

An Analysis of Selected Enhancements to The En Route Central Computing Complex

William Broadley
Harvey Freeman
James Oiesen
Ronald Rutledge
Kenneth Thurber

Transportation Systems Center
Cambridge MA 02142

September 1981
Final Report

This document is available to the public
through the National Technical Information
Service, Springfield, Virginia 22161.



U.S. Department of Transportation
Federal Aviation Administration

Office of Systems Engineering Management
Washington DC 20591

PREFACE

This report summarizes work sponsored by the U.S. Department of Transportation, Federal Aviation Administration, Office of Systems Engineering Management. This work was performed under the leadership of the Transportation Systems Center of the U.S. Department of Transportation.

This report analyzes selected hardware enhancements that could improve the performance of the 9020 computer systems, which are used to provide en route air traffic control services. These enhancements could be implemented quickly, would be relatively inexpensive, and would provide a solution to the short-term but not the long-term problems that the system faces.

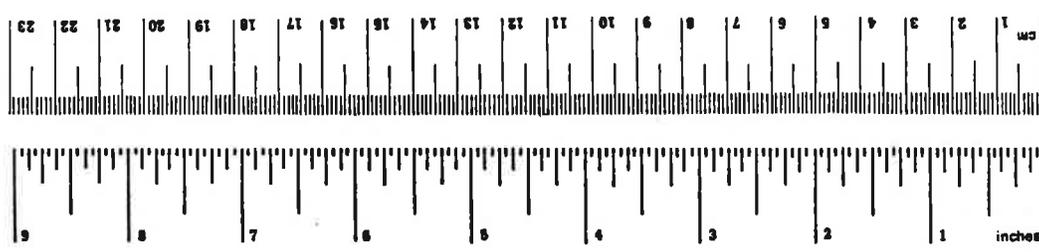
During the course of this study we have received assistance from a number of individuals and organizations, and we would like to express our appreciation for this help. We would like to thank Alfred Cocanower of HH Aerospace, who played a key role in initiating this project. His support, encouragement, and counsel during this work was extremely valuable. Arthur Chantker of the FAA provided invaluable assistance by sharing his expertise and experience on performance, by helping to set up the work at the FAA Technical Center (FAATC), and by giving generously of his time. Especially helpful were personnel in the Data Engineering and Development Division at the FAATC; personnel in the Hardware Engineering Branch and the Software Engineering Branch kindly explained numerous details and otherwise eased our task. Also, OAO personnel helped with the data analysis programs.

Dr. Andres Zellweger, Joseph DeMeo, and Robert Roche of the FAA encouraged and sponsored this work.

The authors of this report are Kenneth Thurber and Harvey Freeman (Architecture Technology Corp.), James Oiesen (HH Aerospace Design Co., Inc.), William Broadley (Pittsburgh Digital Systems, Inc.), and Ronald Rutledge (Transportation Systems Center).

METRIC CONVERSION FACTORS

Symbol	When You Know	Multiply by	To Find	Symbol
LENGTH				
mm	millimeters	0.04	inches	in
cm	centimeters	0.4	inches	in
m	meters	3.3	feet	ft
km	kilometers	0.6	miles	mi
AREA				
cm ²	square centimeters	0.16	square inches	in ²
m ²	square meters	1.2	square yards	yd ²
km ²	square kilometers	0.4	square miles	mi ²
ha	hectares (10,000 m ²)	2.5	acres	ac
MASS (weight)				
g	grams	0.035	ounces	oz
kg	kilograms	2.2	pounds	lb
t	tonnes (1000 kg)	1.1	short tons	st
VOLUME				
ml	milliliters	0.03	fluid ounces	fl oz
l	liters	2.1	pints	pt
l	liters	1.06	quarts	qt
l	liters	0.26	gallons	gal
m ³	cubic meters	36	cubic feet	ft ³
m ³	cubic meters	1.3	cubic yards	yd ³
TEMPERATURE (exact)				
°C	Celsius temperature	9/5 (then add 32)	Fahrenheit temperature	°F



Symbol	When You Know	Multiply by	To Find	Symbol
LENGTH				
in	inches	*2.5	centimeters	cm
ft	feet	30	centimeters	cm
yd	yards	0.9	meters	m
mi	miles	1.6	kilometers	km
AREA				
in ²	square inches	6.5	square centimeters	cm ²
ft ²	square feet	0.09	square meters	m ²
yd ²	square yards	0.8	square meters	m ²
mi ²	square miles	2.6	square kilometers	km ²
ac	acres	0.4	hectares	ha
MASS (weight)				
oz	ounces	28	grams	g
lb	pounds	0.45	kilograms	kg
	short tons (2000 lb)	0.9	tonnes	t
VOLUME				
teaspoon	teaspoons	5	milliliters	ml
Tablespoon	tablespoons	15	milliliters	ml
fl oz	fluid ounces	30	milliliters	ml
c	cups	0.24	liters	l
pt	pints	0.47	liters	l
qt	quarts	0.96	liters	l
gal	gallons	3.8	liters	l
ft ³	cubic feet	0.03	cubic meters	m ³
yd ³	cubic yards	0.76	cubic meters	m ³
TEMPERATURE (exact)				
°F	Fahrenheit temperature	5/9 (after subtracting 32)	Celsius temperature	°C

* 1 in x 2.54 (exactly). For other exact conversions and more detailed tables, see NBS Misc. Publ. 286, Units of Weights and Measures, Price \$2.25, SD Catalog No. C13.10.286.

CONTENTS

	<u>Page</u>
Executive Summary	vii
 <u>Chapter</u>	
1. INTRODUCTION	1
1.1 Purpose and Organization of this Report	1
1.2 The IBM 9020 Computer Systems	2
1.3 Bottlenecks in the 9020A and 9020D Computer Systems	3
2. MEMORY ENHANCEMENTS	9
2.1 Purpose and Organization of this Chapter	9
2.2 Replacement of Memory Boxes	9
2.2.1 Organization of this Section	9
2.2.2 Description of this Enhancement	9
2.2.3 Advantages of this Enhancement	12
2.2.4 Cost	15
2.2.5 Schedule	17
2.2.6 Transition	18
2.2.7 A Variant: Replace All of the 9020D Memory	19
2.3 Replacement of the Memory Stacks in the SE's	19
2.4 Replacement of the Memory Stacks in the IOCE's	22
3. PROCESSOR ENHANCEMENTS	24
3.1 Purpose and Organization of this Chapter	24
3.2 Speed-Up of the 9020A CE Processors	24
3.2.1 Description of this Enhancement.	24
3.2.2 Advantages of this Enhancement	27
3.2.3 Cost and Schedule	30
3.2.4 Transition	31
3.3 Speed-Up of the IOCE Processors	31
3.4 9020A CE Replacement	34
3.4.1 Description of this Enhancement	34
3.4.2 Advantages of this Enhancement	35
3.4.3 Cost and Schedule	35
3.4.4 Transition	35
4. SUMMARY	37
4.1 The Individual Enhancements	37
4.2 Strategies Open to the FAA.	42
 <u>Appendix</u>	
A. THE MODEL OF SYSTEM PERFORMANCE: OVERVIEW	49
B. THE MODEL OF SYSTEM PERFORMANCE: DETAILED EXPOSITION	54
C. NAS REPRESENTATIVE 9020A WORKLOADS	75
D. THE STRATEGIES OPEN TO THE FAA	84
REFERENCES	88

LIST OF FIGURES

<u>Figure</u>	<u>Title</u>	<u>Page</u>
ES-1	Leading Strategies Open to the FAA	xiv
1-1	Simplified 9020A Configuration Diagram	4
1-2	Simplified 9020D Configuration Diagram	5
4-1	Leading Strategies Open to the FAA	45
B-1	The Flow Diagram for the Memory Module Process Module	56
B-2	The Flow Diagram for the Processing Element Process Model	60

LIST OF TABLES

<u>Table</u>	<u>Title</u>	<u>Page</u>
ES-1	Characteristics of the Six Enhancements	xii
1-1	Computer System Configurations for the ARTCC's.	6
1-2	The Critical 9020 Resources	8
2-1	Estimated Cost of Replacing the Memory Boxes	16
3-1	Increased Processing Capacity Due to the CE Speed-Up	29
4-1	Characteristics of the Six Enhancements	39
4-2	Cost and Schedule for Developing the Prototypes	41
B-1	9020A CE Command Model	65
B-2	IOCE Shared Memory Utilization Model	66
B-3	Memory Conflict in 3x2 9020A	68
B-4	Extrapolated Values from Simulation with Various Memory Speedup Options	69
B-5	Overall Performance Improvement Ratios.	69
B-6	Performance Ratios of a 3x2 Multiprocessor with an Eight Port Memory Box	71
B-7	9020A Memory Performance Summary	71
B-8	Effect of Module Number and Size, <u>Not</u> Including Interleaving or Speedup, for a 3x2 Multiprocessor	73
B-9	Effect of Inclusion of a Fourth CE	73
C-1	PE Utilization at 111 Tracks	76
C-2	PE Utilization at 222 and 333 Tracks	78
C-3	Bufferable PE I/O Loading	80
C-4	Memory Loading - Dynamic Buffering.	80
C-5	Memory Loading - No Buffering	82
C-6	I/O Loading	82
C-7	Memory Map	83
D-1	The Possible Strategies Open to the FAA	85

EXECUTIVE SUMMARY

Introduction. The Federal Aviation Administration is now considering ways that the IBM 9020 computer systems, which are used to provide en route air traffic control services, can be upgraded or replaced. The purpose of this report is to give a thorough discussion of some hardware enhancements that could be adopted to upgrade the system. The enhancements discussed in this report fall into the category of actions that could be taken quickly, would be relatively inexpensive, and would provide a solution to the short-term but not the long-term problems that the system faces.

There are three primary short-term problems that the 9020's face. (This report is concerned with two versions of the 9020's, the 9020A and 9020D; there are ten of each in the field.) First, there are potential I/O problems in the areas of bandwidth and device speed for both the 9020A and the 9020D. Second, there is insufficient main memory in both the 9020A and 9020D; moreover, the 9020A has a problem in the area of memory bandwidth. Third, the 9020A has insufficient processing capacity; the 9020D has no problem in this area. In short, these I/O, memory, and processing capacity problems form the context in which any enhancements are to be judged.

This report deals with three memory enhancement and three processor enhancements. Each enhancement is discussed with respect to its description, advantages, risk, cost, schedule, and transition.

Memory enhancements. The first memory enhancement is to replace the 9020 memory boxes, also called storage elements (SE's), with new boxes containing state of the art memory. This enhancement has two main features. First, each system would have enough memory so that all program elements and data would be resident in main memory (with some minor exceptions). Second, the speed of the 9020A's memory would be significantly increased. These features have numerous implications. Because all programs and data would be resident in main memory, buffering would be virtually eliminated. This would decrease I/O activity by 30 to 50 percent, and this would take care of the potential I/O problems. Moreover, having enough main memory to hold almost all program elements and data would also take care of

the memory problems. Therefore, this one enhancement would take care of both the potential I/O and memory problems. Since these are the only problems faced by the 9020D, this one enhancement is sufficient to deal with the 9020D's problems.

This enhancement also deals somewhat with the 9020A's processing capacity problems. The elimination of buffering and the decrease in memory interference due to the faster memory would improve the 9020A's processing capacity by at least 20 percent and perhaps by as much as 60 percent. Further modeling of the 9020A system will be necessary before this estimate can be made more precise. ("Processing capacity" in this report is taken to mean the size of the peak traffic load that the system can handle.) If the increase in 9020A processing capacity yielded by this enhancement is considered adequate, then this enhancement deals with all the problems for both the 9020A and 9020D.

In addition to dealing with these problems, replacing the memory boxes yields three other advantages. First, because there is enough main memory to hold all program elements and data, software maintenance will be made much easier. Currently, the need to deal with the memory constraints greatly complicates and adds to the expense of software maintenance. It could turn out that by easing software maintenance this enhancement could quickly pay for itself.

Second, functional enhancements can be added to the system once the memory constraint is lifted. That is, there are plans to add further capabilities to the system, but these plans are being slowed by the difficulties imposed by the limited memory. With sufficient memory available, these functional enhancements can be implemented more quickly.

Third, system reliability will increase since the new, modern technology memory units would be more reliable than the old.

The cost of replacing the memory boxes at the 23 9020 sites is estimated to be \$8.2 million. Once the FAA places the order for the memory units, 24 months will elapse before the memory replacement is completed at the first

six sites, and 38 months will elapse before the memory replacement is completed at all sites.

This enhancement has virtually no risk. The technical risk is minimal since the memory units being purchased are fairly standard and since there is experience with similar replacements. The financial risk is small since at least six firms are expected to bid; thus, there should be sufficient competition to keep the price down.

The transition when the new units are installed is expected to be smooth since no major changes are anticipated. The system downtime when a memory unit is installed is estimated to be two hours.

The second memory enhancement is to replace not the entire memory boxes but just the memory stacks in the SE's; the memory stacks are the components of the SE's that actually hold the data. Since replacing the stacks would result in the same system performance as replacing the boxes, this enhancement would deal with the 9020's problems and provide the same three advantages as the previous enhancement.

There are five main differences between these two enhancements. First, replacing just the stacks results in a lower cost, i.e., \$5.6 million v. \$8.2 million for memory box replacement, since only the stacks and not the rest of the SE must be purchased. Second, replacing just the stacks is faster, i.e., the first six sites can be enhanced in 8 months v. 24 months for memory box replacement, since only the stacks must be designed and fabricated. Third, the physical installation would be easier with stack replacement since no recabling would be required. Fourth, replacing the stacks does not require that the decision on how many sites are to be enhanced be made in advance, and it does not require long lead time parts, so it gives the FAA more flexibility in deciding how many centers to enhance. Fifth, the memory box replacement would offer the advantage of being a unified design.

The third memory enhancement is to replace the memory stacks in the input-output control elements (IOCE's). This enhancement would allow

program elements to be moved from the 9020's shared memory to the IOCE's memory, and these program elements would then be executed by the IOCE. Further study of this enhancement will be needed before it can be said to what degree it will take care of the 9020's problems; it seems likely, however, that it will increase the processing capacity of the 9020A's by between 10 and 30 percent. To implement this enhancement at the 9020A and 9020D sites would cost an estimated \$3.5 million; it would take 8 months to enhance the first six sites.

Processor enhancements. If it is decided that the memory replacement does not provide a sufficient increase in processing capacity for the 9020A, then there are three processor enhancements that might be adopted to further increase the processing capacity.

The first processor enhancement is to speed up the processors in the 9020A compute elements (CE's). This enhancement consists of replacing the two components of the CE that constrain its speed, the local store and the read only store, with modern, faster components; the CE would then be retuned to take advantage of this faster speed. The gain in processing capacity provided by this enhancement (in conjunction with the memory replacement) is estimated to be between 25 and 100 percent. This enhancement is estimated to cost \$2.0 million; it could be implemented at the first six sites within six months, provided that faster 9020A memory is in place. For this enhancement as well as for the other two CE enhancements, the system downtime during the transition is measured in minutes.

The second processor enhancement is to speed up the processors in the IOCE's. This enhancement would be achieved just as with the CE speed-up; the only difference is that the IOCE's internal memory would need to be replaced with faster memory. The gain in processing capacity provided by this enhancement is estimated to be between 15 and 70 percent (where the basis for comparison is the standard 9020A system). The uncertainty in this estimate would be eliminated once the engineering prototype is completed and its performance is simulated. This enhancement is estimated to cost \$1.6

million if implemented at the 9020A sites and \$2.9 million if implemented at both the 9020A and 9020D sites; it could be implemented at the first six sites within 6 months.

With both of these first two enhancements there is a question as to whether it will be feasible to retune the CE so that the expected gain in performance can be achieved. Current understanding of the CE is not sufficient to say whether there is some complicated timing interaction that would prevent these enhancements from being successful. It would take about \$125,000 and five months to determine whether these enhancements are feasible.

Third, if the speed-up proves infeasible or if it does not provide a sufficient gain in performance, then the 9020A CE's could be replaced by a computer in the one million instruction per second class. This enhancement would provide an increase in processing capacity of between 100 and 200 percent and is estimated to cost \$15.6 million. It would take 24 months to enhance the first six sites. There is virtually no risk associated with this enhancement.

Summary. Table ES-1 summarizes the main characteristics of each of the six enhancements. The first column shows the cost of the enhancement; the cost is shown for implementing the enhancements at both the 9020A and 9020D sites or at just the 9020A sites, depending on what is relevant to each enhancement. The second column shows the increase in processing capacity, and the third gives the estimated probability that this increase can actually be achieved. For example, the enhancement of speeding up the processor in the 9020A CE in conjunction with one of the SE memory enhancements provides an increase in processing capacity of at least 25 percent with probability of 0.98, of at least 50 percent with probability 0.88, and of at least 100 percent with probability 0.49. In order to lower the uncertainty in these estimates, it will be necessary to obtain further data by building an engineering prototype and to do additional simulation modeling. This data-gathering and modeling is also needed for design purposes. The last column in the table shows how long it will take for the enhancement to be implemented at the first six sites once the FAA has placed

TABLE ES-1: CHARACTERISTICS OF THE SIX ENHANCEMENTS

<u>Enhancement</u>	<u>Cost</u> (millions)	<u>Processing Capacity</u> ¹		<u>Schedule</u> (first six sites) (months)
		<u>Increase</u> (%)	<u>Probability</u> ⁴ (%)	
1. Replace SE memory boxes	A&D:\$8.2	A: 20-60 D: 10-30	100 100	24
2. Replace SE memory stacks	A&D: 5.6	A: 20-60 D: 10-30	100 100	8
3. Replace IOCE memory stacks	A: 1.9 A&D: 3.5	A: 10-30 D: 5-15	100 100	8
4. CE Speed-Up ²	A: 2.0	A: 25 A: 50 A: 100	98 88 49	6
5. IOCE Speed-Up ³ memory stacks	A: 1.6 A&D: 2.9	A: 15 A: 30 A: 70 D: 10	98 88 49 88	6
6. CE Replacement ²	A: 15.6	A: 100- 200	100	24

¹ Processing capacity refers to the peak number of tracks that can be handled. This increase is relative to the standard 9020 configuration.

² A prerequisite for this enhancement is replacement of either the memory boxes or the SE memory stacks. The cost of this enhancement excludes the cost of the prerequisite; the increase in processing capacity, however, is the increase that would result from adopting both this enhancement and its prerequisite.

³ A prerequisite for this enhancement is replacement of the IOCE memory stacks. The cost of this enhancement excludes the cost of the prerequisite; the increase in processing capacity, however, is the increase that would result from adopting both this enhancement and its prerequisite.

⁴ These probabilities are best estimates based on a study of the system and on experience; they should not be interpreted as exact probabilities.

the order for the hardware. This time does not include the time needed for design or for building a prototype.

Some of the ways that the FAA could combine these individual enhancements into a comprehensive strategy for dealing with the 9020's potential problems are illustrated in the simplified decision tree in Figure ES-1. The initial decision faced by the FAA is at fork 1 where the FAA would decide whether as a first step in upgrading the 9020's it would be better to replace the SE memory or to upgrade the IOCE's. Suppose that the FAA decides to replace the SE memory; a further choice not shown in this simplified diagram is whether the SE memory should be replaced by replacing the memory boxes or by replacing the memory stacks. Since replacing the SE memory takes care of the memory and I/O problems and provides a modest increase in processing capacity, the FAA at fork 2 might decide that nothing else needs to be done. If, however, the FAA decided that more processing capacity is needed, it can speed up the processors in the 9020A CE's, thus arriving at fork 3. (Not shown in this simplified diagram is the option of increasing processing capacity by replacing the CE's.)

If the FAA is at fork 3 and decides that enough processing capacity has been achieved, then it need do nothing else. If, however, more processing capacity is desired, the FAA can upgrade the IOCE's at the 9020A sites. (Since the SE memory replacement would take care of the 9020D's problems, there would be no need to upgrade the IOCE's at the 9020D sites.) Upgrading the IOCE's means that the IOCE memory stacks are replaced and the IOCE processors are sped up; this simplified diagram does not consider just replacing the IOCE memory stacks.

Suppose now that back at fork 1 the FAA had decided to upgrade the IOCE's instead of replacing the SE memory. This places the FAA at fork 4. If the FAA decides that the IOCE upgrade provides all the needed capabilities, then there would be no need to do anything else. If the IOCE upgrade is not sufficient, then the FAA could further enhance the system by replacing the SE memory and speeding up the processors in the CE's. (Just replacing the SE memory at this stage probably would not be a good idea since the IOCE upgrade would have provided the system with sufficient memory.)

