

ANALYSIS OF EDP PERFORMANCE

Contract NAS2--13758

*IN-61-CR
4569*

FINAL REPORT

14 February 1994

Revised
4 April 1994

Prepared for:
National Aeronautics and Space Administration
Ames Research Center
Moffett Field, California

N94-29945

Unclas

G3/61 0004569

Loral Federal Systems
3700 Bay Area Blvd.
Houston, Texas 77058

(NASA-CR-195726) ANALYSIS OF EDP
PERFORMANCE Final Report (Loral
Space Information Systems) 30 p

CENTER RELEASE FORM

NASA Miscellaneous Publications

(1) INSTRUCTIONS:

- o Block 2 has been completed with available information.
- o Complete Block 3 if there are NO Restrictions on the document and it can be distributed to the public
- OR
- o Complete Block 4 if there are any limitations and the document is only available either FROM a specific office or TO a specific audience (i.e. US Government Agencies only...).
- o Return the completed form to: NASA CASI, Attn: ACQUISITIONS, 800 Elkridge Landing Rd., Linthicum Heights, MD 21090-2934 OR fax form to (301)-621-0134 (Acquisitions Dept).

(2)

Title: Analysis of EDP Performance - Final Report

Author: Loral Federal Systems, Houston, TX

Originating NASA Organization: NASA Ames Research Center

Report Number: NAS2-13759

Publication Date: February 14, 1994 (Revised April 4, 1994)

*IN-61-CR
4569
30P*

(3)

IXI PUBLICLY AVAILABLE

Yuan-Kuan Lin 5/11/94
Signature Date

(4)

LIMITED:

 / Available Only From: _____

 Available Only To (USGA, NASA,...) US Government Agencies til _____

**NOTE: Public Release on _____

Signature Date

(CFxxx 9/93)

Contents

1. Introduction
2. Scope
3. Performance Analysis Factors (SOW Subtask 1)
4. Selection of Benchmarks and Mixes (SOW Subtask 2)
5. Performance versus Power (SOW Subtask 3)
6. Consideration of New Features (SOW Subtask 4)
7. Performance Calculations (SOW Subtask 5)
8. Performance Measurements
9. Observations and Discussions of Results
10. Summary, Conclusions, & Recommendations (SOW Subtask 6)
11. References

1.0 Introduction

The objective of this contract was the investigation of the potential performance gains that would result from an upgrade of the Space Station Freedom (SSF) Data Management System (DMS) Embedded Data Processor (EDP) "386" design with the Intel Pentium (registered trade-mark of Intel Corp.) "586" microprocessor. The Pentium ("586") is the latest member of the industry standard Intel X86 family of CISC (Complex Instruction Set Computer) microprocessors.

This contract was scheduled to run in parallel with an internal IBM Federal Systems Company (FSC) Internal Research and Development (IR&D) task that had the goal to generate a baseline flight design for an upgraded EDP using the Pentium.

2.0 Scope

This final report summarizes the activities performed in support of the referenced contract.

Our plan was to baseline performance analyses and measurements on the latest state-of-the-art commercially available Pentium processor, representative of the proposed space station design, and then phase to an IBM capital funded breadboard version of the flight design (if available from IR&D and Space Station work) for additional evaluation of results.

Unfortunately, the phase-over to the flight design breadboard did not take place, since the IBM Data Management System (DMS) for the Space Station Freedom was terminated by NASA before the referenced capital funded EDP breadboard could be completed. The baseline performance analyses and measurements, however, were successfully completed, as planned, on the commercial Pentium hardware. The results of those analyses, evaluations, and measurements are presented in this final report.

3.0 Performance Analysis Factors (SOW Subtask 1)

The hardware and software factors delineated in *Figure 1* were deemed to be significant as to influence on the instruction execution rate of the Pentium based EDP. These factors were derived from the flow model shown in *Figure 2*, which for analysis purposes, shows the flow direction from main memory to the CPU (the model works with flow in either direction, however).

These factors are shown grouped into an Equation-for-Actual-MIPS (EFAM), shown in *Figure 3 (a)*. The EFAM will yield the upper limit of performance (the Von Neumann Limit-VNL) for the Pentium, as shown in *Figure 3 (b)*, when all factors are set to their maximum/minimum (optimum) values.

The triple path pipelined (superscaler) Pentium will exhibit an absolute VNL for integer performance of "2f" (which is 132 MIPS at 66 Megahertz clocking rate) and an absolute VNL of "f" for operation with all instructions running floating point data (66 MIPS).

Currently published vendor (Intel Corp.) optimized performance results indicate that an efficiency of only 51% is achieved with the SPECint92 benchmark (SPECint92=67.4 MIPS), but an efficiency of 96% with the SPECfp92 (SPECfp92=63.6 MIPS). An examination of the Pentium flow model shows:

- For integer performance--only one flow path is effectively utilized on the average; or that the average instruction is taking approximately two clock cycles,
- For floating-point performance--excellent results are achieved, i.e. the single floating-point path remains essentially full.

HARDWARE (UNITS)

Symbol

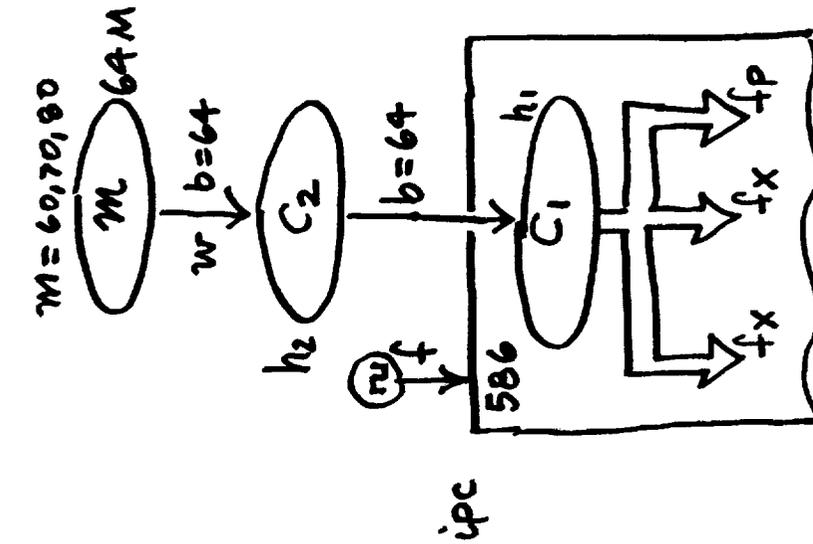
• Clock Frequency (MHz)	f
• Main Memory cycle (nsec)	m
• wait state (#)	w
• BUS width (bits)	b
• L2 cache hit ratio (%)	h_2
• L1 cache hit ratio (%)	h_1
• EDAC delay (nsec)	d
• CPU bus (input) cycle	$m+d$

