REPORT NO. FAA-RD-75-71

ARTS III/PARALLEL PROCESSOR DESIGN STUDY

Vivian J. Hobbs

APRIL 1975

FINAL REPORT

APPROVED FOR FEDERAL AVIATION ADMINISTRATION ONLY. THIS DOCUMENT IS EXEMPTED FROM PUBLIC AVAILABILITY BECAUSE OF THE SENSITIVITY OF COMPARATIVE PERFORMANCE AND COST DATA FROM THREE SPECIFIC DATA PROCESSING SYSTEMS MANUFACTURERS. TRANSMITTAL OF THIS DOCUMENT OUTSIDE THE FEDERAL AVIATION ADMINISTRATION, DEPARTMENT OF TRANSPORTATION MUST HAVE PRIOR APPROVAL OF THE SRDS/FAA/ARD-100.

Prepared for

U.S. DEPARTMENT OF TRANSPORTATION
FEDERAL AVIATION ADMINISTRATION
Systems Research and Development Service
Washington DC 20591

Technical Report Documentation Page

1. Report No.
FAA-RD-75-71

2. Government Accession No.

3. Recipient's Catalog No.

4. Title and Subtitle
ARTS III/PARALLEL PROCESSOR DESIGN STUDY

5. Report Date
April 1975

6. Performing Organization Code

7. Author(s)
Vivian J. Hobbs

8. Performing Organization Report No.
DOT-TSC-FAA-75-2

9. Performing Organization Name and Address
U.S. Department of Transportation
Transportation Systems Center
Kendall Square
Cambridge MA 02142

10. Work Unit No. (TRAIS)
FA 503/R5143

11. Contract or Grant No.

12. Sponsoring Agency Name and Address
U.S. Department of Transportation
Federal Aviation Administration
Systems Research and Development Service
Washington DC 20591

13. Type of Report and Period Covered
Final Report
July 1973-June 1974

14. Sponsoring Agency Code

15. Supplementary Notes

16. Abstract

It was the purpose of this design study to investigate the feasibility, suitability, and cost-effectiveness of augmenting the ARTS III failsafe/failsoft multiprocessor system with a form of parallel processor to accommodate a large growth in air traffic and an increase in automated air traffic control functions.

The major results of the design study are summarized and discussed. Included are: brief descriptions of the proposed parallel processor hardware, approach to interfacing with the ARTS III equipment and software, approaches to failsafe redundancy, partitioning of ATC functions among serial/parallel elements, and provision for support hardware and software. Also presented are: comparative hardware costs, a terminal area traffic projection, and a loading analysis of the ARTS III multiprocessor system.

17. Key Words
Parallel Processor, ARTS III Multiprocessor, Air Traffic Projection, ARTS III Loading, Parallel Processor Cost Estimates

18. Distribution Statement
APPROVED FOR FEDERAL AVIATION ADMINISTRATION ONLY. THIS DOCUMENT IS EXEMPTED FROM PUBLIC AVAILABILITY BECAUSE OF THE SENSITIVITY OF COMPARATIVE PERFORMANCE AND COST DATA FROM THREE SPECIFIC DATA PROCESSING SYSTEMS MANUFACTURERS. TRANSMITTAL OF THIS DOCUMENT OUTSIDE THE FEDERAL AVIATION ADMINISTRATION, DEPARTMENT OF TRANSPORTATION MUST HAVE PRIOR APPROVAL OF THE SRDS/FAA/ARD-100.

19. Security Classif. (of this report)
Unclassified

20. Security Classif. (of this page)
Unclassified

21. No. of Pages
66

22. Price

Form DOT F 1700.7 (8-72)

Reproduction of completed page authorized

PREFACE

A competitive study was conducted to develop a detailed system design for augmenting the ARTS III multiprocessor computer configuration with a parallel processor. Three contractors participated. This report briefly summarizes and comments on each contractor's results in the major areas of technical concern.

The rationale for initiating the study, a terminal area air traffic forecast, and a loading analysis of the ARTS III multiprocessor configuration are also presented.

The original program plan anticipated that a prototype parallel processor would be fabricated, interfaced with an ARTS III multiprocessor system, and evaluated in a test-bed environment. However, the program was cancelled by FAA after the design study phase due to lack of available funding and questionable need.

CONTENTS

Section — Page

1. INTRODUCTION.. 1-1
   - 1.1 Scope of Present Report......................... 1-1
   - 1.2 Background...................................... 1-2
   - 1.3 Approach.. 1-2

2. GENERAL CONCLUSIONS................................. 2-1

3. ARTS III LOADING CONSIDERATIONS..................... 3-1
   - 3.1 Air Traffic Projection.......................... 3-9

4. PARALLEL PROCESSOR COST ESTIMATES................... 4-1

5. DESIGN STUDY RESULTS................................ 5-1
   - 5.1 Parallel Processor Hardware, Interfaces, Failsafe Redundancy, and Their Programming Impact......... 5-1
     - 5.1.1 Goodyear.................................... 5-1
     - 5.1.2 Honeywell................................... 5-2
     - 5.1.3 Texas Instruments........................... 5-3
     - 5.1.4 Comparative Assessment of Parallel Processor Hardware Interfaces, Failsafe Redundancy and Programming Impact......... 5-5
   - 5.2 ATC Functional Software and Executive Interfaces. 5-7
     - 5.2.1 Goodyear.................................... 5-7
     - 5.2.2 Honeywell................................... 5-9
     - 5.2.3 Texas Instruments........................... 5-11
     - 5.2.4 Comparative Assessment - ATC Functions and Executive Interfaces..................... 5-12

6. BENCHMARK EXERCISE.................................. 6-1

7. TEST-BED AND PRODUCTION SYSTEM SUPPORT FACILITIES AND MEASUREMENT PLAN.. 7-1
   - 7.1 Goodyear.. 7-1
   - 7.2 Honeywell....................................... 7-3
   - 7.3 Texas Instruments............................... 7-4

8. SUMMARY... 8-1

CONTENTS (CONTINUED)

Section Page

APPENDIX A.1 - BENCHMARK PROBLEM (INTRODUCTION)............... A-1

A.2 - DEFINITION OF TEXT............................ A-1

A.3 - DEFINITION OF CONFLICT........................ A-2

A.4 BEGINNING TIME OF CONFLICT.................... A-2

A.5 - MINIMUM MISS DISTANCES........................ A-3

A.6 - IMPLEMENTATION OF CONFLICT DETECTION FUNCTION....................................... A-3

A.7 - OUTPUT FORMAT................................. A-5

LIST OF TABLES

Table		Page
Ia	ESTIMATED ARTS III LOAD	3-7
Ib	ARTS III ESTIMATED FAILSAFE SYSTEM CAPACITY LIMITS	3-8
II	AIR TRAFFIC PROJECTION	3-10
III	AIRCRAFT COUNTS	3-12
IV	COST COMPARISON PRODUCTION SYSTEM LOTS (OLD ARTS EQUIPMENT PRICES)	4-3
V	COST COMPARISON PRODUCTION SYSTEM LOTS (NEW ARTS EQUIPMENT PRICES)	4-4
VI	ARTS III EQUIPMENT COST DATA	4-5
VII	SINGLE FAILSAFE SYSTEM COST COMPARISON	4-6
VIII	FAILSAFE SINGLE SYSTEM HARDWARE CONFIGURATION AND COST BREAKDOWN	4-7
IX	PHYSICAL SIZE COMPARISON	5-6
X	ARTS III/PP ATC FUNCTIONAL DISTRIBUTION	5-8
XI	BENCHMARK RESULTS	6-2

1. INTRODUCTION

1.1 SCOPE OF PRESENT REPORT

A general discussion of the results of the ARTS III/ Parallel Processor Design Study is presented. Three contractors were chosen for this study, on the basis of their proposals, to independently develop a detailed system design for augmenting the ARTS III multiprocessor system with a parallel processor. The system design encompassed:

1) Selection of, and justification for, a distribution of terminal ATC functions between the essentially serial ARTS/UNIVAC equipment and the proposed parallel processor
2) Design of a parallel processor architecture based on existing equipment optimized to be cost-effective for the terminal ATC environment
3) Design of reasonable hardware/software interfaces between the ARTS III failsafe/failsoft computer configurations and the parallel processor
4) Provision for real-time failure diagnosis and rapid recovery procedures
5) Support facilities for software development, hardware maintenance and computer system performance measurement.

Each contractor's detailed design documentation fills several large volumes. A final report from each contractor gives an overview of their individual results. This paper will briefly discuss the major points of each approach and highlight some of the strengths and weaknesses. The purposes being (1) to summarize and compare the collective results, (2) to present the rationale for the study, and (3) to document the cost and loading analysis conclusion which evolved.

1.2 BACKGROUND

It was the purpose of this Design Study to investigate the feasibility, suitability, and cost-effectiveness of augmenting the ARTS III failsafe/failsoft multiprocessor systems with a form of parallel processing.

The basic system requirements were 1) to increase the capacity of the terminal automation equipment, via augmentation, without jeopardizing its performance, reliability, and automatic failure diagnosis and failsafe recovery capability; 2) to retain in the overall hardware configuration, at least a minimum failsafe set of ARTS III IOP's and memory modules, and 3) to be cost-effective.

The capacity increase was deemed necessary due to an alarming growth in terminal air traffic, a desire for automation of additional terminal ATC functions, and a lack of data on the saturation thresholds of the enhanced ARTS III configurations, which were still in the development stage at the time that this study began.

The retention of ARTS equipment was deemed necessary to minimize new system development costs, to protect to-date investments in equipment and algorithm development, and to discourage grandiloquent total replacement systems. The thrust of the study was augmentation to extend the life of the third generation ARTS equipment, not replacement with the fourth generation of terminal ATC automation equipment.

