather significant, i.e. $100,000 or better for its purchase, plus a monthly maintenance charge as a DBMS software package is usually proprietary property. The other consideration is that the purchase price may only be for one copy, i.e. for the District Minicomputer Project - 12 copies may have to be purchased.

The largest cost, the one usually overlooked or underestimated, is the database conversion, loading, and establishment of a data dictionary system. There is the cost of loading say a 'flat' COBOL file into the data base. Also there is a cost incurred of not converting or loading COBOL files into a DBMS. For instance, if data cannot be converted or loaded into a data base, then the information is lost. Therefore, the Coast Guard has a decision to make - whether or not to maintain the old procedures for capturing, storing, and manipulating this data or cut with the past and start anew. This maintaining of old ways is costly in programmer maintenance and data entry effort. As such, this decision should be considered carefully. Installing a data dictionary system is not an easy task, and the cost is usually in terms of politics. Individuals consider data to be their own property and not a corporate

entity to be managed as such. The cost is the effort it will take to integrate the data dictionary system into the organization.

These short term costs must be balanced against the long term costs of creating additional files, changing logical structures for different programmers or applications, and the overhead required for data searching and storage capacity. Without a DBMS, maintenance programming becomes prohibitively expensive in just a short time. Changes to data structures or logical programs cannot be done without major disruption to other programs and data. There comes a time where one scarifices performance requirements in order to accommodate new requirements in the application. In the case of reduced requirements, time may have to be degraded as the volume of data to be retrieved or updated reaches the upper limit of the capability of a DBMS.

4. Minimal Redundancy

The importance of this requirement cannot be over emphasized. It is at the core of what database management systems are all about.

> A data base may be defined as a collection of interrelated data stored together without harmful or unnecessary redundancy to serve one or more applications in an optimal fashion; the data are stored so that they are independent of programs which use the data; a common and controlled approach is used in addressing new data and in modifying and retrieving existing data within the data base. One system is said to contain a collection of data bases if they are entirely separate in structure. [Ref. 22]

There are three reasons for keeping data redundancy to a minimum. First, there is the additional cost of storing extra data. At one time this was extremely important, but today with the cost per bit of stored data dropping so rapidly it is becoming less of an issue. The second reason is the requirement to update all the different copies of the data. As a result of duplicate data, computer overhead

increases, and the complexity of the software to keep track of the duplicate data and the access paths increases. The third, and probably the reason most managers distrust computers, is that different copies of the data are in different update states. These different update states produce inconsistant information depending on the access path taken by the DBMS. However, it is important to realize that a certain amount of redundancy may be necessary to improve performance. This is even more so in a distributed type environment than in the traditional centralized method. In cases where redundancy or controlled redundancy is required steps must be taken to insure concurrency control or synchronization of these redundant data items.

5. Query and Report Generation

The DBMS should provide facilities to respond to requests that are not anticipated in any degree of detail - i.e. ad-hoc queries. Since Coast Guard personnel will vary in degree of computer sophistication, it is necessary that the DBMS provide a query or retrieval language that is 'English-like' and relatively easy to use. As an aside, the reasons DBMSs have caught on so well, is their ability to extend the power of the 'bare' machine. It allows both a programmer and non-programmer the ability to manipulate data without having to worry about things like; how is the data stored, what is the access path, is the file indirect sequential or the other many details which do not help the user with his immediate immediate task.

The types of queries can range from very simple to relatively complex. As a minimum, the query must be able to access data in the database through a variety of methods, including direct and Boolean (logical) connectors. It must be capable of interfacing with other system capabilities like application programs and report generators. As a

further requirement it must be able to generate a query against temporary files and support several users concurrently.

The report generator must be flexible enough to support a variety of output formats, and interface with application programs. The language should be easy to use (English-like), logical, and should be able to process single and multiple files with the capability to build and store output for later use.

There exists commercial report generators like FOCUS and RAMIS II which can be procured as an add-on item, if not available with the DBMS software.

6. Integrity Controls

The DBMS system must provide routines to insure that the data in the database are accurate at all times. One can view maintaining the integrity of a database as protecting the data against invalid (not illegal) alteration or destruction.

First, the system must be capable of checking each individual data item value for plausibility, e.g. the hours a vessel is underway during the week cannot exceed 168. In a multi-user system the loss of an update must be guarded against. This can happen when data are shared, and one update is allowed to overwrite the other, thus nullifying the first update. (Note: In some systems, allowing several users to update the same data concurrently is a requirement. An airlines reservation system where it is important not to book a seat more than once, is the classic example.)

In general, loss of integrity will arise as a result of a hardware or software failure such as at a central processor, data channel, or an input/output device. Human error on the part of the terminal user, programming errors in the underlying operating system, or programming errors in

the database application will also adversely affect the integrity of the data base.

Several routines should be provided to support system integrity such as, journaling routines, dump routines, recovery routines, backout routines, checkpoint/restart routines, and detection routines. Journal routines record every operation on the database in a systems log or audit trail. As a minimum the audit trail should contain an identification of the transaction concerned, a timestamp, an identification of the terminal and user concerned, the full text of the input message, and the type of change and the address of the data changed, together with its before and after values. Dump routines are used to make backup copies of the database; usually only selected portions of the database are dumped. Recovery routines are required to restore the database to an earlier state if there is a hardware or software failure. The audit tape will be used as input to this routine. Detection routines will be required to search for any violations of integrity contraints before the database is written on.

7. Security and Privacy

The Coast Guard has a requirement for both security and privacy of data. Information will eventually contain classified data and confidential personnel information.

> Data security refers to protection of data against accidental or intentional disclosure to unauthorized persons, or unauthorized modifications or destruction.
> Privacy refers to the right of individuals and organizations to determine for themselves when, how, and to what extent information about them is to be transmitted to others. [Ref. 23]

The DBMS must be capable of positively identifying users (authentication) before they are allowed to use the DBMS. The software must provide access control to insure ordered accesses are valid. Transactions must be capable of being monitored, and data must be reconstructible from journals.