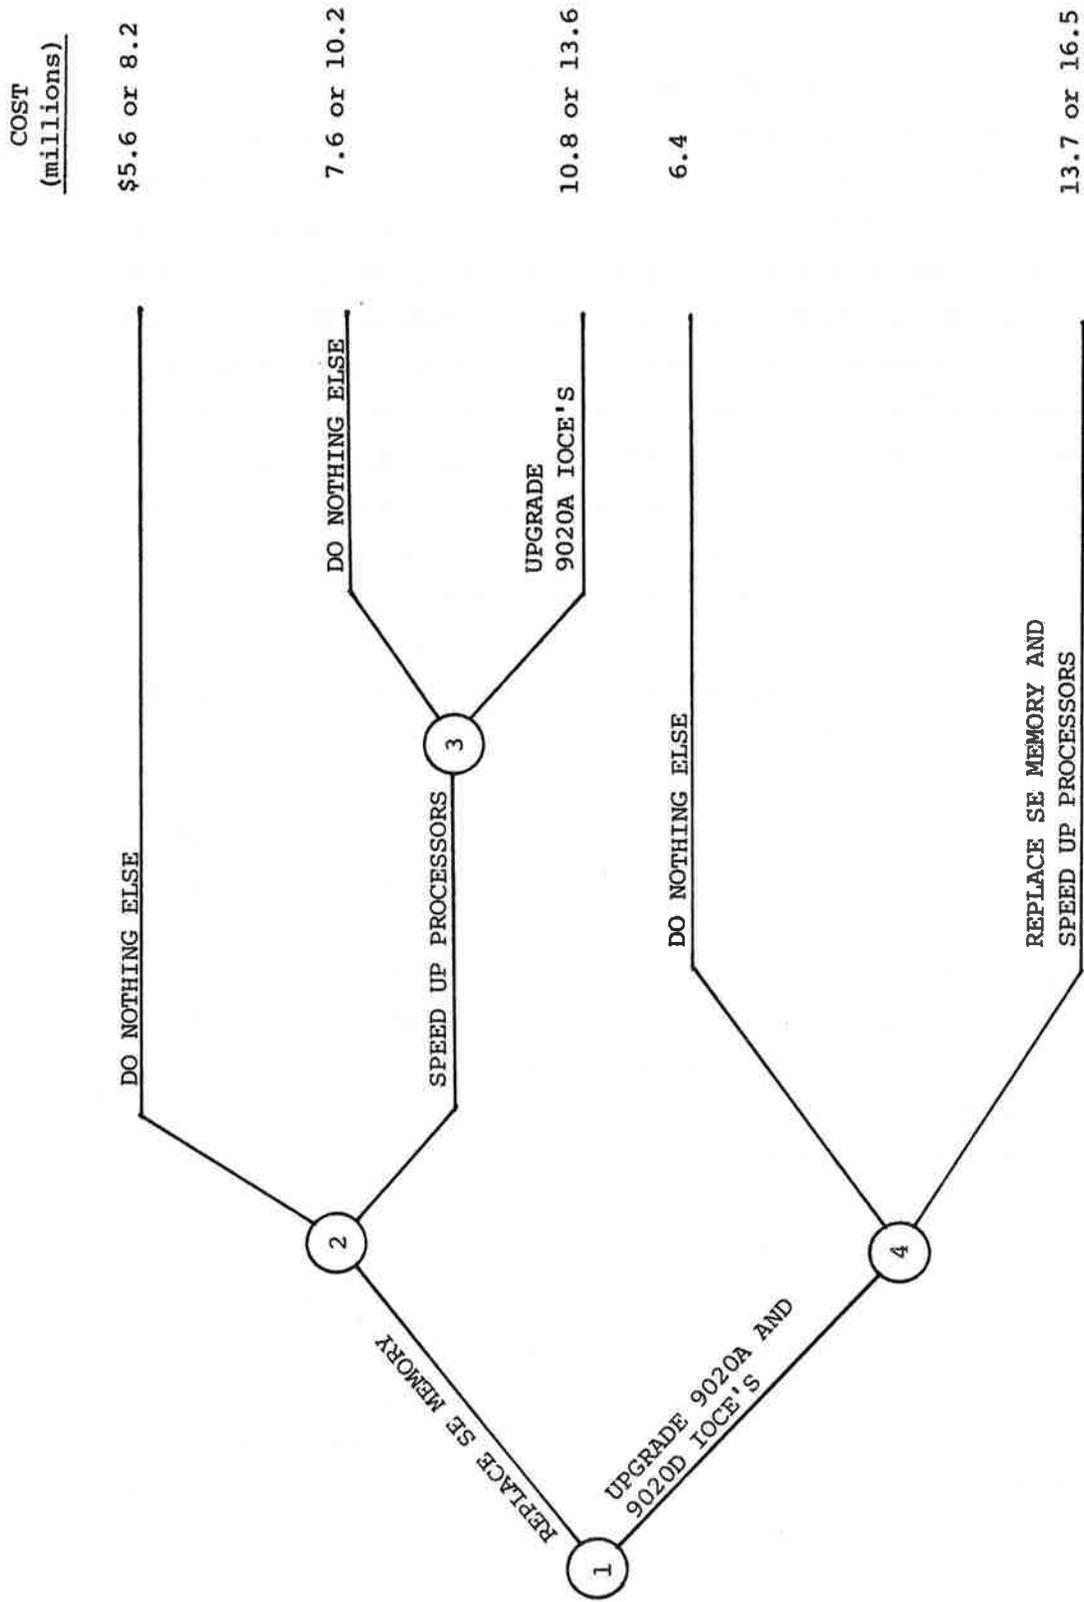


FIGURE ES-1: LEADING STRATEGIES OPEN TO THE FAA

The estimated cost of each strategy is shown in Figure ES-1. This cost reflects the interactions between the various enhancements. Each path that includes "Replace SE memory" has two costs depending on whether the memory stacks or the memory boxes are replaced.

Depending on how much processing capacity is needed, when it is needed, how much each enhancement can provide, and the cost, the FAA can select a path through this decision tree (or perhaps select one of the paths omitted from this simplified diagram) and in this way define a strategy for dealing with the 9020's potential problems.

One all-important point that should be stressed is that the FAA will be in a much better position to decide what combination of enhancements should be adopted once the task of developing working prototypes of the various enhancements is completed; only when the working prototypes are in hand will the FAA know which enhancements are feasible and how much they will contribute to system performance. Since the cost of developing the prototypes is trivial compared to the amounts involved and since the prototype development is critical for providing the information needed as a basis for decisions, proceeding with the prototype development is an immediate step that can make a substantial contribution to dealing with the problems that face the 9020's.

1. INTRODUCTION

1.1 Purpose and Organization of this Report

One of the missions of the Federal Aviation Administration (FAA) is to provide en route air traffic control services. To fulfill this mission the FAA has placed at each air route traffic control center (ARTCC) a computer system that supplies the information that air traffic controllers need; that is, these computer systems keep current the displays that show the location and other characteristics of the aircraft being controlled, and they also print the flight strips that contain detailed information about each flight. These computer systems have been in place and supporting air traffic control (ATC) for about a decade and can be expected to provide effective support for some time to come. These systems, however, will not last forever, and eventually they will need to be upgraded or replaced.

The FAA is considering a number of steps that might be taken to improve the system. These steps range from minor tuning of the system to full-scale replacement. The FAA is currently conducting studies that examine the pros and cons of each step and how the various steps can be fitted together to form a strategy specifying what should be done over the next twenty or thirty years.

The purpose of this report is to discuss some hardware enhancements that can potentially deal with the main problems that the en route computers face over the next ten years, that promise additional advantages, that have a relatively small cost, and that can be quickly implemented. These enhancements fall into the two areas of memory and processor enhancements. Chapter 2 discusses the memory enhancements:

- Replace the memory boxes,
- Replace the memory stacks in the storage elements, and
- Replace the memory stacks in the input-output control elements.

Chapter 3 discusses the processor enhancements:

- Speed up the processors in the compute elements,
- Speed up the processors in the input-output control elements, and
- Replace the compute elements.

Each enhancement is discussed from the following viewpoints.

- Description of the enhancement: What must be replaced, retuned, or otherwise changed?
- Advantages: What are the potential benefits and what is the probability that these benefits will actually be achieved?
- Cost: How much would this enhancement cost?
- Schedule: How long would it take for this enhancement to become operational?
- Transition: What physical modifications would be necessary at each ARTCC and how much system downtime would the enhancement entail?

Chapter 4 shows how the individual enhancements can be combined into strategies for dealing with the potential problems. The rest of this chapter provides background on the current computer system.

1.2 The IBM 9020 Computer Systems

This section describes the computer systems that are now used in providing en route air traffic control services. The computer system at each ARTCC has two parts. First, the central computer complex (CCC)

receives inputs from the radar, flight service stations, controllers, and other sources and then performs the flight data processing and radar data processing. Second, the display channel takes the output from the CCC and uses it to keep each controller's plan view display current. The CCC and display channel together, then, take the raw data that is available, process it, and provide it to the controllers in a way that can be readily grasped and acted on.

There are two different but related computer systems that serve as CCC's, the IBM 9020A and IBM 9020D systems. The main elements in these systems are the compute elements (CE's), storage elements (SE's), input/output control elements (IOCE's), peripheral adapter modules (PAM's), tape units, and disk units. Figures 1-1 and 1-2 show the 9020A and 9020D systems, respectively. These figures show the number of components in each system; the components to the right of the dashed lines are redundant components that are held in reserve in case of a failure. (One additional storage element has been recently added to each 9020A and 9020D and is not shown in these figures.) The CE's and SE's of the 9020A are based on IBM 360/50 engineering; the CE's and SE's in the 9020D are based on IBM 360/65 engineering. The IOCE's, which are identical in the two systems, are based on IBM 360/50 engineering.

There are also two different computer systems that serve as the display channel, the IBM 9020E and the Raytheon 730. The 9020E is almost identical to the 9020D except that some of the storage elements have been replaced by display elements. Since the display channels do not appear to be a bottleneck that degrades system performance, this report will not discuss the display channels.

Table 1-1 shows which versions of the CCC and display channel are present at each ARTCC.

1.3 Bottlenecks in the 9020A and 9020D Computer Systems

This section describes the bottlenecks that are likely to degrade performance of the 9020A and 9020D over the next ten years. This report

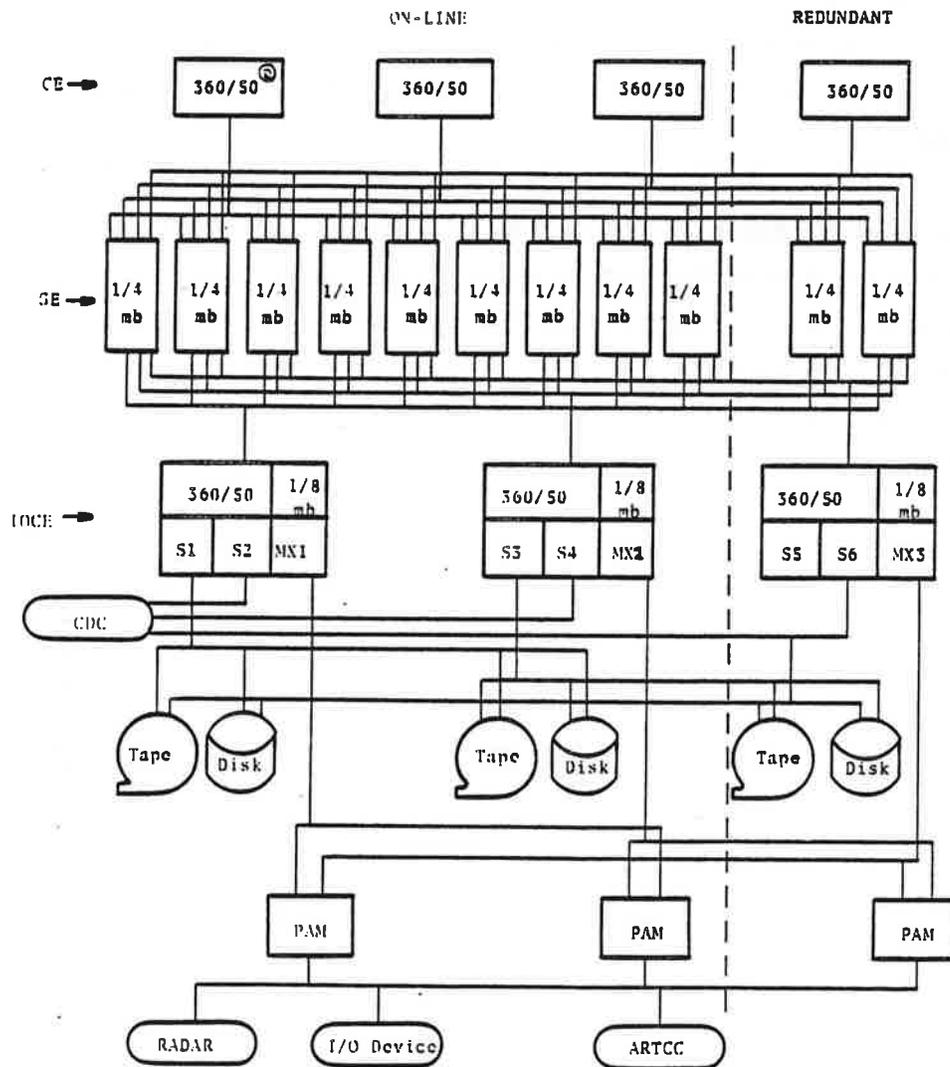


FIGURE 1-1: SIMPLIFIED 9020A CONFIGURATION DIAGRAM

Si - Selector Channel
 MXi - Multiplexor Channel
 PAM - Peripheral Adapter Module
 CDC - Display Channel

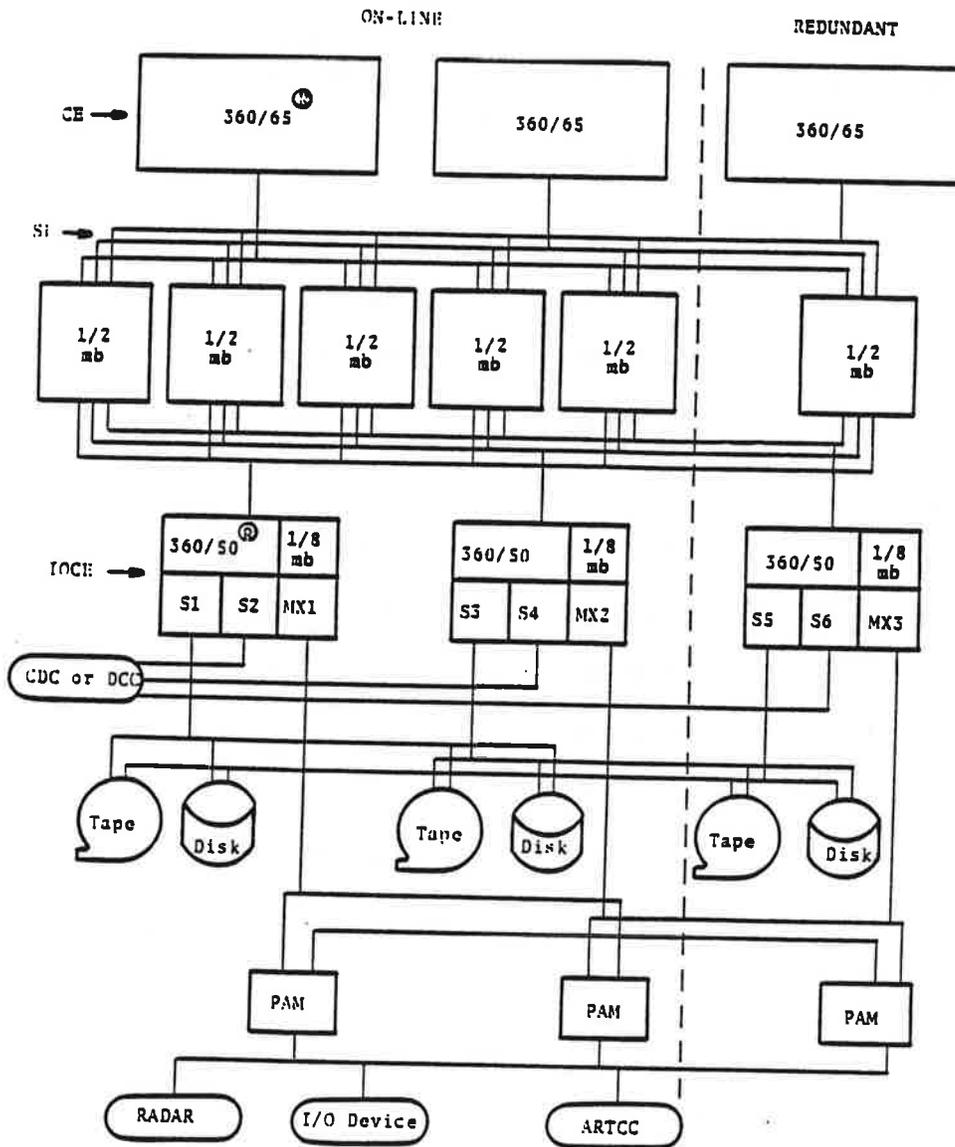


FIGURE 1-2: SIMPLIFIED 9020D CONFIGURATION DIAGRAM

Si - Selector Channel
 MXi - Multiplexor Channel
 PAM - Peripheral Adapter Module
 CDC/DCC - Display Channel

TABLE 1-1: COMPUTER SYSTEM CONFIGURATIONS FOR THE ARTCC'S

<u>Center</u>	<u>CCC</u>	<u>Display</u>
Albuquerque	IBM 9020A	Ray 730
Atlanta	IBM 9020D	Ray 730
Boston	IBM 9020A	Ray 730
Chicago	IBM 9020D	IBM 9020E
Cleveland	IBM 9020D	IBM 9020E
Denver	IBM 9020A	Ray 730
Fort Worth	IBM 9020D	IBM 9020E
Houston	IBM 9020A	Ray 730
Indianapolis	IBM 9020D	Ray 730
Jacksonville	IBM 9020D	Ray 730
Kansas City	IBM 9020D	Ray 730
Los Angeles	IBM 9020D	Ray 730
Memphis	IBM 9020A	Ray 730
Miami	IBM 9020A	Ray 730
Minneapolis	IBM 9020A	Ray 730
New York City	IBM 9020D	IBM 9020E
Oakland	IBM 9020A	Ray 730
Salt Lake City	IBM 9020A	Ray 730
Seattle	IBM 9020A	Ray 730
Washington DC	IBM 9020D	IBM 9020E

will investigate the extent to which the hardware enhancements can eliminate these bottlenecks. In this way one will be able to judge whether the enhancements discussed in this report will provide the needed improvement in system performance.

A study carried out at the Transportation Systems Center [CLAP79, Sec.'s C-4 and C-5] gives a statement of what the bottlenecks are expected to be over the next ten years. This study examined the projected level of activity at the ARTCC's and compared it to the processing capability of the 9020's. The findings are shown in Table 1-2. First, both the 9020A and 9020D are expected to have problems with both I/O bandwidth and I/O device speed. Second, both the 9020A and 9020D are expected to have problems with memory capacity; in addition, the memory bandwidth of the 9020A is another problem area. Third, the 9020A is expected to have inadequate processing capacity; the 9020D is expected to encounter no problems in this area. Processing capacity in this report will be taken to mean the size of the peak traffic load that the system can handle.

In summary, the 9020A and 9020D both have problems with I/O and memory, and the 9020A also has problems with processing capacity. These are problems that are expected to surface over the next few years if nothing is done to avoid them. Solving these problems can be taken to be the minimum that is necessary to preserve satisfactory operation of the 9020's. Therefore, the enhancements discussed in this report will be closely scrutinized to determine how well they deal with these problems.

TABLE 1-2: THE CRITICAL 9020 RESOURCES

Is this resource a bottleneck?

