

2

RE

AD-A242 283

PAGE

Form Approved
OPM No. 0704-0188



Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and reviewing the collection of information, including suggestions for reducing this burden, to Washington Headquarters Service, Directorate for Information Operations and Reports, Washington, DC 20503.

Save time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and reviewing the collection of information, including suggestions for reducing this burden, to Washington Headquarters Service, Directorate for Information Operations and Reports, Washington, DC 20503.

1. AGENCY USE ONLY (Leave Blank)		2. REPORT DATE		3. REPORT TYPE AND DATES COVERED Final: 09 Aug 1991 to 01 Jun 1993	
4. TITLE AND SUBTITLE Alsys, Inc., AlsyCOMP_024, Version 5.3, IBM RISC System 6000, Model 520 under IBM AIX V3.1 (Host & Target), 910809W1.11195				5. FUNDING NUMBERS	
6. AUTHOR(S) Wright-Patterson AFB, Dayton, OH USA				7. PERFORMING ORGANIZATION REPORT NUMBER AVF-VSR-499-0891	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Ada Validation Facility, Language Control Facility ASD/SCEL Bldg. 676, Rm 135 Wright-Patterson AFB, Dayton, OH 45433				8. PERFORMING ORGANIZATION REPORT NUMBER AVF-VSR-499-0891	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Ada Joint Program Office United States Department of Defense Pentagon, Rm 3E114 Washington, D.C. 20301-3081				10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES <i>No software available for distribution per Michelle Rice, ADA file on 11/4/91 ngn</i>					
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited.				12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) Alsys, Inc., AlsyCOMP_024, Version 5.3, Wright-Patterson AFB, OH, IBM RISC System 6000, Model 520 under IBM AIX V3.1 (Host & Target), ACVC 1.11.					
91-15063					
14. SUBJECT TERMS Ada programming language, Ada Compiler Val. Summary Report, Ada Compiler Val. Capability, Val. Testing, Ada Val. Office, Ada Val. Facility, ANSI/MIL-STD-1815A, AJPO.				15. NUMBER OF PAGES	
				16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED		18. SECURITY CLASSIFICATION UNCLASSIFIED		19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	
20. LIMITATION OF ABSTRACT					

02 1104 134

AVF Control Number: AVF-VSR-499-0891
26 August 1991
91-04-22-ALS

Ada COMPILER
VALIDATION SUMMARY REPORT:
Certificate Number: 910809W1.11195
Alsys, Inc.
AlsyCOMP 024, version 5.3
IBM RISC System 6000, Model 520 under IBM AIX v3.1 =>
IBM RISC System 6000, Model 520 under IBM AIX v3.1

Prepared By:
Ada Validation Facility
ASD/SCEL
Wright-Patterson AFB OH 45433-6503

Accession For	
NO. GRADE	5
DATE	
DATE ISSUED	
DATE RECEIVED	
By	
Title	
Activity Code	
Special	
A-1	

Certificate Information

The following Ada implementation was tested and determined to pass ACVC 1.11. Testing was completed on 9 August 1991.

Compiler Name and Version: AlsyCOMP_024, version 5.3

Host Computer System: IBM RISC System 6000, Model 520
under IBM AIX v3.1


Target Computer System: IBM RISC System 6000, Model 520
under IBM AIX v3.1

Customer Agreement Number: 91-04-22-ALS


See section 3.1 for any additional information about the testing environment.

As a result of this validation effort, Validation Certificate 910809w1.11195 is awarded to Alsys, Inc. This certificate expires on 1 June 1993.


This report has been reviewed and is approved.



Ada Validation Facility
Steven P. Wilson
Technical Director
ASD/SCEL
Wright-Patterson AFB OH 45433-6503



Ada Validation Organization
Director, Computer and Software Engineering Division
Institute for Defense Analyses
Alexandria VA 22311



Ada Joint Program Office
Dr. John Solomond, Director
Department of Defense
Washington DC 20301

DECLARATION OF CONFORMANCE

CUSTOMER;
ADA VALIDATION FACILITY: Alsys, Inc.
 Ada Validation Facility (ASD/SCEL)
 Computer Operations Division
 Information Systems and Technology Center
 Wright-Patterson AFB OH 45433-6503

ACVC VERSION: 1.11

ADA IMPLEMENTATION:

COMPILER NAME AND VERSION: ALSYCOMP_024, version 5.3

HOST COMPUTER SYSTEM: IBM RISC System 6000, model 520
 under IBM AIX v3.1

TARGET COMPUTER SYSTEM: IBM RISC System 6000, model 520
 under IBM AIX v3.1

CUSTOMER'S DECLARATION

I, the undersigned, representing Alsys, Inc., declare that Alsys, Inc. has no knowledge of deliberate deviations from the Ada Language Standard ANSI/MIL-STD-1815A in the implementation listed in this declaration.

George Romanski.

Date: 5 AUGUST 1991

George Romanski
Vice President, Engineering
Alsys, Inc.
67 South Bedford Street
Burlington, MA 01803-5152

TABLE OF CONTENTS

CHAPTER 1	INTRODUCTION	
1.1	USE OF THIS VALIDATION SUMMARY REPORT	1-1
1.2	REFERENCES.	1-2
1.3	ACVC TEST CLASSES	1-2
1.4	DEFINITION OF TERMS	1-3
CHAPTER 2	IMPLEMENTATION DEPENDENCIES	
2.1	WITHDRAWN TESTS	2-1
2.2	INAPPLICABLE TESTS.	2-1
2.3	TEST MODIFICATIONS.	2-4
CHAPTER 3	PROCESSING INFORMATION	
3.1	TESTING ENVIRONMENT	3-1
3.2	SUMMARY OF TEST RESULTS	3-1
3.3	TEST EXECUTION.	3-2
APPENDIX A	MACRO PARAMETERS	
APPENDIX B	COMPILATION SYSTEM OPTIONS	
APPENDIX C	APPENDIX F OF THE Ada STANDARD	

CHAPTER 1

INTRODUCTION

The Ada implementation described above was tested according to the Ada Validation Procedures [Pro90] against the Ada Standard [Ada83] using the current Ada Compiler Validation Capability (ACVC). This Validation Summary Report (VSR) gives an account of the testing of this Ada implementation. For any technical terms used in this report, the reader is referred to [Pro90]. A detailed description of the ACVC may be found in the current ACVC User's Guide [UG89].

1.1 USE OF THIS VALIDATION SUMMARY REPORT

Consistent with the national laws of the originating country, the Ada Certification Body may make full and free public disclosure of this report. In the United States, this is provided in accordance with the "Freedom of Information Act" (5 U.S.C. #552). The results of this validation apply only to the computers, operating systems, and compiler versions identified in this report.

The organizations represented on the signature page of this report do not represent or warrant that all statements set forth in this report are accurate and complete, or that the subject implementation has no nonconformities to the Ada Standard other than those presented. Copies of this report are available to the public from the AVF which performed this validation or from:

National Technical Information Service
5285 Port Royal Road
Springfield VA 22161

Questions regarding this report or the validation test results should be directed to the AVF which performed this validation or to:

Ada Validation Organization
Computer and Software Engineering Division
Institute for Defense Analyses
1801 North Beauregard Street
Alexandria VA 22311-1772

INTRODUCTION

1.2 REFERENCES

Reference Manual for the Ada Programming Language, [Ada83]
ANSI/MIL-STD-1815A, February 1983 and ISO 8652-1987.

Ada Compiler Validation Procedures, Version 2.1, [Pro90]
Ada Joint Program Office, August 1990.

Ada Compiler Validation Capability User's Guide, [UG89] 21 June 1989.

1.3 ACVC TEST CLASSES

Compliance of Ada implementations is tested by means of the ACVC. The ACVC contains a collection of test programs structured into six test classes: A, B, C, D, E, and L. The first letter of a test name identifies the class to which it belongs. Class A, C, D, and E tests are executable. Class B and class L tests are expected to produce errors at compile time and link time, respectively.