The desire for cost-effectiveness is self-explanatory and is reinforced by the possibility of installing multiple systems at various ATC terminal sites.

1.3 APPROACH

The approach taken to this investigation was to fund a competitive design study among several qualified contractors. The intention being to use the detailed results of the study in arriving at a decision as to whether an ARTS/parallel processor test-bed development and integrated system evaluation was warranted. The contractor for the test-bed phase would have been chosen from among those participating in the design study.

The necessity of a test-bed evaluation and the possibility of a subsequent production system procurement imposed additional important requirements on the participating contractors which were to be addressed in their overall design effort. These requirements dealt with computer system reliability and support facilities for program development and maintenance, hardware diagnostics and maintenance, and a real-time computer system performance measurement and evaluation. Such facilities were to be defined and provided by the contractor at the FAA test-bed site.

Participation in the design study was limited to contractors who had been active in the previous development of computer hardware of the same general architecture as that being proposed. It was also required that a government-specified benchmark program be coded and executed by the participating contractors on such equipment during the design study period. These constraints were necessary in order to assure a reasonably short-term development and test schedule. Three contracts were awarded: Goodyear Aerospace Inc., Akron, Ohio; Honeywell Inc., Minneapolis, Minn. (subcontractors: ARCON Corp. & General Research Corp.) and Texas Instruments Inc., Austin, Texas.

The design study took place over a six-month period between June and December of 1973. During this time, there were three formal design reviews held separately with each contractor, for a total of nine design reviews. During these reviews, each contractor's progress was assessed, weak areas were identified for further emphasis, and FAA and ARTS III system requirements and restrictions were clarified. The contractors' efforts were focused through these reviews without benefit of specific advice as to how to solve the numerous "decision" problems inherent in the design. This was necessary in order to foster the contractors' individuality of approach and to retain government objectivity and impartiality due to the competitive nature of the study.

This approach has proven worthwhile in that the final designs were significantly better than those presented in proposal form. The originally proposed architectures were all modified during the design study to be more cost-effective and to better fit the problem. Proposed solutions to many problems were better thought-out, and reflected an increased understanding of the overall requirements and the terminal ATC functional environment on the part of the contractors. Cost-estimates were more realistic than originally proposed and reliability estimates were based on a completed design rather than a list of possibilities.

2. GENERAL CONCLUSIONS

The following general conclusions were drawn as a result of this study.

Parallel processing as an augmentation to the ARTS III multiprocessor system:

1) Is feasible (a) within the hardware and software framework of the ARTS III system, and (b) within the current state-of-the-art of parallel processors.
2) Is a reasonable approach in view of the parallel nature of a number of ATC functions.
3) Can provide a significant capacity increase for both air traffic loads and additional automated functions in a modular fashion adjustable to individual sites.
4) Would extend the useful life of the enhanced ARTS III systems, since an ARTS III/parallel processor combination permits modular expansion in both the ARTS III multiprocessor equipment and in the parallel processor equipment.
5) Can be cost-effective (in production quantity lots)for track loads of 450 and up (250 aircraft-dual sensor) and potentially cost-saving for track loads of 875 and up (500 aircraft-dual sensor).
6) Can accommodate a failsafe/failsoft hardware redundancy and recovery philosophy consistent with that envisioned for the enhanced ARTS systems

The ultimate question, however, became one of "how real is the need to augment the ARTS III system?" At the onset of this study, there were no figures available with regard to the capacity threshold of the ARTS III multiprocessor configuration, nor was any reliable data provided as to the air traffic volume at various terminals. It was therefore necessary for TSC to obtain this type of information during the course of the six month design study period.

The subsequent discussion on ARTS III loading will indicate that the final decision on the desirability of proceeding with an ARTS III/ parallel processor test-bed integration and systems evaluation was not cut-and-dried. The projection of future air traffic growth is highly judgemental. The decision, therefore, was subject to considerations other than those of sheer technical merit or potential cost-effectiveness.

The decision was made, by the FAA, to not proceed with the test-bed primarily on the basis of minimal available R&D funds and questionable need. This decision does not reflect any suggestion on the part of the FAA or TSC that parallel processing is not a viable approach to ARTS III augmentation.

3. ARTS III LOADING CONSIDERATIONS

A prime consideration in assessing the urgency of requirements for increasing ARTS III system capacity by any means is the determination of the saturation thresholds of the fully expanded ARTS III multiprocessor configuration. In the absence of any concrete system performance measurement data on these configurations, loading equations developed by UNIVAC[1] from existing program code were used as a sizing estimate. These equations were augmented with data available from an ARTS III system simulation activity[2] and other system-level analyses.[3] Consideration was also given to the effect of altering the existing ARTS III multiprocessor software to increase the amount of parallel activity in the system. This latter effort did not result in a difference in the size of the ARTS III configurations required to support a given system load.

The computer processing requirement for ARTS III is primarily a function of the air traffic volume, the number and type of automated ATC functions in operation at a given site, the number (and % overlap) of sensor inputs to the system, and the number and types of display outputs from the system. For our loading analysis, we considered air traffic volumes of 250, 500, 750, and 1000 aircraft in both single and dual sensor configurations. We assumed a 75% sensor overlap (15 mile separation) in the dual sensor environments after looking at the actual overlaps in several of the dual sensor sites. Overlapped sensor coverage results in a significantly larger processing load, since each sensor is inputting to the computer all that it "sees." There is at least one record in the computer track file for each aircraft each sensor

---

1. "Augmented Radar Beacon Tracking Level System Design Specification," PX7981, March 1973, Final Report, Contract #DOT-FA70WA-2289, Page B13.
2. "Expanded ARTS III Site Configuration Analysis," PX10000, October 1971, Interim Report, Contract #DOT-FA70WA-2289.
3. "Augmented Radar Beacon Tracking Level Test and Evaluation Report," PX10218, October 1973, Final Report, Contract #DOT-FA70WA-2289.

reports. The processing load is, therefore, more directly a function of the track load than of the actual aircraft count. We assumed all digital displays numbering 15, 20, 25, and 30 for aircraft loads of 250, 500, 750, and 1000 respectively. The ATC functions assumed to be operational were as follows:

Radar Input Processing
Beacon Input Processing
All Digital Display Processing
Radar Augmented Tracking
Conflict Prediction
Beacon/Radar Correlation
Critical Data Recording and Collection
Failsafe/Failsoft Executive Processing
System Timeout Processing
Miscellaneous (consisting of Keyboard, Interfacility I/O,
General I/O requirements, and Automatic Offset)

The loading equation used for the above functions were as follows:

Definition of Terms

B_i = Beacon targets sensor i
R_i = Radar targets sensor i
P_i = Predetections sensor i/scan
T_i = Tracks system i
D = Total displays
T_{ui} = Uncontrolled tracks system i
T_{ci} = Controlled tracks system i
U_i = Unused reports system i
PIAC= Peak Instantaneous Aircraft Count
PAS = Peak Aircraft per sensor
P = Number of Processors (IOP's)

Values Assumed for Terms

i = 1, 2 (number of sensors)
PIAC = 250, 500, 750, 1000
PAS = 0.875 PIAC (75% sensor overlap)
P_i = PIAC + 500
U_i = PIAC*3.0/10.0
R_i = 1.2 PAS
B_i = 0.7 PAS
T_i = PAS
T_{ui} = $0.4T_i$
T_{ci} = $0.6T_i$
D = 15, 20, 25, 30 corresponding to PIAC of 250, 500 750, 1000

Equations (Express processing load as % of a single IOP)

NOTE: The use of Display Buffers is assumed.

Beacon Input Processing (BIP) = $0.06B_i = 6.6_i$

Radar Input Processing (RIP) = $0.00375P_i + 0.38R_i + 22.45_i$

All Digital Display Processing = $D\ [1.7068 + 0.00135T_i] + 0.0242T_{ci}$

Radar Augmented Tracking = $2.5 + 0.13T_i + 0.004T_{ui}D + D + 0.04U_i$

Beacon/Radar Correlation = $0.28_i + 0.002B_i + 0.00725R_i$

Conflict Prediction = $[12220 + 636T_i + 5\ \frac{T_i^2 - T_i}{2}]/2X10^4$

Miscellaneous = $0.6D + 1.7_i + 2.8$

System Timeout Processing = 1.8P (P = number of processors)

Failsafe/Failsoft Executive = 7P

Critical Data Recording and Collection ≈ 10% of one IOP

These equations, and the constraints cited below, were used to estimate ARTS III failsafe multiprocessor hardware configurations to support air traffic loads of 250, 500, 750, and 1000.

Configurations Constraints

1) A processor is considered saturated at 80% utilization. This provides a safety margin for unusual loading spikes, overhead due to memory access conflicts, etc.

2) Failsafe configurations include a spare IOP or CPM (or both) and a spare memory module in hot standby mode in the event of an element failure.

3) The maximum ARTS III hardware configuration is physically limited to 8 processors or 12 requestors where:

   IOP = 1 requestor
   CPM = 2 requestors

4) A CMA (Central Memory Access Module) is required if there are more than four processors or 8 memories in the configuration.

5) A CPM has 1.6 times the processing power of an IOP, but no I/O ability.

6) There must be sufficient IOP's in the system to handle I/O requirements in a failsafe manner (redundant connections), i.e., minimum of 3 IOP's in a dual sensor system, 2 in a single sensor system.