<u>Resource</u>	<u>9020A</u>	<u>9020D</u>
I/O Bandwidth	Yes	Yes
I/O Device Speed	Yes	Yes
Memory Capacity	Yes	Yes
Memory Bandwidth	Yes	No
Processing Capacity	Yes	No

Source:[CLAP79, p. C-20]

2. MEMORY ENHANCEMENTS

2.1 Purpose and Organization of this Chapter

The purpose of this chapter is to discuss three enhancements that could be made to the 9020 memories; each enhancement is discussed with respect to its description, advantages, cost, schedule, and transition. Sec. 2.2 discusses the enhancement of replacing the entire memory boxes, i.e., the SE's, with new boxes. A memory box consists primarily of the cabinet, power supply, cooling apparatus, interface to the rest of the machine, and stack (which is what actually holds the data). Sec. 2.3 discusses the enhancement of replacing just the memory stack in the SE, with the rest of the memory box being left intact. Sec. 2.4 discusses the enhancement of replacing the memory stack in the IOCE.

2.2 Replacement of the Memory Boxes

2.2.1 Organization of this Section

2.2.2 describes the enhancement of replacing all of the memory boxes on the 9020A's and some of them on the 9020D's. 2.2.3 explains how this enhancement deals with the problems the 9020's face and how it also provides other advantages. 2.2.4 estimates the cost of this enhancement, and 2.2.5 estimates the schedule according to which it could be implemented. 2.2.6 sketches out what the transition period would be like. Finally, 2.2.7 discusses the variant on this enhancement of replacing all of the memory boxes on the 9020D's instead of just some of them.

2.2.2 Description of this Enhancement

This subsection describes the design decisions the FAA would have to make, the assumed configuration of the enhanced system, the nature of the memory that would be procured, and the changes that this enhancement would imply.

Design decisions. If this enhancement were adopted, the decisions that the FAA would have to make are: How much new memory should each system have? How should the new memory be distributed and interleaved among different boxes? In making these decisions the FAA would be constrained by four factors. First, the 9020A and 9020D can accommodate a maximum of 16 megabytes of main memory (though only the first 10 megabytes can be accessed by the IOCE's). Second, the 9020A is designed for a maximum of 12 memory boxes and the 9020D for a maximum of 10 boxes; these figures include the redundant memory boxes. Third, the 9020A memory is too slow to co-exist with state of the art memory, so the 9020A memory would need to be completely replaced. In contrast, it would be possible to add state of the art memory to the 9020D and to keep the old memory. That is, since the 9020D currently has 7 memory boxes and since it can accommodate as many as 10, it would be possible to add as many as three new boxes without removing any of the old memory. Fourth, so that the advantages of this enhancement can be fully realized, it is necessary for there to be enough main memory to hold all programs and data (except for infrequently used items like pre-stored flight plan data).

Assumed configuration. For concreteness, this report assumes that the new memory boxes would each contain one megabyte; the boxes used on the 9020A and 9020D would be virtually identical. Each would have an eight port switch and be either eight or four bytes wide for the 9020D or 9020A, respectively. It is assumed that all of the 9020A memory boxes are discarded and replaced by six units. It is assumed that the six 9020D memory boxes are retained; three of the new units are added. This means that each 9020 would have six megabytes of shared, main memory. These specific assumptions are made here to illustrate what the enhanced systems might look like and so that the cost estimates can be carried out for a specific system. It should be stressed, however, that additional measurements and simulations are needed in order to determine the optimal configuration of the memory units with respect to the total amount of memory, the number of memory units, and interleaving.

Nature of the new memory. The memory that would be procured would be constructed of solid-state metal oxide semiconductor (MOS) integrated

circuits. Each circuit (or chip) would have either 16,384 or 65,536 bits of memory; in today's market there is no difference in the cost per bit of these two sizes. The memory will contain error checking and correction for single bit errors and detection of double bit errors; this is in addition to the parity bit per byte that the SE stores for the CE. The memory will use the existing uninterruptible power supply.

The speed of the memory would be 750 nanoseconds for an eight byte fetch. (Higher speeds could be obtained by installing a cache memory in each CE.) This speed is chosen because it appears to be the proper trade-off between speed and cost. For the 9020A, a slower memory would make it difficult to achieve the desired increase in processing capacity, and a faster memory would not yield any significant benefit. For the 9020D, the new memory would be about 10 percent slower than the old, but this would not reduce the processing capacity noticeably. (The processing capacity of the 9020D, however, would increase since buffering would be eliminated.)

Implied changes. Essentially, this enhancement would require no major change in the present software. In particular, no change would be required in the application software. There are, however, three minor areas in which some change in the software would be necessary. First, a new system generation would be required to eliminate buffering and to allow for the new memory configuration. This is a function that has been performed many times in the past and is accomplished by changing the appropriate parameters for system generation.

Second, if memory boxes of two different sizes are used, then the dynamic on-line error detection and reconfiguration system would have to be modified so that it recognizes that all memory boxes are not of the same size and, hence, not perfectly substitutable. (This problem would only arise if some of the old boxes on the 9020D are kept.) This modification was done previously by IBM when converting from the 04 to the 08 SE's, so it is already known that the system can accommodate SE's of different sizes without great difficulty. (The 04 SE is an early 9020A SE; the 08 SE is the current 9020A SE.)

Third, all current maintenance programs should run on the new SE's but, because they use solid state technology instead of magnetic cores, the most critical tests, the "worst case pattern" tests, will not be testing the new memories as vigorously as they should. The vendor can either supply worst case diagnostics to run on the system, or he can provide a self-test mode to exercise each SE internally to test for "worst case pattern" failures. Each box would have built in diagnostic functions.

The conclusion drawn from considering the changes in software that this enhancement would require is that the changes are relatively minor and can be carried out at a very small cost and with virtually no risk.

Aside from software, the only other change that this enhancement would require would be to physically connect the new boxes to the system. This cabling would not be major and is described in Sec. 2.2.6.

In summary, the FAA's choice for each 9020 system is to decide how much state of the art memory to add and how to distribute it among different boxes. This choice must satisfy the design constraints of the system, and it should be made so that all programs and data can be resident in main memory throughout the life of the system.

2.2.3 Advantages of this Enhancement

Replacing the current memory with state of the art memory would result in two main effects.

- The 9020A would have a faster memory.
- The 9020A and 9020D would have a larger physical address space that would allow all programs and data to be resident in main memory.

These two features will yield seven advantages. This discussion assumes that only this enhancement is adopted; the additional advantages that would be achieved if faster CE's were used are discussed in the next chapter. When possible the discussion is quantitative; these numerical estimates are

derived from a simulation model of the 9020 systems that is outlined in App. A and is described in detail in App. B.

First, since almost all programs and data will be resident in main memory, buffering can be almost eliminated. This will reduce the I/O load by 30 to 50 percent, and this means that the I/O capacity and bandwidth problems will be dealt with.

Second, the size and speed of the new memory will eliminate the memory capacity and bandwidth problems.

Third, there is an increase in processing capacity. It is estimated that the faster memory in the 9020A will increase capacity by 10 to 40 percent by reducing memory interference. (Increasing capacity by 10 percent means that 10 percent more tracks can be handled at peak load.) Memory interference occurs when two CE's want to access the same memory box at the same time; this means that one of them must wait. With the faster, state of the art memory, the probability of two CE's wanting access to the same box at the same time is smaller. Moreover, when this does occur, because of the faster memory there will be a shorter wait. There is no similar capacity increase for the 9020D since its memory is not slower than (and is, in fact, slightly faster than) the new memory. There will also be an additional increase in processing capacity because, with all programs and data being resident in main memory, buffering will be eliminated. This is estimated to decrease overhead by 10 to 20 percent for the 9020A and by 5 to 10 percent for the 9020D. Therefore, considering the effect of the faster 9020A memory and the elimination of buffering, the increase in processing capacity is expected to be from 20 to 60 percent for the 9020A and from 10 to 30 percent for the 9020D.

Fourth, the 9020A will have a faster response time because of its faster memory, and the 9020A and 9020D will both show a faster response time because buffering is eliminated. The amount by which response time would improve has not been estimated, but it could be estimated using the NAS Systems Model by FEDSIM. The FAA currently uses this model to estimate the performance of the 9020 system.

Fifth, the larger memory would reduce software maintenance cost. Currently at least \$18 million is spent each year on software maintenance [ASI80, p. 6-4], and a considerable portion of this expense is due to the difficulties caused by the shortage of main memory. Because this enhancement would relieve this shortage, a substantial saving in software maintenance cost is expected; in fact, in this way this enhancement could easily pay for itself in a few years.

Sixth, the reliability of the system would be improved. This results from the greater reliability of the state of the art memory. Also, because a significant number of software failures occur during buffering, the elimination of buffering will increase software reliability.

Seventh, because there is a larger memory, more functional enhancements and local adaptation data could be added to the system. This would allow the capabilities of the system to be extended and also allow a greater level of automation to be achieved.

What is the technical risk associated with this enhancement? That is, what is the probability that the new memory will function properly and that these advantages will indeed be obtained? Technically, replacing (or supplementing) the current memory with state of the art memory is straightforward. The procedure is conceptually simple and has been done before in comparable circumstances. Therefore, the conclusion is that there is virtually no risk involved; that is, it is almost certain that the enhanced system would work exactly as described in this report.

In summary, Sec. 1.3 pointed out that if an enhancement is to be of interest, it must be able to deal with I/O, memory, and processing bottlenecks. It is seen that this enhancement does deal with the I/O and memory bottlenecks. It increases processing capacity somewhat, and the FAA would have to judge whether this increase is large enough; if it is not, then one possible course would be to supplement this enhancement with one of the processor enhancements discussed in the next chapter.

2.2.4 Cost

The cost of this enhancement has four components. First, in order to optimize the design of this system it will be necessary to conduct a number of simulations using the model described in App.'s A and B. The estimate of the cost of these simulations is in the range of \$30,000 to \$60,000.

Second, there will be a one-time cost for the engineering that is needed to customize the memory boxes for the 9020 environment. Estimates obtained by phone from Ampex and Intel place this cost in the range of \$200,000 to \$400,000.

Third, there is the cost of the memory boxes. Ampex and Intel estimate that the cost would be \$70,000 for each one megabyte memory unit. Past experience, however, indicates that \$50,000 per one megabyte unit is a realistic cost at final bidding; this lower figure is used here. Ten of the ARTCC's have 9020A's, and ten have 9020D's. There are a 9020A and a 9020D at the FAA Technical Center, and there is a 9020A at the FAA Aeronautical Center. Therefore, there are twelve 9020A's and eleven 9020D's. Since six memory units are needed for each 9020A and three for each 9020D, this means that a total of 105 units would be procured. Throughout this report the amount allotted for spares at each site equals the cost of one unit. At a cost of \$50,000 per unit, then, the cost including spares for the 23 sites is \$6.4 million.

Fourth, even though every effort has been made to make accurate estimates, there might well be unexpected costs. Throughout this report an extra 20 percent will be added to cover contingencies. Therefore, \$1.372 million is allowed for contingencies.

The cost of memory box replacement is shown in Table 2-1. To avoid underestimating the cost, when there is a range the upper limit of the range is used. The measurement and simulation is estimated to cost \$0.06 million, the engineering to cost \$0.4 million, the procurement of the memory units to cost \$6.4 million, and \$1.372 million is allocated for contingencies. The total estimated cost rounded to the nearest hundred thousand is \$8.2

TABLE 2-1: ESTIMATED COST OF REPLACING THE MEMORY BOXES

<u>Component</u>	Cost <u>(millions)</u>
Measurement and simulations	\$0.060
One-time engineering cost	0.400
Memory units	6.400
Contingencies	<u>1.372</u>
Total	\$8.232

million. It should be pointed out that there are some costs that this figure does not include, such as the cost of training technicians to deal with these new units, spare parts, and the cost incurred by the FAA in administering and overseeing the procurement. All of these costs are expected to be minor.

What is the financial risk of this enhancement? That is, what is the chance that this enhancement will cost significantly more than what is estimated here? The main factor in assessing financial risk is that the memory boxes to be procured are standard 360/370 add-on memory and are readily available from a number of sources. Two firms, Ampex and Intel, have bid over the phone, and other firms such as VION/National and Mostek have indicated a high level of interest. From this survey it can be concluded that at least six firms would respond to a request for quotations. Therefore, with this much competition among the bidding firms, the FAA would not have to worry about having to pay an artificially inflated price. The conclusion is that this enhancement entails very little financial risk.

2.2.5 Schedule

The speed with which an enhancement can be implemented is one of the criteria used to evaluate the desirability of that enhancement. So that the enhancements discussed in this report can be seen on a more or less common basis, the zero point on the schedule will be taken to be when the FAA places the order. Therefore, what is of interest is how long various events occur after receipt of order (ARO). It is estimated that the first check-out unit for this enhancement would be delivered twelve months after receipt of order. Initially production would be at the rate of one per month, with the rate rising to one per week by 18 months ARO. Thus, it is estimated that the 105 units would all be delivered by about 38 months ARO. Installation at the six most critical sites could be completed by 24 months ARO.

2.2.6 Transition

The FAA has established the requirement that in any enhancement or replacement of the en route computers, there must be a smooth transition that does not significantly interrupt the provision of air traffic control services. The three main issues are whether there is excessive downtime during installation, whether there is sufficient floorspace, and whether the training requirements can be met. Each issue will be briefly discussed.

Downtime. The cabling on each SE consists of 42 cables (six sets of seven cables), with 14 being short internal cables to an adjacent SE. There are four sets for Data In (lower half word in and out, upper half word in and out) and one set each for Control and Data Out. Only the Data In cable is daisy-chained. Thus, each processor has two cables going to each SE for a total of 26 cables for the processor's memory bus on the 9020A system.

It is estimated that changing a memory box will require 8 man-hours and will result in 2 hours of system downtime. The 08 SE's cabinet can be partially disassembled to allow removal of the SE without moving the cables. This estimate reflects the experience gained on the recent SE additions to the 9020 systems.

Floorspace. A 9020A system when outfitted with the new memory units will take up less space than the system now does, so there would be no floorspace problem. A 9020D system will take up slightly more room since four units will be added, so the ARTCC's will need to be examined for available floorspace; since each unit is quite small, however, it is expected that there will be no floorspace problem.

Training. Since the new memory units would be both conceptually similar to and also simpler than the old memory units, it is expected that the training required would be minimal and would pose no obstacle to a smooth transition.

In summary, because the cabling, floorspace, and training that would be required would be minor, the conclusion is that the transition to the enhanced system can be made without any significant problems.

2.2.7 A Variant: Replace All of the 9020D Memory

This chapter has thus far assumed that three new one megabyte memory boxes would be placed on the 9020D and that the six old 1/2 megabyte units would be retained. A variant on this approach would be to eliminate the old memory and to replace all of it with new memory. For concreteness, assume that six megabytes of new memory are placed on each 9020D. This variant differs from the enhancement discussed in the rest of this chapter in four ways.

First, since all of the 9020A's and 9020D's would have identical memory units, maintenance and logistics would be simplified. Second, the new memory units would be more reliable than the old. Third, since all the 9020D memory is replaced, it would be prudent to procure somewhat faster memory, e.g., memory with a cycle time in the range of 500-600 ns rather than 750 ns. This would raise the cost per box to \$60,000. Fourth, an additional 33 memory boxes would be procured. The cost, which is figured in the same way as in 2.2.4 (except for the greater number of boxes and the higher cost of each box), rises from \$8.2 million to \$12.1 million.

One of the FAA's options not discussed in this report is to upgrade all of the 9020A's to 9020D's. If this is done, it might well be desirable to further upgrade all the systems with the memory replacement discussed in this chapter. The cost of putting six megabytes of state of the art memory on all the systems would be this same figure of \$12.1 million.

2.3 Replacement of the Memory Stacks in the SE's

Sec. 2.2 discussed the possibility of enhancing a SE by replacing the entire memory box; it is possible, however, to enhance an SE by replacing just the memory stack, i.e., the component in the box that actually stores the data. Moreover, it is also possible to enhance the memory in the IOCE's by replacing the memory stacks. These two enhancements, which offer a relatively fast and cheap way to enhance memory, will be discussed in this section and the next, respectively.

Description. The description in 2.2.2 of the enhancement of replacing the memory boxes also applies to this enhancement, except that the FAA would only procure memory boards instead of entire memory boxes. That is, instead of ordering entire boxes from a manufacturer, the FAA would have the new memory designed and have the contractor buy the needed memory chips on the open market and assemble the memory boards. More specifically, the cabinetry, memory interfaces, cable connections, and power supplies would not be replaced; the memory stacks, which will be replaced, consist of everything else, e.g., the line drivers and data planes. Because the 9020A's and 9020D's differ in the word length of memory (36 bit v. 72 bit), in CE speed, and in the interface, the new memory boards for the 9020A would be different from the boards for the 9020D. Since this enhancement does not procure entire boxes, the new memory would not come with built-in diagnostics; new memory diagnostics would have to be written.

Cost. The cost of this enhancement has four components. First, the cost of designing the new memory stacks and building a working, tested, analyzed, and documented engineering prototype for both the 9020A's and the 9020D's is estimated to be \$155,000. (The cost of the design work and the prototype for the 9020A only would be \$95,000 and for the 9020D only would be \$115,000; because of commonality, however, the cost for both is \$155,000.) Second, the estimated cost of writing the new diagnostics is \$100,000, which is \$50,000 for each prototype. Third, the cost of replacing each memory stack with a one megabyte unit is estimated to be \$25,000 for a 9020A SE and \$30,000 for a 9020D SE. Six SE's would be enhanced at each site. At a 9020A site, allowing \$25,000 for spares, the cost of implementing this enhancement is estimated to be \$175,000. At a 9020D site, allowing \$30,000 for spares, the cost is estimated to be \$210,000. Fourth, \$0.933 million is allowed for contingencies. Therefore, the total cost of the design and implementation of this enhancement at the 23 sites is estimated to be \$5.6 million.

Schedule. Once the working prototype is finished (a task which is estimated to take five months), the FAA would be ready to place the order for the parts. The first system could be implemented in 3 months ARO, if parts are in stock. In the worst case, waiting for parts would cause an

additional two month delay, so the first system would be implemented 5 months ARO. (The only long lead time parts are the memory chips, which would cost about \$10,000 for each SE.) It will take about 2 weeks to implement this enhancement at each site. If it takes 5 months to implement the first system, this means that this enhancement could be implemented at the six most critical centers within 8 months ARO.

Transition. It is expected that the stack replacement would be accomplished by installing a small number of boards and by modifying a small number of backplane wires. It is estimated that each stack replacement would take not more than one man-hour. No cable changes would be necessary. The system downtime would only be that necessary for reconfiguring the system, i.e., about 30 seconds for each SE. No additional floorspace would be needed. The amount of training needed by hardware maintenance personnel is expected to be minimal.

Advantages. The seven advantages of replacing the memory boxes described in 2.2.3 would also be obtained from replacing the memory stacks since these advantages stem from the quantity and speed of the memory. Moreover, replacing the memory stacks would, compared to replacing the memory boxes, have four additional advantages. First, the stacks can be procured much faster than the boxes; this is because the cabinet, power supply, and interface need not be designed and manufactured if only the stacks are replaced. The discussion of the schedule implies that the FAA could replace the memory boxes at the first six systems within 8 months after deciding to adopt this enhancement, whereas it would take 24 months if instead the memory boxes were replaced.

Second, the physical installation would be much easier if the stacks are replaced rather than the boxes. The stacks are replaced by substituting a few boards into the cabinet, whereas the boxes are replaced by making a number of cable changes as described in 2.2.6. It would take about 1 man-hour to replace a stack as contrasted with 8 man-hours to replace a box.

Third, it would be cheaper to replace just the memory stacks instead of the entire boxes. For example, the cost of replacing the stacks at the 23

sites is estimated to be \$5.6 million and the cost of replacing the boxes is estimated to be \$8.2 million.

Fourth, there are very short lead times for the parts needed for this enhancement and no significant advantage to buying in quantity. This means that the FAA can try the enhancement at one or more sites and then decide whether to implement it at more sites. The FAA need not commit a large amount of money at the beginning; as the enhancement is put into operation the FAA can gradually decide how many centers should have it without unduly delaying its implementation.

There are four advantages of replacing the boxes rather than the stacks. First, the entire memory box rather than just the stack would contain state of the art components and designs. Second, if the entire boxes were procured, built-in diagnostics would be included. Third, the entire SE would be the responsibility of one vendor. Fourth, if it were later decided to upgrade the 9020A's to 9020D's, then the new memory boxes could be used in the upgrade.