(#: a number; integer)

SOFTWARE

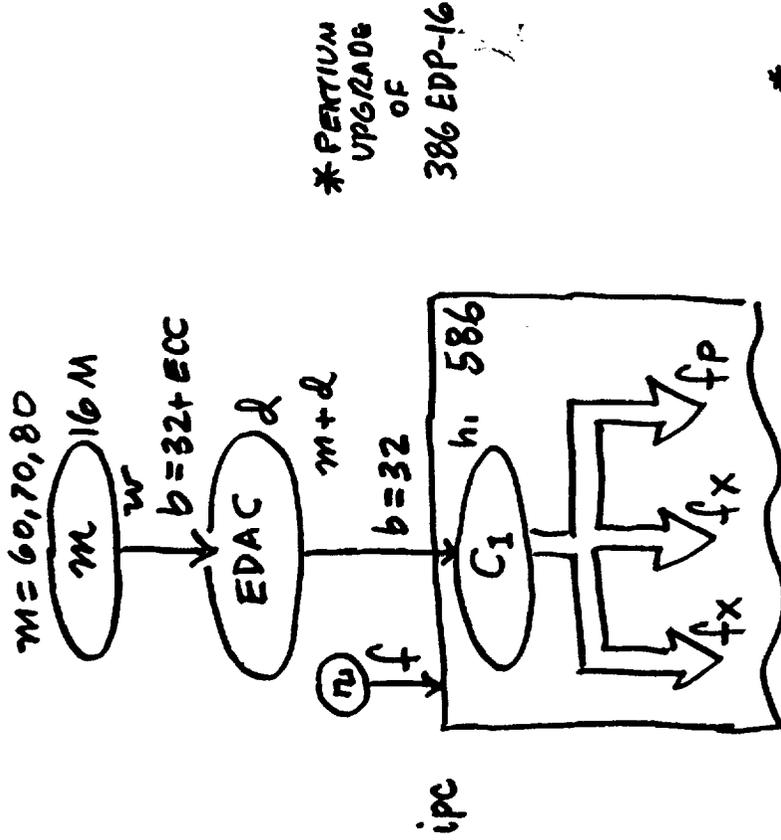
• Compiler optimization/parallelism (ipc - instructions per cycle)	$0 < ipc \leq 2$
• instruction mix (cycles per instr.) (cpi)	$1, 2, \dots, N$
• pipeline (P)	$0 < P \leq 1$
• OS overhead (S)	$0 < S \leq 1$
• Application efficiency (mix) (a)	$0 < a \leq 1$
• I/O, DMA, contention & interference (β)	$0 < \beta \leq 1$

FIGURE 1 = PERFORMANCE ANALYSIS FACTORS



a) Commercially based version

m = memory cycle
 w = wait state(s)
 h_i = cache hit ratio ($i=1,2$)
 f = clock frequency
 d = delay



* PENTIUM
UPGRADE
OF
386 EDP-16

b) Proposed flight design (EDP-16P)

C_1 = on-chip/level 1 cache
 C_2 = off-chip/level 2 cache
 ipc = instructions per cycle
 $EDAC$ = error correction unit
 fx = fixed-point/integer
 fp = floating-point
 b = bus width

FIGURE 2 : PENTIUM (586) FLOW MODEL

ABSORBED INTO ipc

$$a) EFAM = \underbrace{\left(\text{MIX FACTOR} \right) \left(\text{APPL FACTOR} \right) \left(\text{OS OVER HEAD} \right) \left(\text{COMPILER OPTIMIZATION PARALLELISM} \right) \left(\text{PIPE LINE} \right) \left(\text{L1 CACHE} \right) \left(\text{L2 CACHE} \right) \left(\frac{b}{a} \right)}_{\text{SOFTWARE FACTORS}} \underbrace{\left(\text{EDAC DELAY} \right) \left(\text{MAIN MEM} \right)}_{\text{HARDWARE FACTORS}} \left(\text{CLOCK} \right) \left(f \right)$$

fx vs. fp includes I/O

X2 OPTIMIZED
X1 386/486

MEM WORD WIDTH

= 1 if $m+d \leq \text{CPU CYCLE}$

≈ 0.9 OR BETTER

b) • Von Neumann Limit (VNL) = $2f$ WITH ALL FACTORS @ MAX/MIN (OPTIMIZED)