The executable tests are written in a self-checking manner and produce a PASSED, FAILED, or NOT APPLICABLE message indicating the result when they are executed. Three Ada library units, the packages REPORT and SPRT13, and the procedure CHECK FILE are used for this purpose. The package REPORT also provides a set of Identity functions used to defeat some compiler optimizations allowed by the Ada Standard that would circumvent a test objective. The package SPRT13 is used by many tests for Chapter 13 of the Ada Standard. The procedure CHECK FILE is used to check the contents of text files written by some of the Class C tests for Chapter 14 of the Ada Standard. The operation of REPORT and CHECK FILE is checked by a set of executable tests. If these units are not operating correctly, validation testing is discontinued.

Class B tests check that a compiler detects illegal language usage. Class B tests are not executable. Each test in this class is compiled and the resulting compilation listing is examined to verify that all violations of the Ada Standard are detected. Some of the class B tests contain legal Ada code which must not be flagged illegal by the compiler. This behavior is also verified.

Class L tests check that an Ada implementation correctly detects violation of the Ada Standard involving multiple, separately compiled units. Errors are expected at link time, and execution is attempted.

In some tests of the ACVC, certain macro strings have to be replaced by implementation-specific values — for example, the largest integer. A list of the values used for this implementation is provided in Appendix A. In addition to these anticipated test modifications, additional changes may be required to remove unforeseen conflicts between the tests and implementation-dependent characteristics. The modifications required for this implementation are described in section 2.3.

For each Ada implementation, a customized test suite is produced by the AVF. This customization consists of making the modifications described in the preceding paragraph, removing withdrawn tests (see section 2.1), and possibly removing some inapplicable tests (see section 2.2 and [UG89]).

In order to pass an ACVC an Ada implementation must process each test of the customized test suite according to the Ada Standard.

1.4 DEFINITION OF TERMS

Ada Compiler	The software and any needed hardware that have to be added to a given host and target computer system to allow transformation of Ada programs into executable form and execution thereof.
Ada Compiler Validation Capability (ACVC)	The means for testing compliance of Ada implementations, consisting of the test suite, the support programs, the ACVC user's guide and the template for the validation summary report.
Ada Implementation	An Ada compiler with its host computer system and its target computer system.
Ada Joint Program Office (AJPO)	The part of the certification body which provides policy and guidance for the Ada certification system.
Ada Validation Facility (AVF)	The part of the certification body which carries out the procedures required to establish the compliance of an Ada implementation.
Ada Validation Organization (AVO)	The part of the certification body that provides technical guidance for operations of the Ada certification system.
Compliance of an Ada Implementation	The ability of the implementation to pass an ACVC version.
Computer System	A functional unit, consisting of one or more computers and associated software, that uses common storage for all or part of a program and also for all or part of the data necessary for the execution of the program; executes user-written or user-designated programs; performs user-designated data manipulation, including arithmetic operations and logic operations; and that can execute programs that modify themselves during execution. A computer system may be a stand-alone unit or may consist of several inter-connected units.

INTRODUCTION

Conformity	Fulfillment by a product, process, or service of all requirements specified.
Customer	An individual or corporate entity who enters into an agreement with an AVF which specifies the terms and conditions for AVF services (of any kind) to be performed.
Declaration of Conformance	A formal statement from a customer assuring that conformity is realized or attainable on the Ada implementation for which validation status is realized.
Host Computer System	A computer system where Ada source programs are transformed into executable form.
Inapplicable test	A test that contains one or more test objectives found to be irrelevant for the given Ada implementation.
ISO	International Organization for Standardization.
LRM	The Ada standard, or Language Reference Manual, published as ANSI/MIL-STD-1815A-1983 and ISO 8652-1987. Citations from the LRM take the form "<section>.<subsection>:<paragraph>."
Operating System	Software that controls the execution of programs and that provides services such as resource allocation, scheduling, input/output control, and data management. Usually, operating systems are predominantly software, but partial or complete hardware implementations are possible.
Target Computer System	A computer system where the executable form of Ada programs are executed.
Validated Ada Compiler	The compiler of a validated Ada implementation.
Validated Ada Implementation	An Ada implementation that has been validated successfully either by AVF testing or by registration [Pro90].
Validation	The process of checking the conformity of an Ada compiler to the Ada programming language and of issuing a certificate for this implementation.
Withdrawn test	A test found to be incorrect and not used in conformity testing. A test may be incorrect because it has an invalid test objective, fails to meet its test objective, or contains erroneous or illegal use of the Ada programming language.

CHAPTER 2

IMPLEMENTATION DEPENDENCIES

2.1 WITHDRAWN TESTS

The following tests have been withdrawn by the AVO. The rationale for withdrawing each test is available from either the AVO or the AVF. The publication date for this list of withdrawn tests is 2 August 1991.

E28005C	B28006C	C32203A	C34006D	C35508I	C35508J
C35508M	C35508N	C35702A	C35702B	B41308B	C43004A
C45114A	C45346A	C45612A	C45612B	C45612C	C45651A
C46022A	B49008A	B49008B	A74006A	C74308A	B83022B
B83022H	B83025B	B83025D	C83026A	B83026B	C83041A
B85001L	C86001F	C94021A	C97116A	C98003B	BA2011A
CB7001A	CB7001B	CB7004A	CC1223A	BC1226A	CC1226B
BC3009B	BD1B02B	BD1B06A	AD1B08A	BD2A02A	CD2A21E
CD2A23E	CD2A32A	CD2A41A	CD2A41E	CD2A87A	CD2B15C
BD3006A	BD4008A	CD4022A	CD4022D	CD4024B	CD4024C
CD4024D	CD4031A	CD4051D	CD5111A	CD7004C	ED7005D
CD7005E	AD7006A	CD7006E	AD7201A	AD7201E	CD7204B
AD7206A	BD8002A	BD8004C	CD9005A	CD9005B	CDA201E
CE2107I	CE2117A	CE2117B	CE2119B	CE2205B	CE2405A
CE3111C	CE3116A	CE3118A	CE3411B	CE3412B	CE3607B
CE3607C	CE3607D	CE3812A	CE3814A	CE3902B	

2.2 INAPPLICABLE TESTS

A test is inapplicable if it contains test objectives which are irrelevant for a given Ada implementation. Reasons for a test's inapplicability may be supported by documents issued by the ISO and the AJPO known as Ada Commentaries and commonly referenced in the format AI-ddddd. For this implementation, the following tests were determined to be inapplicable for the reasons indicated; references to Ada Commentaries are included as appropriate.

IMPLEMENTATION DEPENDENCIES

The following 201 tests have floating-point type declarations requiring more digits than `SYSTEM.MAX_DIGITS`:

C24113L..Y (14 tests)	C35705L..Y (14 tests)
C35706L..Y (14 tests)	C35707L..Y (14 tests)
C35708L..Y (14 tests)	C35802L..Z (15 tests)
C45241L..Y (14 tests)	C45321L..Y (14 tests)
C45421L..Y (14 tests)	C45521L..Z (15 tests)
C45524L..Z (15 tests)	C45621L..Z (15 tests)
C45641L..Y (14 tests)	C46012L..Z (15 tests)

The following 20 tests check for the predefined type `LONG_INTEGER`; for this implementation, there is no such type:

C35404C	C45231C	C45304C	C45411C	C45412C
C45502C	C45503C	C45504C	C45504F	C45611C
C45613C	C45614C	C45631C	C45632C	B52004D
C55B07A	B55B09C	B86001W	C86006C	CD7101F

C35713D and B86001Z check for a predefined floating-point type with a name other than `FLOAT`, `LONG_FLOAT`, or `SHORT_FLOAT`; for this implementation, there is no such type.

C45531M..P and C45532M..P (8 tests) check fixed-point operations for types that require a `SYSTEM.MAX_MANTISSA` of 47 or greater; for this implementation, `MAX_MANTISSA` is less than 47.

C45536A, C46013B, C46031B, C46033B, and C46034B contain length clauses that specify values for `'SMALL` that are not powers of two or ten; this implementation does not support such values for `'SMALL`.

C45624A..B (2 tests) check that the proper exception is raised if `MACHINE_OVERFLOW` is `FALSE` for floating point types and the results of various floating-point operations lie outside the range of the base type; for this implementation, `MACHINE_OVERFLOW` is `TRUE`.

B86001Y uses the name of a predefined fixed-point type other than type `DURATION`; for this implementation, there is no such type.

C96005B uses values of type `DURATION`'s base type that are outside the range of type `DURATION`; for this implementation, the ranges are the same.