8. Three-Level Hierarchy

In the previous chapter there was a reference to the American National Standards Institute ANSI X3 SPARC study group's interim report regarding the three-level hierarchy view of a DBMS. A DBMS must support this conceptual view of a database, i.e. internal schema - (physical view), conceptual schema - (logical view), and external schema - (programmer's view). Each schema should be independent from the other schemas; in other words:

1. The programmer utilizing the external schema should be provided (by the conceptual schema) all the information required to use the external schema and nothing more.
2. The data base administrator utilizing the conceptual schema should be provided (by both the external schema and the internal schema) all the information required to logically design the database and nothing more.
3. The implementor of the internal schema should be provided (by the conceptual schema) all the information required to complete the implementation of the internal schema and nothing more.

This type of construction in the DBMS provides flexibility, and ensures maintainability and reliability. Flexibility is provided because changes in the size of the data base, storage of the data base, and the number of users, can be accommodated up to the limits of the technology imposed upon the system i.e. memory size, speed of CPU, etc. The entire DBMS should be more maintainable as those items that are likely to change, e.g. the storage structure, corporate view of the data and programmer use of data, are isolated from one another. Hence, changes are not propogated throughout the entire system. Because the entire DBMS is more maintainable, the reliability should follow as a natural consequence.

C. COAST GUARD INFORMATIONAL GOALS

Today the Coast Guard is not facing the complex problems of other organizations who have many years of data processing experience and who have continuously used the state-of-the-art technology over a long period of time. Basically the Coast Guard is just beginning to use some of the newer technology available today - database management systems, networks i.e. TELENET, 16-bit microcomputers, etc. This lack of experience is a handicap; it provides a very small knowledge base upon which to draw experienced personnel. On the other hand, the Coast Guard can learn from the organizations that have already introduced the new technology, like DBMSs and distributed processing. The first thing that is discernible from organizations that have implemented database technology, is that the requirements for a DBMS are easily met by many of the commercially, available packages. This is not particularly surprising, especially as the requirements, although difficult to implement, provide handsome returns in the market sector of society. The two most acute problems stem from managerial decisions. One concerns the implementation of database management technology, while the other, concerns the policies to keep current with computer technology. First, an organization will invest a great deal of resources, developing applications using COBOL type files. The organization finds out later, they cannot take advantage of new technology, e.g. DBMS. The reason is there is no <u>economical</u> way to convert the data and procedures of the COBOL oriented system to a DBMS. The second problem affects the federal government probably more than the private sector. This is concerned with the upward compatibility of resources, i.e. computers, operating systems, DBMSs, and secondary storage devices. Management invests heavily in current computer

technology only ending up having to scrap the previous effort to take advantage of the newer technology. Today, management is no longer allowing this to happen. The marketing strategies of the computer industry recognize this fact, and now provides for upward compatibility of new items. An exception is the microcomputer industry. Thus, the problem reduces to determining a migration path for expanded computer system requirements. Therefore the goals of a DBMS strategy are:

1. To provide for the protection of the organization's intellectual development and
2. To provide a migration path for expanded requirements.

What is of concern in regards to goals is the ability to not only integrate database management technology within the organization, but to insure its future and to reduce the turbulence that comes from follow-on changes.

1. Protection of Intellectual Development

Most entities utilize a DBMS because it provides easier access to data while allowing the programmer to spend his time writing programs to manipulate that data. Although this is a very desirable feature of DBMSs, it is not the main concern. In fact, because DBMSs have abstracted the data so that it is closer to the way people think, an entity's data base can be allowed to grow willy-nilly. For example, an application for law enforcement is determined to be extremely valuable for the organization. To help with the development of a management information system, the designers and implementors decide to use a DBMS to organize the data. However, the data going into the data base is defined and utilized by only a small fragment of the organization even though the data is of benefit to other entities within the organization. Without too much trouble, one can

picture what happens. The other entities within the organization develop their own management information system, defining data in their own way and using different DBMSs. Before long, the organization has a proliferation of data with much duplication and inconsistency. But the worst thing is that there is no way to protect the data and procedures that have been developed when the organization either tries to consolidate the redundant data, to reduce bottlenecks in performance, or to take advantage of new technology, like a back-end processor to handle DBMS functions.

An obvious solution is to centralize everything and have all the organization's data placed into one big corporate data base that will provide the manager with everything there is to know about the organization. Many companies and agencies have tried this only to fail. It is just too big an undertaking for most large organizations. An approach taken from software engineering is to modularize the process and take advantage of what the DBMS software provides. First, try to minimize the effects of change. For example where the data bases are developed by application, a change in a data element resulting from a law or regulation will have a ripple effect throughout the organization. However, if data bases are developed by subject area the effects of change are reduced. Therefore, one way to help insure protection of an organization's intellectual development is to develop data bases by subject area versus application areas.

To help implement subject area data bases, and maintain consistency throughout the organization, the use of a data dictionary system is essential. If the Coast Guard requires a three-level hierarchy for a DBMS in each subject area then, surely at the conceptual schema level, one must know what the entities and characteristics of those entities

are that one is managing. For one person to manage 500 or more data elements without some kind of software tools is unthinkable. This is a conservative estimate of the number of data elements that will comprise a subject area data base [Ref. 10: p. 3-4]. In short what a data dictionary system provides is the capability for:

1. Specifying the type of an entity, such as a form, a data element, or a computer file.
2. Uniquely naming an entity and describing it in appropriate terms, such as the range of values of a data element or a narrative description of its meaning.
3. Specifying the flow and the storage locations of data entities within the organization or within the computer installation.
4. Specifying associations and relationships among the data entities; for example, appearance on the same form, or derivation of an entity from another.
5. Specifying and producing reports about the data dictionary content, such as a listing of all data elements or a cross-reference listing of all entities. [Ref. 24]

The authors believe that the use of a data dictionary must be a requirement and the burden a user must pay in order to obtain the benefits of a DBMS.