2.4 Replacement of the Memory Stacks in the IOCE's

Description. Each IOCE currently has 1/8 megabyte of memory, called MACH memory, that can be accessed only by that IOCE. One possible enhancement is that the memory stack in each IOCE could be replaced with up to 6 megabytes of state of the art memory; for concreteness it is here assumed that the new stacks contain 2 megabytes. The replacement memory would be generally the same as that described in Sec. 2.3.

Advantages. If this enhancement were followed by moving program elements into the enlarged MACH memory, some of the processing load could then be shifted to the IOCE. The potential increase in 9020A processing capacity is estimated to be between 10 and 30 percent. Since, however, replacing the IOCE memory stacks makes the most sense when the IOCE processor is sped up, the discussion of the advantages of this enhancement is postponed to Sec. 3.3 where the advantages of jointly implementing these two enhancements are discussed.

Cost. The cost of this enhancement for the 9020A's has four components. First, the cost of designing the new memory stack and building the prototype is estimated to be \$105,000. (This cost figure assumes that the 9020A SE memory stack replacement prototype is not built; if it is built, then the additional cost of the IOCE memory stack replacement prototype would be \$20,000.) Second, the estimated cost of writing the new diagnostics is \$50,000. Third, the cost of the 2 megabytes of new memory for each IOCE is estimated to be \$30,000. Allowing \$30,000 for spares, the cost of this enhancement at each center is estimated to be \$120,000. Fourth, allow 0.319 million for contingencies. Therefore, the total cost of this enhancement at the 12 9020A sites is estimated to be \$1.9 million. If the IOCE memory stacks are also replaced at the eleven 9020D sites, the additional cost is \$1.320 million for parts and installation and \$0.264 million for contingencies. Therefore, the cost of replacing the IOCE memory stacks at the eleven 9020D sites is \$1.6 million, and the cost at all 23 sites is \$3.5 million.

Schedule. The schedule for this enhancement is the same as that for replacing the stacks in the CE's; the first six systems would be upgraded within 8 months ARO.

3. PROCESSOR ENHANCEMENTS

3.1 Purpose and Organization of this Chapter

Chapter 2 has described several memory enhancements that can provide some relief in the areas of I/O, memory, and processing capacity where the 9020's face potential problems. If the FAA decides that these memory enhancements alone are not sufficient to deal satisfactorily with the 9020's problems, then the FAA might decide to supplement the memory enhancements with one or more processor enhancements. The purpose of this chapter is to describe three possible processor enhancements that can be considered for adoption.

This chapter is organized as follows. Sec. 3.2 discusses the enhancement of speeding up the processors in the 9020A CE's by replacing selected components. Sec. 3.3 discusses the enhancement of speeding up the processors in the IOCE's. Either of these enhancements would provide a significant increase in computing capacity if it proved to be feasible. Unfortunately, study of this problem has not yet progressed to the stage where it can definitely be said whether the speed-up is feasible. Therefore, Sec. 3.4 discusses the fall-back option of replacing the 9020A CE's. This enhancement would provide the needed increase in computing capacity, and it would be suitable for adoption if the speed-up proves to be infeasible or too risky or for some reason undesirable.

3.2 Speed-Up of the 9020A CE Processors

3.2.1 Description of this Enhancement

The CE speed-up enhancement is accomplished by replacing two of the subsystems of the 9020A CE that are bottlenecks limiting CE speed. One subsystem to be replaced is the local store, which contains the CE's registers. The other subsystem to be replaced is the read only store (ROS), which contains the microinstructions for the processor. Each can be replaced by an integrated circuit system that would be smaller, take less

power, be more reliable, and run from 5 to 8 times faster. The CE would need to be retuned to take advantage of these faster components. This enhancement would not require any changes in software or in any other part of the system. (One minor exception to this statement is the diagnostics, which are mentioned below.) A prerequisite for this enhancement is a faster memory; therefore, this enhancement assumes either that the memory boxes or the SE memory stacks have been replaced. The rest of this subsection describes in more detail the subsystems to be replaced and the installation procedure to be followed.

Local store. The local store is a 0.5 microsecond, 64 word by 32 bit, linear select, core memory system which contains the general purpose registers, the floating point registers, and several internal registers. It is wholly contained on a single card and lends itself very well to implementation with the random-access memory (RAM) now available.

There are several 4 x 256 bipolar RAM chips available with access times in the 50 nanosecond range. (1000 nanoseconds equals 1 microsecond.) Nine of these chips would constitute the memory array, and an additional 20 chips would provide the interface to IBM's solid logic technology (SLT) and would perform various control functions.

Read only store. The ROS contains 2,816 90-bit words in a 0.5 microsecond, read only capacitative memory. It is physically very large, comprising about 15 percent of the total processor. It also is well contained and could be readily replaced by a state of the art subsystem that would be one-tenth the size and 8 times as fast as the old subsystem.

The new memory array would be constructed of 66 8x512 programmable read only memories (PROM's) if the current size of 2,816 words were retained. It would be possible, however, to increase the size to 4,096 words by using 88 PROM's. In either case these PROM's would be mounted on three separate boards with supporting circuitry.

Retuning the CE. Once the new, faster components are installed in the CE, it will need to be retuned to take advantage of them. The following

discussion gives a general idea of what this retuning will consist of. The microcycle is the basic unit of time that the processor uses; any particular task that the processor carries out is allotted some number of microcycles. For the 9020A the microcycle time is 500 nanoseconds. In order to reference the 9020A memory, 5 microcycles are currently needed; this is called the storage timing ring. Therefore, the processor can be sped up by decreasing the number of microcycles in the storage timing ring and by reducing the microcycle time. The idea behind this enhancement is that the faster memory on the 9020A and the new components in the CE will allow the number of microcycles in the storage timing ring and the length of each microcycle to be reduced; this is referred to as retuning the CE.

Installation procedure. The modifications to reduce the storage timing ring would require some modified modules and back plane wiring changes. Although these changes would be minor, it might be advantageous to replace the affected modules with modules made from standard integrated circuits to minimize the conversion time and reduce the chance of error in changing the module for maintenance reasons.

The local store and ROS upgrades would replace whole motherboards with their load of modules with a printed circuit board with integrated circuits mounted directly on the board. The technology would be Schottky TTL (LS, S, ALS, AS, and/or F series) with Schmidt trigger inputs and discrete output drivers to interface with IBM's SLT modules. The local store upgrade would be a replacement of one motherboard with one printed circuit board. The ROS upgrade would replace five motherboards with three printed circuit boards.

The CE speed-up modifications would not change the characteristics of the IBM diagnostics, but whenever they indicate a defective module in the ROS or local store, a separate chart would indicate which card to replace. These charts could be decals affixed to the panels that a maintenance engineer would normally approach to replace the indicated defective module. In the case of modified modules, care must be taken that the modified module is replaced by a similarly modified unit. Again the judicious use of labels as well as the general awareness of the maintenance engineer should suffice to make the correct replacements.

3.2.2 Advantages of this Enhancement

If this enhancement were adopted, it would result in six advantages. First, it is judged that with a probability of 0.98 the storage timing ring could be decreased from 5 to 4 microcycles. (This probabilistic judgment and the ones below are based on experience with the System/360 architecture and with making similar changes to other processors.) Since the current microcycle time is 0.5 microseconds, this would reduce the storage cycle time from 2.5 to 2.0 microseconds. This reduction would be made possible by the faster memory. According to the simulation described in App. A, a reduction of 0.5 microseconds in the storage timing loop would result in a 21 percent increase in performance. Because of memory interference and other considerations, however, not all memory references would benefit from this faster cycle time and the actual increase in performance would be somewhat less than 21 percent. A sampling of the microcode indicates that approximately 75 percent of the memory references would benefit from this shorter storage timing loop; thus, there is a 15 percent increase in processing capacity. This figure, however, only reflects the increase due to faster memory and reduced memory interference; it does not include the increase due to having more memory. This latter increase is estimated to be at least 10 percent and perhaps as much as 30 percent. Therefore, the increase in processing capacity by reducing the number of microcycles in the storage timing ring is estimated to be 25 percent. (The standard IBM 360/50 CPU uses four 500 nanosecond microcycles. The 9020A CE is essentially model 360/50 memory; the main difference is that the 9020A CE has an eight port switch. The delay in this switch is about 100 nanoseconds. Since the microcycle time cannot be varied in the 360/50, the presence of this switch required that a full microcycle be added to the storage timing ring for the 9020A.)

Second, this enhancement will allow the microcycle time to be decreased. The reasoning behind this judgment is as follows. The three main CE subsystems that currently are major bottlenecks on performance are the local store, the ROS, and the 32-bit adder. This enhancement replaces the old, 500 nanosecond local store with a new, 50 nanosecond component. It also replaces the old, 500 nanosecond ROS with a new, roughly 62.5

nanosecond component. A series of measurements made of an IOCE executing a full, 32-bit add and carry indicates that the worst case timing is 120 nanoseconds though the specification is 360 nanoseconds (i.e., 360 nanoseconds are currently allowed in the timing sequence but only 120 are needed). Thus, it appears that it will not be necessary to replace the adder even with a microcycle time of 300 nanoseconds. (If it turns out that the adder is slower than these measurements indicate, then replacing the adder might be considered. The adder's functions are scattered on various boards, and it would be the most difficult of the three subsystems to replace. The difficulty of replacing the adder has not been fully evaluated since replacement does not appear necessary.)

How much would this enhancement allow the microcycle time to be reduced? This question cannot at the present be answered because the reduction that could be achieved depends on timing interactions and on other complicated and not fully understood factors. The best estimates of the probabilities with which various microcycle times could be achieved are that the current time of 500 nanoseconds could be reduced to 400 with probability 0.9, to 300 with probability 0.5, to 250 with probability 0.2. It is judged that a 200 nanosecond cycle time could not be achieved.

These first two sources of an increased processing capacity are summarized in Table 3-1. Consider the second row of this table. Suppose that the storage timing ring is decreased from 5 to 4 microcycles and that the microcycle time is decreased from 500 to 400 nanoseconds. Then the storage cycle time is reduced from 2500 to 1600 nanoseconds. This yields an increase in processing capacity of at least 50 percent. The probability that this 50 percent increase will be achieved is 0.88, which is 0.98 (the probability that the storage timing ring can be decreased from 5 to 4 microcycles) times 0.9 (the probability that the microcycle time can be decreased to at least 400 nanoseconds). The third row in this table shows that there is a 0.49 probability that processing capacity can be increased by at least 100 percent.

Third, if the ROS is expanded beyond the current 2,816 word size, this would allow a further increase in computing capacity. That is, a sequence

TABLE 3-1: INCREASED PROCESSING CAPACITY DUE TO THE CE SPEED-UP

<u>Storage Cycle Time (ns)</u>	<u>Capacity Increase (%)*</u>	<u>Probability of Achieving</u>
5x400 = 2000	25	0.98
4x400 = 1600	50	0.98 x 0.9 = 0.88
4x300 = 1200	100	0.98 x 0.5 = 0.49
4x250 = 1000	-	0.98 x 0.2 = 0.20

* These estimates are conservative estimates of the total increase in processing capacity due to all factors.

of instructions that is commonly used could, in effect, be made into a single instruction and coded into the ROS; the sequence would then execute much faster. In order to achieve this advantage, it would be necessary to identify the frequently used sequences and then to code them. Therefore, this additional increase in computing capacity would not happen automatically when the ROS is enlarged; it would require additional work before it were realized.

Fourth, the CE's would be made substantially more reliable since the local store and the ROS are being replaced by modern technology components, which are perhaps an order of magnitude more reliable than the old components. This is especially significant for the ROS, which uses a great deal of power, comprises a large portion of the CPU, and is the most unreliable portion of the CPU.

Fifth, since the new ROS would use much less power and would dissipate less heat, the cooling of the CE's would be improved.

Sixth, the ease of installation would contribute to a smooth transition. That is, other options that the FAA is considering would require laying new cables and making many new connections, and this can be a difficult job because of the confusing mass of cables in the ARTCC's. This enhancement avoids these possible problems since no cable changes or disconnects are needed.

3.2.3 Cost and Schedule

There are three components to the cost of this enhancement. First, measurements and simulations need to be done to determine how the speed-up is to be accomplished and to complete the engineering prototype. This stage has begun; to finish it will cost an additional \$125,000 (plus support from the Technical Center) and will take five months. (This cost would be cut to \$20,000 if the IOCE processor speed-up were carried out before the CE processor speed-up.) Second, the modification that speeds up the CE's must be implemented. Each speed-up kit is estimated to cost \$25,000. At each 9020A site, then, the cost is estimated to be the cost of speeding up four

CE's (\$100,000) plus another \$25,000 for modifying spares, for a total of \$125,000 per site. Since there are twelve 9020A sites, the cost is estimated to be \$1.5 million. Third, allow 0.325 million for contingencies. Therefore, the total cost of this enhancement is estimated to be \$2.0 million.

Since delivery of the speed-up kits could start three months ARO and since one site could be sped up every two weeks, the first six sites could be sped up within 6 months ARO.

This discussion assumes that it does prove possible to speed up the CE's. If it turns out that this effort is not successful, then it is estimated that \$50,000 would be lost. The remaining \$75,000 would be applicable to the CE replacement and to the memory stack replacements.

3.2.4 Transition

It is expected that the CE speed-up would be accomplished by replacing four boards and by replacing or modifying a small number of modules and backplane wires. It is estimated that a conversion of the four CE's would take four hours. No cable changes would be necessary. The system downtime would only be that necessary for reconfiguring the system, i.e., about 30 seconds for each CE. No additional floorspace would be needed. The amount of training needed by hardware maintenance personnel is expected to be minimal.

3.3 Speed-Up of the IOCE Processors

The next processor enhancement to be discussed is to speed up the processors in the IOCE's; a prerequisite for this enhancement is the IOCE memory stack replacement discussed in Sec. 2.4. Since the processors in the IOCE's are virtually identical to the processors in the CE's, this enhancement is in many ways quite similar to the CE speed-up enhancement just discussed; the differences between these two enhancements will now be discussed.

Description. The main difference between speeding up the IOCE processor and the CE processor is that if the SE memory is not replaced with faster memory, then the IOCE must reference memory with two different speeds. That is, the IOCE processor would reference the new faster IOCE memory and also the old, slower SE memory. This can be dealt with by providing a different timing sequence for the references made to the SE memory.

For this enhancement to provide its main advantages, some software changes would need to be made. Selected program elements (PE's) would be removed from the shared memory and made resident in the IOCE's memory; tables would be left in shared memory. If the IOCE is executing a PE in MACH storage, only operand fetches in data tables in shared memory would generate memory contention; all instruction fetches would be contention free and faster. The software changes that would be required are not discussed in this report.

Advantages. There are five main advantages that are obtained if the IOCE memory stacks are replaced and the IOCE processors are sped up.

First, because the sped-up processors execute the program elements that have been placed in the IOCE memory, the processing power of the system increases. It is estimated that this increase in processing power for the 9020A's is at least 15 percent with probability 0.98, at least 30 percent with probability 0.88, and at least 70 percent with probability 0.49. This increased processing power will not all be realized immediately but only as program elements are moved into the IOCE's.

Second, because the PE's moved to the IOCE's need no longer be executed from main memory, this deals somewhat with the memory capacity problem. The degree to which the lack of shared memory is taken care of depends on the size and number of PE's that are moved to the IOCE's.

Third, insofar as the memory capacity problem is taken care of, there will be less need to buffer programs and data on disk. Therefore, swapping in and out of main memory will be decreased, and this will at least partly deal with the I/O problems.

It is seen that these IOCE enhancements can deal partially and perhaps fully with the three main problem areas of processing capacity, memory, and I/O. The degree to which these enhancements deal with these problems cannot presently be answered; the answers can only be provided once further studies are done of these enhancements and once the FAA specifies the improvements that are needed. The additional advantages of these enhancements will now be discussed.

Fourth, this enhancement would speed up the channels. This would allow the current peripherals (e.g., disk drives) to be replaced with faster and more reliable modern peripherals.

Fifth, if the new ROS that is installed in the sped-up IOCE processor is enlarged, this would allow the IOCE to recover the floating point and decimal instructions that are now lacking because of ROS space limitations. This would require either that IBM furnish the needed microcode or that the microcode be obtained from the microstore of a 9020A CE.

Besides these advantages, the other advantage obtained by replacing the memory in the SE's that are described in Sec. 2.2.3. would be obtained: lower response time, reduced software maintenance cost, increased reliability, and more scope for functional enhancements. Whether these advantages would be obtained in the same degree depends on the size of the PE's moved to the IOCE's.

Cost. The cost of speeding up the IOCE processors at the 9020A sites has three components. First, the cost of designing the converted processor and building the prototype is estimated to be \$125,000. (This cost would be cut to \$20,000 if the CE processor speed-up were carried out first. That is, the prototypes for both processor speed-ups could be built for \$145,000.) Second, the cost of speeding up each IOCE processor is \$25,000. With three IOCE's at each site, and allowing another \$25,000 for spares, the cost for each site is \$100,000, and the cost for the 12 9020A sites is \$1.2 million. Third, add \$0.265 million to cover contingencies. Therefore, the total cost of speeding up the IOCE processors at the 12 sites is estimated to be \$1.6 million. If the IOCE's are also sped up at the 11 9020D sites,

this adds \$1.1 million plus \$0.220 million to cover contingencies, for a total of \$2.9 million for speeding up the IOCE processors at all sites. These cost estimates do not include the cost of the required software changes; a preliminary investigation indicates that the cost of these software changes will not be significant.

3.4 9020A CE Replacement

3.4.1 Description of this Enhancement

If the two speed-up options described in Sec.'s 3.2 and 3.3 prove to be infeasible or to provide an insufficient increase in processing capacity, then the fall-back option is to replace each 9020A CE by a machine with capabilities similar to an IBM 4341. That is, the new machine would be able to execute perhaps one million instructions per second and would have cache and internal main memories with a 300 nanosecond access time. The machine would require hardware and firmware modifications to allow it to work in the 9020A environment, e.g., a modification to the ROS would be necessary to enable it to execute the 9020A's special instructions. It is assumed that this CE replacement is preceded by the memory replacement described in Ch. 2. The main question is how the different memories are to be used; the three different memories involved are the memory shared by all the processors, the main memory of each processor, and the cache memory of each processor.

The method that at this time seems best is to use the shared memory and each processor's cache memory but not to use each processor's main memory. In this scheme the system would operate in much the same way as the present system except that a cache memory is added. For cache memory to work properly, only instruction fetches can be cached.

An alternate method, which probably would not be needed, would be to use all three levels of memory. The program elements would be stored in each processor's main memory and transferred from there to the cache as needed. The shared memory would contain only the tables and software flags. It is thought that this method would not be desirable because it would require

extensive software changes and because the first method would probably provide the desired increase in processing capacity.

3.4.2 Advantages of this Enhancement

The advantages of this enhancement are for the most part the same as the advantages of speeding up the CE's that are discussed in 3.2.2. The main difference is that replacing the CE's would at least double (and possibly triple) the processing capacity of the system, which is a larger possible gain than can be attained by speeding up the CE's. This doubling of processing capacity could, it is estimated, be obtained with a 95 percent probability if the method using only the shared memory and cache memory is adopted. If the more elaborate method using all three levels of memory is adopted, then the doubling of capacity could be obtained with a 100 percent probability.

3.4.3 Cost and Schedule

The cost of this enhancement has three components. First, there is a one-time engineering cost that will fall somewhere in the interval from \$0 to \$1.0 million; the best estimate is \$1.0 million. Second, the cost per processor is estimated to be from \$100,000 to \$300,000 per processor; the best estimate is \$200,000. With four processors per site, and adding in \$200,000 to cover spares, the cost per site is \$1.0 million; the cost of the new processors for the twelve 9020A sites is then \$12.0 million. Third, \$2.6 million is added for contingencies. Therefore, the total cost of this enhancement is \$15.6 million.

It is estimated that the first processor would be delivered twelve months ARO, and the rate at which processors are delivered would gradually rise until they are being delivered at the rate of one per week 18 months ARO. This means that delivery will be completed 27 months ARO. The first six sites would be enhanced within 24 months ARO.

3.4.4 Transition

Replacing a processor would result in two outages lasting five minutes each while the interprocessor cable is disconnected and connected; other cables can be handled while the system is active. Processor swap time, which mostly consists of physically moving cabinets, is estimated at four hours. Replacing one processor a day would allow a twenty hour shakedown period of the last processor before the next processor is installed.