The following configurations resulted:

250 aircraft, single sensor
- 2 IOP's
- 2 CPM's
- 0 CMA's
- 5 MM's (Memory Modules)
- 1 RFDU (Reconfiguration & Fault Detection Unit)

250 aircraft, dual sensor
- 5 IOP's
- 0 CPM's
- 1 CMA
- 6 MM's
- 1 RFDU

500 aircraft, single sensor

- 4 IOP's
- 3 CPM's
- 1 CMA
- 7 MM's
- 1 RFDU

500 aircraft, dual sensor

- 4 IOP's
- 3 CPM's
- 1 CMA
- 7 MM's
- 1 RFDU

750 aircraft, dual sensor* and 1000 aircraft dual sensor

cannot be handled in a single ARTS III multiprocessor configuration

Table Ia shows the % processing load results for dual sensor environments of 250, 500, 750, and 1000. We estimate from this sizing exercise that the ARTS III multiprocessor system under the given function loading may handle 500 aircraft, dual sensor environments but runs out of processing power somewhere before it reaches a 750 aircraft, dual sensor environment. Alternatively, if we consider the largest (processor wise) single systems that we can configure with the ARTS III multiprocessor equipment and apply the same loading equations, we obtain the results in Table Ib. We can see from this table that the ARTS III multiprocessor system should be able to handle any reasonable traffic load in a single sensor site with the functions that we have assumed. However, in a dual sensor environment, 600 aircraft appears to be the threshold.

Since the equations are inexact, the functions are nominally selected, memory access interference is not a known quantity, and sensor overlap is not always 75% but sometimes greater, a safer

*Note: 750 and 1000 aircraft would not be reasonable in a single sensor environment.

estimate of maximum air traffic load might be around 500 aircraft (in a dual sensor environment this is roughly equivalent to a trackload of 875).

TABLE Ia ESTIMATED ARTS III LOAD (Including Redundant Elements in Failsafe Configuration)
(All ARTS III Configurations Assume 80% Saturation Threshold)

FUNCTIONS	250 TOTAL AIRCRAFT 15 DISPLAYS PERCENT IOP	500 TOTAL AIRCRAFT 20 DISPLAYS PERCENT IOP	750 TOTAL AIRCRAFT 25 DISPLAYS PERCENT IOP		1000 TOTAL AIRCRAFT 30 DISPLAYS PERCENT IOP	
BEACON INPUT PROCESSING	29	47	66		84	
RADAR INPUT PROCESSING	70	92	114		136	
TRACKING (RADAR/BEACON)	83	168		260	359	
ALL DIGITAL DISPLAY PROCESSING WITH DISPLAY BUFFERS	41	70		106		147
CONFLICT PREDICTION	26	76		150		247
MISCELLANEOUS (KEYBOARD, INTERFACILITY, ETC.)	19	22	25		28	
IOP REQUIREMENT (LESS EXECUTIVE)	268	475	205	516	607	394
IOP REQUIREMENT (WITH EXECUTIVE)	296	517	303	558	656	429
REQUIRED CONFIGURATION (IOP/CPM/CMA/MM)	4/0/0/5	4/2/1/6	3/0/0/4	3/3/1/5	4/3/1/6	3/2/0/5
IOP CAPABILITY	320	576	240	624	704	448
FAILSAFE CONFIGURATION (IOP/CPM/CMA/MM)	5/0/1/6	4/3/1/7	4/0/0/5	3/4/1/6	4/4/1/7	3/3/1/6

TABLE Ib ARTS III ESTIMATED FAILSAFE SYSTEM CAPACITY LIMITS (ARBTL System with Conflict Prediction)

SYSTEM	CONFIGURATION IOP/CPM	MAXIMUM AIRCRAFT	MAXIMUM TRACKS	MAXIMUM AIRCRAFT/SENSOR
SINGLE SENSOR	8/0	850	850	850
	2/5	900	900	900
	4/4	950	950	950
DUAL SENSOR*	8/0	525	915	457
	2/5	(not failsafe - no redundant IOP)		
	4/4**	600	1050	525

* ASSUMES 15 MILE SEPARATION OF SENSORS (75% OVERLAP)

** APPARENT INCREASE IN TRACKFILE IS QUESTIONABLE (DUE TO CONSTANTS IN LOADING EQUATIONS)
DUAL SENSOR: 550 A/C, TRACKS -962, AIRCRAFT/SENSOR -481 }
TRIPLE SENSOR: 350 A/C, TRACKS -912, AIRCRAFT/SENSOR -306 } ARE SAFER ESTIMATES

3.1 AIR TRAFFIC PROJECTION

The question of the necessity of ARTS III computer augmentation, then, became one of whether 500 instantaneous aircraft is a load which will be reached in the terminal area during the life of the third generation ARTS equipment. This is a question to be answered by the FAA and was beyond the scope of this design study, which began with the assumption that augmentation was necessary. We did, however, attempt to address this question in a limited fashion in order to get a feel for the schedule we would be confronting for production versions of parallel processor systems.

Our approach was to use the "Federal Aviation Administration Eastern Region Airport Activity" report for Fiscal Year 1973 to obtain peak aircraft counts for some of the busiest terminals. These counts were then increased by 6% per year and projected to 1984. (The increase in air traffic over the last ten years has been approximately 6% per year.) Near the end of our design study, however, the energy crises became severe and it was felt as a result that a 6% per year growth in air traffic would probably not be achieved over the next ten years. In the absence of any firm guidelines as to the impact of the energy crises on air traffic, we postulated a possible 3% per year growth. These projections are shown in Table II. The results are shown as trackload rather than aircraft count since the counts in some of the dual sensor sites are the sum of the counts from both antennae and include redundant aircraft. It is not possible to determine the exact aircraft count from the data as presented in the Airport Activity Report.

At first glance these projections are not alarming and seemingly fall within the saturation thresholds which we estimated for the ARTS III multiprocessor configurations. We hasten to point out, however, a number of flaws.

TABLE II AIR TRAFFIC PROJECTION

HUB	*DUAL ASR PIAC	EQUIVALENT TRACK LOAD	TRACK LOAD FY'84 PROJECTION 3%	6%
NYCIFRR	*423	423	510	682
LOS ANGELES	*220	440**	530	708
CHIC/O'HARE	*216	216	260	347
LONG BEACH	*157	157	193	252
DALLAS/FT. WORTH	*144	144	174	233
WASH., D.C.	*128	128	154	206
SAN FRAN/OAK	*103	103	125	166
SACRAMENTO	* 94	94	114	152
JACKSONVILLE	* 87	87	105	140
MIAMI	* 82	82	99	132
ST. LOUIS	154	154	186	249
DETROIT	146	146	176	234
PHOENIX	144	144	174	233
SAN DIEGO	135	135	163	218
BALTIMORE	129	129	155	208
SEATTLE	122	122	147	197
MINN/ST. PAUL	110	110	133	177
PITTSBURGH	107	107	129	172
PHILADELPHIA	98	98	118	158
HOUSTON	97	97	117	156
CLEVELAND	96	96	115	154
KANSAS CITY	80	80	96	128
BOSTON	71	71	86	115
LAS VEGAS	61	61	74	98
TAMPA	61	61	74	98
NEW ORLEANS	59	59	72	97
SAN JUAN	56	56	66	88
DENVER	54	54	66	88
HONOLULU	36	36	43	57

**Add Traffic from March Field "Seen" by Lax Radar.

1. The reported ASR PIAC is a count of beacon-equipped aircraft only. It is not a count of all aircraft in the control area.
2. The count is taken only in the control area. It does not include all aircraft "hit" by the radar and consequently processed by the computer. (The range of the radar is typically 60nm.) If there are airports peripheral to the terminal control area, their aircraft activity will be picked up by the radar and will undergo a certain amount of processing by the computer before being filtered out (thus increasing computer processing requirements).
3. The procedure for taking the ASR PIAC count seems to vary from site to site, from an exact count of targets on the display screen to a count of one or two sectors, which result is then used to estimate the total of all sectors.
4. The sensor overlap varies in the dual sensor sites.

Table III illustrates these problems. Our conclusion is only that the data is quite shaky and that our projected track loads should be construed as minimum expectations, with a low level of confidence. In that light, 1984 appears as a year of concern. Clearly, better data is needed to suggest whether the 1984 air traffic load at a given site will exceed a PIAC of 500. We present this data only to highlight the pitfalls.

The preceding loading analysis was performed in order to establish a performance baseline against which a parallel processor augmentation to the ARTS system might be compared; and the air traffic projections were made to estimate a working schedule for potential augmented system evaluation and implementation. It must be pointed out that the addition of sensors to a site, or the addition, deletion, or modification of automated ATC functions could radically alter the conclusions drawn with regard to ARTS III saturation thresholds. It must also be emphasized that the conclusions drawn were not based on actual system performance measurements.

TABLE III AIRCRAFT COUNTS

DUAL SENSOR HUB	REPORTED PIAC	PIAC	SENSORS COUNTED	AIRCRAFT PER SENSOR	TRACKLOAD	RANGE OF COUNT	RANGE OF SENSOR N.M.	% A/C NOT COUNTED	SENSOR SEPA-RATION
NY CIFRR	423		2		423+	JURIS-DICTION	60	?	
CHICAGO MIDWAY	150 56		2	150+ 56		30 20	60 60	? ?	23
LOS ANGELES	220	220	1	220	440+	30	60	?	1.5
LONG BEACH (EL TORO)	157		1	157		30	60	?	20
WASHINGTON, D.C. DCA ADW	69 59		2			30 30	55 60	∼35-40	7
DALLAS/ FORT WORTH BSW DAL	144		1			30	60	FEW	12

FORECASTS SHOULD BE BASED ON TRACKLOAD, OR A COMBINATION OF RADAR TARGETS AND TRACKLOAD + ADD TRACKS FOR % OF A/C TURNING AND/OR FOR AIRCRAFT SEEN BY SENSOR BUT NOT COUNTED

It is recommended that an effort be made to actually measure the performance characteristics of the ARTS III multiprocessor configurations under various ATC function loads and aircraft loads; and that a closer look be taken at the definition and projection of the peak loads placed on the computer facility at the busy hubs, especially those with multiple sensors and peripheral airports.