In conjunction with a data dictionary system there exists the requirement for a database administrator. The design of the conceptual schema will normally be done by the database administrator and a staff. Also this person will be responsible for the implementation of the conceptual schema as a physical data base. Once implemented the database administrator performs these functions:

1. The creation of subschemas for external views.
2. The granting of authorization to use the data base or certain parts of it.
3. Modification of the conceptual schema should the original design prove faulty or the requirements of the organization change.
4. Modification of the implementation of the conceptual schema by the physical schema, should past usage of the data base indicate another organization of the data base be more efficient.

AD-A127 364 US COAST GUARD ALTERNATIVES FOR DISTRIBUTED DATA BASE MANAGEMENT SYSTEMS(U) NAVAL POSTGRADUATE SCHOOL MONTEREY CA E M FIEGEL ET AL. DEC 82 2/2

UNCLASSIFIED F/G 5/1 NL

END
DATE
FILMED
5-83
DTIC

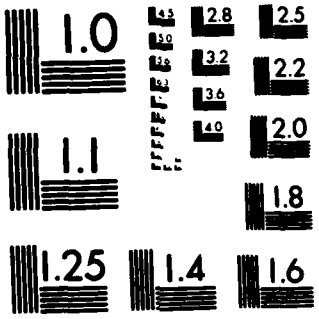

MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

5. Making backup copies of the data base and repairing damage (e.g. 'trashed pointers') to the data base. [Ref. 25]

One should not underestimate the difficulty of this task. Defining a conceptual schema to accommodate the external views of multiple users is no easy task. In fact, most personnel hired as database administrators leave out of frustration. One cannot expect this position to be filled by a junior person within the organization. The position requires the talents of a project manager. He must have the respect of the organization, and a fairly in-depth technical background in computer science. Although this is an important job, traditionally, it has been either underestimated in its importance, or ignored. In either case, the effective use of database management technology has been minimal. The Coast Guard must be willing to establish this position and fill it with highly qualified personnel at those offices utilizing database management technology.

While the development of data bases by subject area, and the use of data dictionary systems will help provide some continuity, it will not insure a high degree of protection of developed data bases. If those data bases are allowed the choice of selecting DBMSs without any overall central policy, then the efforts of integrating a data dictionary and data base administrator, will fail. Because the policy and selection of equipment for the District Minicomputer Project, is a Headquarters controlled undertaking, they will help insure the protection of the intellectual development of data bases at different District Offices. Because the District Offices are somewhat independent, responsible to Program Managers at the Headquarters level and responsible for field units in the District's jurisdiction, there is now little need to distribute database management systems among the different districts. In fact there is little need at the present time to distribute

the data between Districts. Although the distribution of data may be desirable in some cases, the communication costs, protocols, control, and lack of technically knowledgeable personnel make this a high risk undertaking at the present time. The risks of trying to distribute both the data and the processing among districts at the present time outweigh any of the benefits. However, at a later time when the distribution of processing and data are deemed desirable, then it may be accomplished without wholesale disruption of past efforts. This can happen only if the Districts are required to use predetermined data bases, based upon subject areas, and if the Districts are required to use the same commercially available DBMS package. Therefore, the control of the selection of DBMSs and the broad definition of subject area data bases should be a Headquarters level function. It is important to note that development, implementation, and use of data bases is a decentralized function, i.e. the responsibility of the District Office or the Program Manager.

2. Growth

There are two fundamental facts about data processing, change and growth. Costs are dropping, machine structures are changing, new architectures are appearing. What the Coast Guard thinks its requirements are today, and the technology to handle those requirements, will not be sufficient tomorrow. The Coast Guard will need to move to new systems because they are more cost-effective; however, the Coast Guard cannot re-write the multitude of old programs because there will not be enough programmers. The computer industry itself has recognized these two facts for years. They realize they cannot design, manufacture, and market equipment so that it stays continuously current. The manufacturers have developed a strategy to allow for both

change and growth. Both operating systems and the 'bare' machine (processor) on which the operating system runs are developed using the manufacturer's strategy for handling change and growth. An initial version of an operating system will be released, and as new technology is added, newer versions of the operating system are available. This usually happens without major disruption to the utilities and applications utilizing the operating system. (Of course this is not true if an organization has modified the operating system.) Hardware is also planned to be upwardly compatible. Prime Computer, Inc.'s 50 series is an example of a compatible hardware family. With this marketing strategy, an organization can move up to a more powerful processor without major re-writes or disruptions. The computer manufacturers, in short, plan long-range migration paths for their customers. Some examples are the facilities in a new operating system may lead the way to future hardware change, or a software architecture is planned to accommodate future machines. Therefore, the Coast Guard must adopt a similar strategy.

The initial environment the Coast Guard will be dealing with is a stand-alone, large minicomputer including the operating system, utilities such as compilers, sort routines, etc, and DBMS software. For discussion purposes, it is assumed there is a front-end processor to handle the terminal communication protocols. In the District Offices, compatible machines will reside with compatible software. This cannot be an assumption with regards to the other computer systems within the Coast Guard. For example, the Operational Computer Center (OCC) uses Prime 750s, while a Headquarter's Office may be use a VAX-11/780. There are three migration paths available. The first is a migration up to larger processors, the second is creation of a network of minicomputers, and the third is use of a smaller minicomputer as a back-end database machine.

If the path chosen were migration to a larger, more powerful processor, then a requirement for not having to re-write, the application programs and library routines, would most certainly be mandatory. To avoid this re-writing evolution, one must have compatibility of hardware, operating system, and DBMS software. Incompatibility in or among any of these features will require some re-writes. Maintaining the compatibility of hardware and operating systems is a slim possibility in a federal agency because of the procurement restrictions of the Brook's Act. Providing for a competitive procurement and requiring compatible hardware, operating system, and DBMS software is a task with little chance of success. If this migration path were chosen, the Coast Guard should be prepared to do some re-writes of the application programs, and in some cases 'major' re-writes.