CD1009C checks whether a length clause can specify a non-default size for a floating-point type; this implementation does not support such sizes.

CD2A53A checks operations of a fixed-point type for which a length clause specifies a power-of-ten `TYPE'SMALL`; this implementation does not support decimal `'SMALLs`. (See section 2.3.)

IMPLEMENTATION DEPENDENCIES

CD2A84A, CD2A84E, CD2A84I..J (2 tests), and CD2A84O use length clauses to specify non-default sizes for access types; this implementation does not support such sizes.

BD8001A, BD8003A, BD8004A..B (2 tests), and AD8011A use machine code insertions; this implementation provides no package MACHINE_CODE.

CE2401H, EE2401D, and EE2401G use instantiations of DIRECT_IO with unconstrained array and record types; this implementation raises USE_ERROR on the attempt to create a file of such types.

The tests listed in the following table check that USE_ERROR is raised if the given file operations are not supported for the given combination of mode and access method; this implementation supports these operations.

Test	File Operation	Mode	File Access Method
CE2102E	CREATE	OUT FILE	SEQUENTIAL IO
CE2102F	CREATE	INOUT FILE	DIRECT IO
CE2102J	CREATE	OUT FILE	DIRECT IO
CE2102N	OPEN	IN FILE	SEQUENTIAL IO
CE2102O	RESET	IN FILE	SEQUENTIAL IO
CE2102P	OPEN	OUT FILE	SEQUENTIAL IO
CE2102Q	RESET	OUT FILE	SEQUENTIAL IO
CE2102R	OPEN	INOUT FILE	DIRECT IO
CE2102S	RESET	INOUT FILE	DIRECT IO
CE2102T	OPEN	IN FILE	DIRECT IO
CE2102U	RESET	IN FILE	DIRECT IO
CE2102V	OPEN	OUT FILE	DIRECT IO
CE2102W	RESET	OUT FILE	DIRECT IO
CE3102F	RESET	Any Mode	TEXT IO
CE3102G	DELETE	-----	TEXT IO
CE3102I	CREATE	OUT FILE	TEXT IO
CE3102J	OPEN	IN FILE	TEXT IO
CE3102K	OPEN	OUT FILE	TEXT IO

The tests listed in the following table check the given file operations for the given combination of mode and access method; this implementation does not support these operations.

Test	File Operation	Mode	File Access Method
CE2105A	CREATE	IN FILE	SEQUENTIAL IO
CE2105B	CREATE	IN FILE	DIRECT IO
CE3109A	CREATE	IN FILE	TEXT IO

IMPLEMENTATION DEPENDENCIES

CE2203A checks that WRITE raises USE_ERROR if the capacity of an external sequential file is exceeded; this implementation cannot restrict file capacity.

CE2403A checks that WRITE raises USE_ERROR if the capacity of an external direct file is exceeded; this implementation cannot restrict file capacity.

CE3304A checks that SET_LINE_LENGTH and SET_PAGE_LENGTH raise USE_ERROR if they specify an inappropriate value for the external file; there are no inappropriate values for this implementation.

CE3413B checks that PAGE raises LAYOUT_ERROR when the value of the page number exceeds COUNT'LAST; for this implementation, the value of COUNT'LAST is greater than 150000, making the checking of this objective impractical.

CE3202A expects that function NAME can be applied to the standard input and output files; in this implementation these files have no names, and USE_ERROR is raised. (See section 2.3.)

2.3 TEST MODIFICATIONS

Modifications (see section 1.3) were required for 20 tests.

The following tests were split into two or more tests because this implementation did not report the violations of the Ada Standard in the way expected by the original tests.

B23004A	B24007A	B24009A	B28003A	B32202A	B32202B
B32202C	B37004A	B61012A	B91004A	B95069A	B95069B
B97103E	BA1101B	BC2001D	BC3009C		

BA2001E was graded passed by Evaluation Modification as directed by the AVO. The test expects that duplicate names of subunits with a common ancestor will be detected as compilation errors; this implementation detects the errors at link time, and the AVO ruled that this behavior is acceptable.

EA3004D was graded passed by Evaluation and Processing Modification as directed by the AVO. The test requires that either pragma INLINE is obeyed for a function call in each of three contexts and that thus three library units are made obsolete by the re-compilation of the inlined function's body, or else the pragma is ignored completely. This implementation obeys the pragma except when the call is within the package specification. When the test's files are processed in the given order, only two units are made obsolete; thus, the expected error at line 27 of file EA3004D6M is not valid and is not flagged. To confirm that indeed the pragma is not obeyed in this one case, the test was also processed with the files re-ordered so that the re-compilation follows only the

IMPLEMENTATION DEPENDENCIES

package declaration (and thus the other library units will not be made obsolete, as they are compiled later); a "NOT APPLICABLE" result was produced, as expected. The revised order of files was 0-1-4-5-2-3-6.

When run as is, the implementation fails to detect an error on line 27 of test file EA3004D6M. This is because the pragma INLINE has no effect when its object is within a package specification. However, the results of running the test as-is do not confirm that the pragma had no effect, only that the package was not made obsolete. By re-ordering the compilations so that the two subprograms are compiled after file D5 (the re-compilation of the "with"ed package that makes the various earlier units obsolete), we create a test that shows that indeed pragma INLINE has no effect when applied to a subprogram that is called within a package specification. The test must execute and produce the expected NOT APPLICABLE result (as though INLINE were not supported at all). The recommended re-ordering of EA3004D test files is: 0-1-4-5-2-3-6.

CD2A53A was graded inapplicable by Evaluation Modification as directed by the AVO. The test contains a specification of a power-of-10 value as 'SMALL for a fixed-point type. The AVO ruled that, under ACVC 1.11, support of decimal 'SMALLs may be omitted.

CE3202A was graded inapplicable by Evaluation Modification as directed by the AVO. This test applies function NAME to the standard input file, which in this implementation has no name; USE ERROR is raised but not handled, so the test is aborted. The AVO ruled that this behavior is acceptable pending any resolution of the issue by the ARG.

CHAPTER 3

PROCESSING INFORMATION

3.1 TESTING ENVIRONMENT

The Ada implementation tested in this validation effort is described adequately by the information given in the initial pages of this report. For technical and sales information about this Ada implementation, contact:

George Romanski, Vice-President, Engineering
Alsys, Inc.
67 South Bedford Street
Burlington, MA 01803-5152
(617) 270-0030

Testing of this Ada implementation was conducted at the customer's site by a validation team from the AVF.

3.2 SUMMARY OF TEST RESULTS

An Ada Implementation passes a given ACVC version if it processes each test of the customized test suite in accordance with the Ada Programming Language Standard, whether the test is applicable or inapplicable; otherwise, the Ada Implementation fails the ACVC [Pro90].

For all processed tests (inapplicable and applicable), a result was obtained that conforms to the Ada Programming Language Standard.

PROCESSING INFORMATION

The list of items below gives the number of ACVC tests in various categories. All tests were processed, except those that were withdrawn because of test errors (item b; see section 2.1), those that require a floating-point precision that exceeds the implementation's maximum precision (item e; see section 2.2), and those that depend on the support of a file system — if none is supported (item d). All tests passed, except those that are listed in sections 2.1 and 2.2 (counted in items b and f, below).

a) Total Number of Applicable Tests	3794
b) Total Number of Withdrawn Tests	95
c) Processed Inapplicable Tests	80
d) Non-Processed I/O Tests	0
e) Non-Processed Floating-Point Precision Tests	201
f) Total Number of Inapplicable Tests	281
g) Total Number of Tests for ACVC 1.11	4170

3.3 TEST EXECUTION

A magnetic tape containing the customized test suite (see section 1.3) was taken on-site by the validation team for processing. The contents of the magnetic tape were loaded directly onto a VAX system (from a standard 1/2 inch, 9-track VMS BACKUP tape) and then transferred to the IBM RISC System 6000 via ethernet.

After the test files were loaded onto the host computer, the full set of tests was processed by the Ada implementation.

The tests were compiled, linked, and executed on the host computer system, as appropriate. The results were captured on the computer system.