4. PARALLEL PROCESSOR COST ESTIMATES

As part of the design study, each contractor was required to develop a set of cost estimates for production quantities of his proposed parallel processor hardware. The production quantities were defined to consist of four distinct cases:

Case 1 - 5 dual sensor systems to be installed in 1978

one, 1000 aircraft system
three, 500 aircraft systems
and one, 250 aircraft system

Case 2 - 10 dual sensor systems - Case 1 plus the following to be installed in 1979

two, 250 aircraft systems
three, 500 aircraft systems

Case 3 - 15 systems - Case 2 (dual sensor) plus the following single sensor systems to be installed in 1980

five, 250 aircraft systems

Case 4 - 17 systems - Case 3 plus the following single sensor systems to be installed in 1981

two, 250 aircraft systems

These cases were selected on the basis of an air traffic projection performed early in the design study.

The basic ground rules outlined for the development of costs were as follows:

1. The cost of production systems must be based on 1973 dollars for labor.
2. Any projected cost-reductions for components in the 1978-1981 time-frame must be substantiated by supporting documentation.

3. The development and breakdown of all costs must be shown.
4. Each case (as defined above) must be considered as a separate entity.

The details of the results of these cost estimates are given in each contractor's final study documentation. For the purpose of comparison, summaries have been prepared and are shown in Tables IV-VII. The parallel processor hardware costs in Tables IV and V were arrived at by computing a single system cost (Table VII) for a 250 aircraft system (both single and dual sensor), a 500 aircraft system dual sensor, and a 1000 aircraft dual sensor system based on a total production quantity of 10 systems; this basic cost was then multiplied by the number of systems of that particular size contained in each production lot case. Thus, the figures are not an exact replica of those provided by the contractor, but a very slightly simplified version. The ARTS equipment costs in Table IV are those which were originally provided by the FAA and given to the participating contractors.

Table V was prepared because the FAA felt that the ARTS equipment costs were out-of-date by the end of the design study and therefore a new set of ARTS cost estimates was provided.* The new costs were not made known to the contractors, since they were obtained after the completion date of the design study. The new ARTS III equipment costs used in Table V were the ones used to determine system cost-effectiveness.

Table VI is a list of the ARTS III hardware components along with the old costs, used in Table IV, and provided to the participating contractors; and the new costs used in computing the figures in Table V.

Table VIII shows the hardware configurations proposed by the three contractors to handle the various aircraft loads. These are the single systems whose costs are given in Table VII and used to compute Tables IV and V.

* New ARTS cost estimates supplied by ARD-120.

TABLE IV COST COMPARISON PRODUCTION SYSTEM LOTS (OLD ARTS EQUIPMENT PRICES)

DESIGN STUDY PRODUCTION LOTS (17 SYSTEMS) (1984 LOADS)	ARTS III ALONE	ARTS & PP GOODYEAR	ARTS & PP HONEYWELL		ARTS & PP TEXAS INSTRUMENTS
				** SAVINGS RVD	
CASE 1 1978 1 - 1000 A/C DS 3 - 500 A/C DS 1 - 250 A/C DS	5224K*	3768K Δ = -1456K	3406K Δ = -1818K	360K	4675K Δ = -549K
CASE 2 1979 CASE 1 + 2 - 240 A/C DS 3 - 500 A/C DS	9435K	7353K Δ = -2082K	6566K Δ = -2869K	760K	8681K Δ = -754K
CASE 3 1980 CASE 2 + 5 - 250 A/C SS	12,200K	10,028K Δ = -2172K	9231K Δ = -2969K	960K	11,736K Δ = -464K
CASE 4 1981 CASE 3 + 2 - 250 A/C SS	13,306K	11,098K Δ = -2208K	10,297K Δ = -3009K	1040K	12,958K Δ = -348K

DS = Dual Sensor
SS = Single Sensor
Δ = Difference in price between ARTS III alone and ARTS III + Parallel Processor

* Since our analysis indicates that we cannot configure a single ARTS III multiprocessor system to handle a 1000 aircraft dual sensor environment, the ARTS III alone price for the 1000 A/C system is based on two ARTS systems configured to share the load.

** Potential RVD savings since Honeywell does target detection in PP rather than RVD. These savings were <u>not</u> included in the totals.

TABLE V COST COMPARISON PRODUCTION SYSTEM LOTS (NEW ARTS EQUIPMENT PRICES)

DESIGN STUDY PRODUCTION LOTS (17 SYSTEMS) (1984 LOADS)	ARTS III ALONE	GOODYEAR	HONEYWELL		TEXAS INSTRUMENTS
				SAVINGS RVD**	
CASE 1 1 - 1000 A/C DS 1978 3 - 500 A/C DS 1 - 250 A/C DS	3400K*	2818K Δ = -582K	2546K Δ = -854K	360K	3677K Δ = +277K
CASE 2 CASE 1 + 2 - 250 A/C DS 1979) 3 - 500 A/C DS	6131K	5453K Δ = -678K	4871K Δ = -1260K	760K	6733K Δ = +602K
CASE 3 CASE 2 + 1980) 5 - 250 A/C SS	7881K	7643K Δ = -238K	6826K Δ = -1055K	960K	8963K Δ = +82K
CASE 4 CASE 3 + 1981) 2 - 250 A/C SS	8581K	8519K Δ = -62K	7608K Δ = -973	1040K	10,003K Δ = +1422K

DS = Dual Sensor
SS = Single Sensor
Δ = Difference in price between ARTS III alone and ARTS III + Parallel Processor

* Since our analysis indicates that we cannot configure a single ARTS III multiprocessor system to handle a 1000 aircraft dual sensor environment, the ARTS III alone price for the 1000 A/C system is based on two ARTS systems configured to share the load.

** Potential RVD savings since Honeywell does target detection in PP rather than RVD. These savings were not included in the totals.

TABLE VI ARTS III EQUIPMENT COST DATA

	OLD	NEW	DECREASE
IOP	75	50	33%
CPM	64	45	30%
MM	45	22	51%
CMA	50	50	0
RFDU	50	50	0

The "new" ARTS III equipment costs represent the FAA's best estimate of future production quantity costs of this equipment based on ARTS III implementation experience and using 1973 dollars.

TABLE VII SINGLE FAILSAFE SYSTEM COST COMPARISON (New ARTS III EQUIPMENT PRICES)

ATC LOAD	ARTS III	ARTS III/PP		
		GOODYEAR	HONEYWELL	TEXAS INSTRUMENTS
250 A/C SINGLE SENSOR	350K	438K Δ = +88K	391K Δ = +41K RVD -40K	446K Δ = +96K
250 A/C DUAL SENSOR	482K	515K Δ = +33K	441K Δ = -41K RVD -80K	520K Δ = +38K
500 A/C DUAL	589K	535K Δ = -54K	481K Δ = -108K RVD -80K	672K Δ = +83K
750 A/C DUAL	922K	610K Δ = -312K	581K Δ = -341K RVD -80K	867K Δ = -55K
1000 A/C DUAL	1151K	698K Δ = -453K	662K Δ = -489K RVD -80K	1141K Δ = -10K

TABLE VIII FAILSAFE SINGLE SYSTEM HARDWARE CONFIGURATION AND COST BREAKDOWN (NEW ARTS COSTS)

ATC LOAD	ARTS III ALONE 1 RFDU ASSUMED FOR ALL SYSTEMS	ARTS + PP *		
		GOODYEAR	HONEYWELL	TEXAS INSTRUMENTS
250 A/C SINGLE SENSOR	2/2/0/5 2 IOP'S 2 CPM'S 0 CMA'S 5 MM'S = 350K TOTAL 350K	ARTS 2/0/1/4 = 283K PP 1(1024)/24 1-1024 BIT-64 WORD ARRAY 24K CONTROL MEMORY = 155K TOTAL 438K	ARTS 2/0/0/4 = 238K PP 2/9 (1S) 2 CONTROL UNITS 9 PROCESSING ELEMENTS (INCLUDES 1 SPARE PE) = 153K TOTAL 391K	ARTS 2/0/0/5 = 260K PP 3/4 3 SP/CU (CONTROL ELEMENTS) 4 PE'S (PROCESSING ELEMENTS) = 186K TOTAL 446K
250 A/C DUAL SENSOR	5/0/1/6 = 482K	3/0/1/5 = 360K 1(1024)/24 = 155K TOTAL 515K	3/0/0/4 = 288K 2/9 (1S) = 153K TOTAL 441K	3/0/0/5 = 310K 3/6 (1S) = 210K TOTAL 520K
500 A/C DUAL SENSOR	4/3/1/7 = 589K	3/0/1/5 = 360K 2(1024)/32 = 175K TOTAL 535K	3/0/0/4 = 288K 2/18(2S) = 193K TOTAL 481K	3/0/0/5 = 310K 3/15(2S) = 362K TOTAL 672K
750 A/C DUAL SENSOR	4/0/0/5 + 3/4/1/6 = 922K	3/0/1/5 = 360K 4(2048)/32 = 250K TOTAL 610K	4/0/1/4 = 388K 2/18(2S) = 193K TOTAL 581K	3/0/0/5 = 310K 3/25(2S) = 555K TOTAL 867K
1000 A/C DUAL SENSOR	4/4/1/7 + 3/3/1/6 = 1151K	3/0/1/5 = 360K 5(2048)/32 = 338K TOTAL 698K	4/0/1/4 = 388K 2/36(4S) = 274K TOTAL 662K	4/0/1/6 = 432K 3/33(2S) = 709K TOTAL 1141K

* 1 RFDU at 50K is included in all system totals.