4. SUMMARY

4.1 The Individual Enhancements

This report has argued that the types of problems that the 9020's will face over the next few years primarily lie in the areas of processing capacity, memory capacity, and I/O capacity. The six enhancements that have been proposed as possible building blocks to use to construct a strategy that will deal with these problems are:

Memory Enhancements

- Replace the SE memory boxes
- Replace the SE memory stacks
- Replace the IOCE memory stacks

Processor Enhancements

- Speed up the CE processors
- Speed up the IOCE processors
- Replace the CE's.

Replacing the SE memory boxes or replacing the SE memory stacks would solve the memory and I/O problems for both the 9020A and 9020D systems, as Chapter 2 has shown. Replacing the IOCE memory stacks would deal with these problems somewhat, but it cannot at present be said to what degree this enhancement would take care of these problems. All three of these memory enhancements would, moreover, provide some increase in processing capacity. Whether this increase in processing capacity is sufficient to take care of the 9020A's processing capacity problem depends on how much of an increase the 9020A's need and on how much these enhancements can provide; both of these are open questions. If it is decided that enhancing the memory will not provide the needed increase in processing capacity, then one of the processor enhancements could be adopted.

The enhancements of speeding up the processors in the CE's and of speeding up the processors in the IOCE's are attractive because of their relative inexpensiveness and the speed with which they can be implemented. Either of these enhancements could be adopted, or both could be adopted if that were necessary to achieve the desired increase in processing capacity. One problem with the processor speed-up is that it is currently not known for certain whether it is feasible. A \$125,000 study will be needed to determine whether it is feasible. If it is infeasible, or if these enhancements cannot provide the needed increase in processing capacity, or if these enhancements prove to be unsuitable for some other reason, then the fall-back option of replacing the CE's could be adopted.

Table 4-1 summarizes the main information about each enhancement.

Replacing the SE memory boxes would cost an estimated \$8.2 million. This would increase processing capacity by between 20 and 60 percent for the 9020A's and by between 10 and 30 percent for the 9020D's; there is full confidence that these increases can be attained. This enhancement could be implemented at the first six sites within 24 months after receipt of order (ARO).

Replacing the SE memory stacks would cost an estimated \$5.6 million. This would increase processing capacity by between 20 and 60 percent for the 9020A's and by between 10 and 30 percent for the 9020D's; there is full confidence that these increases can be attained. This enhancement could be implemented at the first six sites within 8 months ARO.

Replacing the IOCE memory stacks only at the 9020A sites would cost an estimated \$1.9 million; replacing the stacks at both the 9020A and 9020D sites would cost an estimated \$3.5 million. This would increase processing capacity by between 10 and 30 percent for the 9020A's and by between 5 and 15 percent for the 9020D's. This enhancement could be implemented at the first six sites within 8 months ARO.

Speeding up the CE processors at the 9020A sites would cost an estimated \$2.0 million. When combined with an SE memory enhancement, this enhancement

TABLE 4-1: CHARACTERISTICS OF THE SIX ENHANCEMENTS

<u>Enhancement</u>	<u>Cost</u> (millions)	<u>Processing Capacity</u> ¹		<u>Schedule</u> (first six sites) (months)
		<u>Increase</u> (%)	<u>Probability</u> ⁴ (%)	
1. Replace SE memory boxes	A&D: \$8.2	A: 20-60 D: 10-30	100 100	24
2. Replace SE memory stacks	A&D: 5.6	A: 20-60 D: 10-30	100 100	8
3. Replace IOCE memory stacks	A: 1.9 A&D: 3.5	A: 10-30 D: 5-15	100 100	8
4. CE Speed-Up ²	A: 2.0	A: 25 A: 50 A: 100	98 88 49	6
5. IOCE Speed-Up ³ memory stacks	A: 1.6 A&D: 2.9	A: 15 A: 30 A: 70 D: 10	98 88 49 88	6
6. CE Replacement ²	A: 15.6	A: 100- 200	100	24

¹ Processing capacity refers to the peak number of tracks that can be handled. This increase is relative to the standard 9020 configuration.

² A prerequisite for this enhancement is replacement of either the memory boxes or the SE memory stacks. The cost of this enhancement excludes the cost of the prerequisite; the increase in processing capacity, however, is the increase that would result from adopting both this enhancement and its prerequisite.

³ A prerequisite for this enhancement is replacement of the IOCE memory stacks. The cost of this enhancement excludes the cost of the prerequisite; the increase in processing capacity, however, is the increase that would result from adopting both this enhancement and its prerequisite.

⁴ These probabilities are best estimates based on a study of the system and on experience; they should not be interpreted as exact probabilities.

would increase processing capacity by at least 25 percent with probability 0.98, by at least 50 percent with probability 0.88, and by at least 100 percent with probability 0.49. This enhancement could be implemented at the first six sites within 6 months ARO.

Speeding up the IOCE processors only at the 9020A sites would cost an estimated \$1.6 million; speeding them up at both the 9020A and 9020D sites would cost an estimated \$2.9 million. When combined with the replacement of the IOCE memory stacks, this enhancement would increase the 9020A processing capacity by at least 30 percent with probability 0.88 and by at least 70 percent with probability 0.49. This enhancement could be implemented at the first six sites within 6 months ARO.

Replacing the CE's at the 9020A sites would cost an estimated \$15.6 million. This would increase the 9020A processing capacity by between 100 and 200 percent; we can have full confidence that the increase will at worst fall into this range. This enhancement could be implemented at the first six sites within 24 months ARO.

Information about the cost and schedule of developing engineering prototypes for the enhancements that involve a memory stack replacement or a processor speed-up is of special interest since there is uncertainty about whether these enhancements are feasible and about exactly how much of an increase in processing capacity they would provide. The upper part of Table 4-2 shows for the four relevant enhancements the cost of developing the prototype under the assumption that the prototype is built for only this enhancement. Also shown is the estimated time it would take; this prototype would need to be completed before the FAA placed the order for the hardware. The lower part of Table 4-2 shows the cost and schedule for combinations of enhancements where there is an interaction. For example, building the prototype just for the 9020A CE processor speed-up costs \$125,000, and building the prototype just for the IOCE processor speed-up also costs \$125,000; both prototypes, however, could be built for \$145,000.

The considerations that arise when trying to devise a combination of these enhancements to deal with the 9020's problems are discussed in the next section.

TABLE 4-2: COST AND SCHEDULE FOR DEVELOPING THE PROTOTYPES

<u>Enhancement</u>	<u>Cost</u>	<u>Schedule (months)</u>
Replace SE memory stacks	A: \$ 95,000 D: 115,000 A&D: 155,000	5
Replace IOCE memory stacks	105,000	5
CE Speed-Up	125,000	5
IOCE Speed-Up	125,000	5
Replace A&D memory stacks and IOCE memory stacks	175,000	6
CE Speed-Up and IOCE Speed-Up	145,000	6

4.2 Strategies Open to the FAA

Choosing among strategies. It seems unlikely that the FAA will be able to deal with the 9020's problems by adopting a single enhancement; the FAA will probably need to combine two or more enhancements in order to form a workable strategy. This section will sketch out some of the relevant considerations and lay out some of the strategies that the FAA might adopt.

In choosing among the six enhancements, there are two sets of constraints that should be observed. First, some of the enhancements have prerequisites. Speeding up the CE's or replacing the CE's requires that the memory boxes or the SE memory stacks be replaced. Speeding up the processors in the IOCE's requires that the IOCE memory stacks be replaced. Second, it would not make sense to replace both the memory boxes and the SE memory stacks, and it would not make sense to both speed up the CE's and replace the CE's.

Even after these constraints are taken into account, one can still construct 20 strategies from combinations of the 6 enhancements; these 20 strategies are exhibited in Appendix D. Since this is too many strategies to discuss individually, three further simplifications will be made.

Simplifications. First, consider the choice between upgrading the 9020's shared memory by buying new memory boxes or by replacing the SE memory stacks. There are four relative advantages of buying new memory boxes. First, the entire memory box would contain state of the art components and designs. Second, built-in diagnostics would be included. Third, the entire SE would be the responsibility of one vendor. Fourth, if it were later decided to upgrade the 9020A's to 9020D's, the new memory boxes could be used in this upgrade.

There are four relative advantages to replacing the memory stacks rather than the entire boxes. First, replacing just the stacks is cheaper, i.e., \$5.6 million v. \$8.2 million. Second, replacing just the stacks is much faster; it would take about 8 months to upgrade the first six systems compared to 24 months if the memory boxes were replaced. Third, replacing

just the stacks is physically easier and less prone to problems since no recabling is required. Fourth, with memory stack replacement the decision on whether to upgrade at any particular center could be made on a case by case basis since there is no advantage to buying the components in bulk and since there is a short lead time. In contrast, if the memory boxes were replaced, the number of centers at which this enhancement is to be implemented should be decided when the contract for the boxes is let. Therefore, replacing just the memory stacks gives the FAA more flexibility in deciding how many centers will be upgraded and when.

In sum, these two memory enhancements differ mainly not in performance but in other ways. The decision which is preferred would depend on how the appeal of replacing the entire boxes as a single unit is weighed against the time and cost savings and the flexibility of replacing just the memory stacks. To simplify the discussion, these two memory enhancements will be lumped together as the enhancement of "replace SE memory;" this enhancement will stand for either replacing the SE's or replacing the memory stacks.

The second simplification to be made lies in the choice between achieving an increase in processing capacity by replacing the CE's or by speeding up the CE processors. The relative advantage of replacing the CE's is that with very little uncertainty the processing capacity of the 9020A can be doubled or tripled. There are two relative advantages of speeding up the processors. First, the increase in processing capacity can be achieved much faster, i.e., 6 months v. 24 months for the first 6 systems if the CE's are replaced. Second, speeding up the processors is much cheaper, i.e., \$2.0 million v. \$15.6 million for the 12 systems. Since the speed-up is so much faster and cheaper than the replacement, for purposes of discussion it will be assumed that the speed-up is preferred to the replacement. It should be emphasized that this assumption is made only to simplify the exposition.

The third simplification concerns the enhancements to the IOCE's. While it is possible that the IOCE memory stacks might be replaced without speeding up the IOCE processor, this seems like an unlikely event.

Therefore, these two IOCE enhancements will be grouped together under the title of "IOCE upgrade."

Decision tree. Now consider the simplified decision tree in Figure 4-1, which shows some of the choices facing the FAA. At fork 1 the FAA would decide whether as a first step in upgrading the 9020's it would be better to replace the SE memory or to upgrade the IOCE's at the 9020A and 9020D sites. The cost and schedule of these two enhancements are not dramatically different, so the choice between them would be made on the basis of the four differences between them. First, replacing the SE memory involves more hardware changes. That is, if the IOCE's are upgraded, changes need be made only in the three IOCE's; if the SE memory is replaced, all the SE's would be affected, and if it is followed by speeding up the processors, all the CE's would be affected. Therefore, upgrading the IOCE's would entail less change to the hardware. Second, upgrading the IOCE's involves more software changes. Replacing the SE memory would require no significant software changes, whereas upgrading the IOCE's would require that program elements be moved from shared memory to the MACH memory. Third, replacing the SE memory would immediately take care of the 9020A and 9020D memory and I/O problems. In contrast, upgrading the IOCE's provides relief only insofar as the needed software changes are made, and it is not yet clear how difficult it will be to make these changes. Fourth, these enhancements differ in their potential increase in processing capacity. Replacing the SE memory would yield an increase in processing capacity for the 9020A of between 20 and 60 percent; if the processors in the CE's are then sped up, the total increase in processing capacity is between 25 and 100 percent. Upgrading the IOCE, in contrast, would provide an increase in processing capacity of between 15 and 70 percent.

Suppose that at fork 1 the FAA decides to replace the SE memory. The FAA then has the further decision, not shown in Figure 4-1, of whether this should be done by replacing the memory boxes or stacks; the relative advantages of each are discussed above. Suppose now that the FAA is at fork 2. Since replacing the SE memory takes care of the memory and I/O problems and provides a modest increase in processing capacity, the FAA might decide that nothing else needs to be done. If, however, the FAA decides that more

COST
(millions)

\$5.6 or 8.2

7.6 or 10.2

10.8 or 13.6

6.4

13.7 or 16.5

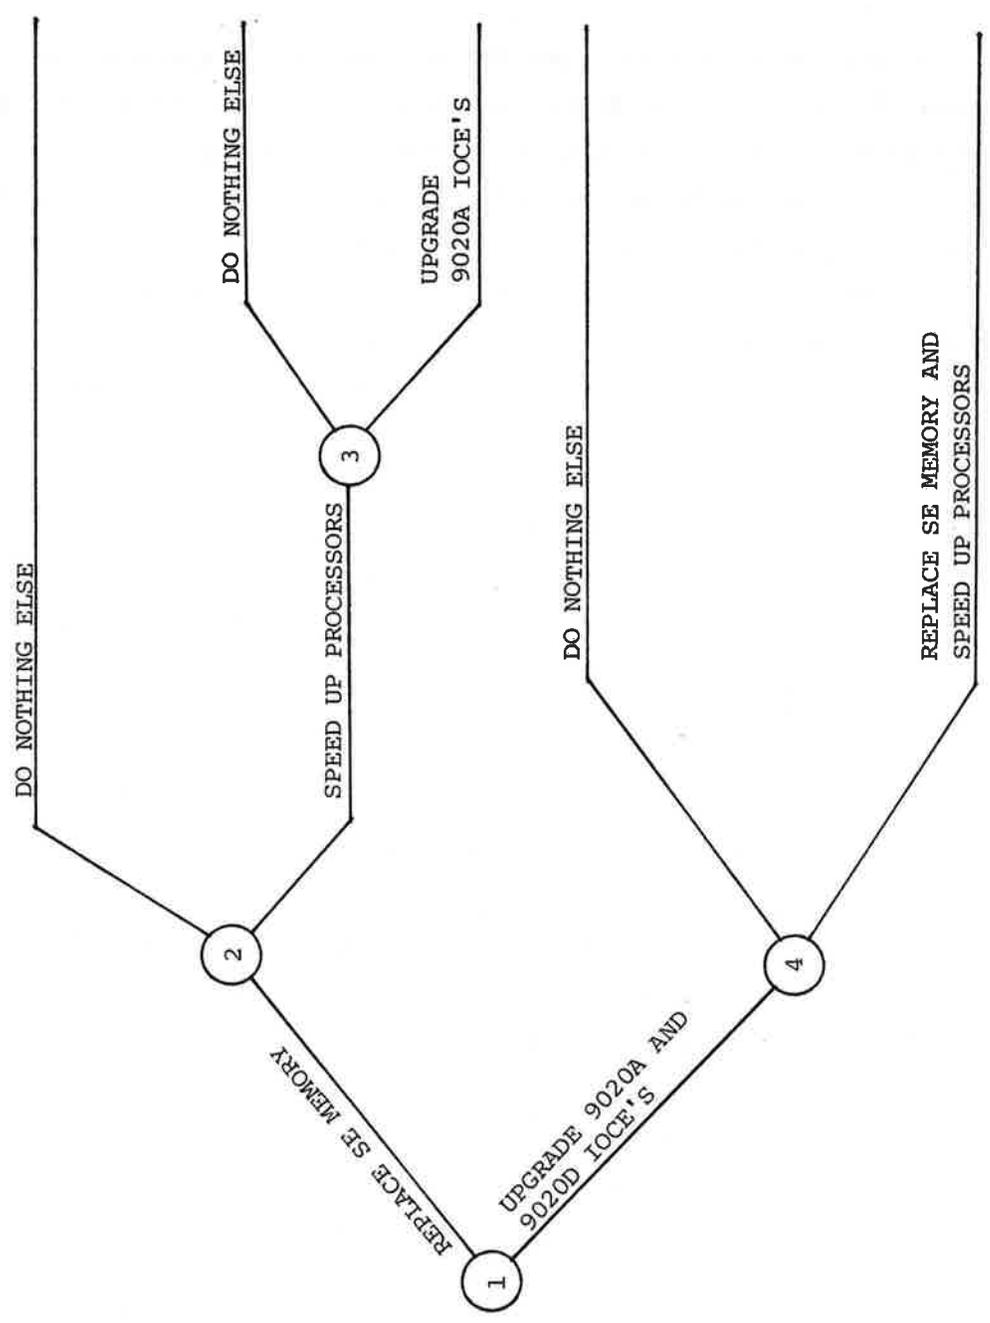


FIGURE 4-1: LEADING STRATEGIES OPEN TO THE FAA

processing capacity is needed, it can speed up the processors in the CE's, thus arriving at fork 3.

If the FAA is at fork 3 and decides that enough processing capacity has been achieved, then it need do nothing else. If, however, more processing capacity is desired, the FAA can upgrade the IOCE's at the 9020A sites. (Since the SE memory replacement would take care of the 9020D's problems, there would be no need to upgrade the IOCE's at the 9020D sites.)

Suppose now that back at fork 1 the FAA had decided to upgrade the IOCE's instead of replacing the SE memory. This places the FAA at fork 4. If the FAA decides that the IOCE upgrade provides all the needed capabilities, then there would be no need to do anything else. If the IOCE upgrade is not sufficient, then the FAA could further enhance the system by replacing the SE memory and speeding up the processors in the CE's. (Just replacing the SE memory at this stage probably would not be a good idea since the IOCE upgrade would have provided the system with sufficient memory.)

The estimated cost of each strategy is shown in Figure 4-1. This cost reflects the reduction in prototype development cost that occurs because of interaction between the enhancements, and it also reflects the resulting reduction in the amount allowed for contingencies. Each path that includes "Replace SE memory" has two costs depending on whether the memory stacks or the memory boxes are replaced.

Exactly which path, if any, through this tree is chosen depends on how much of an increase in processing power is needed, when it is needed, and how much each enhancement can provide. Two comments about these choices should be made. First, the times at which the decisions are made are not specified in the tree. On the one hand, the FAA might make all the decisions at one time. On the other hand, the FAA might make the decisions sequentially. That is, the FAA might implement one enhancement and then only decide whether to implement another enhancement after seeing how well the first enhancement works, what the projected need is for processing capacity, and how quickly the 9020 replacement program is proceeding.

Second, this decision tree does not take into account the possibility of replacing the CE's. A somewhat different tree would need to be drawn to reflect this enhancement.

Summary. The five strategies depicted in Figure 4-1 are:

1. Replace the SE memory,
2. Replace the SE memory and speed up the 9020A CE processors,
3. Replace the SE memory, speed up the 9020A CE processors, and upgrade the IOCE's at the 9020A sites,
4. Upgrade the IOCE's at the 9020A and 9020D sites, and
5. Upgrade the IOCE's at the 9020A and 9020D sites, replace the 9020A SE memory, and speed up the 9020A CE processors.

The choice among these strategies depends on the increase in processing capacity that is needed, when it is needed, how much each enhancement provides, and on the perceived difficulty of the hardware and software modifications that the various enhancements require.

In brief, there are a number of hardware enhancements to the 9020's that the FAA could potentially adopt. By developing the requirements that the 9020's must fulfill over the next few years and by studying the characteristics of these enhancements, the FAA will be able to combine selected enhancements into a strategy for dealing with the 9020's potential problems.

In closing, one important point that must be stressed is that if the FAA wants to know quickly and with precision the magnitude of the advantages yielded by these enhancements, then it should complete the development of the engineering prototypes. Since there are only minor differences between the CE processor speed-up and the IOCE processor speed-up, one prototype development lasting about five months will provide the needed information

about both these enhancements. Similarly, one prototype development lasting about five months would provide the needed information about the memory stack replacement enhancements. These prototype studies should proceed for three reasons. First, these studies will provide information needed if the FAA is to decide which strategy best meets its needs. Currently, it is not known whether the processor speed-up is feasible, and it is not known with precision how much of an increase in processing capacity each enhancement would provide; this information can only be obtained by completing the prototypes. Second, the rapid implementation times quoted in this report assume that the working prototype has been developed. That is, the CE processor speed-up can be implemented at the first six sites in eight months, but only if the prototype has already been developed; if it has not been developed, then another five months must be added to this schedule. Third, compared to the amounts of money at stake, the prototype studies involve a trivial cost. In sum, immediate development of these prototypes is suggested since this will provide at a low cost the information that the FAA can use to decide what strategy is best and since this will bring closer the time when the strategy that is eventually chosen can be implemented.