Testing was performed using command scripts provided by the customer and reviewed by the validation team. See Appendix B for a complete listing of the processing options for this implementation. It also indicates the default options. The options invoked explicitly for validation testing during this test were:

Compiler Options:

TEXT => NO	Do not show source code in listing (used for all but the B tests)
TEXT => YES	Show source code in listing (used for the B tests)

PROCESSING INFORMATION

SHOW => NO	Do not show header nor error summary in listing.
WARNING => NO	Do not include warning messages.
ERRORS => 999	Maximum number of compilation errors permitted before terminating the compilation.
CALLS => INLINED	This option allows insertion of code for subprograms inline and must be set for the pragma <code>INLINE</code> to be operative.
GENERICCS => STUB	This option places code of generic instantiations in separate subunits.

The default options were used for all binds.

Test output, compiler and linker listings, and job logs were captured on magnetic tape and archived at the AVF. The listings examined on-site by the validation team were also archived.

APPENDIX A

MACRO PARAMETERS

This appendix contains the macro parameters used for customizing the ACVC. The meaning and purpose of these parameters are explained in [UG89]. The parameter values are presented in two tables. The first table lists the values that are defined in terms of the maximum input-line length, which is the value for \$MAX_IN_LEN--also listed here. These values are expressed here as Ada string aggregates, where "V" represents the maximum input-line length.

Macro Parameter	Macro Value
\$MAX_IN_LEN	255
\$BIG_ID1	(1..254 => 'A', 255 => '1')
\$BIG_ID2	(1..254 => 'A', 255 => '2')
\$BIG_ID3	(1..127 => 'A') & '3' & (1..127 => 'A')
\$BIG_ID4	(1..127 => 'A') & '4' & (1..127 => 'A')
\$BIG_INT_LIT	(1..252 => '0') & "298"
\$BIG_REAL_LIT	(1..250 => '0') & "690.0"
\$BIG_STRING1	"" & (1..127 => 'A') & ""
\$BIG_STRING2	"" & (1..127 => 'A') & '1' & ""
\$BLANKS	(1..235 => ' ')
\$MAX_LEN_INT_BASED_LITERAL	"2:" & (1..250 => '0') & "11:"
\$MAX_LEN_REAL_BASED_LITERAL	"16:" & (1..248 => '0') & "F.E:"
\$MAX_STRING_LITERAL	"" & (1..253 => 'A') & ""

MACRO PARAMETERS

The following table lists all of the other macro parameters and their respective values.

Macro Parameter	Macro Value
\$ACC_SIZE	32
\$ALIGNMENT	4
\$COUNT_LAST	2147483647
\$DEFAULT_MEM_SIZE	2**32
\$DEFAULT_STOR_UNIT	8
\$DEFAULT_SYS_NAME	RS_6000
\$DELTA_DOC	2#1.0#E-31
\$ENTRY_ADDRESS	FCNDECL.ENTRY_ADDRESS
\$ENTRY_ADDRESS1	FCNDECL.ENTRY_ADDRESS1
\$ENTRY_ADDRESS2	FCNDECL.ENTRY_ADDRESS2
\$FIELD_LAST	255
\$FILE_TERMINATOR	' '
\$FIXED_NAME	NO_SUCH_FIXED_TYPE
\$FLOAT_NAME	NO_SUCH_FLOAT_TYPE
\$FORM_STRING	""
\$FORM_STRING2	"CANNOT_RESTRICT_FILE_CAPACITY"
\$GREATER_THAN_DURATION	75_000.0
\$GREATER_THAN_DURATION BASE LAST	131_073.0
\$GREATER_THAN_FLOAT_BASE LAST	1.0E+39
\$GREATER_THAN_FLOAT_SAFE LARGE	1.0E+38

MACRO PARAMETERS

\$GREATER_THAN_SHORT_FLOAT_SAFE_LARGE
 1.0E+38

 \$HIGH_PRIORITY 10

 \$ILLEGAL_EXTERNAL_FILE_NAME1
 /this/is/a/junk/directory/name

 \$ILLEGAL_EXTERNAL_FILE_NAME2
 THIS-FILE-NAME-IS-TOO-
 LONG-FOR-MY-SYSTEM

 \$INAPPROPRIATE_LINE_LENGTH
 -1

 \$INAPPROPRIATE_PAGE_LENGTH
 -1

 \$INCLUDE_PRAGMA1 PRAGMA INCLUDE ("A28006D1.TST")
 \$INCLUDE_PRAGMA2 PRAGMA INCLUDE ("B28006D1.TST")

 \$INTEGER_FIRST -2147483648

 \$INTEGER_LAST 2147483647

 \$INTEGER_LAST_PLUS_1 2147483648

 \$INTERFACE_LANGUAGE C

 \$LESS_THAN_DURATION -75_000.0

 \$LESS_THAN_DURATION_BASE_FIRST
 -131_073.0

 \$LINE_TERMINATOR ASCII.LF

 \$LOW_PRIORITY 1

 \$MACHINE_CODE_STATEMENT
 NULL;

 \$MACHINE_CODE_TYPE NO_SUCH_TYPE

 \$MANTISSA_DOC 31

 \$MAX_DIGITS 15

 \$MAX_INT 2147483647

 \$MAX_INT_PLUS_1 2_147_483_648

 \$MIN_INT -2147483648

MACRO PARAMETERS

\$NAME	SHORT_SHORT_INTEGER
\$NAME_LIST	S370,I80X86,I80386,MC680X0,VAX, TRANSPUTER,RS_6000,MIPS
\$NAME_SPECIFICATION1	/alsys/acvc/X2120A
\$NAME_SPECIFICATION2	/alsys/acvc/X2120B
\$NAME_SPECIFICATION3	/alsys/acvc/X3119A
\$NEG_BASED_INT	16#F000000E#
\$NEW_MEM_SIZE	2**32
\$NEW_STOR_UNIT	16
\$NEW_SYS_NAME	RS_6000
\$PAGE_TERMINATOR	ASCII.LF & ASCII.FF
\$RECORD_DEFINITION	NEW_INTEGER
\$RECORD_NAME	NO_SUCH_MACHINE_CODE_TYPE
\$TASK_SIZE	32
\$TASK_STORAGE_SIZE	1024
\$TICK	0.1
\$VARIABLE_ADDRESS	FCNDECL.VARIABLE_ADDRESS
\$VARIABLE_ADDRESS1	FCNDECL.VARIABLE_ADDRESS1
\$VARIABLE_ADDRESS2	FCNDECL.VARIABLE_ADDRESS2
\$YOUR_PRAGMA	INTERFACE

APPENDIX B

COMPILATION SYSTEM OPTIONS

The compiler options of this Ada implementation, as described in this Appendix, are provided by the customer. Unless specifically noted otherwise, references in this appendix are to compiler documentation and not to this report.

Compiler Options:

```
COMPILE (SOURCE => source_name | INSTANTIATION,
         LIBRARY => library_name,
         OPTIONS =>
           (ANNOTATE => character_string,
            ERRORS => positive_integer,
            LEVEL => PARSE | SEMANTIC | CODE | UPDATE,
            CHECKS => ALL | STACK | NONE,
            GENERICS => STUBS | INLINE,
            FLOAT => OPTIMIZED | IEEE,
            MEMORY => number_of_kbytes),
        DISPLAY =>
           (OUTPUT => SCREEN | NONE | AUTOMATIC | file_name,
            WARNING => YES | NO,
            TEXT => YES | NO,
            SHOW => BANNER | RECAP | ALL | NONE,
            DETAIL => YES | NO,
            ASSEMBLY => CODE | MAP | ALL | NONE),
        ALLOCATION =>
           (CONSTANT => positive_integer,
            GLOBAL => positive_integer),
        IMPROVE =>
           (CALLS => NORMAL | INLINED,
            REDUCTION => NONE | PARTIAL | EXTENSIVE,
            EXPRESSIONS => NONE | PARTIAL | EXTENSIVE,
            OBJECT => NONE | PARTIAL | EXTENSIVE),
        KEEP =>
           (COPY => YES | NO,
            DEBUG => YES | NO,
            TREE => YES | NO));
```

COMPILATION SYSTEM OPTIONS

OPTION/SWITCH	EFFECT
SOURCE	=> source_name INSTANTIATION
LIBRARY	=> library_name
ANNOTATE	=> character_string
ERRORS	=> positive_integer
LEVEL	=> PARSE SEMANTIC CODE UPDATE
CHECKS	=> ALL STACK NONE
GENERIC	=> STUBS INLINE
FLOAT	=> OPTIMIZED IEEE
MEMORY	=> number_of_kbytes
OUTPUT	=> SCREEN NONE AUTOMATIC file_name
WARNING	=> YES NO
TEXT	=> YES NO
SHOW	=> BANNER RECAP ALL NONE
DETAIL	=> YES NO
ASSEMBLY	=> CODE MAP ALL NONE
CONSTANT	=> positive_integer
GLOBAL	=> positive_integer
CALLS	=> NORMAL INLINED
REDUCTION	=> NONE PARTIAL EXTENSIVE
EXPRESSIONS	=> NONE PARTIAL EXTENSIVE
COPY	=> YES NO
DEBUG	=> YES NO
TREE	=> YES NO

LINKER OPTIONS

The linker options of this Ada implementation, as described in this Appendix, are provided by the customer. Unless specifically noted otherwise, references in this appendix are to linker documentation and not to this report.