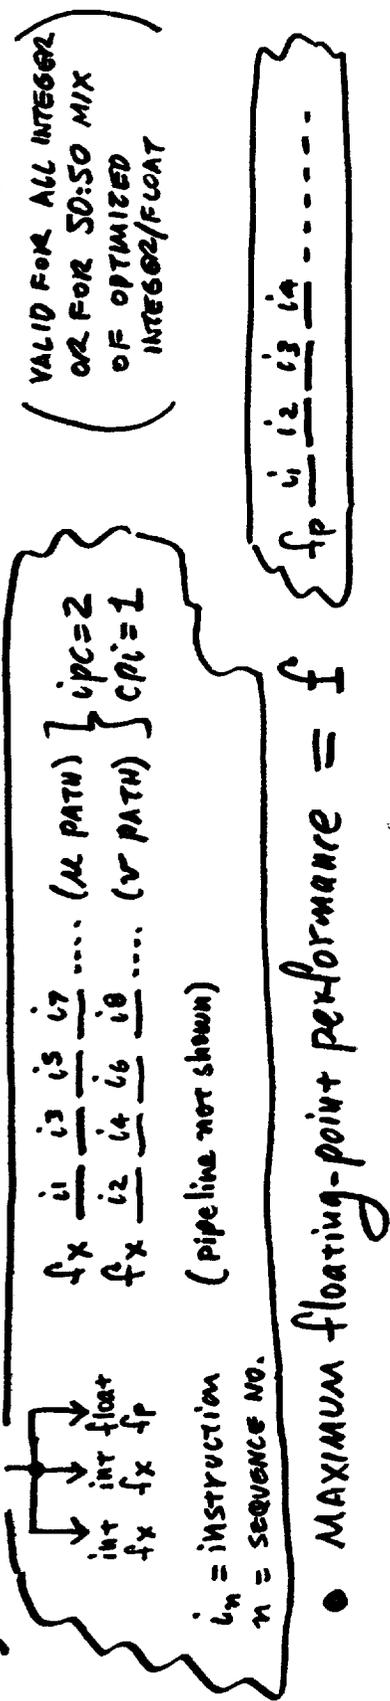


FIGURE 3 : EQUATION FOR ACTUAL MIPS (EFAM)

An examination of the Pentium architecture at other clock rates reveals consistency for the integer performance. At 90 Megahertz, the vendor optimized results are 90 SPECint92 MIPS; and for the P54C the vendor optimized results at 100 Megahertz are 100 SPECint92 MIPS. Both of these yield an efficiency of 50% when compared with the Von Neumann Limit of "2f".

The floating-point performance, based on published vendor optimized results, however tends to fall in efficiency at the higher frequencies. The 90 Megahertz version exhibits 72 SPECfp92 MIPS, while the 100 Megahertz P54C is reported at 81 SPECfp92 MIPS. These performance values are approximately 80% efficient--down by 16% from the 96% efficiency of the 66 Megahertz version under test and evaluation in this project. This efficiency, however, is still quite high, and may merely indicate that the floating-point efficiency for the Pentium superscaler design should be realistically considered in the 80 to 90 per-cent range (maximum).

The EFAM and the Pentium architecture will be referenced and discussed again in the section comparing theoretical performance with actual measured results.

4.0 Selection of Benchmarks and Mixes (SOW Subtask 2)

Various benchmarks and mixes, as listed in *Figure 4 (a & b)*, were examined for consideration as "standard" measures of performance for the EDP processor family, i.e. from 386 versions--through 586 versions--onto future versions utilizing the Intel X86 instruction set architecture. A unified means of measurement is considered key for understanding the performance gains from generation to generation and for understanding the performance differences that may potentially result from EDP variations, such as application

a) MIXES

- ADDS — a quick indication of typical perf.
- 80/20 — a general purpose "rule-of-thumb" representative of aerospace throughput
- GIBSON — an early aerospace mix
- DAIS — developed to evaluate MIL-STD-1750A & similar aerospace/MIL ISAs
- STATION — REPRESENTATIVE OF SPACE STATION DATA MANAGEMENT SYSTEM S/W

b) BENCHMARKS

- WHETSTONE — floating point math/scientific
- DHRYSTONE — INTEGER TEST OF PROCESSOR AND COMPILER EFFICIENCY
- LINPACK — HEAVY SCIENTIFIC APPLICATIONS
- SPEC MARK — COMPREHENSIVE INTEGER AND FLOATING-POINT SUITES.

FIGURE 4: MIXES & BENCHMARKS

dependencies, card types, etc. Definitions of the two measurement sets are as follows:

- **Mixes**--*Figure 4 (a)*--Generally computed for a processor by hand, using instruction timing data from the hardware reference manual. The "mix" is typically controlled by the number or the percentage of the instruction types used. Mixes are useful as "rules-of-thumb" for computers, and are usually easy to compute (in machine language). Mixes are not often available as software packages. Mixes are usually designed to represent a particular problem/application type--such as, integer, scientific, aerospace (GN&C), floating point, etc. Mixes can be as simple as one instruction, as in the case for the integer ADD mix. The strategy behind a mix is to quickly represent the machine performance with a number that is relatively easy to calculate, but actually yields a rating that is commensurate with realistic throughput for that application. The best historical example of this is the 80/20 mix that was used in the 60's & 70's for aerospace computers. The 80% "add" portion would realistically represent the application, and the 20% "multiply" portion would typically cover the rest of the machine overhead from executive software to input-output operations, etc.
- **Benchmarks**--*Figure 4 (b)*--Benchmarks are usually larger and more comprehensive pieces or groups of actual software, designed to represent a particular application, mode of operation, language, operating system, or some combination of these. Benchmarks, as with mixes, can also be designed to exercise architectural features of the computer, such as cache memories, pipelining, floating point, memory management, etc. Due to the fact that benchmarks are software packages, they may be transported from machine to machine, and many are becoming essentially standardized. Many of these software based benchmarks (SPECMARKS, Whetstone,

Dhrystone, etc.) are utilized by the industry for comparisons between models in an architectural family and for comparisons between competing architectures.

Criteria used for evaluation of the benchmarks and mixes included:

- Does it realistically and/or accurately represent real throughput?
- How easy is it to use--or how easy are result numbers generated? Is it controlled by software means? Does it require hand calculations, assembly, or any hand coding?
- How transportable and consistent is the benchmark/mix across various platforms?
- How standard is the benchmark/mix within the industry?
- Are relative comparisons between computers meaningful?
- How adaptable is the benchmark/mix to architectural and/or parameter changes?
- Is the benchmark/mix "fooled" by architectural tricks or features, such as pipelining, caching, etc.?
- How tailored is the benchmark/mix--will it reflect a bias for a particular application or architecture?