The second option is to create a network of minicomputers, where the procurement of a new minicomputer represents a new node in a new network structure. Certainly this will increase the 'computing power' of the system, but new problems are introduced with this approach. If the configuration is to have both processors access the same data, then linking the two processors together, and providing for synchronization of reads and writes will be required. At the present time this is not a common approach and the experience base is not large. It is the author's understanding that Battelle Labs of Columbus, Ohio, who are developing the Marine Safety Information System (MSIS), are proposing an architecture similar to this. The experience they are gaining in this effort of linking several Prime 750 computers together should be very valuable. Another approach within this option is the creation of a network of minicomputers by adding another layer of software - the network operating system. This is similar to IBM's System

Network Structure (SNA) which allows many processors to be hooked together. The problem here is that the addition of the network operating system and another processor may not relieve the burden of the DBMS functions on the first machine. Also a number of these networks like SNA support only the manufacturers equipment. Thus the Coast Guard must again contend with the Brook's Act in the procurement arena.

The third option is the use of a back-end database machine. As discussed in Chapter IV, this is an emerging technology where one is truly working within the state-of-the-art. The advantages of a back-end database machine is its ability to handle a heterogeneous environment easier than the other two approaches. However, loading of the data from the host to the back-end may be extremely difficult if the host say supports a CODASYL approach and the database machine supports a relational approach. This migration path at present is the most expensive and difficult to achieve of the three, and the percent usage of the host to DBMS functions will determine if this migration path is even cost-effective.

In summary option one, migration to a larger processor, is easiest to implement, and there is more knowledge and experience about this migration path than any other. The second solution is harder than the first, and the database machine is possibly the hardest because it is current state-of-the-art and there does not seem to be much experience to date with this alternative.

D. HETEROGENEOUS VERSUS HOMOGENEOUS SYSTEMS

Distributed database technology requires a hard look at the problems one encounters when examining heterogeneous systems - heterogeneous systems are interconneted processors, software, and data structures that are dissimilar.

There are three (3) main problems when dealing with heterogeneous systems at the database management level. First, different types of database software are incompatible (even without considering distributed databases). Second, file structures are expensive to convert to other file structures or database structures. Finally, even if all the software is compatible, problems may arise from incompatible data fields and data structures due to inadequate data administration in an organization. Although there is at present little need to distribute data bases in the Coast Guard, later efforts to do so will be inhibited, if there is not an early commitment to database technology which requires the data bases in an organization to be compatible. Therefore, the emphasis should be on developing homogeneous systems - homogeneous systems are interconnected processors, software, and data structures that are similar.

One realizes that the Coast Guard survives in an environment of regulations that strongly discourages the development of homogeneous systems. The Federal Procurement Regulations requiring procurements to be competitively bid is one such regulation. There are some exceptions to this rule, but they are such that the Coast Guard cannot in general apply these exceptions. The approach taken by the District Minicomputer Project is one solution to providing for a high degree of compatibity. A similar effort must be undertaken at the Headquarters level if one expects to distribute data bases to meet higher goals. At the present time, the cost and degree of technological sophistication is more than the Coast Guard can handle in distributing heterogeneous systems. Therefore, emphasis should be placed on developing homogeneous systems, not on trying to distribute present-day heterogeneous systems.

E. SUMMARY

The Coast Guard is committed to investing in the architecture of the future today, so that the Coast Guard may become an 'information corporation' of the 1990's. The consequence of this is an exponential growth in computer systems over the next few years. An important aspect of these computer systems will be database management system technology. A Coast Guard DBMS has the following requirements - provide for multiple views of data, provide for acceptable performance, minimize cost, minimize redundant data items, support query and report generation, provide for integrity of the database, provide for both security and privacy access control, and support the three-level hierarchy. The Coast Guard must protect its intellectual development during this period by developing data bases by subject area instead of by application, requiring a data dictionary system, and committing itself to provide qualified database administrators. Orderly growth must be determined by a migration path now instead of just letting it happen. In most cases the simplest solution is to use the more advanced technology provided by the original vendor; however, the creation of a minicomputer network may be feasible as the Coast Guard gains more experience in the computer field. Finally the Coast Guard should strive now to develop homogeneous systems instead of trying to distribute present-day heterogeneous systems.

VII. CONCLUSIONS

The following are the conclusions which the researchers have formulated as a result of their thesis research.

First, the Coast Guard should design its information systems for change. The rearchers have discussed the need for system designers to be conscious of the effect which technological change has had, and will continue to have in the computer industry. The Coast Guard does not want to lock itself out of any promising technological advances which will occur, and it must also recognize the need for growth and change in user requirements. Therefore the Coast Guard should develop flexible systems which are amendable to future change.

Second, the researchers have pointed out that since the Coast Guard is relatively new at procuring information systems, the need to have sophisticated computer networks, such as IBM's SNA, is not warranted at this time. This policy is in accordance with prudent evolution, and leads to the conclusion that there will not be a great deal of data sharing required among various Coast Guard organizational elements initially, although it must be recognized that in the future the Coast Guard may need to create these sophisticated computer networks (or distributed systems). In order to support future networking possibilities, it is recommended that the Coast Guard strive for homogeneous systems as much as possible so that this interconnection is much easier to accomplish. The Coast Guard's scarce resources dictate that connecting heterogeneous systems be postponed at this time, and its efforts directed toward the development of homogenous systems.

Third, the authors have concluded that the Coast Guard should require good documentation of all of the systems which are developed and installed. Being a military organization implies a great deal of personnel rotation, and therefore in order that new personnel coming aboard understand, operate, and maintain the system, good documentation is essential becuase there probably won't be anyone around who was present at installation after a few years.