We can conclude from these tables and the loading analysis in section 3 that, given the automated ATC functions which we have assumed, a parallel processor augmentation to ARTS III:

1) is not necessary nor cost effective for a 250 peak aircraft single sensor environment,
2) is not necessary for a 250 aircraft dual sensor environment or 500 aircraft single sensor environment, but is not greatly different in cost from the ARTS III equipment configuration for that size system,
3) is potentially cost-effective, and possibly necessary in a 500 aircraft dual sensor system,
4) is both cost-effective and required in either a 750 or 1000 aircraft dual sensor environment.

When we say "necessary" and "required" in the preceding discussion, we mean that the ARTS III multiprocessor equipment in its present form at its maximum single system configuration probably cannot handle the indicated air traffic loads. Parallel processor augmentation appears to be a feasible and potentially cost-effective augmentation which would accommodate those and even larger air traffic loads. It is, of course, entirely possible that other reasonable alternatives could be postulated.

5. DESIGN STUDY RESULTS

As mentioned earlier, the results of each contractor's design study fills several large volumes. In addition, each contractor's final report contains a summary of their results. Our purpose here is to comment on those results in some of the major areas touched on in the design study.

5.1 PARALLEL PROCESSOR HARDWARE, INTERFACES, FAILSAFE REDUNDANCY, AND THEIR PROGRAMMING IMPACT

5.1.1 Goodyear

Goodyear proposed an "associative processor" architecture, with 1-bit serial arithmetic and 64-bit parallel search capability. The basic array size is 1024 bit by 64 words with provision for expansion to a 2048 bit X 64 word array for large aircraft loads. ECL is used for logic and array. The number of arrays in a given configuration would vary from one to five depending on aircraft load. The control memory is solid state (N-channel MOS) and varies in size from 24 - 32K. This is used for program, data, and microprogram storage. Additionally, a 16-word high-speed cache memory is proposed to improve instruction fetch time for the 1000 aircraft configurations.

Failsafe redundancy is achieved by duplicating the entire AP in a system configuration. Each AP is interfaced directly to an ARTS III memory port, to two IOP's and the RFDU logic. A simple switching arrangement is used to switch to the redundant hot standby in the event of an AP failure.

The Goodyear Associative Processor is a familiar machine. The proposed AP is basically a standard STARAN with some tailoring specifically for the ATC application. The tailoring resulted in a lower cost architecture with enough computer power for the task, but with a high usage efficiency as well. In addition, most of the existing software would still be applicable, thus the changes provide little technical risk. The fact that direct multiplication is not possible (an APPLE subroutine is used) makes the implementation of certain algorithms and processes awkward. Positive features include loop instructions, pointers and length counters to facilitate variable block/field operations. The concept of "folds" of data in the larger air traffic loads as opposed to data roll-in from control memory in the smaller configurations introduces a slight inconsistency in the software from small to large systems, but this is a minor price to pay for the overall cost-savings resulting from the use of smaller and fewer arrays.

Goodyear's approach to failsafe redundancy is sound and appropriate to their architecture; its major shortcoming being the expense of a completely redundant AP.

5.1.2 Honeywell

Honeywell proposed a Parallel Element Processing Ensemble (PEPE) type of architecture, called HAPP - Honeywell Associative Parallel Processor. The basic machine consists of a control unit which contains separate logic for sequential and parallel operations, linked to a set of 8 to 36, 32-bit processing elements (PE's). The PE's are equipped with floating-point and local indexing capabilities. Each PE contains 4K of memory. The ensemble of PE's is directly manipulated by the control unit, which contains a 16K control memory and 8-16K SCL data memory, and provides for a high degree of overlapped operation in its design. The participation of a given PE in a parallel operation can be determined by tag setting or accumulator content, thus providing an "associative" capability. All program code is stored in control unit memory, and most data is contained in PE memory. All memories employ 1024 X 1 bit bipolar RAM chips.

HAPP interfaces with ARTS equipment via a CMA memory port connection and channel connections with two IOP's. In addition, a direct interface (with 8K memory), is also proposed with RDAS.

Failsafe redundancy is achieved through use of dual control units, dual interfaces, spare PE's, spare CU program memory, and triply redundant critical PE status lines interconnected through a complex switching network and controlled by a configuration control unit.

HAPP is a complex piece of machinery which appears on the surface to be more than required for the job. On closer examination, however, many of the apparently superfluous features are used in the development of functional software and in the integration with ARTS equipment. All elements of the system appear to the programmer as conventional processors with conventional looking instructions. Design and coding of the ATC system on HAPP should be relatively straightforward because of the understandable nature of the machine and the plethora of features with which is has been provided.

The interface with the RDAS could result in a cost-saving and more efficient operation than accepting RDAS data through an IOP, but the justification for this is not strongly presented.

A great deal of effort is made to provide redundancy, but the last link in the memory interface, the CMAA, is a single point failure. In addition, the multiple-level switching design for failsafe purposes implies a large amount of wires, gates, drivers, receivers, and connectors which may off-set the reliability gain sought by this very redundancy switching scheme.

5.1.3 Texas Instruments

T.I. proposes to use their commercial 16-bit 980A mini-computers as basic Processing Elements (PE's). These mini-computers are implemented in TTL and dynamic MOS memory with standard commercial packaging. The PE ensemble (8-64 PE's) is interfaced to a bank of global memory modules via a global bus. Synchronization of parallel activity is achieved through special logic in the form of a MRRL (multiple request resolver logic). There are no control units, per se. A PE becomes known as a control unit (SP/CU) by virtue of the fact that is has special priority on the bus, interfaces with an ARTS IOP, and special-purpose executive-like code will be run on it and not on other PE's. Each PE can have from 8-24K of memory. An SP/CU will normally also be distinguished by having more memory than the average PE in a given configuration. Global memory would be provided in 32K increments from 32-192K. The addressing scheme used by the hardware to access global memory requires that a given task not cross a module (32K) boundary. The PE's are capable of independent operation (not in parallel mode), and PE memory can contain both program and data, although the proposed system design does not utilize these features.

The SIMDE II (as this ensemble is called by T.I.) interfaces with three IOP's via their channels and to ARTS memory via an available port. There is one SP/CU for each IOP channel interface. Thus, failsafe redundancy is achieved through multiple SP/CU's, multiple interfaces, spare PE's, and bus system redundancy.

The strength of T.I.'s approach lies in the flexibility of the processing mode, i.e., a portion of the PE ensemble may execute parallel tasks in "lock-step" out of global memory, while other PE's may do independent multiprocessing via SPLIT and MERGE procedures. The idea of using commercial minicomputers is also appealing. However, these very features contribute to the weaknesses of the system. The software system design proposed in this study did not make use of the independence of the PE's, and on close examination of the procedures and tools necessary to do so, one understands why. Designing, implementing, and debugging code to take advantage of this feature would be a programmer's nightmare. It is also not clear whether such a feature would be exploitable in the ATC application.

Although the use of minicomputers is unique and clever, the resulting architecture is unwieldy. A very elaborate bus interconnect scheme is employed which raises questions as to reliability, since connectors are a major offender in system failures.

The lack of an instruction set to control and manipulate the ensemble is judged to be a major shortcoming. In addition, the block data transfer between PE local memory and the global memory is clumsy.

5.1.4 Comparative Assessment of Parallel Processor Hardware Interfaces, Failsafe Redundancy and Programming Impact

The Goodyear and Honeywell approaches provide sufficient and expandable computer configurations for the design study designated ATC loads and functions. The interfaces to the ARTS equipment are clean and straightforward. The hardware designs of the processors themselves are very good and can accomodate the ATC applications. The hardware is realizable using the technology proposed in each case and generally represents low technical risk.

Programming for either machine would be reasonable, although the HAPP is more desirable from an algorithm design and coding point-of-view since it has more sophisticated arithmetic/logic capability.

The approach to failsafe redundancy taken by all three contractors is adequate and compatible with the ARTS recovery scheme. Goodyear's, however, seems to be the least likely to encounter reliability problems in this area due to their simple hot-standby approach.

The Texas Instruments interface approach is clean and realizable with current technology. The hardware architecture, however, presents an unwieldy programming situation, and a confusing and clumsy bus interconnect scheme which raises questions of reliability. This PP is, as a result, judged considerably less desirable than either Goodyear's or Honeywell's.

A comparison of the physical size and space requirements for each contractor's equipment is given in Table IX.

TABLE IX PHYSICAL SIZE COMPARISON

Including maintenance equipment and peripherals but not measurement equipment nor off-line card repair equipment.

		VOLUME		FLOOR SPACE	
250 A/C SYSTEM	RACKS	CU. IN.	CU. FT.	SQ. FT.	
T.I.	6 19"	120X26X70	126	27.7	plus external card reader, printer, 2 silent keyboards
HONEYWELL	5 24"	120X24X72	120	20	plus external card reader display, line printer
GOODYEAR	4 28"	112X32X79	165	25	plus external dec writer, disc, printer, card reader
(NAFEC) 500 A/C SYSTEM					
T.I.	8 19"	160X26X70	168	29	plus external card reader, printer, 2 silent keyboards
Honeywell	6 24"	144X24X72	144	24	plus external card reader display, line printer
GOODYEAR	4 28"	112X32X79	165	25	plus external dec writer, disc, printer, card reader
1000 A/C SYSTEM					
T.I.	13 19"	260X26X70	274	47	plus external card reader, printer, 2 silent keyboards
HONEYWELL	8 24"	192X24X72	192	32	plus external card reader display, line printer
GOODYEAR	4 28"	112X32X79	165	25	plus external.dec writer, disc, printer, card reader

5.2 ATC FUNCTIONAL SOFTWARE AND EXECUTIVE INTERFACES

Each contractor was expected to do an analysis to determine which of the automated ATC functions were suitable to parallel processing in general, and appropriate to execution on their specific architectures in particular, while considering the flow of data from function to function and machine to machine. The overall serial/parallel distribution which resulted in each case was, not surprisingly, similar although the extent and quality of analysis presented by each contractor varied considerably. Table X shows the functions selected by each contractor for execution on his equipment. Those functions without an X beside them were assigned to the serial ARTS equipment.