APPENDIX A. THE MODEL OF SYSTEM PERFORMANCE: OVERVIEW

A.1 Purpose and Organization of this Appendix

In order for the FAA to decide which of the enhancements discussed in this report should be adopted, it is desirable to have estimates of the gain in performance that each enhancement would yield. To provide these estimates a model of 9020A system performance has been constructed. In addition to estimating possible gains in performance, this model can also be used to help design the enhancements. This appendix gives a high-level discussion of the model and its main features. App. B then gives a detailed discussion of the model and the results that have been obtained from it.

Sec. A.2 describes the model inputs, i.e., the parameters that can be varied between runs of the model to reflect the different enhancements and work loads. Sec. A.4 describes the model outputs, i.e., the information about system performance that the model yields. Sec. A.3 describes the model logic, which tells how the model views the process being modeled; that is, the model logic tells how the model goes about transforming inputs into outputs. To increase the usefulness of this model, more data is needed to serve as input and to validate the model; Sec. A.5 lists the measurements that should be taken to provide this data.

This appendix only gives a general discussion of the model designed to acquaint the reader with its main features; for a more detailed understanding, the reader should consult App.'s B and C.

A.2 Model Inputs

Each run of the model simulates a different scenario; various scenarios differ in the characteristics of the computer system or in the workload that is placed on the computer system. A scenario is characterized by choosing values for the model's inputs, and the goal is to choose values that represent a scenario of interest. The inputs that can be varied between runs of the model fall into three areas.

First, there are characteristics of the 9020 system that, while they could be changed, typically are not changed between runs because they are unaffected by the enhancements discussed in this report. Examples of these inputs are a list of the program elements (PE's), the instruction mix for each PE, and for each PE the average number of instructions executed each time it is activated.

Second, there are the characteristics of the 9020 system that typically are changed between runs because they are affected by the enhancements discussed in this report. The primary inputs that fall into this category are:

- memory cycle time,
- execution time of every instruction (not counting the memory cycle time),
- number of memory units,
- memory map, which shows where all programs and data are stored.

For any one run of the model, values are chosen for these inputs that describe the particular enhancement that is being considered. For example, for the memory replacement enhancement, the memory cycle time drops because the memory cycle falls from five to four microcycles; the number of memory units increases from seven to ten in the 9020D and decreases from eleven to seven in the 9020A; the memory map changes significantly since buffering is eliminated. When memory replacement is supplemented with a CE enhancement, this decreases the microcycle time, which is reflected in the inputs by reducing the memory cycle time and the instruction execution time.

Third, there are the inputs that reflect the workload that is placed on the system. The main input describing workload is the number of times each PE is activated per hour.

Once these inputs have been specified, the model is ready to run; the model logic then uses those inputs to determine how the system performs.

A.3 Model Logic

The model logic describes how the 9020 system operates; the simulation, by tracing out this operation, can determine how the enhanced system would perform with various enhancements.

Start by considering a single processor that has an instruction to execute. This processor follows an eight step cycle.

- 1) A random number is drawn to determine what the specific instruction is. (This depends not only on the random number but also on the PE being executed.) Once the specific instruction is determined, then various things are known, e.g., how many, if any, references to memory must be made.
- 2) If no reference to memory is made, go to step 6); if a reference is made to memory, go to step 3).
- 3) Determine which memory unit must be accessed.
- 4) This processor goes to the relevant memory unit; if the unit is tied up serving another request, then this processor queues up until it is given access to this memory unit.
- 5) The processor receives the desired information from memory.
- 6) The processor executes the instruction.
- 7) To obtain the next instruction to be executed, the processor determines which memory unit must be accessed, queues up if necessary at that unit, and eventually receives the next instruction to be executed.
- 8) Go back to step 1).

The memory replacement enhancement causes an increase in performance since there is less memory interference at steps 4) and 7) and since the memory cycle time (for the 9020A) is faster in steps 5) and 7). The CE enhancement, which decreases the microcycle time, causes an increase in performance at steps 5), 6), and 7).

While this eight step procedure is the heart of the model, it is not the entire model. The model logic also governs the order in which PE's are executed and how PE's are allocated among the processors.

A.4 Model Output

During the simulation, statistics are kept that describe what happens during the simulation. The primary output of the model is the amount of simulated time that it takes for the specified workload to be carried out. That is, given a workload, the model predicts how long it would take the enhanced 9020 system to dispose of that workload. The performance figures cited in the text refer to the decrease in time it would take for the enhanced 9020 system to perform a set task.

A.5 Needed Data

This report gives estimates of the performance gains that each enhancement would yield (e.g., Table 4-1). These estimates were obtained by running the model with the best available data, but confidence in the model's results could be improved if new measurements were made to obtain the data that is most critical to the model. The measurements that are needed to provide input data and to validate the model are as follows.

1. Memory References Per Time Unit
 - Each CE
 - Each IOCE
2. Peripheral Utilization
 - Disks
 - Tapes
 - Selector and Multiplexor Channel Activity
3. PE Activity
 - Number Activations Per Time Unit
 - Time Active For Activation
 - Number Memory References Per Activation
 - SE Number For Each Activation
 - Start Time For Each Activation

4. Number of SVC's Per Time Unit
5. Dispatcher Activations and Time Active
6. I/O Interrupt Processor Activation and Time Active
7. SVC Handling Activations and Time Active
8. Non-PE Activity in CE-Number Activations and Times
9. Subprogram RIN Memory References Per Time Unit
10. Number of Tracks Active Per Time Unit
11. Total Number of Instructions Executed Per Time Unit
12. Sequence of CE's Requesting SE Access
13. Number of Proposed Tracks Per Time Unit
14. Wait-On-CE Delay for PE's
15. I/O Delay for PE's
16. Lock Delay for PE's

APPENDIX B. THE MODEL OF SYSTEM PERFORMANCE: DETAILED EXPOSITION

B.1 Purpose and Organization of this Appendix

This appendix details a simulation model used to analyze the 9020A system. The model is adapted from a model used in [PATT73] and uses the techniques described in [FRAN77]. The technique used has previously been used successfully by members of the staff of Architecture Technology to model for Navy Real-Time environments the AN/UYK-7 and for BMD Site Defense environments the CDC 7700, 3X2 CDC 7600 configurations, H6000 Series multiprocessors and Univac 1100 Series multiprocessors.

The general structure of the model is described in Sec. B.2. The model was definitized by parameters and suitable modifications until it accurately represented the 9020A. The results obtained by running the model are described in Sec. B.3.

B.2 The Model

The model, which is implemented by a SIMULA program, represents a system consisting of two types of entities. These are processing elements and memory modules. Processing elements are parameterized to represent either a CPU (processor) or an IOC within the 9020A system. Memory modules are established to service requests which result from the operation of the processing elements in the system.

The memory modules in the model are instances of a SIMULA process class. This means that each individual memory module is modeled by a process which interacts with other elements of the system. At certain points in the action of this process, simulated time is used to allow for the proper interactions. The device for this interaction is represented by a switch of ports through which requests for service can be made by various processing elements. These entries model the switch connections which can be made with memory modules in the 9020A. Bus connections are represented by assigning each processing element having use of specified ports into the memory module process.

The action of the memory module process is described by a cyclic acknowledgement of processing requests. Since the simulation model does not detail the content of memory references, no information transfers are represented. A request exists because of the action of a processing element in simulated time. The request is represented by a flag in the port entry. The memory module process services the request by simply clearing the flag; the result is a simple synchronization exchange. As long as requests to be serviced remain in the array of ports for an individual memory module process, it cycles to service requests. Each cycle involves locating the request, and signaling completion of service for the request. Therefore, the first port is given highest priority, etc. A complete description of this process is given in the flow chart in Figure B-1.

As long as requests exist, a memory module process remains active. If no more requests remain at the beginning of a new cycle, then the process passivates. Entry of a new request into a port of an individual memory module process restarts the passivated processes as required.

The action of the processing element is also cyclic. However, the possible paths during a cycle are greater in number and the decision points are controlled by pseudo-random draws from given distributions. The results of the cycling of the processing element process are requests to various memory modules for service. As indicated above, these requests are represented as synchronization exchanges. Accordingly, the progress is partly controlled by the memory modules.

Each processing element contains two sources of requests to memory modules. These are the instruction word reference, denoted iref, and the operand reference, denoted oref. Associated with each of these sources is a dedicated bus assignment represented as a port ordinal. This ordinal is an integer from one to eight. Only one request source may be assigned a given port ordinal or bus number. The lower numbered busses have the higher priority. To be consistent with 9020A characteristics, the iref source of an individual processing element process should be assigned a lower bus number than the oref source. However, the model itself does not require this.

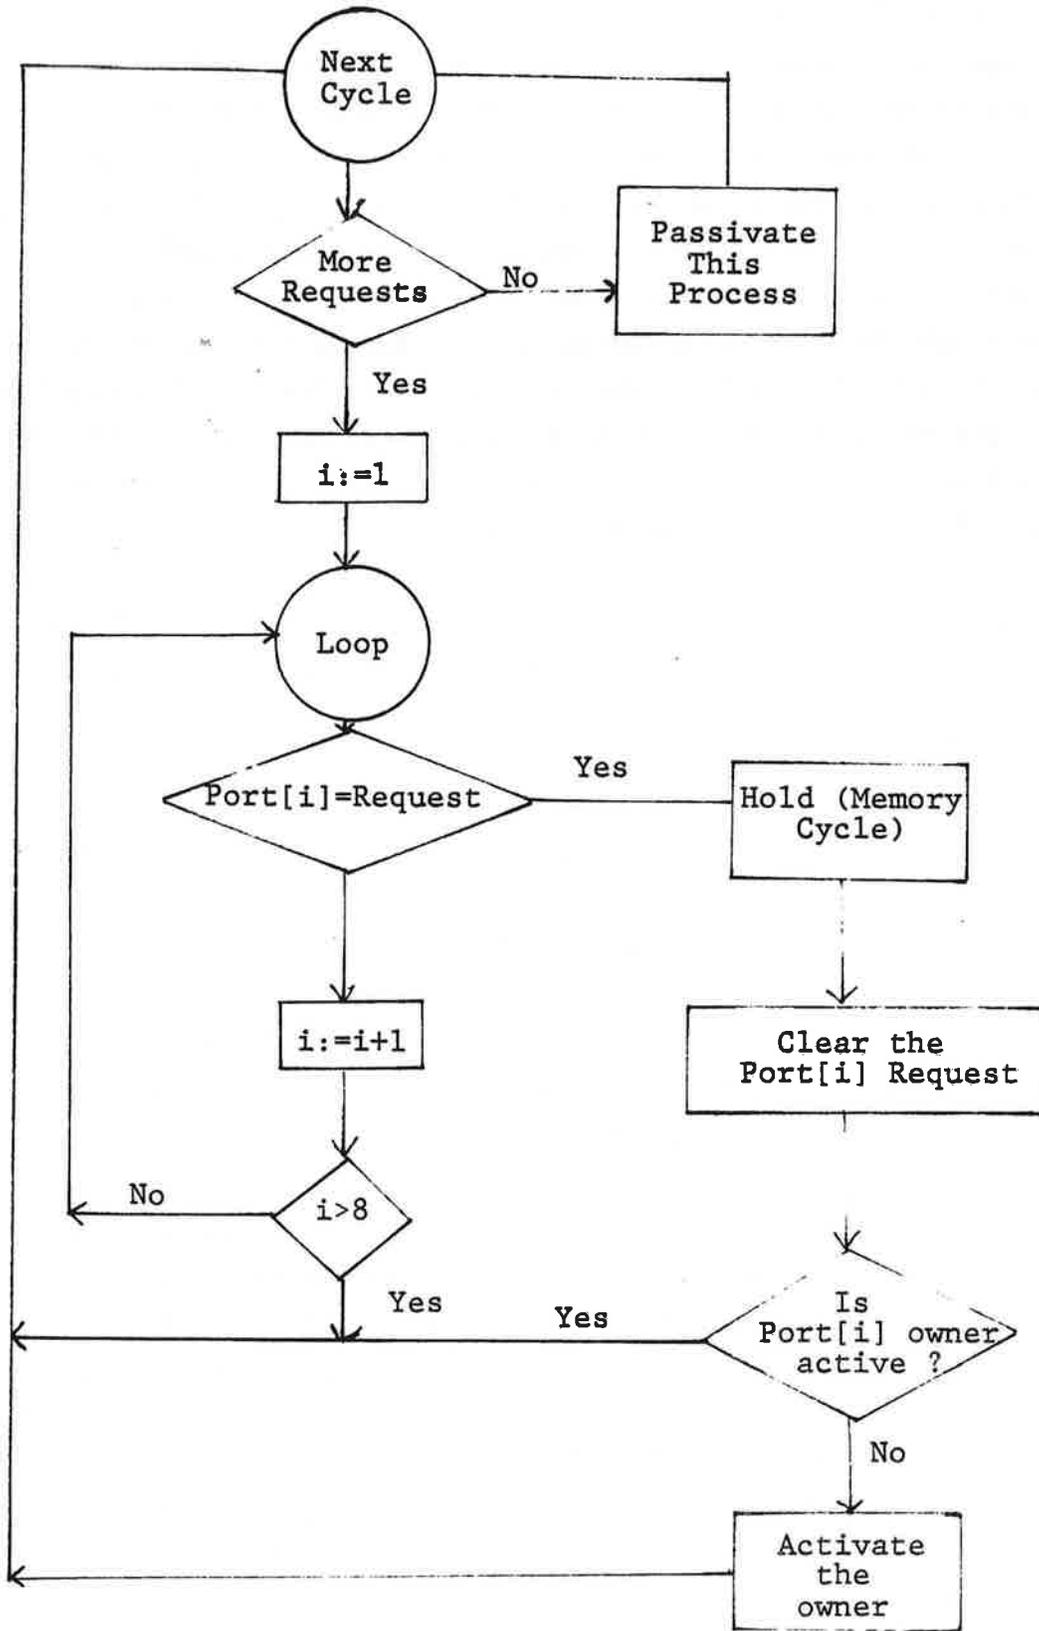


FIGURE B-1: THE FLOW DIAGRAM FOR THE MEMORY MODULE PROCESS MODULE

In what follows, we describe the flow of the processing element process cycle. Where the decision points depend on a pseudo-random draw, the word DRAW will indicate this. The paragraphs below identify the various distributions required and how they influence the behavior of the processing element model.

The selection of a command within the processing element cycle and the determination of its characteristics is controlled by four vectors. These are the instructions probability vector, denoted INSTPROB, the instruction cost vector, denoted INSTCOST, the instruction type vector, denoted ISTYPE, and the instruction length vector denoted CLENGTH. Each of these vectors are of size N where N represents the number of command orders to be simulated. Specifically, INSTPROB [I] represents the probability that a command will be order I. Given that the command order is I, INSTCOST[I] represents the time cost for any execution portion of the command, which may be zero. ISTYPE[I] is either 1, 2, or 3 and indicates if the command is a jump, no operand, or operand command, respectively. CLENGTH[I] specifies the amount of the current instruction word utilized by the command. Notice that the units of CLENGTH need only be consistent with IPW.

References to the memory modules are generated as integers specifying which memory module must service the request. This action operates essentially as a Markov process. A current state for instruction and operand reference is maintained as PREG and QREG, respectively. Each reference is then a transition from the current state PREG (or QREG) to the next state which becomes the new value of PREG (or QREG). The memory reference is also simulated along with each state transition. The simulation model identifies four distinct transition types within the framework of the processing element model. These transitions correspond to instruction reference on sequential references, instruction references on branching references (jumps), operand references first kind, and operand references second kind.

Before dealing with the specific interpretation of these four transitions, we should develop the notational machinery a bit more. Formally, a reference transition can be represented as $p=F(T,p)$ where p represents the old state, T a transition matrix, and F a function operating

on T and p . The transition matrix T is an $m \times m$ matrix where m is the number of memory modules which are addressable. For $T = [p_{ij}]$, p_{ij} represents the probability of the next reference to memory going to module j given the last reference to i . For convenience, a mode can be associated with certain special cases of the matrix T . These are listed as follows:

uniform $p_{ij}=k$ for all i,j

banked $\begin{cases} p_{ij}=1 & \text{for } i=j \\ p_{ij}=a & \text{for } i \neq j \end{cases}$

phased $F(T,p) = \begin{cases} p+1 & p < M \\ 1 & p = M \end{cases}$

Through trivial extensions to the SIMULA program, additional specialized transitions could be defined. However, the simulation program has the option for defining the access to memory by an explicit statement of the transition matrices for each of the four transitions.

The first transition involves the memory reference for the next sequential instruction word. This is termed the read next instruction (RNI) sequence. Basically, the transition defined for RNI is a specification of how sequential addresses are mapped to the memory modules. More likely than not, this transition will be a function of hardware configuration than of software organization.

The second transition concerns the branch or jump command. Since the occurrence of a jump command is a break in the sequential behavior of the RNI operations, an alternate transition matrix (or mode) is in order. This would usually depend more heavily on software organization since jump instructions may cross certain hardware partitions, etc. Alternatively, the degenerate case of jump references involving the same transition probabilities as RNI can be easily handled by establishing the same definition for both.

The operand reference transitions are of two kinds; this splitting is arbitrary from a hardware or architectural point of view. Operand memory references occur relative to a last reference state QREG. Ordinarily, one might think that operand references would address memory modules independent

of the instruction referencing state; however, this is not true. Programs dealing with vector or array numerical structures display behavior very dependent on program control. Rather than attempting to model this factor in terms of the operand to program dependence, the possibility of two types of operand transitions were allowed. Which of the two kinds occur is controlled by a Boolean draw based on a probability γ . What the two kinds of transitions are and how they differ is then supplied as part of the definition of the model. An example would be to allow operand references to uniformly random with probability 0.5 (kind 1 is uniform, $\gamma=0.5$) and phased through the sequential module number with probability 0.5 (kind 2 is phased, $\gamma=0.5$). A further discussion of the utility of the dual operand availability is contained in the section reporting the results of examining program behaviors.

Figure B-2 shows the basic flow of the processing element model. This figure does not detail the dual operand transition, but does show how the command order identification interacts with memory reference transitions. This interaction provides for realistic statistical dependence between the command order distributions and the memory module addressing distributions. Complete models of the 9020A system along with an appropriate workload can be provided in terms of these parameters and command distributions.

The description of the 9020A simulation program in detail is in terms of the CONTROL DATA implementation of SIMULA. The reader may refer to CDC publication number 60234800 for the SIMULA reference manual; however, the description provided below will contain minimal dependence on the details of the CDC SIMULA implementation.