A summary of the advantages, disadvantages, and suitability of the particular benchmarks and mixes are shown in *Figure 5*. The final trade matrix results are then summarized numerically in *Figure 6* for all of the benchmark and mix sets considered.

CANDIDATE	PROS	CONS	SUITABILITY
• ADDS	EASY TO COMPUTE	UNREALISTIC FOR PENTIUM	POOR
• BO/20	EASY TO COMPUTE	NOT REPRESENTATIVE FOR SUPERSTAR	POOR
• GIBSON	GENERAL PURPOSE MIX	MANUALLY COMPUTED	FAIR
• DAIS	AEROSPACE MIX	DIFFICULT TO COMPUTE	FAIR
• STATION	REPRESENTS PROPOSED STATION DMS	HEAVY FLOAT-POINT MANUALLY CALCULATED	GOOD
• WHETSTONE	TRANSPORTABLE	HEAVY FLOAT	FAIR
• DHRYSTONE	TRANSPORTABLE	INTEGER ONLY	POOR
• LINPACK	TRANSPORTABLE	NOT REPRESENTATIVE	POOR
• SPEC MARK	TRANSPORTABLE STANDARDIZED	COMPILED ONLY	BEST

FIGURE 5: MIX & BENCHMARK SUITABILITY

CRITERIA	MIXES						BENCHMARKS			
	ADD	80/20	GIB	DAIS	STA		WHET	DHRY	LIN	SPEC
REALISTIC?	0	7	3	3	10		8	5	5	10
EASE OF USE	10	10	7	5	10		9	7	7	9
TRANSPORTABLE?	5	5	5	5	5		10	10	10	10
STANDARD	5	10	7	7	5		10	10	10	10
RELATIVE COMPARISONS MEANINGFUL?	7	5	5	3	9		9	7	8	10
ADAPTABLE	2	4	3	1	5		8	7	6	9
NOT FOOLED	10	10	10	5	10		7	5	8	9
NOT TAILORED	7	7	7	5	5		8	5	7	9
	46	58	47	34	59		69	56	61	76

BEST BENCHMARK

BEST MIX

ALSO GOOD RULE-OF-THUMB

FIGURE 6: MIX / BENCHMARK SUMMARY TRADE MATRIX

After these considerations of the benchmark and mix candidates it was determined that the SPECMARK suite would be used (as indicated in the August 4, 1993, kickoff meeting between IBM and NASA Ames for this project), since it is becoming a universally accepted *de-facto* "standard".

5.0 Performance versus Power (SOW Subtask 3)

When implemented with the low-power version of the Pentium (P54C), the upgraded flight version of the EDP-16 (to be designated as EDP-16P) will potentially have a lower total power dissipation. The removal of the 386 and 387 chips from the EDP-16 will typically reduce the overall page power by 6 to 8 watts. The addition of the P54C will add back approximately 4 watts when running at full clock. The P54C is anticipated to be designed with on-chip power switching for the various functional areas--which will permit power to automatically vary from a fraction of a watt to full power, as a function of performance. Unused portions of the chip will be permitted to drop into the quiescent mode, dissipating essentially "near-zero" power, when not in use by the application (software) running. The range of power anticipated for the EDP-16P flight version will potentially be:

- At full capability (for the ≥ 40 MIPS goal),

$$P_{d_{max}} = \underbrace{(EDP-16)}_{28} - \underbrace{(386)}_{3.5} - \underbrace{(387)}_{3.5} + \underbrace{(P54C)}_2 = 23 \text{ WATTS} \\ @ 33\frac{1}{3} \text{ MHz}$$

- At idle (in standby mode),

$$P_{d_{min}} = \underbrace{(P_{d_{max}})}_{(23)} - \underbrace{(P54C)}_{(2)} = 21 \text{ WATTS}$$

These values are based on the current power dissipation estimate for the EDP-16 (with its 386 and 387 chips).

6.0 Consideration of New Features (SOW Subtask 4)

New features available with the Pentium were surveyed from a performance point of view in order to determine if any should be candidates for consideration in the EDP-16P design. The new features of the Pentium examined included the following:

- The triple superscaler pipeline (two integer plus one floating point path),
- Automatic on-chip power switching,
- Self-checking functions (redundancy and fault-tolerance aids),
- Performance monitoring and control,
- New instructions,
- New architectural features.

Figure 7 provides an overview of the results of the consideration of new features in the Pentium from a performance point of view. The most significant new feature that can potentially benefit throughput is the three parallel flows within the superscaler architecture--two paths for integer instructions and one path for floating-point instructions. This feature will theoretically permit the number of instructions executed per clock cycle to approach two (IPC=2). The architecture attempts to automatically fill up to two of the pipes at any given time, but is highly dependent upon the instruction flow presented by the compiler. For this reason, existing 386 and 486 code, unless recompiled, will execute in the Pentium with $IPC=1+\Delta$, where Δ will

PENTIUM NEW FEATURE	DIRECT PERFORMANCE BENEFIT	APPLICABLE TO EDP-16P DESIGN CURRENT VERSION	FUTURE VERSIONS
<ul style="list-style-type: none"> • 3 way superscaler pipeline (2 int + 1 float path) 	<p>YES (yields ipc = 2 for optimized code)</p>	YES	YES
<ul style="list-style-type: none"> • AUTO ON-CHIP POWER SW 	NO	YES (WITH PS4C)	YES
<ul style="list-style-type: none"> • SELF-CHECK FUNCTIONS - dual chip 	NO	PARTIALLY (SELECTED) (TBD BY ADDITIONAL TRADES)	YES (RED-MGT & FAULT-TOL & RAD-TOL)
<ul style="list-style-type: none"> • PERFORMANCE MONITORING - M & V signals 	YES (an aid to real-time optimization)	NO	TBD
<ul style="list-style-type: none"> • NEW INSTRUCTIONS 	NO	YES	YES
<ul style="list-style-type: none"> • NEW ARCHITECTURE - superscalar 	NO	NO (NOT SUPPORTED BY CURRENT DMS SOFTWARE)	YES (FOR RED-MGT, PM2-MGT, & FAULT-TOL, & RAD-TOL.)