For each function selected for PP execution, a complete algorithm description was to be provided, including data mapping descriptions, timing analyses, and memory requirements. Intersystem data transfer volume and frequency of transfer were to be considered. These results can be found in the design study documents delivered by each contractor. The following are general comments on the overall results.

5.2.1 Goodyear

The algorithms proposed by Goodyear were well developed, presented in good detail, and drew heavily on their previous experience in Knoxville and TRANSPO. Every attempt has been made to take advantage of the parallelism of the AP. These algorithms, however, differ in some respects from those currently in use in the ARTS system. Parallel processor time and memory requirements were well covered for all algorithms.

The tracking algorithm proposed is a new development and has never been evaluated in a terminal environment.

Conflict prediction is treated as a dual sensor problem, but intersystem conflicts are not addressed.

TABLE X ARTS III/PP ATC FUNCTIONAL DISTRIBUTION

FUNCTION	GOODYEAR	HONEYWELL	TEXAS INST.
EXECUTIVE SCHEDULING			-
BEACON I/O			
CONFLICT PREDICTION	X	X	X
R/B TRACKING	X	X	X
RADAR/BEACON CORRELATION	X	X	X
BEACON TARGET PROCESSING	X	X	X
DISPLAY PROCESSING	X		
DISPLAY OUTPUT PROC			
RADAR I/O		X	
R. TARGET PREDETECTION		X*	X
R. TARGET FINAL DETECTION	X	X	X
TARGET DECLARATION	X	X	X
R. QUANTIZER & CLUTTER			X
MISCELLANEOUS	TRACK RELATED		

14 FUNCTIONS SUMMARY
6 FUNCTIONS ALWAYS IN PP
3 FUNCTIONS ALWAYS IN ARTS
DIFFERENCES IN (1) RADAR HANDLING
(2) DISPLAY PROCESSING

* USES NO HARDWARE PREDETECTION

Goodyear relies on the RVD-4 and IOP for radar predetection, and on the IOP for display output processing. In the Goodyear system design there is only one copy of the trackfile, and it is resident in the AP. Therefore, Goodyear also does all display processing, which involves the trackfile, in the AP.

The implementation of large and small systems would be somewhat different in the Goodyear approach since data files are stored differently depending on system size (in control memory on small systems and in array folds on large systems).

Goodyear has apparently given much thought to optimizing the flow of data. The average data transfer rate is well within system capacity and peak data transfers can be met. The favorable system data flow was achieved through good functional partitioning.

Executive interface is an important consideration which Goodyear handled adequately. The Executive interface uses small control programs inbedded in the lattice structure to schedule AP programs. This is a reasonable approach with one particular shortcoming. The IOP which schedules the AP control program is apparently kept waiting until the AP has completed the requested task. This inefficiency seems unnecessary.

The overall functional flow in Goodyear's system design is a structured sequence of tasks paced by an input data interrupt which initiates the final target detection task (~ every 100 milliseconds).

5.2.2 Honeywell

Honeywell performed an extensive functional analysis to determine the amount of parallelism to be exploited within each ATC function when implementing it on their proposed architecture. As a result, they chose to perform radar target predetection in HAPP rather than use the RVD-4 predetection hardware. They propose this as a potential hardware cost saving since the expensive RVD predetection hardware would not be necessary. Their proposed software predetection algorithm is efficient and clever but needs some work since at least one counter-example exists.

In general, the algorithms are well developed and presented in sufficient detail. Emphasis was wisely placed on parallel implementation of existing algorithms rather than invention of new ones. This approach reduced technical risk. The conflict prediction algorithm was featured since it is a new function. Honeywell proposed a different air sectorization scheme for use with the conflict prediction function to optimize parallel processing efficiency. The scheme is clever and easily mapped from the standard ARTS air sectorization scheme used for all other functions. They did not, however, address the problem of intersystem conflicts in their approach.

Honeywell's timing and memory estimates are particularly well developed-- based on coding of major instruction sequences. System dynamics (task elapsed time and flow of data) was well thought out and sizing of the system was supported by detailed analysis.

Executive interfaces were thoroughly addressed. The design relies strictly on the use of the existing multiprocessor executive and utilizes the executive service request (ESR) to exercise control over the parallel processor. The software on the parallel processor side of the interface operates in a slave mode. A PP function is invoked and controlled by an IOP program through the use of special subroutines. The IOP program can elect to wait until the PP has finished the function, or it may proceed with other IOP operations concurrent with the PP. The ARTS executive POPUP feature is also used. Multiple simultaneous calls from various ARTS processors are handled on the ARTS end of the interface via a first-in first-out queue. Only one request is sent to the PP at a time.

The Honeywell instruction set has all the provisions necessary for coding re-entrant routines. A multi-level PP interrupt structure is therefore realizable in the PP code and used in the design.

5.2.3 Texas Instruments

T.I.'s distribution of functions appears to have been based largely on intuition and generally accepted speculation rather than any real analysis, although the results are acceptable.

T.I. chose to implement the ATC functions as nearly as possible to the manner in which they are done in the ARTS system, thus reducing technical risk. The conflict prediction algorithm, in particular, seems to be well thought out, matching the task and the processor especially well. They were especially careful in their design of the data updating facility so that there is no chance that one processor will update a record while the other is using it. The processing elements are assigned one track per sector, which is very good when traffic is uniform across sectors. They have, however, also provided for "sector overloading" in an acceptable manner.

Data tables through which the IOP and PP software interface are duplicated in whole or in part in the two systems. A system data base management function has been added to pass data back and forth between the duplicate tables. Some re-design of the IOP's copy of the Central Track Store seems necessary.

T.I.'s 16-bit word architecture seems to complicate some of the algorithms, notably radar target detection where the IOP sends a 30-bit word to the PP. Each PE then operates on a 30-bit word as two 15-bit segments.

T.I. uses the RVD hardware predetection results and then proceeds with software predetection from that point. It appears that the use of hardware predetection is marginally beneficial in their approach. Except when there are no targets in the sweep, all PE's wait for the busy processor anyway.

Although flow charts and descriptions were provided for each PP algorithm, it was not always clear what detailed steps would be involved in the parallel implementation, nor how the data flows between global and local memory. A detailed timing analysis was presented along with memory usage estimates.

T.I. uses the ARTS multiprocessor executive to schedule PP tasks via two new ESR's which will allow the IOP to continue processing concurrent with the SIMDE II PP. A message processor on the PP side of the interface enters the PP function request into one of three priority queues (background, planned task, or POPUP task). These queues are serviced by a small PP executive routine. T.I.'s interfaces with the ARTS executive and inter-system communication and control approach are by far the most sophisticated, clear, and flexible of the three contractors. T.I. demonstrated a thorough grasp of both the ARTS III multiprocessor executive concepts and multiprocessing in general.

5.2.4 Comparative Assessment - ATC Functions and Executive Interfaces

All three contractors arrived at a reasonable distribution of ATC functions between the serial and parallel portions of the system; both Honeywell and Goodyear provided convincing arguments for such distributions.

Implementation of the algorithms appeared to be easiest on the Honeywell PP due to its 32-bit word, powerful arithmetic/logic capability, and extremely useful instruction set--especially the ensemble control instructions.

Both Honeywell and T.I. stayed as close as possible to the proven ARTS algorithms, while Goodyear made some modifications to fit the functions to their equipment and to "improve" performance. Goodyear's approach would require that algorithm performance be viewed extra closely in the test-bed. T.I. did not attempt to exploit the PE independence, which they emphasized as an architectural advantage, in their software design.
T.I.'s approach to executive interface and inter-system communication and control was outstanding, Honeywell's was good and Goodyear's was acceptable but not as good as it could have been.

The timing and storage analyses presented by each contractor was adequate, and sufficient to indicate that all loads could be easily handled. Honeywell's analysis was the most credible, however, since major function blocks were flowcharted and timed in detail.

6. BENCHMARK EXERCISE

A benchmark exercise was required of each contractor for the purposes of (1) demonstrating the availability of a parallel processor similar to the one being proposed for the ATC application, and (2) assuring that personnel were assigned to the study who understood the equipment and could code a problem on it. The primary objective of specifying the benchmark exercise was to limit the competitors to those who had previously been involved in the construction and use of parallel processors and thus eliminate "paper designs." Timing was not a real consideration in this exercise since the equipment would not be identical to that proposed, except to establish the fact that parallel execution of the algorithm was an improvement over serial execution. This proved to be the case on all three machines. Timing results are not reported here because they were not intended to be used for comparison nor where they derived in each case in the same manner and are therefore not amenable to comparison.

The benchmark problem (a conflict prediction algorithm) was fully specified by the government in the form of equations, parameters, and limitations. Three sets of input data (on tape) representing air traffic loads of 50, 100, and 200 instantaneous aircraft were also provided. The input data was not given to the contractors in advance of the demonstration. A sample tape was provided in advance, however, to allow for debugging of the tape input routines. The results of the Benchmark exercise were documented by each contractor as part of their design study deliverables. Table XI summarizes the discrepancies found in each contractor's results when compared against the baseline output, which was coded in PL/I; by MITRE Corp. and run on an IBM 360.