Binder Options:

```

BIND (PROGRAM => unit_name,
      LIBRARY => library_name,
      OPTIONS =>
        (LEVEL      => CHECK | BIND | LINK,
         OBJECT     => AUTOMATIC | file_name,
         UNCALLED   => REMOVE | KEEP,
         SLICE      => NO | positive_integer,
         BLOCKING   => YES | NO | AUTOMATIC),
      STACK =>
        (MAIN      => positive_integer,
         TASK      => positive_integer,
         HISTORY   => MAIN | TASK | ALL | NONE),
      HEAP  =>
        (SIZE      => positive_integer,
         INCREMENT => natural_integer),
      INTERFACE =>
        (DIRECTIVES => options_for_linker,
         MODULES    => file_names,
         SEARCH     => library_names),
      DISPLAY =>
        (OUTPUT    => SCREEN | NONE | AUTOMATIC | file_name,
         DATA     => BIND | LINK | ALL | NONE,
         WARNING   => YES | NO),
      KEEP  =>
        (DEBUG     => YES | NO));

```

COMPILATION SYSTEM OPTIONS

OPTION/SWITCH		EFFECT
PROGRAM	=> main_program_name	Ada name of main subprogram.
LIBRARY	=> library_name	Name of program library.
LEVEL	=> CHECK BIND LINK	Specify level of bind.
OBJECT	=> AUTOMATIC file_name	Name of object module.
UNCALLED	=> REMOVE KEEP	Removes or keeps uncalled subprograms.
SLICE	=> NO positive_integer	Define time-slice.
BLOCKING	=> YES NO AUTOMATIC)	Does I/O block other tasks?
MAIN	=> positive_integer	Size of main stack.
TASK	=> positive_integer	Size of task stack.
HISTORY	=> YES NO	Generate stack traces for unhandled exceptions?
SIZE	=> positive_integer	Size the initial heap allocation.
INCREMENT	=> positive_integer	Size subsequent heap allocations.
DIRECTIVES	=> options_for_linker	Supplies linker directives.
MODULES	=> file_names	Additional object modules.
SEARCH	=> library_names	Additional object libraries.
OUTPUT	=> SCREEN NONE AUTOMATIC file_name	Location of binder listing.
DATA	=> BIND LINK ALL NONE	Amount of data to be included in binder listings.
WARNING	=> YES NO	Generate warning messages?
DEBUG	=> YES NO	Generate debug information?

APPENDIX C

APPENDIX F OF THE Ada STANDARD

The only allowed implementation dependencies correspond to implementation-dependent pragmas, to certain machine-dependent conventions as mentioned in Chapter 13 of the Ada Standard, and to certain allowed restrictions on representation clauses. The implementation-dependent characteristics of this Ada implementation, as described in this Appendix, are provided by the customer. Unless specifically noted otherwise, references in this Appendix are to compiler documentation and not to this report. Implementation-specific portions of the package STANDARD, which are not a part of Appendix F, are:

package STANDARD is

type INTEGER is range -2_147_483_648..2_147_483_647;

type FLOAT is digits 6 range
-2#1.111_1111_1111_1111_1111_1111#E+127..
2#1.111_1111_1111_1111_1111_1111#E+127;

type DURATION is delta 2#0.000_000_000_000_01# range
-131072.00000..131071.99994;

type SHORT_INTEGER is range -32768..32767;

type SHORT_SHORT_INTEGER is range -128..127;

type SHORT_FLOAT is digits 6 range
-2#1.111_1111_1111_1111_1111_1111#E+127..
2#1.111_1111_1111_1111_1111_1111#E+127;

type LONG_FLOAT is digits 15 range
-2#1.1111_1111_1111_1111_1111_1111_1111_1111_1111_1111_1111_1111_1111_1111_1111#E+1023
..
2#1.1111_1111_1111_1111_1111_1111_1111_1111_1111_1111_1111_1111_1111_1111_1111#E+1023;

...

end STANDARD;

BETA DOCUMENTATION - MAY 1991

**Alsys Ada Software Development Environment
for RS/6000**

APPENDIX F

Version 5

Copyright 1991 by Alsys

All rights reserved. No part of this document may be reproduced in any form or by any means without permission in writing from Alsys.

Printed: May 1991

Alsys reserves the right to make changes in specifications and other information contained in this publication without prior notice. Consult Alsys to determine whether such changes have been made.

Alsys, AdaWorld, AdaProbe, AdaXref, AdaReformat, and AdaMake are registered trademarks of Alsys.

Unix is a registered trademark of AT&T.

386/ix is a registered trademark of Interactive Systems Corporation.

Microsoft is a registered trademark of Microsoft Corporation.

IBM, PC AT and PS/2 are registered trademarks of International Business Machines Corporation.

INTEL is a registered trademark of Intel Corporation.

TABLE OF CONTENTS

APPENDIX F	1
1 Implementation-Dependent Pragmas	3
1.1 INLINE	3
1.2 INTERFACE	3
1.3 INTERFACE_NAME	4
1.4 INDENT	5
1.5 Other Pragmas	5
2 Implementation-Dependent Attributes	7
2.1 P'IS_ARRAY	7
2.2 P'RECORD_DESCRIPTOR, P'ARRAY_DESCRIPTOR	7
2.3 E'EXCEPTION_CODE	7
3 Specification of the package SYSTEM	9
3.1 Specification of the package SYSTEM	9
4 Support for Representation Clauses	15
4.1 Enumeration Types	16
4.1.1 Enumeration Literal Encoding	16
4.1.2 Enumeration Types and Object Sizes	16
4.2 Integer Types	18
4.2.1 Integer Type Representation	18
4.2.2 Integer Type and Object Size	18
4.3 Floating Point Types	20
4.3.1 Floating Point Type Representation	20
4.3.2 Floating Point Type and Object Size	21
4.4 Fixed Point Types	21

	4.4.1 Fixed Point Type Representation	21
	4.4.2 Fixed Point Type and Object Size	22
4.5	Access Types and Collections	24
4.6	Task Types	24
4.7	Array Types	25
	4.7.1 Array Layout and Structure and Pragma PACK	25
	4.7.2 Array Subtype and Object Size	28
4.8	Record Types	29
	4.8.1 Basic Record Structure	29
	4.8.2 Indirect Components	33
	4.8.3 Implicit Components	37
	4.8.4 Size of Record Types and Objects	37
5	Conventions for Implementation-Generated Names	39
6	Address Clauses	41
6.1	Address Clauses for Objects	41
6.2	Address Clauses for Program Units	41
6.3	Address Clauses for Interrupt Entries	42
7	Unchecked Conversions	43
8	Input-Output Packages	45
8.1	Introduction	45
8.2	The FORM Parameter	46
	8.2.1 File Protection	47
	8.2.2 File Sharing	49
	8.2.3 File Structure	50
	8.2.4 Buffering	52
	8.2.5 Appending	52
	8.2.6 Blocking	52
	8.2.7 Terminal Input	53

9	Characteristics of Numeric Types	55
9.1	Integer Types	55
9.2	Floating Point Type Attributes	55
9.3	Attributes of Type DURATION	56
10	Other Implementation-Dependent Characteristics	57
10.1	Characteristics of the Heap	57
10.2	Characteristics of Tasks	57
10.3	Definition of a Main Subprogram	59
10.4	Ordering of Compilation Units	59
11	Limitations	61
11.1	Compiler Limitations	61
11.2	Hardware Related Limitations	61
	INDEX	63

APPENDIX F

Implementation - Dependent Characteristics

This appendix summarizes the implementation-dependent characteristics of the Alsys Ada Software Development Environment for RS/6000. Appendix F is a required part of the *Reference Manual for the Ada Programming Language* (called the RM in this appendix).