FIGURE 7: NEW FEATURES CONSIDERATIONS - PERFORMANCE RELATED

vary from zero to about 0.1 to 0.2, depending on the code layout and the ability of the hardware to sort out the flow in real-time. As will be seen, the results of our performance measurements will show essentially an IPC=1 when compared with vendor results that were typically optimized for IPC=2.

7.0 Performance Calculations (SOW Subtask 5)

Performance calculations and estimates for the selected benchmark and mix set are shown in *Figure 8*. The data is tabulated into two side-by-side columns for comparison of the 20 MHz 386 EDP-16 with the projected 586 EDP-16P (flight version). The data in the chart for the 586 EDP-16P is shown at 60 MHz, which was the speed of the commercial version used for performance measurements. This will permit meaningful comparisons to be made.

The data for the "station-mix" clearly indicates that the 40 MIPS goal for the EDP-16P is achievable and that the clock rate can be reduced to 33 1/3 MHz to reach this level of throughput. At this set-point a realistic IPC of 1.2 is utilized--i.e., allowing for 20 % utilization of the multiple integer and floating-point paths through the chip. This level of utilization does not place severe demands on the compiler, and should be reached also with 386 and 486 code that is carried forward to the 586 without a recompile. This should be beneficial to projects where conservation of existing software without additional rework or where budget limitations are key. *Figure 8* shows these parameters for the 40 MIPS goal. *Figure 9* shows a graphical derivation of the set-point that yields the 40 MIPS result with the 33 1/3 MHz clock at the IPC of 1.2 (20 % multipath utilization).

For reference purposes, the "station-mix" is listed in *Figure 10* for the X86 instruction set architecture. This mix will yield an in-line flow for the 386, and can be manipulated to utilize the parallel

MIX	EDP-16 (20 MHz / 386/387)	EDP-16P (60 MHz / 586)
<ul style="list-style-type: none"> • ADDS • 80/20 • GIBSON • DAIS • STATION 	10 MIPS (MAX) 5 MIPS (integer) 2.71 MIPS 2.5 MIPS 4.15 MIPS EDP-16P GOAL	30 MIPS (ipc=1) / 60 MIPS (ipc=2) 16.2 MIPS (ipc=1) / 30 MIPS (ipc=2) 42 MIPS (ipc=1) 68.6 MIPS (ipc=2) 45.6 MIPS (ipc=1) 118 MIPS (ipc=1.97) 40 MIPS @ 33 1/3 MHz ipc=1.2
BENCHMARK <ul style="list-style-type: none"> • WHETSTONE • DHRystone • LINPACK • SPEC MARK 	1.0 TO 1.6 MWIPS 4000 TO 7500 DMIPS 0.2 MFLOPS 2 TO 4 integer 1 TO 2 float	19.1 (EST) 92,400 DMIPS/SEC (EST) 2.6 MFLOPS (EST) 59-61 integer (ipc=2) 52-58 float (ipc=2) [VENDOR CALCULATIONS/SIMULATIONS]