TABLE XI BENCHMARK RESULTS

COMPANY	A/C VALUES	# OF MISSED CONFLICTS	# OF DECLARED CONFLICTS NOT IN BASELINE**	# OF A OR B DISCREPANCIES AT MARGIN VALUES (ALT 500) HORIZ DIST 3.00 2.99)			# OF A OR B DISCREPANCIES INVOLVING BAD INPUT DATA***			UNACCOUNTABLE DISCREPANCIES REMAINING
				A	B	TOT	A	B	TOT	
GOODYEAR	50	4	12	3	12	15	0	0	0	1
	100	8	13	8	11	19	0	0	0	2
	200	29	64	20	56	76	9	13	22	5
	TOTAL	41	89	31	79	110	9	13	22	8
HONEYWELL	50	4	0	1	0	1	0	0	0	3
	100	1	0	0	0	0	0	0	0	1
	200	15	20	0	0	0	2	0	2	33*
	TOTAL	20	20	1	0	1	2	0	2	37
TEXAS INSTRUMENTS	50	1	4	0	4	4	0	0	0	1
	100	1	21	0	21	21	0	1	1	1
	200	20	106	2	87	89	9	25	34	12
	TOTAL	22	131	2	112	114	9	26	35	14

* DURING HONEYWELL'S 200 AIRCRAFT RUN, SEVERAL OF THE PE'S ON THE PEPE PROTOTYPE FAILED.

** IT IS NOT CLEAR WHETHER THESE CONFLICTS SHOULD HAVE APPEARED IN THE BASELINE RUN. WE THEREFORE CANNOT REALLY CONSIDER THEM TO BE ERRORS. IT IS ENTIRELY POSSIBLE THAT THE BASELINE WAS NOT ADEQUATE TO CATCH THEM.

*** THE INPUT TAPES CONTAINED SOME AIRCRAFT WHOSE VELOCITIES WERE OUT OF THE ACCEPTABLE RANGE AND/OR INFINITELY LARGE.

7. TEST-BED AND PRODUCTION SYSTEM SUPPORT FACILITIES AND MEASUREMENT PLAN

The ability to make efficient use of any computer is greatly influenced by the quality of the support facilities provided for that computer (i.e., compilers, assemblers, debugging tools, linkers, loaders, file manipulation utilities, peripheral equipment, hardware diagnostic aids, etc.). Indeed, without adequate support facilities, a computer is virtually useless. We were, therefore, quite concerned in this design study with the available or proposed support for each contractor's parallel processor.

In addition, one of the purposes of the ARTS/PP test-bed was to be to evaluate the overall performance of an ARTS multiprocessor system augmented by a parallel processor against an ARTS multiprocessor system alone. It was, therefore, also required that each contractor develop a method for obtaining meaningful performance measurements on the ARTS/PP system, and to define a set of meaningful experiments which would ultimately demonstrate that his particular ARTS III/PP system design provided the sought after capacity increase and performance improvement. The following are comments on each contractor's support facilities and measurement plan.

7.1 GOODYEAR

Goodyear's support system for the AP consists of a PDP-11/20 facility with 24K of core, a keyboard printer, line printer, paper tape reader/punch, disk drive and card reader. No magnetic tapes are suggested. The bulk of Goodyear's support software has been developed for STARAN and would require relatively minor modification for use with their ATC AP. Their APPLE MACRO assembler is particularly useful in dealing with the weak arithmetic capabilities of the AP. An assembler, linker, loader, and macro pre-processor, file utility program, text editor, AP control module, AP debug module, and AP program supervisor are proposed. Of these, the latter three would require development. Production system on-site support would consist of a subset of these facilities developed especially to run on an IOP. This software, of course, is not currently developed either.

In general, Goodyear's proposed facilities are acceptable and reasonable. The fact that the existing PDP-11/20 STARAN software can be modified for ATC AP use makes this a low-risk approach and a favorable one.

Goodyear's measurement plan includes the use of a hardware monitor augmented by software-derived measurement data. They provided an exhaustive list of parameters to be measured but did not seem to have a handle on which were key parameters to be measured or what subset of these parameters would be required for various tests. Their plan lacked organization and was rather lopsided. That is, there were meaningful experiments described with regard to determination of algorithm accuracy; however, no experiments per se were outlined for gaining insight into system performance. There was no indication as to which of the exhaustive list of measurement data will be extracted simultaneously (hopefully not all data would be extracted for every experiment), nor any plans as to how such data should be combined and assessed to view some particular aspects of system performance.

Software data extraction was to be done through modification of an existing (ARTS) data collection program. This is a rather clumsy approach since the ARTS program was not designed for system performance measurement but to collect controller and traffic related data. The modifications to the ARTS program were not spelled out and would seem to indicate that this had not been examined thoroughly. Problems peculiar to measurement data extraction and analysis in a multiprocessor environment were not addressed.

There were a number of data reduction problems which would arise as a result of Goodyear's overall measurement plan which were not addressed in their design, i.e., different record formats on output tapes from the hardware monitor and the software monitor, different word sizes in the IOP, AP, and PDP-11.

The plan, in general, seems clumsy and not thoroughly worked out as to the overall logistics of measurement data extraction and reduction. The problem of assessing system performance seems to have escaped sufficient attention.

7.2 HONEYWELL

In support of HAPP, Honeywell proposed to develop an assembler and loader which would run on the ARTS IOP, backed up by a small cross assembler and loader which would run on an H-716 (to be installed at the test-bed site primarily to support hardware diagnosis). Initial program development would be done on a UNIVAC-1108 whose MACRO assembler closely resembles that of the ARTS ULTRA assembler. Debugging facilities would be provided through the H-716.

This plan is pleasing in that the HAPP would be supportable on the ARTS equipment and, therefore, fit into the NAFEC test-bed operation nicely, but it carries an element of risk. None of the proposed support software is extant. There would be, therefore, a significant development effort required to put together the requisite software on three separate and distinct computers during the overall program. The risks would be encountered in meeting software development schedules and in operating with entirely new support software. The problems involved in coping with three machines with three different word sizes to support a fourth machine with still another word size were not adequately addressed. (IOP-30 bit word, 1108-36 bit word, H-716-16 bit word, Happ-32 bit word.)

The support software specifications, per se, were thorough and reasonable-providing a workable test-bed development environment.

Honeywell's measurement plan included the use of a hardware monitor, a hybrid monitor (internal to the PP) and a software monitor. The overall measurement data to be extracted in order to assess system performance was succinctly defined. A number of very meaningful experiments and air traffic scenarios were described. These experiments were not comprehensive but were definitely germane.

The details of implementing software probes on the ARTS side of the system were conspicuously absent, and the data reduction plan was sketchy particularly with regard to the use of ARTS software measurements. It was clear, however, that Honeywell planned to develop data reduction programs peculiar to the test-bed objectives. They demonstrated an excellent understanding of the goals of measurement and evaluation and the ARTS/PP assessment in particular. There seemed to be concise knowledge of which data to collect rather than using the approach of collecting everything imaginable from the system.

The overall plan leaned a little too heavily on PP measurements when details were described giving it a certain uneveness. The problems peculiar to multiprocessor measurement data collection and reduction were not adequately adressed.

Experiments with regard to determination of algorithm accuracy were sketchy.

7.3 TEXAS INSTRUMENTS

Support facilities proposed by T.I. included a 980A minicomputer with disk, magnetic tape, card reader, line printer and ASR terminal. The existing support software for this minicomputer would be modified and expanded to support the SIMDE II. Since the 980A support software was not designed to handle an ensemble of processors, the modification effort would be much greater than that of Goodyear's. There are, however, a disk operating system and file management routines available providing a good foundation system for software development activity. The SIMDE II debugging facilities would have to be designed from scratch. The provisions for loading the SIMDE II are clumsy and time-consuming. It is necessary to create the load tape (casette) on the support equipment and then load it via a separate ASR attached to SIMDE II. It is not clear how much time it would take to load the PP in this manner, but it seems to be on the order of an hour or so which is rather undesirable. The overall description of support facilities was vague and gave the general impression that this area was not given adequate attention.

T.I.'s approach to measurement and evaluation emphasizes software monitoring with a hardware monitor used only to gain information not suitable to software monitoring. One PE in the test-bed SIMDE II configuration would be dedicated to extraction of measurement data. The types of data proposed for extraction are germane. T.I. exhibited a good concern for the procedures and logistics necessary to coordinate measurement experiments and to log and track various output files. There was, however, no discussion of how software monitoring would be implemented on the ARTS equipment nor what data would be extracted from the ARTS software. Consequently, there was also no consideration of measurement problems peculiar to multiprocessor operation.

In general, T.I. seems to have understood the general objectives of the test-bed measurement and evaluation effort, although there were a number of holes in the description of their overall plan. They were, in particular, missing any presentation on data reduction reports specifically developed for the ARTS/PP performance evaluation.

8. SUMMARY

Our loading analysis indicates that the ARTS III multiprocessor system can probably handle up to 500 instantaneous aircraft in a dual sensor environment (75% overlap) operating with the functions described in Section 3 (basically a radar augmented beacon tracking level system with conflict prediction, and all digital displays utilizing display buffers).

Our air traffic forecasting effort was hampered by ambiguous data, but seems to indicate that some large hubs could begin to experience loading problems as early as 1984.

Our design study for augmentation of the ARTS III multiprocessor system with a form of parallel processor indicates that such augmentation:

1) Is feasible (a) within the hardware and software framework of the ARTS III system, and (b) within the current state-of-the-art of parallel processors.

2) Is a reasonable approach in view of the parallel nature of a number of ATC functions.