The sections of this appendix are as follows:

1. The form, allowed places, and effect of every implementation-dependent pragma.
2. The name and the type of every implementation-dependent attribute.
3. The specification of the package SYSTEM.
4. The description of the representation clauses.
5. The conventions used for any implementation-generated name denoting implementation-dependent components.
6. The interpretation of expressions that appear in address clauses, including those for interrupts.
7. Any restrictions on unchecked conversions.
8. Any implementation-dependent characteristics of the input-output packages.
9. Characteristics of numeric types.
10. Other implementation-dependent characteristics.
11. Compiler limitations.

The name *Alsys Runtime Executive Programs* or simply *Runtime Executive* refers to the runtime library routines provided for all Ada programs. These routines implement the Ada heap, exceptions, tasking control, and other utility functions.

General systems programming notes are given in another document, the *Application Developer's Guide* (for example, parameter passing conventions needed for interface with assembly routines).

Section 1

Implementation-Dependent Pragmas

1.1 INLINE

Pragma `INLINE` is fully supported; however, it is not possible to inline a subprogram in a declarative part.

1.2 INTERFACE

Ada programs can interface with subprograms written in Assembler and other languages through the use of the predefined pragma `INTERFACE` and the implementation-defined pragma `INTERFACE_NAME`.

Pragma `INTERFACE` specifies the name of an interfaced subprogram and the name of the programming language for which parameter passing conventions will be generated. Pragma `INTERFACE` takes the form specified in the RM:

```
pragma INTERFACE (language_name, subprogram_name);
```

where,

- *language_name* is `ASSEMBLER`, `ADA`, or `C`.
- *subprogram_name* is the name used within the Ada program to refer to the interfaced subprogram.

The only language names accepted by pragma `INTERFACE` are `ASSEMBLER`, `ADA` and `C`. The full implementation requirements for writing pragma `INTERFACE` subprograms are described in the *Application Developer's Guide*.

The language name used in the pragma `INTERFACE` does not have to have any relationship to the language actually used to write the interfaced subprogram. It is used only to tell the compiler how to generate subprogram calls; that is, what kind of parameter passing techniques to use. The programmer can interface Ada programs with

subroutines written in any other (compiled) language by understanding the mechanisms used for parameter passing by the Alsys Ada Software Development Environment for RS/6000 and the corresponding mechanisms of the chosen external language.

1.3 INTERFACE_NAME

Pragma `INTERFACE_NAME` associates the name of the interfaced subprogram with the external name of the interfaced subprogram. If pragma `INTERFACE_NAME` is not used, then the two names are assumed to be identical. This pragma takes the form:

```
pragma INTERFACE_NAME (subprogram_name, string_literal);
```

where,

- *subprogram_name* is the name used within the Ada program to refer to the interfaced subprogram.
- *string_literal* is the name by which the interfaced subprogram is referred to at link time.

The pragma `INTERFACE_NAME` is used to identify routines in other languages that are not named with legal Ada identifiers. Ada identifiers can only contain letters, digits, or underscores, whereas the UNIX Linker allows external names to contain other characters, for example, the dollar sign (\$) or commercial at sign (@). These characters can be specified in the *string_literal* argument of the pragma `INTERFACE_NAME`.

The pragma `INTERFACE_NAME` is allowed at the same places of an Ada program as the pragma `INTERFACE`. (Location restrictions can be found in section 13.9 of the RM.) However, the pragma `INTERFACE_NAME` must always occur after the pragma `INTERFACE` declaration for the interfaced subprogram.

The *string_literal* of the pragma `INTERFACE_NAME` is passed through unchanged, including case sensitivity, to the Unix object file. There is no limit to the length of the name.

The user must be aware however, that some tools from other vendors do not fully support the standard object file format and may restrict the length or names of symbols. For example, most Unix debuggers only work with alphanumeric identifier names.

The *Runtime Executive* contains several external identifiers. All such identifiers begin with either the string "ADA_" or the string "ADAS_". Accordingly, names prefixed by "ADA_" or "ADAS_" should be avoided by the user.

If `INTERFACE_NAME` is not used, the default link name for the subprogram is its Ada name converted to all upper case characters.

Example

```
package SAMPLE_DATA is
  function SAMPLE_DEVICE (X: INTEGER) return INTEGER;
  function PROCESS_SAMPLE (X: INTEGER) return INTEGER;
private
  pragma INTERFACE (ASSEMBLER, SAMPLE_DEVICE);
  pragma INTERFACE (ADA, PROCESS_SAMPLE);
  pragma INTERFACE_NAME (SAMPLE_DEVICE, "DEVIO$GET_SAMPLE");
end SAMPLE_DATA;
```

1.4 INDENT

Pragma `INDENT` is only used with *AdaReformat*. *AdaReformat* is the Alsys reformatter which offers the functionalities of a pretty-printer in an Ada environment.

The pragma is placed in the source file and interpreted by the Reformatter. The line

```
pragma INDENT(OFF);
```

causes *AdaReformat* not to modify the source lines after this pragma, while

```
pragma INDENT(ON);
```

causes *AdaReformat* to resume its action after this pragma.

1.5 Other Pragmas

Pragmas `IMPROVE` and `PACK` are discussed in detail in the section on representation clauses and records (Chapter 4).

Pragma `PRIORITY` on Unix systems is accepted with the range of priorities running from 1 to 10 (see the definition of the predefined package `SYSTEM` in Section 3). Undefined priority (no pragma `PRIORITY`) is treated as though it were less than any defined priority value.

In addition to `pragma SUPPRESS`, it is possible to suppress all checks in a given compilation by the use of the Compiler option `CHECKS`. (See Chapter 4 of the *User's Guide*.)

Section 2

Implementation-Dependent Attributes

2.1 P'IS_ARRAY

For a prefix P that denotes any type or subtype, this attribute yields the value TRUE if P is an array type or an array subtype; otherwise, it yields the value FALSE.

2.2 P'RECORD_DESCRIPTOR, P'ARRAY_DESCRIPTOR

These attributes are used to control the representation of implicit components of a record. (See Section 4.8 for more details.)

2.3 E'EXCEPTION_CODE

For a prefix E that denotes an exception name, this attribute yields a value that represents the internal code of the exception. The value of this attribute is of the type INTEGER.

Section 3

Specification of the package SYSTEM

The implementation does not allow the recompilation of package SYSTEM.

3.1 Specification of the package SYSTEM

package SYSTEM is

-- The order of the elements of this type is not significant.

```
type NAME is (S370,
              I80X86,
              I80386,
              MC680X0,
              VAX,
              TRANSPUTER,
              RS_6000,
              MIPS);

SYSTEM_NAME : constant NAME := RS_6000;
STORAGE_UNIT : constant      := 8;
MAX_INT      : constant      := 2**31 - 1;
MIN_INT      : constant      := - (2**31);
MAX_MANTISSA : constant      := 31;
FINE_DELTA   : constant      := 2#1.0#E-31;
MAX_DIGITS   : constant      := 15;
MEMORY_SIZE  : constant      := 2**32;

TICK         : constant      := 1.0;
```

subtype PRIORITY is INTEGER range 1 .. 10;

type ADDRESS is private;
NULL_ADDRESS : constant ADDRESS;

-- Converts a string to an address. The syntax of the string and its
-- meaning are target dependent.

--

-- For the 8086, 80186 and 80286 the syntax is:

-- "SSSS:0000" where SSSS and 0000 are a 4 digit or less hexadecimal
-- number representing a segment value and an offset.

-- The physical address corresponding to SSSS:0000 depends
-- on the execution mode. In real mode it is $16 * SSSS + 0000$.

-- In protected mode the value SSSS represents a segment
-- descriptor.