FIGURE 8: PERFORMANCE CALCULATIONS - OVERVIEW

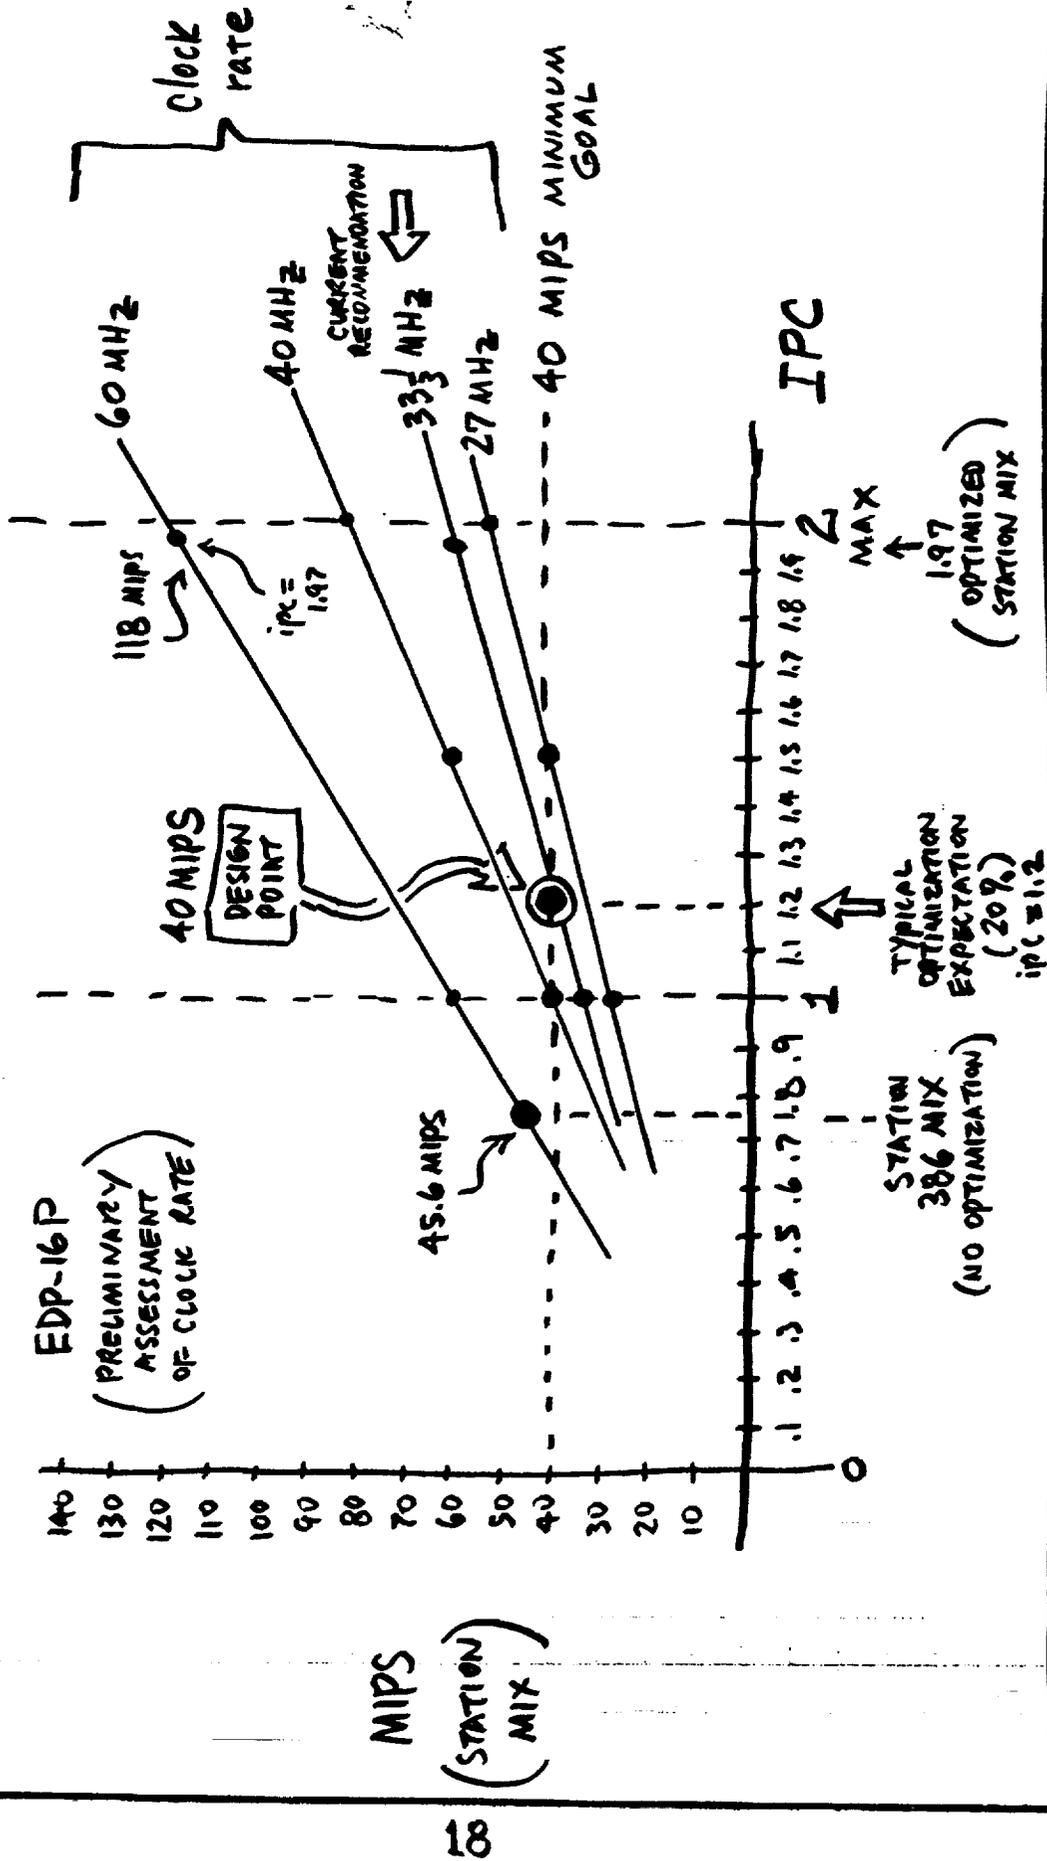


FIGURE 9 : CLOCK RATE RANGE/SELECTION FOR EDP-16P

INSTRUCTION MNEMONIC	FREQUENCY OF OCCURRENCE
MOVE R, M	16.5
MOVE M, R	12.8
PUSH R	9.9
MOVE R, R	9.4
MOVE R, I	7.0
JCC TAKEN	2.5
JCC FAIL	3.5
ALU2 R, R	6.0
POP R	3.3
JMP	2.0
ALU2 R, M	2.0
ALU2 M, I	1.5
CALL	2.1
SHIFT R	3.8
ALU2 R, I	3.9
RET	2.0
STRING	0.7
ALU1 R	3.3
LD POINTER	0.6
ALU2 M, R	0.3
ALU1 M	0.3
PUSH M	0.3
NO OPERAND	1.3
PUSH SR	1.8
MOVE SR, I	1.2
MOVE M, I	1.5
POP SR	0.2
MULW X	0.1
OTHERS	0.2
total	100.0 %

%

FIGURE 10 : SPACE STATION MIX

multipath flow for the 586 (integer instructions flowing in parallel "U" and "V" paths when permitted by the alignments; and floating-point following its path).

8.0 Performance Measurements

8.1 Hardware Configuration

The hardware configuration used for the measurements included a commercial IBM PS/2 Server Model 95-560 and an IBM PS/2 Model 80. The Model 95-560 machine was equipped with an Intel Corp. Pentium central processor unit running at 60 Megahertz, 64 Megabytes of main memory and a 256 kilobyte level-2 cache. The PS/2 Model 80 was equipped with a 386 central processor unit running at 20 Megahertz, a 387 floating-point coprocessor, and 16 Megabytes of main memory. The Model 80 was chosen because it was the commercial design baseline for the Space Station DMS Embedded Data Processor (EDP).

In addition to the two PS/2 machines, an IBM RISC/6000-730 workstation was available and was used for early familiarization with the benchmark software before the PS/2 95-560 became available. Although not part of this task, the benchmark data obtained from the RISC/6000 is included in this report merely for additional comparison and references purposes.

8.2 Software Configuration

AIX Version 1.3, a UNIX based operating system, was used on both PS/2 platforms since the execution of the SPEC benchmark software assumes the presence of a UNIX based system. In addition, AIX Version 1.3 also supports the X86 family of processors, including the Pentium (586).