3) Can provide a significant capacity increase for both air traffic loads and additional automated functions in a modular fashion adjustable to individual sites.

4) Would extend the useful life of the enhanced ARTS III systems, since an ARTS III/parallel processor combination permits modular expansion in both the ARTS III multiprocessor equipment and in the parallel processor equipment.

5) Can be cost-effective (in production quantity lots) for <u>track</u> loads of 450 and up (250 aircraft-dual sensor) and potentially cost-saving for track loads of 875 and up (500 aircraft-dual sensor).

6) Can accommodate a failsafe/failsoft hardware redundancy and recovery philosophy consistent with that envisioned for the enhanced ARTS system.

Our conclusions with regard to the three ARTS III/PP system designs and attendant support facilities are as follows:

All three designs are feasible and reasonable. The Goodyear and Honeywell architectures are based on existing parallel processor architectures, while Texas Instruments has broken new ground in linking a set of minicomputers into an ensemble.

The Goodyear AP has a 1-bit serial arithmetic/logic capability with no direct multiply features. This seems to require special attention to detailed implementation of algorithms for reasonable performance. Honeywell's PE ensemble has a powerful arithmetic/logic capability (32-bit parallel) and a good ensemble-control instruction set, thus appearing to be the architecture which requires the least programming effort. T.I.'s independent processor architecture lacks a simple ensemble control and synchronization mechanism, thereby introducing programming and software design complexity. In addition, their complex bus interconnect structure and data flow seems clumsy.

The approach to failsafe redundancy taken by all three contractors is adequate and compatible with the ARTS recovery scheme. Goodyear's, however, seems to be the least likely to encounter reliability problems in this area due to their simple hot-standby approach.

The distribution of ATC functions between serial and parallel computers was similar in each of the three designs. Differences were primarily in the handling of radar input and display processing (See Table X). Goodyear has proposed an "improved" tracking algorithm - made reasonable to implement through the parallelism of the AP, while both Honeywell and T.I. stayed as close as possible to the ARTS algorithms.

Conflict Prediction, a new function, was implemented in an acceptable manner by each contractor. Honeywell, however, proposed a unique sectorization scheme for that function to optimize its performance on their PP.

Executive interfaces were satisfactory in all three designs. Honeywell's and T.I.'s approach were similar, but T.I.'s was exceptionally good. Goodyear's approach was workable, but seemingly inefficient in its lack of provision for concurrent activity in IOP and PP.

Adequate support facilities were proposed by each contractor. Most of Goodyear's support facilities have already been developed making them the least risky in this area, although they defer a considerable development effort to the production phase where they propose to design a subset of support facilities for operation on the ARTS equipment. Honeywell's approach is most attractive since they are most compatible with the NAFEC operation. They carry the largest risk, however, since virtually none of their proposed support software has been developed. T.I.'s approach, like Goodyear's, builds on existing support facilities, but they would still have a considerable development effort to support the SIMDE II configuration. In addition, their proposed SIMDE II loading scheme is awkward and time-consuming.

The measurement and evaluation plan proposed by Goodyear leaves a lot to be desired and gives the general impression that they did not have a good handle on the objectives of this task, nor the methods necessary to carry it out effectively. Both Honeywell and T.I. did an acceptable job of putting together a preliminary measurement plan, but Honeywell clearly had the best grasp of both the requirements and the tools and techniques necessary to meet the objectives.

APPENDIX

A.1 BENCHMARK PROBLEM (INTRODUCTION)

The benchmark problem will be used to demonstrate that the contractor has a working machine, to show the programming ease of his machine, and to verify in a very limited way the performance one would expect from his machine. It will not be the primary means of evaluating the effectiveness of his overall design. For the purpose of running the benchmark, a simplified version of the conflict prediction problem has been selected.

A.2 DEFINITION OF TEXT

The contractor will be supplied a tape containing the track positions and velocities of a varying number of aircraft in a 60-mile range and 0 to 10,000 foot altitude layer. He shall code a conflict prediction program which will project the positions of all aircraft once every 4 seconds and will determine whether any pair of aircraft are in conflict as defined in paragraph A3. The contactor's program will output specific information in specific format as described in paragraph A7 each time it determines a conflict.

The above problem is deterministic and has been chosen so that a common solution can be obtained. The contractor is free to use any technique which he desires in order to perform the problem. A real-time test is not required. A minimum of processing should take place in the host machine. The host machine should only be used to load the PP and, if necessary, control the PP's execution. The test will be run in the presence of FAA/TSC evaluators and will include the following documentation to be delivered at the time of the benchmark run:

1. Description and the listing of the programs in the host machine and in the parallel processor.

2. Raw measurement data and descriptive documentation specifying methods used for measurements, format of output, resolution of clocks, etc.
3. Processed performance measures indicating the percent of time that the I/O interface between the host and the parallel processor was busy, memory storage used in the host and in the PP (indicate program and data storage separately), and the total elapsed time of each case in the test, and the total test.
4. The output described in paragraph A7 for the predicted conflicts.

A.3 DEFINITION OF CONFLICT

The contractor's benchmark program is to process for each scan all the aircraft data contained within a scan to determine which aircraft pairs are in a conflict state. An aircraft pair is in a conflict state when the straight line projection of aircraft positions (using the position and velocity data provided) over a specified look-ahead time interval of duration t_w results in the separation of the aircraft at some time over the interval being less than d_h in the horizontal plane and simultaneously less than d_p in the vertical dimension. The constants t_w, d_h, and d_p will be 30 seconds, 3 miles, and 500 feet, respectively for the purpose of this benchmark.

A.4 BEGINNING TIME OF CONFLICT

The time that conflict is projected to begin is required as output data. This is the number of seconds from the present time until the above projection of aircraft positions first enter the conflict state. This should be a non-negative number, if the aircraft are presently in a conflict state, a zero value should be printed.

A.5 MINIMUM MISS DISTANCES

For those aircraft pairs which are projected to be in the conflict state during the look-ahead interval, the projected minimum separations over that interval in the horizontal plane and the vertical dimension are required as output.

A.6 IMPLEMENTATION OF CONFLICT DETECTION FUNCTION

For the purpose of providing a frame of reference for the benchmark, the above defined conflict detection function was implemented on an IBM 360/50 computer. The computations performed in that program are summarized here to clarify any ambiguity in the preceding defintions.

Define the following variables for the i^{th} aircraft:

X_i, Y_i -- location in horizontal plane (n.mi.)

Z_i -- altitude (100 feet per unit)

V_{xi},V_{yi} -- velocity in horizontal plane (n.mi./sec.)

V_{zi} -- velocity in vertical dimension (ft./sec.)

The computations are as follows for aircraft 1 and 2:

$$\Delta Z = Z_1 - Z_2$$

$$\Delta V_z = V_{z_1} - V_{z_2}$$

$$\Delta X = X_1 - X_2$$

$$\Delta Y = Y_1 - Y_2$$

$$\Delta V_x = V_{x_1} - V_{x_2}$$

$$\Delta V_y = V_{y_1} - V_{y_2}$$

$$V_c = \Delta X \Delta V_x + \Delta Y \Delta V_y$$

$$V^2 = \Delta V_x^{\,2} + \Delta V_y^{\,2}$$

$$t_m = -Vc/V^2$$

$$d_o^{\,2} = \Delta X^2 + \Delta Y^2$$

$$d_m^2 = d_o^2 + V_c t_m$$

$$\Delta t = 1/V \sqrt{d_h^2 - d_m^2}$$

$$t_s = t_m - \Delta t$$

$$t_e = t_m + \Delta t$$

$$t_m' = -100 \; \Delta Z/\Delta V_z$$

$$t_s' = t_m - 100 |d_p / \Delta V_z|$$

$$t_e' = t_m' + 100 | d_p / \Delta V_z |$$

The above computations produce the intervals (t_s, t_e) and (t_s', t_e') during which the separation is projected to be less than the separation standards in the horizontal plane and the vertical dimension, respectively. If the intersection of these intervals and the interval (o, t_w) is null, then there is no conflict. Otherwise, the aircraft are in the conflict state. That is,

if $(t_s, t_e) \wedge (t_s', t_e') \wedge (0, t_w) = \phi$ (null)

then there is no conflict.

For those aircraft that are in conflict, the projected starting time of the conflict T_s is

$$T_s = \max (o, t_s, t_s')$$

To compute the projected minimum horizontal and vertical separations D_m and D_m' during the interval (o, t_w) the following procedure is used.

$$D_m = \begin{cases} d_m, & o<t_m \geq t_w \\ d_o, & t_m \leq o \\ \sqrt{d_m^{\,2} + V^2 \, (t_w - t_m)^2}, & t_m > t_w \end{cases}$$

$$D_m{}' = \begin{cases} |100 \; \Delta Z| - |\Delta V_z \; t_w| , & t_m{}' \geq t_w \\ o , & o<t_m{}' \leq t_w \\ |100 \; \Delta Z|, & \Delta V_z = o \text{ or } t_m{}' \leq 0 \end{cases}$$

A.7 OUTPUT FORMAT

When a conflict is detected, the following data is to be printed out on a single line, with two blank spaces between data items:

Variable	Format	Units
Time	F (5)	Seconds
Aircraft ID1	A (8)	-
Aircraft ID2	A (8)	-
Projected Beginning time of conflict (T_s)	F (5)	Seconds
Projected Minimum horizontal miss distance (D_m)	F (5,2)	n.mi.
Projected Minimum vertical miss distance (D_m')	F (4)	feet

The third through the seventh characters of the aircraft I.D. are numerics, and the output should be printed such that for each line the aircraft with the lowest numeric value in this field is printed first. A blank line should be inserted after the conflicts reported for each 4-second scan.