-- Example:

-- "0014:00F0"

--

-- For the 80386 the syntax is:

-- "00000000" where 00000000 is an 8 digit or less hexadecimal number
-- representing an offset either in the data segment or in the
-- code segment.

-- Example:

-- "00000008"

--

-- The exception CONSTRAINT_ERROR is raised if the string has not the
-- proper syntax.

function VALUE (LEFT : in STRING) return ADDRESS;

-- Converts an address to a string. The syntax of the returned string
-- is described in the VALUE function.

```

subtype ADDRESS_STRING is STRING(1..8);

function IMAGE (LEFT : in ADDRESS) return ADDRESS_STRING;

-- The following routines provide support to perform address
-- computation. The meaning of the "+" and "-" operators is
-- architecture dependent. For example on a segmented machine
-- the OFFSET parameter is added to, or subtracted from the offset
-- part of the address, the segment remaining untouched.

type OFFSET is range 0 .. 2**28 -1;

-- On a segmented architecture the function returns true if the
-- two addresses have the same segment value. On a non segmented
-- architecture it always returns TRUE.

function SAME_SEGMENT (LEFT, RIGHT : in ADDRESS) return BOOLEAN;

-- The exception ADDRESS_ERROR is raised by "<", "<=", ">", ">=", "-"
-- if the two addresses do not have the same segment value. This
-- exception is never raised on a non segmented machine.
-- The exception CONSTRAINT_ERROR can be raised by "+" and "-".

ADDRESS_ERROR : exception;

function "+" (LEFT : in ADDRESS; RIGHT : in OFFSET) return ADDRESS;
function "+" (LEFT : in OFFSET; RIGHT : in ADDRESS) return ADDRESS;
function "-" (LEFT : in ADDRESS; RIGHT : in OFFSET) return ADDRESS;

-- The exception ADDRESS_ERROR is raised on a segmented architecture
-- if the two addresses do not have the same segment value.

function "-" (LEFT : in ADDRESS; RIGHT : in ADDRESS) return OFFSET;

```

```

-- Perform an unsigned comparison on addresses or offset part of
-- addresses on a segmented machine.

function "<=" (LEFT, RIGHT : in ADDRESS) return BOOLEAN;
function "<" (LEFT, RIGHT : in ADDRESS) return BOOLEAN;
function ">=" (LEFT, RIGHT : in ADDRESS) return BOOLEAN;
function ">" (LEFT, RIGHT : in ADDRESS) return BOOLEAN;

function "mod" (LEFT : in ADDRESS; RIGHT : in POSITIVE) return NATURAL;

-- Returns the given address rounded to a specific value.

type ROUND_DIRECTION is (DOWN, UP);

function ROUND (VALUE      : in ADDRESS;
                DIRECTION : in ROUND_DIRECTION;
                MODULUS    : in POSITIVE) return ADDRESS;

-- These routines are provided to perform READ/WRITE operation
-- in memory.
-- Warning: These routines will give unexpected results if used with
-- unconstrained types.

generic
  type TARGET is private;
function FETCH_FROM_ADDRESS (A : in ADDRESS) return TARGET;

generic
  type TARGET is private;
procedure ASSIGN_TO_ADDRESS (A : in ADDRESS; T : in TARGET);

-- Procedure to copy LENGTH storage unit starting at the address
-- FROM to the address TO. The source and destination may overlap.

```

```
-- OBJECT_LENGTH designates the size of an object in storage units.

type OBJECT_LENGTH is range 0 .. 2**31 -1;

procedure MOVE (TO      : in ADDRESS;
               FROM    : in ADDRESS;
               LENGTH  : in OBJECT_LENGTH);

private

...

end SYSTEM;
```


Section 4

Support for Representation Clauses

This section explains how objects are represented and allocated by the Alsys Ada Software Development Environment for RS/6000 and how it is possible to control this using representation clauses. Applicable restrictions on representation clauses are also described.

The representation of an object is closely connected with its type. For this reason this section addresses successively the representation of enumeration, integer, floating point, fixed point, access, task, array and record types. For each class of type the representation of the corresponding objects is described.

Except in the case of array and record types, the description for each class of type is independent of the others. To understand the representation of array and record types it is necessary to understand first the representation of their components.

Apart from implementation defined pragmas, Ada provides three means to control the size of objects:

- a (predefined) pragma PACK, applicable to array types
- a record representation clause
- a size specification

For each class of types the effect of a size specification is described. Interactions among size specifications, packing and record representation clauses is described under the discussion of array and record types.

Representation clauses on derived record types or derived tasks types are not supported.

Size representation clauses on types derived from private types are not supported when the derived type is declared outside the private part of the defining package.

4.1 Enumeration Types

4.1.1 Enumeration Literal Encoding

When no enumeration representation clause applies to an enumeration type, the internal code associated with an enumeration literal is the position number of the enumeration literal. Then, for an enumeration type with n elements, the internal codes are the integers 0, 1, 2, ... , $n-1$.

An enumeration representation clause can be provided to specify the value of each internal code as described in RM 13.3. The compiler fully implements enumeration representation clauses.

As internal codes must be machine integers the internal codes provided by an enumeration representation clause must be in the range $-2^{31} .. 2^{31} - 1$.

An enumeration value is always represented by its internal code in the program generated by the compiler.

4.1.2 Enumeration Types and Object Sizes

Minimum size of an enumeration subtype

The minimum possible size of an enumeration subtype is the minimum number of bits that is necessary for representing the internal codes of the subtype values in normal binary form.

A static subtype, with a null range has a minimum size of 1. Otherwise, if m and M are the values of the internal codes associated with the first and last enumeration values of the subtype, then its minimum size L is determined as follows. For $m \geq 0$, L is the smallest positive integer such that $M \leq 2^L - 1$. For $m < 0$, L is the smallest positive integer such that $-2^{L-1} \leq m$ and $M \leq 2^{L-1} - 1$. For example:

```
type COLOR is (GREEN, BLACK, WHITE, RED, BLUE, YELLOW);  
-- The minimum size of COLOR is 3 bits.
```

```
subtype BLACK_AND_WHITE is COLOR range BLACK .. WHITE;  
-- The minimum size of BLACK_AND_WHITE is 2 bits.
```

subtype BLACK_OR_WHITE is BLACK_AND_WHITE range X .. X;
-- Assuming that X is not static, the minimum size of BLACK_OR_WHITE is
-- 2 bits (the same as the minimum size of its type mark BLACK_AND_WHITE).

Size of an enumeration subtype

When no size specification is applied to an enumeration type or first named subtype, the objects of that type or first named subtype are represented as signed machine integers. The machine provides 8, 16 and 32 bit integers, and the compiler selects automatically the smallest signed machine integer which can hold each of the internal codes of the enumeration type (or subtype). The size of the enumeration type and of any of its subtypes is thus 8, 16 or 32 bits.

When a size specification is applied to an enumeration type, this enumeration type and each of its subtypes has the size specified by the length clause. The same rule applies to a first named subtype. The size specification must of course specify a value greater than or equal to the minimum size of the type or subtype to which it applies:

```
type EXTENDED is  
  (-- The usual ASCII character set.  
  NUL, SOH, STX, ETX, EOT, ENQ, ACK, BEL,  
  ...  
  'x', 'y', 'z', '{', '|', '}', '-', DEL,  
  -- Extended characters  
  C_CEDILLA_CAP, U_UMLAUT, E_ACUTE, ...);
```

```
for EXTENDED'SIZE use 8;  
-- The size of type EXTENDED will be one byte. Its objects will be represented  
-- as unsigned 8 bit integers.
```

The compiler fully implements size specifications. Nevertheless, as enumeration values are coded using integers, the specified length cannot be greater than 32 bits.

Size of the objects of an enumeration subtype

Provided its size is not constrained by a record component clause or a pragma PACK, an object of an enumeration subtype has the same size as its subtype.

4.2 Integer Types

There are three predefined integer types in the Alsys implementation for I80386 machines:

```
type SHORT_SHORT_INTEGER      is range -2**07 .. 2**07-1;
type SHORT_INTEGER            is range -2**15 .. 2**15-1;
type INTEGER                  is range -2**31 .. 2**31-1;
```

4.2.1 Integer Type Representation

An integer type declared by a declaration of the form:

```
type T is range L .. R;
```

is implicitly derived from a predefined integer type. The compiler automatically selects the predefined integer type whose range is the smallest that contains the values L to R inclusive.