8.3 Benchmark Software

The SPEC benchmark software was used to measure the performance of the Pentium PS/2 and the 386 PS/2. The SPEC-92 benchmark release consisted of two separate suites to measure integer operations and floating-point operations. The integer suite, CINT92, consisted of six integer intensive benchmark programs, written in C language. The floating-point suite, CFP92, consisted of 14 benchmark programs, written in C and FORTRAN; stressing both single and double precision floating-point operations.

Table 1 presents a list of the benchmark programs that make up each suite with a brief description of each.

The SPEC benchmark procedure was to run each program on the system to be measured with only one process active. The resulting run time was then recorded and normalized by dividing by a reference time supplied by SPEC. This reference time corresponds to the time necessary to run the same program on a DEC VAX 11/780 computer, as measured by SPEC. The resulting number constituted the SPECratio for that particular benchmark routine. After all of the programs in a suite were executed, the geometric mean of all the individual program SPECratios were computed. This result was the single metric for each suite, known as the SPECint92 (for the integer results), and the SPECfp92 (for the floating-point results).

8.4 Benchmark Execution

The SPEC benchmark software was originally intended for measuring the performance of different workstation platforms. In order to facilitate this, the software is provided in a source code format that must be compiled for the system under analysis. The need for compilation results in the benchmark metrics reflecting not only the raw processor performance, but also the performance of the

Table 1. SPEC Benchmark programs

Program	Language	Description
INTEGER SUITE:		
008.espresso	C	A tool for generation and optimization of Boolean expressions
022.li	C	A LISP interpreter that solves the nine queens problem
023.eqntott	C	Translates a logical expression of a Boolean equation to a truth table
026.compress	C	A text compression and decompression utility using an adaptive Lempel-Ziv algorithm
072.sc	C	A spread-sheet that calculates budgets and 15-year amortization schedules
085.gcc	C	GNU C compiler compiling preprocessed source files
FLOATING POINT SUITE:		
013.spice2g6	Fortran	An analog circuit simulation application
015.doduc	Fortran	A thermohydraulic simulation of a nuclear reactor
034.mdljdp2	Fortran	A double precision program that solves the equations of motion for a model of 500 atoms
039.wave5	Fortran	Solves Maxwell's equations and a particle's equations of motion on a Cartesian mesh.
047.tomcatv	Fortran	A vectorized mesh generation program
048.ora	Fortran	Traces rays through an optical system
052.alvinn	C	Trains a neural network using back propagation
056.ear	C	A simulation of the human ear
077.mdljsp2	Fortran	A single precision version of 034.mdljdp2
078.swm256	Fortran	Solves a system of shallow water equations using finite differences approximations
089.su2cor	Fortran	A quantum physics mass computation
090.hydro2d	Fortran	Solves hydrodynamical Navier-Stokes equations to compute galactical jets
093.nasa7	Fortran	A collection of seven kernels that perform matrix multiplications, vortex solutions, matrix inversions ,etc.
094.fpppp	Fortran	A Quantum Chemistry program that measures the performance of a two electron integral derivative

compilers. Since the goal of this project was the comparison of the Pentium performance to that of the 386-EDP, and not differences in compilers, the benchmark programs were compiled (but not optimized for the multiple execution paths) in the Pentium PS/2 and then copied (as load modules) to the 386 PS/2. This eliminated any compiler effects from the SPEC metrics--including any optimization for multiple path execution for the 586 (the case where IPC equals 2). The compilers used were an IBM AIX PS/2 C, version 1.2, and an IBM AIX PS/2 VS FORTRAN compiler, version 1.1.1.

8.5 Elimination of Two Metrics--Problems Encountered

Many of the SPEC software benchmarks have been derived from publicly available applications and have evolved to a point where they can be ported to many different platforms. A significant number of tools and procedures are provided with the software to aid the users in running the benchmark for their particular architecture. Regardless of this, however, individual hardware and compiler idiosyncrasies may still affect the success of some program executions. This was the case with 085.gcc and 039.wave5 routines. The programs ran to completion, but comparison with the expected SPEC results were not successful. Consequently, these two routines were eliminated from the benchmarks.

8.6 Software Compatibility

The Pentium (586) instruction set was advertised to be fully compatible with the 386 instruction set, even though five new instructions were added. There were no incompatibility issues noted during the execution of the SPEC benchmark software suites or with the other commercial operating systems and compilers used.

8.7 Performance Measurement Results and Comparisons

SPEC benchmark results are shown in **Table 2** (Integer Suite--SPECint92) and **Table 3** (Floating-Point Suite--SPECfp92). As can be seen in the tables, the Pentium processor provides, as expected, a significant performance improvement over the 386 processor for both integer and floating-point programs. Since the objective of the project was to compare realistic applications, no efforts were made to optimize the performance through manipulation of code, procedures, compiler parameters, etc. Therefore, the resulting data presented should reflect viable performance results that may be compared to determine performance ratios between the two processors. The Pentium results, when compared with SPECMARK values published by the chip vendor, indicate that our 60 MHz version is running at approximately unity IPC--the vendor results being about twice our values, indicating that their code was optimized for IPC=2.

Tables 2 and **3** also include, as a courtesy, the additional results obtained for the IBM RISC 6000-730 (running AIX version 3.1).