Fourth, the researchers have concluded that the Coast Guard should stress the importance of human needs in the systems which it develops and acquires. By this the researchers are pointing to the need for systems to be easy to learn as well as easy to use in order that personnel be more compelled to use the systems to assist in their job activities, and thereby the system be more responsive to user requirements. Other human factors which should be considered by the Coast Guard include training and education and involving end users in the design and development of systems.

Fifth, the authors have concluded and recommend that the Coast Guard should establish database administrator positions throughout the organization, and it should seek to fill those positions with talented professionals. The duties of the database administrator and staff have been discussed in previous chapters, and the researchers consider that this position is vital to the success of the proper management of any organization's data base(s).

Sixth, a key tool which the database administrator should have is a data dictionary. The data dictionary is essential to assist in implementing subject area databases, to maintain consistency throughout the organization, to identify sources and uses of the data resources, and to construct standards and procedures for those data resources. The data dictionary must be a requirement for the Coast Guard to be successful in implementing database technology.

Seventh and finally, a Coast Guard DBMS should include the following requirements: provide for multiple views of data, provide for acceptable performance, minimize cost, minimize redundant data items, support query and report generation, provide for integrity of the database, provide for both security and privacy access control, and support the three-level hierarchy.

APPENDIX A
RESOURCES

1. Personnel Resources. These are approximate figures, subject to change.

a. Headquarters (does not reflect interim reassignments):

	CO	WO	ENL	CIV	TOT
G-T	2	0	1	2	5
G-TEE	41	13	9	28	91
G-TIS	7	0	4	42	53
G-TTM	9	4	22	13	48
G-TFP	6	0	6	1	13
G-TNR	6	0	8	0	14
TOTAL	71	17	50	86	224

b. District Offices:

Dist	otm					eee					fds					irm				TOTAL
	CO	WO	EN	CV	TOT	CO	WO	EN	CV	TOT	CO	WO	EN	CV	TOT	CO	WO	EN	CV	TOT	
1	2	3	16	0	21	3	3	10	5	21	0	0	0	7	7	2	1	0	0	3	52
2	1	1	10	0	12	2	1	0	0	3	0	0	0	5	5	2	1	0	0	3	23
3	2	3	24	0	29	3	3	5	5	16	0	0	1	10	11	2	1	0	0	3	59
5	2	2	15	2	21	2	3	1	4	10	0	0	0	8	8	2	1	0	0	3	42
7	2	2	17	0	21	3	3	3	3	12	0	0	0	5	5	2	1	0	0	3	41
8	2	3	14	1	20	2	4	2	1	9	0	0	0	5	5	2	1	0	0	3	37
9	2	2	13	0	17	2	2	2	5	11	0	0	0	6	6	2	1	0	0	3	37
11	1	2	10	0	13	2	2	1	3	8	0	0	0	5	5	4	1	0	0	5	31
12	2	2	17	1	22	2	3	1	2	8	0	0	0	6	6	2	1	0	0	3	39
13	2	2	12	1	17	3	3	3	3	12	0	0	0	6	6	2	1	0	0	3	38
14	2	3	14	1	20	3	2	4	3	12	0	0	0	6	6	2	1	0	0	3	41
17	2	3	21	1	27	5	4	4	1	14	0	0	0	5	5	2	1	0	0	3	49
TOT	22	28	183	7	240	32	33	36	35	136	0	0	1	74	75	26	12	0	0	38	489

	CO	WO	ENL	CIV	TOTAL
DISTRICT TOTALS	80	73	220	116	489

c. Area Offices:

	CO	WO	ENL	CIV	TOTAL
Atlantic	7	3	3	0	13
Pacific	6	3	4	0	13
	13	6	7	0	26

d. Communications Stations and Radio Stations

Boston	1	3	64	0	68
Portsmouth	2	2	82	3	89
Miami	1	1	30	1	33
San Juan	0	1	18	0	19
New Orleans	0	2	43	0	45
San Francisco	2	3	87	1	93
Guam	0	1	27	0	28
Honolulu	1	1	49	0	51
Kodiak	2'	2	81	0	85
TOTAL	9	16	481	5	511

e. Laboratories

EECEN	17	6	87	25	135
Sta Alexandria	8	5	54	22	89
	25	11	141	47	224

f. Electronic Shops (ES, ESM, EST, ESMT, Yard)

	1	18	414	112	545

g. Other (Vessels, LORAN Stations, etc.)

	91	110	1475	0	1676
GRAND TOTAL	290	251	2788	366	3695

APPENDIX B
SDD-1

INTRODUCTION

SDD-1 [Ref. 26], was developed by Computer Corporation of America and supported by the Defense Advanced Research Project Agency of the U.S. Department of Defense. It is a system for managing data bases whose storage is distributed