The simulation program manipulates three files or datasets. Two of these are the datasets INPUT and OUTPUT. The third dataset is called DATA. The dataset INPUT must contain cards describing the identification of the dataset DATA as a SCOPE operating system file. This file will contain the input data for the descriptions of the simulation runs. The cards must be of the form

```
DATASET,DATA= Lfn  
DATASET,END
```

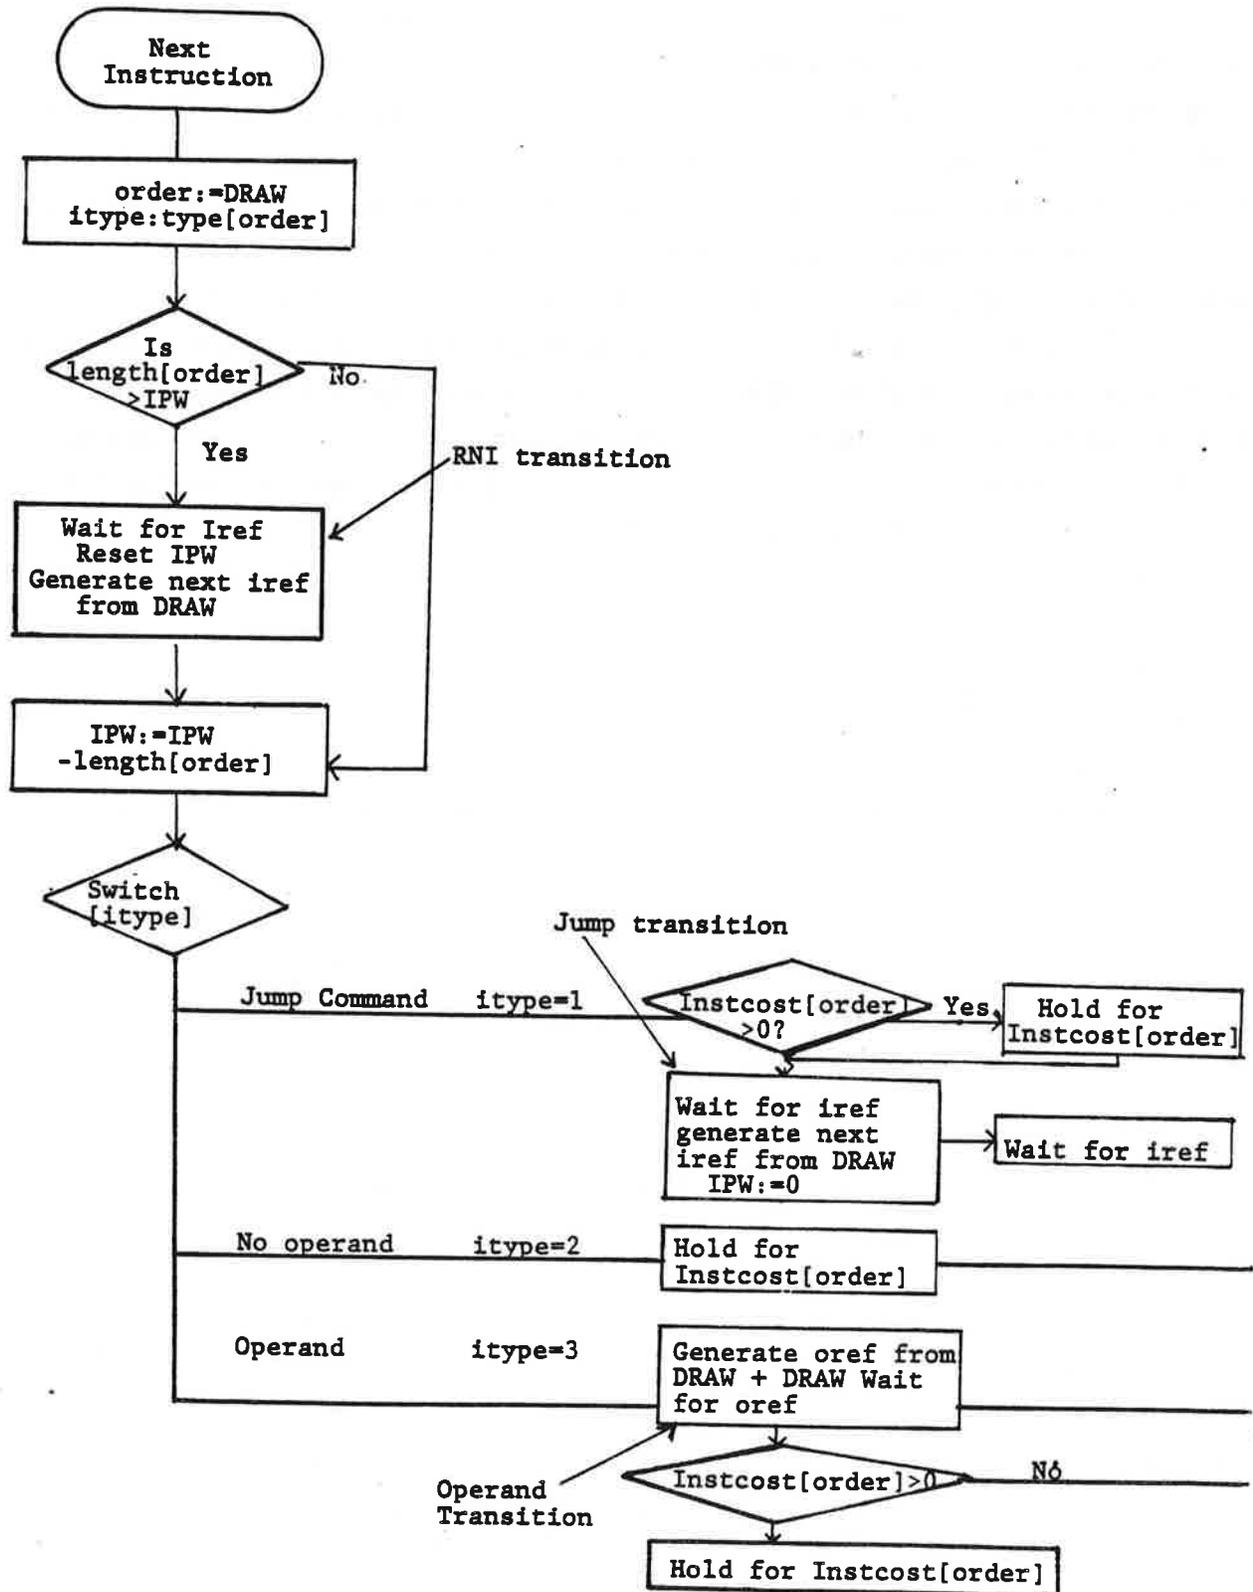


FIGURE B-2: THE FLOW DIAGRAM FOR THE PROCESSING ELEMENT PROCESS MODEL

where Lfn is the SCOPE file name. For example, if the file name is SAM, then the first card above would be

```
    DATASET,DATA=SAM
```

In any case, the file provided as the dataset DATA must be rewound.

The general format of the information provided on the dataset DATA is a sequence of problems each describing a simulation run. The SIMULA program reads the information for each problem, executes the simulation, prints the results, and proceeds to the next problem. This action is halted by an EOF condition on the dataset DATA.

The format of the information read for each problem is consistent with the SIMULA free form input/output conventions. Because this implies a sequential ordering dependence on the entire set of parameters for a problem, various keyword fields have been introduced for the sake of redundancy to prevent errors. Each keyword must begin in column 1 of a data card. In the explanation below, <string> will denote a keyword given by the indicated string of characters. The description of a simulation run as a problem is headed by the following information:

```
<PROBLEM>      m      g      mc      rt      rb      rn
```

The parameters are m, the number of memory modules; g, the number of processor groups or types; mc, the memory cycle time; rt, the run time of the simulation after the initial bias run; rn, the number of runs of length; rt and rb, the time of simulation for purposes of removing initial bias. Note that the initial bias period is followed by the clearing of all statistics and counters followed by the running of rn simulation periods of length rt. Also note that all times are of arbitrary units. However, the problem description must be consistent thus the natural unit of time would be microseconds.

The above information provides the general framework of the problem. Specific information about each group of the g groups must follow on the

dataset DATA. Accordingly, the SIMULA program expects to find on the dataset DATA g sets of items of the form:

```

<GROUP>   p   ni
<GAMMA>   δ
<IPW>     ipw
<UMASTER> u
<RNI>     rmode          [TR]
<JUMPTM>  jmode          [TJ]
<OPERAND1TM> omodel      [T01]
<OPERAND2TM> omode        [T02]
<INSTPROB> (Prob[i], i=1,ni)
<INSTCOST> (C[i], i-1,ni)
<INSTTYPE>
<INSTLENGTH>
<PE>      preg      greg      ibus      obus
.
.
.
<PE>      preg      greg      ibus      obus

```

P items

The number of processors is p . The number of instructions or commands in the processor workloads is ni . δ is the probability that an operand is operand1 rather than operand2. ipw is the number of instruction units per word. The interaction between ipw and the instruction lengths described by $L[i]$ determine the rate and distribution of instruction word memory references. u is an integer seed from which all random number streams within the processor group are started. For the mode values, the following integer values can be used.

- 1 indicates matrix to be used.
- 2 uniformly random references to all modules.
- 3 all references to same module.
- 4 references to sequential modules.

If the mode values are other than 1, then these values completely describe the discipline for memory references. If a mode value is 1, then a $m \times m$

matrix must follow. These values must be real and represent the rows of the transition matrix. Each row will be the probability density function, not necessarily normalized, for a reference to a memory module. The previous reference determines which row is used.

The next four keywords, <INSTPROB> through <INSTLENGTH>, determine the instruction mix for this group of processors. Prob [i] is a vector of the probability density function determining the instruction distribution. C[i] is a vector giving the execution time factor for the instruction selection, L[i] is a vector giving the number of instruction units this command uses in the current instruction word, and T[i] is a vector giving the instruction type. There are three types of instructions.

- 1 A jump command. C[i] is executed any outstanding instruction reference completed, a new instruction reference generated, and the new reference completed.
- 2 No operand command. C[i] is executed.
- 3 Operand required. An operand reference is processed followed by the execution of C[i].

Finally, for each of the processing elements in the group, the following information is obtained from DATA.

preg the memory module for the initial instruction reference.
greg the memory module for the initial operand reference.
ibus the bus number for instruction references (1<ibus<8).
obus the bus number of operand references (1<obus<8).

Note that two processors cannot share a bus. Accordingly, the limit of 8 busses restricts the total number of processing elements in the total system model to eight. To model a 3x2 9020A configuration three simulated processing elements would be used to simulate the three CPUs and two individual processing elements would be used to simulate the two IOCs. The model is extendable to allow handling multiprocessor configurations that drive more than eight addresses in parallel.

B.3 Results

A memory conflict analysis was done on the 9020A as a 3.2 (three CE, two IOCE) multiprocessor using a SIMULA coded model on the University of Minnesota CDC Cyber 74. This model is an instruction level simulator. The program is sufficiently general that it will handle any number processors of a variety of command structures and functional specialization in the system subject to the limit of a maximum of eight prioritized memory bus connections. The specific memory conflict model employed here is based on a general performance limited model of system function in a shared main memory multiprocessor. The model enables a close study of system performance limitation due to CE and IOCE conflicts at the memory bus or memory interface. By selective variation of parameters the user can relax constraints that cause performance limitation due to processor contention for the shared memory resource and "tune" the system at an architectural level. In this appendix we will describe the results together with implications.

Table B-1 presents the 9020A CE model input for the SIMULA program; these statistics were derived from Table 4-2 Instruction Mix and Execution Times, [KELL77]. This model organizes the commands executed in that sample into sixteen categories by instruction execution time including operand fetch (if any) from memory. The sixteen categories are further divided into three types: Type 1 are jump commands, assumed to occur 20 percent of the time, type 2 are register to register commands that do not require an operand from memory, and type 3 are the main sequential memory references for both instruction and operand. This mi is used in the workload given in Sec. B.2 to drive the simulator.

The 9020A IOCE shared memory utilization model shown in Table B-2 was derived from experience with similar configurations of similar machines in tactical real-time radar data processing applications, because a dynamic workload was not available from any of our sources for the 9020A IOCE's. The 9020A IOCE is significantly different from others previously studied, however, in that it has local memory. The most conservative modeling choice in this case was to assume that both IOCE's were fully occupied performing input/output functions for the three CE's. This assumption leads to a worst

TABLE B-1: 9020A CE COMMAND MODEL

Instruction Category	Instruction Frequency	Instruction Timing*	Type	Length
1	102768	2.5	3	2
2	16122	3.0	3	2
3	35747	3.5	3	2
4	25549	4.0	3	2
5	12131	5.5	3	2
6	4260	12.75	3	2
7	1683	14.25	3	2
8	21057	14.9	3	2
9	8764	21.0	3	2
10	7352	1.0	1	2
11	15120	1.3	1	2
12	37719	4.2	1	2
13	22248	2.5	1	2
14	8445	0.5	2	1
15	34857	0.75	2	1
16	10223	1.25	2	1

* Instruction times do not include the 2.5 microsecond fetch time for the instruction itself.

TABLE B-2: IOCE SHARED MEMORY UTILIZATION MODEL

Ratio*	Percent <u>INSTPROB</u>	Total I/O Memory Load us	<u>INSTCOST</u> I/O Memory Load Per IOCE
5:1	10	12.5	6.25
4:1	20	10.0	5.0
3:1	30	7.5	3.75
2:1	20	5.0	2.5
1:1	9	2.5	1.25
1:2	5	1.25	0.87
1:3	3	0.84	0.42
1:4	2	0.64	0.32
1:5	1	0.50	0.25

*Ratio of instructions executed to data words input or output

case conflict situation for the three CE's since they have lower bus priorities and thus will be more frequently shut out than they would have been should a light I/O load have been estimated.

The results of the simulation run on the baseline data in Tables B-1 and B-2 are given in Table B-3. The influence of the assumption of high I/O demand on shared memory plus the high priority of the IOCE's can be seen in the table. Lowering the I/O demand from 100 kops to 50 would increase one CE memory service level from 94 to about 144 but would not change the total. These results do not account for careful memory mapping to reduce CE conflict and level CE loading on memory. Since we have found previously that memory mapping reduces first order CE memory conflicts we can set an upper bound for its effectiveness as being equal to the effect of interleaving memory. Assuming nearly perfect memory mapping then allows us to take as effective memory bandwidth 465 thousand memory references per second (kmrps) rather than the 272 kmrps computed by the SIMULA model which does not account for mapping. Comparing this result with the theoretical maxima indicates that the effective memory bandwidth is far less than the possible maximum and the actual instruction rate for the CE's is less than the rate three independent CE's could sustain at an AIET of 6.23 usec or 160.5 kops per CE. If mapping is as effective as interleaving, then the system sustains a rate of 301 kops which is considerably less than the potential rate of 160.5 per CE and 50 per IOCE. This reduction must not only be understood as a consequence of sharing the main memory resource but also as a tradeoff in favor of enhanced system availability.

Table B-4 extends the baseline SIMULA results for a number of memory speedup options. The table shows memory speed of 2.0, 1.5, 1.0, 0.8, and 0.5 μ sec beyond the current or baseline value of 2.5 μ sec. The value 0.8 μ second was chosen because that is state of the art for large main memory and the value 0.5 was chosen as a first approximation to the effect of a cache memory in the 9020A system. These extensions of one baseline SIMULA results allow comparison of the speedings effect of each option with and without two way and four way memory interleaving as shown in Table B-5. For this data to be valid, the CE must be modified to allow for asyndmonous operation with respect to the memory at these specified rates. If shared memory is the

TABLE B-3: MEMORY CONFLICT IN 3x2 9020A

Memory Speed	2.5 μ sec
Max. Instruction Time per CE	5.0 μ sec
Corresp. Instruction rate per CE	200 kops
Average Instruction Execution Time (AIET)	6.23 μ sec
Corresponding Conflict free Instruction rate per CE	160.5 kops

MEMORY CONFLICT MODEL RESULTS

	<u>Three CEs</u>		<u>Two IOCEs*</u>		<u>System Totals</u>	
	<u>kops</u>	<u>kmrps</u>	<u>kops</u>	<u>kmrps</u>	<u>kops</u>	<u>kmrps</u>
No interleave	94	172	100	100	194	272
2 Way interleave	201	365	100	100	301	465
4-way inter.	316	569	100	100	416	669

* Assumes IOCE's priority 1 and 2 with CE's 3, 4, and 5.

Also assumes IOCE has local memory thus loading shared memory at a constant level at full I/O load capability.

TABLE B-4: EXTRAPOLATED VALUES FROM SIMULATION
WITH VARIOUS MEMORY SPEEDUP OPTIONS

3x2 9020A with Memory Speedup						
Memory Speed	2.5	2.0	1.5	1.0	0.8	0.5
The Max. inst/op time	5.0	4.0	3.0	2.0	1.60	1.0
The Max. inst/op rate	200 kops	250	333	500	625	1000
AIET	6.23	5.34	4.59	3.97	3.77	3.46
Conflict free inst. rate	160.5	187.3	219.3	251.9	265.3	289.0
System Totals	<u>kops</u>					
No intl.	194	234	260	289	305	329
2-way	301	366	406	457	482	521
4-way	416	509	563	631	671	728
	<u>kmrps</u>					
No int.	272	331	366	411	433	469
2-way	465	572	629	671	752	817
4-way	669	827	807	1034	1088	1184

TABLE B-5: OVERALL PERFORMANCE IMPROVEMENT RATIOS

Memory Speed in Micro Seconds						
	2.5	2.0	1.5	1.0	0.8	0.5
No intl.	1.00	1.21	1.34	1.49	1.57	1.69
2-way	1.55	1.89	2.09	2.35	2.48	2.69
4-way	2.14	2.62	2.90	3.25	3.46	3.75

critical resource in the system, why is the performance improvement not greater than shown in this table. The model shows primarily improvement due to conflict reduction which shows diminishing return with further speedup options not only because memory speedup can only reduce processor wait time to the basic memory independent rate of the processor. The 9020A CE has many instructions that run much longer than the 5.0 μ sec turn around time (i.e., instruction plus operand fetch times) of its current memory. Thus Table B-5 encourages the conclusion that state of the art memory would improve the performance of the 9020A 3x2 multiprocessor by 57 percent. This result is not all good news, however since the memory banking issue has not yet been considered. State of the art memory technology is not only faster but encourages larger memory modules. This if replacing the current nine banks 2.5 μ sec 9020A memory one would probably use only 4 much larger banks of 0.8 μ sec memory.

Table B-6 shows the simulation results of varying number of memory banks and degree of interleave in a 3x2 multiprocessor capable of five simultaneous memory requests through an eight port memory switch. This table indicates that reduction of eight to four banks of memory with either no interleave or two way interleave results in a reduced performance level 1.87/1.54 or 1.89/1.54 or about 82 percent. This reduction applied to the 1.57 times improvement of memory speedings reduces the potential gain to $1.57 \times .82 = 1.29$ or 29 percent over the current state, however there is some gain due to the larger memory size. A comparison of Tables C-4 and C-5 in Appendix C shows an overage improvement of six percent for elimination of program overlays by memory size sufficient to store all of the program. This improvement due to the combination of fewer but faster memory banks above for combined total of 35 percent.

If the current memory rate is 465 kmprs (thousands of memory references per second) as discussed above then the system performance improvement due to four banks of 0.8 μ sec memory is shown in Table B-7. The current memory loads for 111,222 and 333 tracks is taken from Table C-4 by converting from kmprh (thousands of memory references per hour) to kmprs. The current 9020A 3x2 system shows 78 percent memory saturation on the table. Option A is installation of four banks of 0.8 μ second memory, which the table indicates will handle the 222 track case but certain low priority functions will have

TABLE B-6: PERFORMANCE RATIOS OF A 3x2 MULTIPROCESSOR
WITH AN EIGHT PORT MEMORY BUS

DEGREE OF INTERLEAVING	NUMBER OF MEMORY MODULES					
	1	2	4	6	8(9)	10
none	1.00	1.20	1.54	1.75	1.87	1.93
2-way	-----	1.28	1.58	1.79	1.89	1.96
4-way	-----	-----	1.74	-----	1.95	-----

TABLE B-7: 9020A MEMORY PERFORMANCE SUMMARY

Load in No. of Tracks	Current Memory Load in KMRPS	Percent Memory Saturation		
		Current	Option A	Option B
111	361	78	58	38
222	633	136	100	67
333	949	204	151	100

Option A. careful memory mapping to reduce conflict with 0.8 micro second

Option B. two way interleaving of four 0.8 (micro) sec memory banks.

to be suspended to allow processing 333 tracks. The Option B column shows the additional performance gain due to interleaving by two ways the four banks of large memory. In this case 333 track can be processed without suspending any secondary functions.

The performance gain of interleaved memory in a real-time computer system must be traded off against reduced availability. For example, without interleaving the four banks, if one fails the system can reduce to a casualty mode based on reconfiguring into the remaining three banks. With two-way interleaving loss of one bank means loss of two (i.e. the faulted one plus its interleaved partner) and casualty mode becomes problematic with only half the memory. With four-way interleave of only four banks a single memory fault reduces to complete system outage and casualty mode, if any, must invoke another facility or backup means.

Table B-5 relates memory speedup possibilities with memory interleave alternatives and Table B-6 relates the latter to number of memory banks. Two other factors that are not analyzed quantitatively but are none the less important are memory size and the application of cache technology to the 9020A. Larger main memory can be employed in the system to advantage first by eliminating the need for overlays to gain a 6 percent advantage independent of other means. Beyond this advantage is the possibility of having sufficient main memory that critical programs shared by numerous processes could be replicated in each memory bank as required to further reduce conflict. This improvement possibility is not completely independent of other conflict reduction techniques. In applying memory size advantage it is best to increase the number of memory modules rather than merely to increase the size of each module only. Table B-8 relates the performance improvement due to the combined size per module and number of module factors. The improvement shown in this table is due to two factors, one enabled by memory size and one by conflict reduction as the number of independent module increases to (and slightly beyond) the number of simultaneous memory requests. The two latter factors are, first, reduction of memory demand if overlays are not required, and, second, reduction in memory conflict if routines that may be called simultaneously by different processors can be shared in each memory bank. The overall improvement for

TABLE B-8: EFFECT OF MODULE NUMBER AND SIZE, NOT INCLUDING INTERLEAVING OR SPEEDUP, FOR A 3x2 MULTIPROCESSOR

Memory Size/Module	Number of Memory Modules						
	1	2	4	6	8	(9)	10
512k bytes (Current System)	--	--	--	--	0.9	1.0	1.08
1024k bytes	--	--	0.74	1.39	1.64	--	--
2048k bytes	--	0.6	1.23	1.39	1.64	--	--
4096k bytes	0.5	0.97	1.23	1.39	1.64	--	--

TABLE B-9: EFFECT OF INCLUSION OF A FOURTH CE

	<u>3x2 System Totals</u>		<u>4x2 System Totals</u>		Approx. Improvement (percent)
	kops	kmrps	kops	kmrps	
No interleave	194	272	243	341	25
Interleaved	301	465	379	589	26
4-way intl.	416	669	528	850	27

four banks of 4096 k memory the same speed as is in current use would be about 23 percent.