9.0 Observations and Discussions of Results

The performance of the Pentium microprocessor is clearly superior to that offered by its 386 predecessor. The following observations of the measurements and evaluations were made:

- The performance ratio between the 20 MHz 386 and the 60 MHz 586 was found to be approximately 10:1 for the non-optimized (IPC=1) case, running the station-mix;
- Pentium performance can achieve an additional 2:1 boost over the above ratio through optimization of the parallel flow (IPC=2),

Table 2. SPEC Benchmark Comparisons - Integer Suite (SPECint92)

Benchmark	SPEC Reference Time (Sec)	RS/6000-730		PS/2-80 (386) (20 MHz)		PS/2-95 (PENTIUM) (60 MHz)	
		Elapsed Time (Sec)	SPECratio	Elapsed Time (Sec)	SPECratio	Elapsed Time (Sec)	SPECratio
008.espresso	2270	110.0	20.64	1261.3	1.80	64.1	35.41
022.li	6210	292.2	21.25	2328.3	2.67	127.3	48.78
023.eqntott	1100	42.4	25.94	719.4	1.53	46.6	23.61
026.compress	2770	130.7	21.19	1540.0	1.80	140.0	19.79
072.sc	4530	218.5	20.73	2060.8	2.20	130.5	34.71
Geom Mean: SPECint92			21.31		1.96		30.87

Reference SPEC time= Execution time measured in a VAX 11-780 processor

SPECratio = Execution time of system being measured divided by the Reference SPEC time

SPECint92 = Geometric mean (nth root of the product) of the individual values in the SPECratio column

Table 3. SPEC Benchmark Comparisons - Floating Point Suite (SPECfp92)

Benchmark	SPEC Reference Time (Sec)	RS/6000-730		PS/2-80 (386) (20 MHz)		PS/-95 (PENTIUM) (60 MHz)	
		Elapsed Time (Sec)	SPECratio	Elapsed Time (Sec)	SPECratio	Elapsed Time (Sec)	SPECratio
013.spice2g6	24000	815.1	29.44	13352.1	1.80	849.0	28.27
015.doduc	1860	61.9	30.05	1735.0	1.07	58.7	31.69
034.mdljdp2	7090	157.3	45.07	8735.7	0.81	212.5	33.36
047.tomcatv	2650	32.0	82.81	2963.5	0.89	134.6	19.69
048.ora	7420	183.9	40.35	5118.0	1.45	152.8	48.56
052.alvinn	7690	110.2	69.78	5626.2	1.37	164.2	46.83
056.ear	25500	463.0	55.08	21388.8	1.19	508.4	50.16
077.mdljdp2	3350	172.4	19.43	10539.8	0.32	172.0	19.48
078.swm256	12700	627.3	20.25	16647.7	0.76	687.6	18.47
089.su2cor	12900	159.8	80.73	7995.9	1.61	460.5	28.01
090.hydro2d	13700	284.6	48.14	15882.2	0.86	541.0	25.32
093.nasa7	16800	206.5	81.36	18582.2	0.90	1115.8	15.06
094.fppp	9200	188.1	48.91	8015.9	1.15	167.6	54.89
Geom Mean: SPECfp92			43.18		1.01		29.73

Reference SPEC time= Execution time measured in a VAX 11-780 processor

SPECratio = Execution time of system being measured divided by the Reference SPEC time

SPECfp92 = Geometric mean (nth root of the product) of the individual values in the SPECratio column

- When clock rates are equalized for 386 and 586, the performance ratio running the station-mix (with IPC=1) remains at approximately 4:1 (12 MIPS versus 45 MIPS for both at 60 MHz); with IPC=2, this ratio grows back to 10:1 (12 MIPS vs. 118 MIPS).

10.0 Conclusions, and Recommendations (SOW Subtask 6)

An upgrade of the Space Station 386-EDP-16 to the Pentium based 586-EDP-16P is deemed to be feasible, based on the results of our evaluations. The 40 MIPS goal (for station-mix) appears to be adequately achievable with a 33 1/3 MHz clock--allowing for 20 percent fill of the parallel paths (IPC=1.2). This EDP-16P configuration should then yield approximately 17 MIPS for both SPECint92 and SPECfp92, at IPC=1, and approximately 34 MIPS peak for both suites with IPC=2.

The following recommendations for additional analysis, evaluation and testing are included for your consideration:

- The actual flight implementation of the EDP-16P should be made with the low-power P54C version of the Pentium; lower power dissipation results from the smaller geometry, and on-chip power switching permits near quiescence,
- A parallel EDAC implementation should be considered for the EDP-16P design to eliminate the in-line time delay between main memory and the on-chip level one cache (to permit lower cost/slower main memory chips). The analysis and results presented in this work assumed that the EDAC and memory cycle time was short enough to present zero wait conditions to the CPU,
- The dual chip capability of the Pentium architecture should be investigated for operation in both fault-tolerant modes and in radiation-tolerant modes. The master/slave cross compare error

detection capability permits near 100 % coverage of soft/transient errors,

- Additional measurements of the EDP-16P are needed to isolate and determine the remainder of the hardware and software performance parameters, as indicated in *Figure 1*,
- The superstate for power management should be investigated for utilization in redundancy management and fault-tolerance applications; additional software residing at this level can perform operations transparent to the conventional application software running,
- Enhanced performance monitoring and control (in real-time) can be achieved with the Pentium, to provide advanced AGE capabilities; in addition, two special signals (IU and IV) indicate the progress through the integer pipelines--potentially useful for dynamically controlling the instruction stream for achievement of maximum throughput.

11.0 References

1. Performance Analysis and Kernel Size Study of the Lynx real-time Operating System; Liu, Gibson, & Fernquist; Computing-in-Aerospace 9; San Diego, October 1993.
2. Operation of Commercially-Based Microprocessor Technology in a Space Radiation Environment; Yelverton; Computing-in-Aerospace 9; San Diego, October 1993.
3. Technique and Applications of Device Driver Dynamic Loading; Liu, Gibson, & Fernquist; Computing-in-Aerospace 9; San Diego, October 1993.

4. P5 Engineering Data Books--Architecture and Instruction Set; Intel Corporation; 1992
5. EDP-16P Design Notes; Hai Vu; IBM Federal Systems, Houston; November 1993
6. IBM PS/2 Pentium Design Notes; IBM Boca Raton; 1992
7. Performance Calculation Notes; Yelverton; IBM Federal Systems; Houston, 1993.
8. Performance Measurements Notes; Alari; IBM Federal Systems; Houston, 1993.