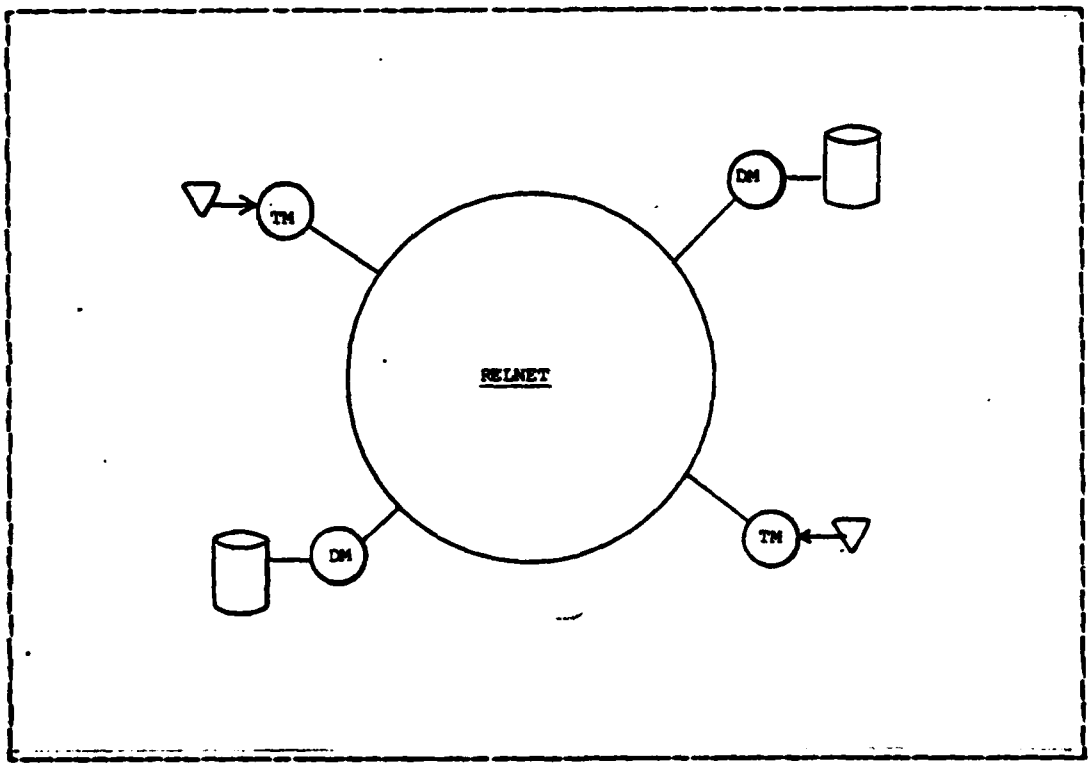

Figure B.1 SDD-1 Architecture.

over a network of computers. The general configuration of SDD-1 is given in Figure B.1 and consists of three virtual machines; Transaction Module (TM), Data Module (DM), and a

Reliable Network (RELNET). More about these 'virtual' machines later. SDD-1 is based upon relational data structures because complex data operations can be simpler and more precise than when data structures are used which are not two-dimensional. It is believed by the creators of SDD-1 that this type of architecture is appropriate for activities requiring access to a single pool of information distributed over a wide geographical area. It permits decentralized processing for performance, reliability and flexibility of function reasons.

Users interact with SDD-1 in a high level language called Datalanguage which is in fact a general purpose programming language. The data stored at the nodes of SDD-1 are portions of a relation which may be either a vertical subset composed of specified fields of the relation or a horizontal subset defined by one or more expressions e.g. Value of Field Id # = WAGB. Figure B.2 represents a 'horizontal' fragment and Figure B.3 represents a 'vertical'

SHIP	NAME	ID	HOMEPORT	STATUS
ICE_1A	GLACIER	WAGB-4	San Francisco	Ops
ICE_1B	NORTHWIND	WAGB-282	Wilmington	Charlie
ICE_1C	WESTWIND	WAGB-281	Milwaukee	Ops
ICE_2A	POLAR STAR	WAGB-11	Seattle	Ops
ICE_2B	POLAR SEA	WAGB-10	Seattle	Charlie

Figure B.2 Horizontal Fragment of a Relation.

fragment. A fragment is either completely present or completely absent at any node (DM) and any fragment may be stored redundantly at more than one Data Module (DM).

SHIP	NAME	ID	HOMEPORT	STATUS
ICE_1A	ICE_1A.1		ICE_1A.2	ICE_1A.3
ICE_1B	ICE_1B.1			ICE_1B.2
ICE_1C	ICE_1C.1	ICE_1C.2		ICE_1C.3
ICE_2A	ICE_2A.1	ICE_2A.2		
ICE_2B	ICE-2B.1			ICE_2B.2

ICE_1A.1 = (Name = GLACIER;Id = WAGB-4)
ICE_1A.2 = (Homeport = San Francisco)
ICE_1A.3 = (Status = Ops)
etc.

Figure B.3 Vertical Fragment of a Relation.

The implementation of a distributed database system presents three fundamental problems; concurrency control, distributed query processing, and reliable posting of updates. Each one of these problems can be associated with one of three functions for distributed database technology. For example, concurrency control can be associated with the functions of database management, likewise distributed query processing with management of distributed transactions, and reliable posting of updates with distributed data base management system reliability. The authors of SDD-1 have placed each one of these problems in a distinct processing phase - Read, Execute, and Write. Furthermore, each processing phase is contained in one of the 'virtual' machines mentioned earlier, i.e. Data Module (DM), Transaction Module (TM), and Reliable Network (RelNET). The functions of the Data Module are to respond to:

1. READ part of the Data Module's data base into a local workspace at that Data Module.

2. MOVE part of a local workspace from this Data Module to another Data Module.

3. MANIPULATE data in a local workspace at the Data Module.

4. WRITE part of the local workspace into the permanent data base stored at the Data Module. [Ref. 26: p. 5]

The functions of the Transaction Module plan and control the distributed execution of transactions. It performs the following tasks.

1. FRAGMENTATION: The Transaction Module translates queries on relations into queries on logical fragments and decides which instance of stored fragments to access.

2. CONCURRENCY CONTROL: The Transaction Module synchronizes the transaction with all other active transactions in the system.

3. ACCESS PLANNING: The Transaction Module compiles the transaction into a parallel program which can be executed cooperatively by several Data Modules.

4. DISTRIBUTED QUERY EXECUTION: The Transaction Module coordinates execution of the compiled access plan, exploiting parallelism whenever possible. [Ref. 26: p. 5]

Lastly, the Reliable Network interconnects the Data Modules and Transaction Modules to provide for:

1. GUARANTEED DELIVERY; allowing messages to be delivered even if the recipient is down at the time the message is sent, and even if the sender and receiver are never up simultaneously.

2. TRANSACTION CONTROL; a mechanism for posting updates at multiple Data Modules guaranteeing that either all Data Modules post the update or none do.