A configurational alternative would be to apply the redundant fourth CE to the workload. As Table B-9 shows, the use of the redundant machine is about the same as employing redundant memory. In general this approach will not be fruitful, as K. J. Thurber points out in Large Scale Computer Architecture, pp. 307-311. If the total number of CE's and IOCE's in a multiprocessor system drive more addresses simultaneously than the number of memory banks, then performance is degraded. In the case of the 9020A one could drive up to eight addresses in parallel before this conflict situation would cause serious performance loss. In this section of his book Thurber also shows how a secondary memory multiprocessor experiences less performance loss due to memory conflicts than a primary memory multiprocessor like the 9020A. Isolation of the shared memory resource could be provided in the 9020A by providing each CE with a small buffer memory or cache. This approach could produce a potential gain of 58 percent; however, this value must be reduced by the hit rate of the cache. If the cache is very small, for example only a few words, then the hit rate will be about 80 percent (assuming that every fifth instruction is a jump or change in sequence). If the cache is 4096 bytes or larger, then the system could attain a hit rate as high as 94 percent. In the first case the improvement could be as large as 6 percent and in the second case no larger than 54 percent.

APPENDIX C. NAS REPRESENTATIVE 9020A WORKLOADS

This appendix describes representative workloads and their derivations for the existing National Airspace System (NAS) 9020A Computer Complex. A "simplified" configuration diagram is shown in Figure 1-1 and consists of nine 1/4 mb memories, three 360/50 compute elements, two 360/50 IOCE's, two 2314 disk units, and two 2401 tape units [NIEL77]. Only this primary or nonredundant portion of Figure 1-1 was considered in deriving the representative workloads. Actual measurements as reported by several organizations, theoretical calculations, and program descriptions and specifications as reported in the documents listed in the references were used in preparing these workloads. The workloads, derived for three cases - 111, 222, and 333 tracks -- are termed "representative" because there has not been a complete set of measurements made for any one version of the NAS Program. Versions NAS A3d2.1, 2.2, 2.3, 2.4, 2.7, and 2.9 were all used to gather the necessary statistics that in turn were used to derive the workloads. Because the purpose of constructing a workload is to drive the model to examine 9020A memory interference, this representative workload appears to offer a fairly accurate picture of NAS Program activity. Information derived from the Workload Tables C-1 through C-6 compared favorably with material in the references that were not previously used in the workload derivations.

Several groups have measured NAS activity either in actual operation or at the FAA Technical Center and have found that approximately 26 program elements (PE) account for approximately 90% of the processor activity [KELL77, NOPAR77]. These PE's and their size are shown in Table C-1. Also given is whether they are permanently resident in memory or they are dynamically loaded when needed [NOPAR77].

Three traffic load cases -- 111, 222, and 333 tracks -- were used in deriving the workload. 111 tracks, for which measurements using various monitoring devices have been made [NOPAR77], is representative of the typical non-saturated case. Table C-1 lists the measured number of activations per hour per selected PE and the associated percentage of one computer element utilization. There were not any counts for four of the

TABLE C-1: PE UTILIZATION AT 111 TRACKS

#	PE	Size (Bytes)	Buffer- able	Activations Per Hour	% CE Utilization	Memory Loading (10 ⁶ Refs/Hour)
1	FTM	10,376	Y	601	2.25	22.0
2	COP	6,665	N	715	1.19	11.7
3	CRU	11,328	N	391	1.27	12.4
4	CSF	18,160	N	560	3.14	30.8
5	DAM	22,728	N	279	-----	-----
6	RAT	7,424	Y	600*	2.21	21.7
7	DUZ	18,136	N	337	4.25	41.6
8	MOR	528	N	23,656	5.91	57.9
9	RDA	17,216	N	3,605	10.18	99.8
10	PDE	3,544	N	1,495	1.12	11.0
11	JQN	15,800	Y	2,066	1.38	13.5
12	HTI	30,848	N	2,012	20.23	198.3
13	HHM	3,880	N	3,602	2.49	24.4
14	RSL	464	N	3,850	-----	-----
15	RTG	11,904	N	3,606	11.19	109.7
16	MRM	11,264	Y	600*	5.80	56.8
17	RCD	28,112	Y	908	3.46	33.9
18	CNN	23,816	Y	848	-----	-----
19	CSS	584	N	523	-----	-----
20	PNA	1,824	N	1,295	1.17	11.5
21	JTU	2,056	Y	600*	-----	-----
22	CRJ	11,280	N	1,115	-----	-----
23	CBC	18,832	N	1,215	1.02	10.0
24	RRA	9,488	Y	300*	3.32	32.5
25	RFA	26,232	Y	301	1.09	10.7
26	FWR	3,200	N	601	1.29	12.6
TOTAL		315,788		55,681	84.0	822.8

*estimated

PE's in the original measurements but they were easily estimated because of their periodicity. Table C-1 shows that these 26 PE's consumed 84% of the resources of one processor. The final set of numbers in Table C-1 is the number of memory references per hour for each of the selected PE's. The total load from just these PE's is 822.8 million memory references per hour. A 2.5 microsecond memory or storage element (SE) as used in the 9020A configuration, has a bandwidth of 1440 million memory references per hour.

The memory reference figures of Table C-1 were derived from instruction times listed in Kelley's report[KELL77]. Kelley found that the average instruction execution time for the NAS Program was 6.23 microseconds. From Kelley's instruction times and counts charts, it was determined that 30.4% of the executed instructions involve one memory reference and 69.6% involved two memory references (instruction and operand fetch). Therefore, there are 1.696 memory references per instruction. The number of memory references per hour for a PE is then found from the equation:

$$\frac{\text{PE reference}}{\text{hour}} = \frac{3600 \times 10^6 / \text{sec}}{\text{hour}} \times 1 / (6.23 \times 12.696 \times \% \text{ CE utilization}).$$

The 222 track case is a saturated system case. Although actual measurements had been made for this case [NOPAR77], it was noted that some of the numbers were suspect because the system was saturated. 222 tracks are handled in actual operation today by removing some of the operational PE's as the system approaches saturation[SENA80]. The figure for the total number of activations per hour and the % CE utilization were derived from some of the actual measurements[NOPAR77] and by estimating the PE's operation[PDSI78, PDSII79]. Some 9020D measurements for 222 and 444 tracks [NOPAR77] were used as guidelines in determining ratios between PE activity at various track sizes. The memory load from these PE's would saturate the memory if all could operate as in the case with 111 tracks.

The 333 track case was selected because it is a load that is well into system memory saturation that could possibly be moved to the non-saturated region by increasing the memory speed or size, or interleaving references, or using a cache. The total activations and CE utilizations for the 333 track case were extrapolated from the previous sets of numbers to obtain the figures listed in Table C-2.

TABLE C-2: PE UTILIZATION AT 222 AND 333 TRACKS

	Total Total Activations Per Hour	% CE Utilization	Memory Loading (10 ⁶ Refs/ Hour)	Total Total Activations Per Hour	% CE Utilization	Memory Loading (10 ⁶ Refs Hour)
#PE	<u>222 Tracks</u>	<u>222 Tracks</u>	<u>222 Tracks</u>	<u>333 Tracks</u>	<u>333 Tracks</u>	<u>333 Tracks</u>
1 FTM	599	2.25	22.0	600	2.25	22.0
2 COP	1,400	2.4	23.5	2,100	3.6	35.3
3 CRU	780	2.6	25.5	1,170	3.9	38.2
4 CSF	1,100	6.0	58.8	1,650	9.0	88.2
5 DAM	560	1.4	13.7	840	2.1	20.6
6 RAT	600	4.4	43.1	600	6.6	64.7
7 DUZ	660	8.5	83.3	990	12.8	125.4
8 MOR	40,000	11.0	107.8	60,000	17.0	166.6
9 RDA	3,550	15.0	147.0	3,600	22.0	215.6
10 PDE	3,000	2.2	21.6	4,500	3.3	32.3
11 JQN	2,000	1.4	13.7	2,100	1.5	14.7
12 HTI	4,000	30.0	294.0	6,000	40.0	392.0
13 HHM	3,508	3.5	34.3	3,600	4.5	44.1
14 RSL	8,000	2.0	19.6	1,200	3.0	29.4
15 RTG	3,551	22.0	215.6	3,600	33.0	323.4
16 MRM	600	11.6	113.7	600	17.4	170.5
17 RCD	1,200	6.0	58.8	1,500	9.0	88.2
18 CNN	1,800	1.4	13.7	2,700	2.1	20.6
19 CSS	1,000	1.2	11.8	1,500	1.8	17.7
20 PNA	1,800	1.4	13.7	2,100	1.6	15.7
21 JTU	900	1.0	9.8	1,200	1.2	11.8
22 CRJ	2,200	1.4	13.7	3,300	2.1	20.6
23 CBC	1,500	1.3	12.7	1,800	1.6	15.7
24 RRA	300	3.0	29.4	300	3.0	29.4
25 RFA	298	1.2	11.8	300	1.3	12.7
26 FWR	601	2.6	25.5	600	3.9	38.2
TOTAL	85,507	148.8	1,438.1	119,250	209.6	2,053.6

The total load on memory is not due only to PE activity but also to Operating System (OS), I/O (disks and tapes), and IOCE activity. The dynamic buffering of PE's affects the OS and I/O activity. Therefore, Table C-3 was prepared to determine how many memory references or words per hour were used in loading these PE's into core memory (SE's) from disk storage.

The Operating System or Monitor loading was derived from measurements made at the 9020A Memphis ARTCC site [NIEL77]. The following items and % CE utilization compare the OS loading:

Dispatcher - 4%	(actual PE dispatching)
SVC - 2%	
I/O interrupt processor - 2%	
Load module relocate subroutine - 2.6%	
TAR generation - 10%	
Pool management subroutines - 3.7%	
Other monitor services - 6%	

Total	30.3%
-------	-------

Using the same equation as for PE loading, the OS loading was determined and is listed in Table C-4. For larger memories, therefore, eliminating the need for buffering, 6.3% of the OS load (Load module relocate subroutine and Pool management subroutines) can be removed. The OS load without buffering is shown in Table C-5.

The I/O load on main memory was assumed due to the transfer of disk and tape information. Table C-6 lists the peripheral parameters used and calculates the number of memory references per hour based on utilization rates found by LOGICON [NIEL77]. Table C-4 lists the memory loading for the I/O for the three cases. Because eliminating dynamic buffering eliminates the need to transfer the buffered PE's from disk, the I/O loads for the non-buffered cases were determined by reducing the I/O loads in Table C-4 by the totals in Table C-3 and are shown in Table C-5.

TABLE C-3: BUFFERABLE PE I/O LOADING

# PE	Size (Words)	Total Activations Per Hour	Memory Loading 10 ⁶ Refs/Hr	Total Activations Per Hour	Memory Loading 10 ⁶ Ref/Hr	Total Activations Per Hour	Memory Loading 10 ⁶ Refs/Hr
		111 Tracks		222 Tracks		333 Tracks	
1 FTM	2,594	601	1.6	599	1.6	600	1.6
6 RAT	1,856	600	1.1	600	1.1	600	1.1
11 JQN	3,950	2,066	8.2	2,000	7.9	2,100	8.3
16 MRM	2,816	600	1.7	600	1.7	600	1.7
17 RCD	2,028	908	6.4	1,200	8.4	1,500	10.5
18 CNN	5,954	848	5.0	1,800	10.7	2,700	16.1
21 JTU	514	600	0.3	900	0.5	1,200	0.6
24 RRA	2,372	300	0.7	300	0.7	300	0.7
25 RFA	6,558	301	2.0	298	2.0	300	2.0
Totals	33,642		27.0		34.6		42.6

TABLE C-4: MEMORY LOADING - DYNAMIC BUFFERING

Component	111 Tracks	222 Tracks	333 Tracks
	10 ⁶ memory Refs/Hr	10 ⁶ Memory Refs/Hr	10 ⁶ Memory Refs/Hr
PE	822.8	1438.1	2053.6
OS	296.9	481.0	822.5
I/O	171.3	342.6	513.9
IOCE	9.8	18.6	29.4
TOTAL	1300.8	2280.3	3419.4

The final memory loading component is due to PE's and other software executing in the IOCE. RIN is the most active PE in the IOCE but as with other programs in the IOCE, it executes out of IOCE local memory. The only additional load on main memory due to the IOCE, then, is the transfer of information to main memory tables as the result of IOCE PE activity. Because no measurement of this type of reference could be found, an estimate was made based on table size, information transfer, and frequency of activation[PDSI78, PDSII79]. The result is listed in Tables C-4 and C-5 and is the same with or without buffering.

The total memory loading is then calculated by summing the loadings for the four components -- PE, OS, I/O, and IOCE. Tables C-4 and C-5 list these totals for 111, 222, and 333 tracks, both buffered and with no buffering. These totals then served as input to the simulation model to investigate memory interference problems and possible performance improvement approaches.

In order to check some of the assumptions made for bufferable PE activity, a memory map (Table C-7) was constructed. Using sizing figures for NAS 3d2.4 [NOPAR77], resident PE's were optionally placed such that subsequent PE's in the processing flow chain do not reside in the same memory module. This chart illustrates that PE's can be placed in memory such that interference from processing simultaneous tracks is kept to a minimum and the buffering of non-resident PE's can be uniform throughout the nine SE's.

Thus, a representative workload for three different cases of air traffic activity was developed. Based on actual measurements, simulation data, specifications, and extrapolations, the workload figures reflect a reasonable driving function for the 9020A EnRoute System Configuration Model.

TABLE C-5: MEMORY LOADING - NO BUFFERING

<u>Component</u>	111 Tracks	222 Tracks	333 Tracks
	<u>10⁶ memory Refs/Hr</u>	<u>10⁶ Memory Refs/Hr</u>	<u>10⁶ Memory Refs/Hr</u>
PE	822.8	1438.1	2053.6
OS	235.2	431.9	579.4
I/O	144.3	308.0	471.3
IOCE	9.8	18.6	29.4
TOTAL	1212.1	2196.6	3233.7

TABLE C-6: I/O LOADING

<u>Unit</u>	<u>Transfer Rate (Kb/sec)</u>	<u>% Utilization</u>	<u>10⁶ Memory Refs Per Hour</u>
2314 Disk ₁	312	21	59.0
2314 Disk ₂	312	32	89.9
2401-II Tape	60	16	8.6
2401-III Tape	90	17	13.8
Total			171.3

TABLE C-7: MEMORY MAP

<u>SE #</u>	<u>Resident PE's</u>	<u>Total PE Resident KBytes</u>	<u>Remaining Resident KBytes</u>	<u>Total Resident KBytes</u>	<u>Total Resident Tables, & Misc. KBytes</u>	<u>Total Dynamic Buffer Area KBytes</u>	<u>Total Memory Size KBytes</u>
1	CRU PDE	16	62	78	111	73	262
2	CSF PNA	31	47	78	111	73	262
3	DUZ FWR	23	55	78	111	73	262
4	HTI	31	47	78	111	73	262
5	CSS MOR	2	76	78	111	73	262
6	RDA RSL	19	59	78	111	73	262
7	RTG HHM	16	62	78	111	73	262
8	DAM CRJ	35	43	78	111	73	262
9	COP CBC	26	52	78	111	73	262
	Total	198	504	702	996	657	2358

APPENDIX D. THE STRATEGIES OPEN TO THE FAA

Chapters 2 and 3 discussed the six enhancements that the FAA might adopt, and Chapter 4 discussed how the enhancements can be combined into strategies for upgrading the 9020's. Chapter 4 only explained the strategies that now seem to be most attractive; as conditions change or as the appreciation of the problem deepens, however, it might be that other strategies gain in appeal. Therefore, this appendix exhibits all the strategies that can be constructed from the six enhancements and explains how the strategies highlighted in the decision tree in Chapter 4 were chosen.

The following five constraints must be observed in forming strategies from the six enhancements.

1. Replacing the memory boxes and replacing the SE memory stacks are not both adopted.
2. Speeding up the CE's and replacing the CE's are not both adopted.
3. Speeding up the CE's can only be done if either the SE memory boxes or the SE memory stacks are replaced.
4. Replacing the CE's can only be done if either the memory boxes or the SE memory stacks are replaced.
5. Speeding up the IOCE processors can only be done if the IOCE memory stacks are replaced.

Any combination of the six enhancements that does not violate one of these constraints is considered to be a strategy. There are 20 possible strategies, and these are shown in Table D-1. Each row of this table represents one strategy; the X's in a row show which enhancements constitute the strategy. For example, strategy 17 consists of replacing the memory boxes, replacing the IOCE memory stacks, speeding up the CE's, and speeding up the IOCE's.

TABLE D-1: THE POSSIBLE STRATEGIES OPEN TO THE FAA

<u>Enhancements</u>						
<u>Strategy Number</u>	<u>Replace SE Memory Boxes</u>	<u>Replace SE Memory Stacks</u>	<u>Replace IOCE Memory Stacks</u>	<u>Speed Up CE's</u>	<u>Speed Up IOCE's</u>	<u>Replace CE's</u>
1	X					
2		X				
3			X			
4	X		X			
5		X	X			
6	X			X		
7	X					X
8		X		X		
9		X				X
10			X		X	
11	X		X	X		
12	X		X		X	
13	X		X			X
14		X	X	X		
15		X	X		X	
16		X	X			X
17	X		X	X	X	
18	X		X		X	X
19		X	X	X	X	
20		X	X		X	X

The relationship between the strategies in Table D-1 and the strategies in the decision tree in Figure 4-1 is as follows.

The first path through the tree corresponds to strategies 1 and 2; this one path corresponds to two strategies since the decision tree does not distinguish between replacing the memory boxes and replacing the stacks. The second path through the tree corresponds to strategies 6 and 8. The third path corresponds to strategies 17 and 19. The fourth path also corresponds to strategies 17 and 19; the difference between these two paths lies in the timing of the decisions and in whether the IOCE's are upgraded in just the 9020A's or also in the 9020D's. The fifth path corresponds to strategy 10.

It now must be explained why strategies 3, 4, 5, 7, 9, 11, 12, 13, 14, 15, 16, 18, and 20 were omitted from the tree. Strategy 3 was omitted since replacing the IOCE memory stacks and doing nothing else probably will not deal with the 9020's short-run problems. Strategies 4 and 5 were omitted since once the SE memory is replaced, the additional memory gained by replacing the IOCE memory stacks and doing nothing else does not yield much of an advantage. Strategies 11 and 14 were omitted since once the SE memory is replaced and the CE is sped up, the additional memory gained by replacing the IOCE memory stacks and doing nothing else apparently offers no significant advantage. Strategies 12 and 15 were omitted since once the IOCE is upgraded, replacing the SE memory, though it would increase the available memory, would probably not yield much more performance. Strategies 7, 9, 13, 16, 18, and 20 were omitted since, as Sec. 4.2 explains, the enhancement of replacing the CE's is tentatively assumed to be undesirable since it is both more expensive and more time-consuming than speeding up the CE's. It should be emphasized that these 13 omitted strategies are omitted because, given our current understanding of the problem, they appear to be relatively unattractive and because of the desire to keep Figure 4-1 as simple as possible.

In summary, this appendix has exhibited all 20 of the strategies that can be constructed from the 6 enhancements and has explained why the strategies appearing in the decision tree in Figure 4-1 were selected as the

leading strategies. It is quite possible, however, that the relative attractiveness of these strategies will change over time as the situation evolves, so this should by no means be taken as a definitive demonstration of the undesirability of these 13 strategies.