3. SITE MONITORING; to keep track of which sites have failed, and to inform sites impacted by failures.

4. NETWORK CLOCK; a virtual clock kept approximately synchronized at all sites. [Ref. 26: p. 5]

The rest of this appendix will be concerned with a description of how SDD-1 solves the three problems mentioned earlier.

CONCURRENCY CONTROL:

In a centralized database management system concurrency problems are handled by simply locking the database until the write operation has been completed. This has been the approach taken by most implementors of a distributed database technology; however, locking in a distributed database system with a high level of updates can cause substantial performance degradation because there is no activity allowed on the data until the update is complete. On a distributed system where the data has many copies the performance problem becomes even more severe.

A system will provide for concurrency control if the system can provide for serializability. The property of serializability is that the interleaved operation of the system is equivalent to one in which the transactions are run to completion, one at a time, serially. Hence serializability requires that whenever transactions execute concurrently, their effect must be identical to some serial (i.e. non-interleaved) execution of those same transactions. Two major types of conflicts occur to prevent serializability; update conflicts and read conflicts. If transactions are grouped into classes based on which Data Module executes the transaction and what data the Transaction Module uses then potential conflicts between classes can be analyzed. This is exactly what is done by the data base administrator at design time, and is known as conflict analysis.

With conflict graph analysis the Data Base Administrator (DBA) defines transaction classes which are named groups of commonly executed transactions. These transactions are defined by its name, read-set, write-set and Transaction Module at which it runs. Figure B.4 is an example of an update conflict. Through conflict graph analysis four

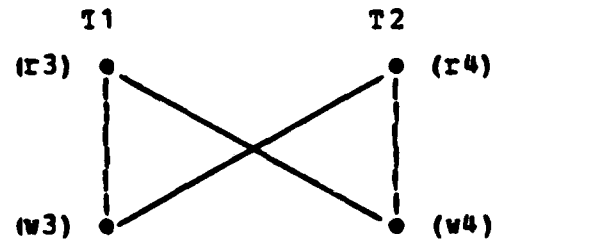

Let T1, T2 = Transaction Modules
Let r3, r4 = Read-sets
Let w3, w4 = Write-sets

Figure B.4 Conflict Graph.

different situations may arise each requiring different synchronization protocols. They are; no conflict, update conflict, read conflict with one Data Module, and read conflict with more than one Data Module.

To solve the update conflict the two transactions must be synchronized, i.e. one must be run first and the other delayed until the first transaction is completed. This order of processing transactions is determined by the total ordering of transactions induced by timestamps. This is accomplished by piping which requires each Transaction Module to send its Write commands in timestamp order. One should note that the 'virtual' clock does not have to be absolutely synchronized at each site as the Transaction Module identifier is attached to each timestamp, thus ordering the timestamps globally.

Thus through conflict graph analysis a DBA can loosen the restrictions on the synchronization protocols for those transactions that may not conflict. This in turn increases

system performance. The technique of piping achieves consistency and serializability throughout the system.

DISTRIBUTED QUERY PROCESSING

In processing a query in a distributed mode one can take the straight forward approach and move all the transactions' read-sets to a single Data Module and then execute the transaction. This however, is quite prohibitive in terms of communication costs for very large read-sets and no use is made of parallel processing. SDD-1 attempts to solve this problem by splitting query processing into two distinct phases. The first phase reduces the read-set as much as possible without changing the transaction's answer. This reduces the amount of data to be transmitted between sites and makes efficient use of parallel processing as a transaction may be worked on concurrently at more than one Transaction Module at once. The final phase transmitts the reduced read-set to the final Data Module where the transaction is executed. A temporary file (buffer) is created at the final Data Module so that the answer to the transaction can be written to the database or displayed whichever was requested.

RELIABLE WRITING

There are three potential problems that must be solved to insure that updates are reliable. The first is to guard against a failure of a receiving Data Module, i.e. must have reliable delivery. The second is to take appropriate measures when a failure occurs at the point of origin, i.e. transaction control. Finally the system must insure updates to different transactions are installed in the same 'effective' order at all Data Modules.

The first problem is handled by a mechanism the creators called spooler. It is a function/process with access to secondary storage devices that serve as first-in, first out (FIFO) message queues for a failed site. Thus any message sent to a site which has failed will be re-routed to the spooler instead. Once the failed site comes back on-line updates from the spooler are processed by the failed Data Module.

The second problem - transaction control - is handled in a more subtle way with a technique that employes a variant of the two-phase commit. During phase 1 the file 'F' is segregated into 'n' number of files corresponding to the receiving Data Modules. For example file 'F' will be split into files F(1),...,F(n) to be sent to Data Modules DM(1),...,DM(n). The receiving Data Modules do not install these files yet. This completes phase 1. During phase 2 the Data Module that originated the transaction sends a commit message to each DM(1),...,DM(n) where upon each DM(i) installs F(i). If some DM(k) has received F(k), but not a commit, and the Data Module that originated the transaction has failed, then the DM(k) consults other DM's affected by the transaction. If any of the affected DM's have received a commit then DM(k) proceeds to install F(k). If no other affected DM has received a commit the transaction is aborted.

Problem three, insuring updates to different transactions are installed in the same 'effective' order at all Data Modules, is handled quite easily. Since every physical data item in the data base is timestamped with the most recent updated transaction, and each update (Write command) carries the timestamp of the Transaction Module that generated it; the following write rule is applied: For each data item, 'x' in the update, the value of 'x' is modified at the Data Module if and only if x's stored timestamp is less than

the timestamp of the write command. The overhead of timestamps on each data item is reduced to acceptable levels by caching the timestamps according to the authors of SDD-1.

SUMMARY

SDD-1 relies heavily on the use of directories to maintain the location of data throughout the system. For SDD-1 to be effective, the management of these directories must be efficient and flexible. The approach taken by the authors is to treat the directories just like data, hence directory management is a design issue. Therefore, the directories can be fragmented, distributed with arbitrary redundancy and updated from arbitrary Transaction Modules.

The development of SDD-1 uses a relational database management system called Datacomputer and runs on PDP-10 equipment using the ARPANET and the communication network. The different modules require a fair amount of main memory for the object code, but not an unreasonable amount. There does however, exist two fundamental problems with SDD-1 for the Coast Guard. The first, the equipment is all homogeneous, hence complexities are reduced somewhat. Here the Coast Guard cannot expect to ever be in the position of having homogeneous systems in the near future. Secondly, the Data Base Administrator (DBA) has a difficult job just to create a logical database design. In fact there are few examples of corporations that have had great success with this fundamental approach to the use of database technology. The Data Base Administrator's job becomes ever more complex by having to perform the conflict graph analysis and splitting transactions into different classes. Although the Coast Guard has many talented people, to find this required talent would be a herculean task.

LIST OF REFERENCES

1. THE COAST GUARD CAPITAL INVESTMENT PLAN, perspective: FY 1977-1986, CG 380 Goverment Printing office, 1976.

2. "Location of the Coast Guard Within the Federal Government," The Bulletin - U.S. Coast Guard Academy Alumni Association Vol. 43, No 4, New London, Ct, July/August 1981.

3. A Conversation with Admiral John B. Hayes - Cost Effectiveness in the Coast Guard. American Enterprise Institute for Public Policy Research, Washington, D.C., 1981

4. Command, Control and Communications Support Program Plan FY 82-92, Office of Command, Control and Communication, U.S. Coast Guard, Washington, D.C., 1981.

5. Commandant Instruction M5230.8A Automated Data Processing (ADP) Plan, Draft, (Undated - Received August 1982).

6. Nicholson, C.M., LCDR, USCG Coast Guard Information Architecture Concept Paper (DRAFT) Office of Command, Control and Communications (G-T) Information Architecture Project, 11 February 1982.

7. "DBMS for Mini-computers," EDP Analyzer, Vol. 19 No. 3, March, 1981.

8. Lorin, H. Aspects of Distributed Computer Systems. John Wiley and Sons, 1980.

9. "The Coming Impact of New Technology," EDP Analyzer, Vol. 19, No. 1, March, 1981.

10. Contract Number: DTCG-23-81-C-30079, ASSESSMENT AND ANALYSIS OF DATA AND ANALYSIS OF DATA AND INFORMATION PROCESSING REQUIREMENTS AT U.S. COAST GUARD DISTRICT OFFICES, Final Report, 2 July 1982.

11. Hussain, Donna & K. M., Information Processing Systems for Management. Richard D. Irwin, Inc., Homewood, Illinois, 1981.

12. Smith, L. B., "The Use of Interactive Graphics to Solve Numerical Problems," Communications of the ACM, Vol. 13, No. 10, October, 1970.

13. Bray, O.H., Freeman, H.A., Data Base Computers, Lexington Books, 1981.

14. Martin, James, Design and Strategy for Distributed Data Processing, Prentice-Hall, Inc., 1981.

15. The CODASYL Systems Committee, Proceedings, National Computer Conference, Distributed Data Base Technology: An Interim Report of the CODASYL Systems Committee, 1978.

16. Bray, O. H. Distributed Database Management Systems, Lexington Books, 1981.

17. Champine, George A., "Four Approaches to a Data Base Computer", Datamation, December 1978.

18. Olle, T. William, The Codasyl Approach to Data Base Management, John Wiley & Sons, 1978.

19. Upham, David, "Commercially Available Relational DBMS Compared," Computerworld, Vol. XVI, No. 2, p. 51-52, 11 January 1982.

20. DataPro Research Corporation, DataPro 70 the EDP Buyer's Bible, Delran, New Jersey, p. 70E-132-01b - 70E-132-01c, April 1980.

21. Miller, Robert B., Response Time in Man-Computer Conversational Transactions, Fall Joint Computer Conference Proceedings, 1968.

22. Martin, James, Computer Data-Base Organization, Prentice-Hall, Inc., 1975.

23. Martin, James, Security, Accuracy, and Privacy in Computer Systems, Prentice-Hall, Inc., 1973.

24. Department of Commerce, National Bureau of Standards, Federal Information Processing Standards (FIPS) Publication 76, Specifications for Guideline for Planning and Using a Data Dictionary System, 20 August 1980.

25. Ullman, Jeffrey D., Principles of Database Systems, Computer Science Press, 1980.

26. Rothnie Jr., J.B., Bernstein, P.A., Fox, S., Goodman, N., Hammer, M., Landers, T.A., Reeve, C., Shipman, D.W., and Wong, E., "Introduction to a System for Distributed Databases (SDD-1)," ACM Transactions on DataBase Systems, Volume 5 Number 1, March 1980.

INITIAL DISTRIBUTION LIST

		No. Copies
1.	Defense Technical Information Center Cameron Station Alexandria, Virgina 22314	2
2.	Library, Code 0142 Naval Postgraduate School Monterey, California 93940	2
3.	Department Chairman, Code 54 Administrative Sciences Department Naval Postgraduate School Monterey, California 93940	1
4.	Curricular Office, Code 37 Computer Technology Naval Postgraduate School Monterey, California 93940	1
5.	Associate Professor Norman R. Lyons, Code 54Lb Department of Administrative Sciences Naval Postgraduate School Monterey, California 93940	1
6.	LCDR John R. Hayes, Code 54Ht Department of Administrative Sciences Naval Postgraduate School Monterey, California 93940	1
7.	Lt. Mark Fiegel 115 North Meadow Drive Glen Burnie, Maryland 21061	1
8.	CDR Robert Williams c/o Commandant (G-TPP-2) 2100 Second Street S.W. Washington, D.C. 20593	1
9.	Lt Stephen H. Goetchius c/o Commander (dt) First Coast Guard District 150 Causeway St. Boston, Massachusetts 02114	2
10.	Commander (dt) First Coast Guard District 150 Causeway St. Boston, Massachusetts 02114	2