Binary code is used to represent integer values. Negative numbers are represented using two's complement.

4.2.2 Integer Type and Object Size

Minimum size of an integer subtype

The minimum possible size of an integer subtype is the minimum number of bits that is necessary for representing the internal codes of the subtype values in normal binary form.

For a static subtype, if it has a null range its minimum size is 1. Otherwise, if m and M are the lower and upper bounds of the subtype, then its minimum size L is determined as follows. For $m \geq 0$, L is the smallest positive integer such that $M \leq 2^{L-1}$. For $m < 0$, L is the smallest positive integer that $-2^{L-1} \leq m$ and $M \leq 2^{L-1} - 1$. For example:

```
subtype S is INTEGER range 0 .. 7;
-- The minimum size of S is 3 bits.
```

subtype D is S range X .. Y;
-- Assuming that X and Y are not static, the minimum size of
-- D is 3 bits (the same as the minimum size of its type mark S).

Size of an integer subtype

The sizes of the predefined integer types `SHORT_SHORT_INTEGER`, `SHORT_INTEGER` and `INTEGER` are respectively 8, 16 and 32 bits.

When no size specification is applied to an integer type or to its first named subtype (if any), its size and the size of any of its subtypes is the size of the predefined type from which it derives, directly or indirectly. For example:

type S is range 80 .. 100;
-- S is derived from `SHORT_SHORT_INTEGER`, its size is
-- 8 bits.

type J is range 0 .. 255;
-- J is derived from `SHORT_INTEGER`, its size is 16 bits.

type N is new J range 80 .. 100;
-- N is indirectly derived from `SHORT_INTEGER`, its size is
-- 16 bits.

When a size specification is applied to an integer type, this integer type and each of its subtypes has the size specified by the length clause. The same rule applies to a first named subtype. The size specification must of course specify a value greater than or equal to the minimum size of the type or subtype to which it applies:

type S is range 80 .. 100;
for S'SIZE use 32;
-- S is derived from `SHORT_SHORT_INTEGER`, but its size is
-- 32 bits because of the size specification.

type J is range 0 .. 255;
for J'SIZE use 8;
-- J is derived from `SHORT_INTEGER`, but its size is 8 bits
-- because of the size specification.

```
type N is new J range 80 .. 100;
-- N is indirectly derived from SHORT_INTEGER, but its
-- size is 8 bits because N inherits the size specification
-- of J.
```

Size of the objects of an integer subtype

Provided its size is not constrained by a record component clause or a pragma PACK, an object of an integer subtype has the same size as its subtype.

4.3 Floating Point Types

There are two predefined floating point types in the Alsys implementation for I80x86 machines:

```
type FLOAT is
  digits 6 range  $-(2.0 - 2.0^{**(-23)}) \cdot 2.0^{**127} .. (2.0 - 2.0^{**(-23)}) \cdot 2.0^{**127}$ ;
```

```
type LONG_FLOAT is
  digits 15 range  $-(2.0 - 2.0^{**(-52)}) \cdot 2.0^{**1023} .. (2.0 - 2.0^{**(-52)}) \cdot 2.0^{**1023}$ ;
```

4.3.1 Floating Point Type Representation

A floating point type declared by a declaration of the form:

```
type T is digits D [range L .. R];
```

is implicitly derived from a predefined floating point type. The compiler automatically selects the smallest predefined floating point type whose number of digits is greater than or equal to D and which contains the values L to R inclusive.

In the program generated by the compiler, floating point values are represented using the IEEE standard formats for single and double floats.

The values of the predefined type FLOAT are represented using the single float format. The values of the predefined type LONG_FLOAT are represented using the double float format. The values of any other floating point type are represented in the same way as the values of the predefined type from which it derives, directly or indirectly.

4.3.2 Floating Point Type and Object Size

The minimum possible size of a floating point subtype is 32 bits if its base type is `FLOAT` or a type derived from `FLOAT`; it is 64 bits if its base type is `LONG_FLOAT` or a type derived from `LONG_FLOAT`.

The sizes of the predefined floating point types `FLOAT` and `LONG_FLOAT` are respectively 32 and 64 bits.

The size of a floating point type and the size of any of its subtypes is the size of the predefined type from which it derives directly or indirectly.

The only size that can be specified for a floating point type or first named subtype using a size specification is its usual size (32 or 64 bits).

An object of a floating point subtype has the same size as its subtype.

4.4 Fixed Point Types

4.4.1 Fixed Point Type Representation

If no specification of `small` applies to a fixed point type, then the value of `small` is determined by the value of `delta` as defined by RM 3.5.9.

A specification of `small` can be used to impose a value of `small`: The value of `small` is required to be a power of two.

To implement fixed point types, the compiler uses a set of anonymous predefined types of the form:

```
type SHORT_FIXED is delta D range (-2.0**7-1)*S .. 2.0**7*S;  
for SHORT_FIXED'SMALL use S;
```

```
type FIXED is delta D range (-2.0**15-1)*S .. 2.0**15*S;  
for FIXED'SMALL use S;
```

```
type LONG_FIXED is delta D range (-2.0**31-1)*S .. 2.0**31*S;  
for LONG_FIXED'SMALL use S;
```

where `D` is any real value and `S` any power of two less than or equal to `D`.

A fixed point type declared by a declaration of the form:

type T is delta D range L .. R;

possibly with a small specification:

for TSMALL use S;

is implicitly derived from a predefined fixed point type. The compiler automatically selects the predefined fixed point type whose small and delta are the same as the small and delta of T and whose range is the shortest that includes the values L to R inclusive.

In the program generated by the compiler, a safe value V of a fixed point subtype F is represented as the integer:

$V / F'BASE'SMALL$

4.4.2 Fixed Point Type and Object Size

Minimum size of a fixed point subtype

The minimum possible size of a fixed point subtype is the minimum number of binary digits that is necessary for representing the values of the range of the subtype using the small of the base type.

For a static subtype, if it has a null range its minimum size is 1. Otherwise, s and S being the bounds of the subtype, if i and I are the integer representations of m and M, the smallest and the greatest model numbers of the base type such that $s < m$ and $M < S$, then the minimum size L is determined as follows. For $i \geq 0$, L is the smallest positive integer such that $I \leq 2^{L-1}$. For $i < 0$, L is the smallest positive integer such that $-2^{L-1} \leq i$ and $I \leq 2^{L-1}-1$.

type F is delta 2.0 range 0.0 .. 500.0;
-- The minimum size of F is 8 bits.

subtype S is F delta 16.0 range 0.0 .. 250.0;
-- The minimum size of S is 7 bits.

subtype D is S range X .. Y;
-- Assuming that X and Y are not static, the minimum size of D is 7 bits
-- (the same as the minimum size of its type mark S).

Size of a fixed point subtype

The sizes of the predefined fixed point types SHORT_FIXED, FIXED and LONG_FIXED are respectively 8, 16 and 32 bits.

When no size specification is applied to a fixed point type or to its first named subtype, its size and the size of any of its subtypes is the size of the predefined type from which it derives directly or indirectly. For example:

type S is delta 0.01 range 0.8 .. 1.0;
-- S is derived from an 8 bit predefined fixed type, its size is 8 bits.

type F is delta 0.01 range 0.0 .. 2.0;
-- F is derived from a 16 bit predefined fixed type, its size is 16 bits.

type N is new F range 0.8 .. 1.0;
-- N is indirectly derived from a 16 bit predefined fixed type, its size is 16 bits.

When a size specification is applied to a fixed point type, this fixed point type and each of its subtypes has the size specified by the length clause. The same rule applies to a first named subtype. The size specification must of course specify a value greater than or equal to the minimum size of the type or subtype to which it applies:

type S is delta 0.01 range 0.8 .. 1.0;
for S'SIZE use 32;
-- S is derived from an 8 bit predefined fixed type, but its size is 32 bits
-- because of the size specification.

type F is delta 0.01 range 0.0 .. 2.0;
for F'SIZE use 8;
-- F is derived from a 16 bit predefined fixed type, but its size is 8 bits
-- because of the size specification.

type N is new F range 0.8 .. 1.0;
-- N is indirectly derived from a 16 bit predefined fixed type, but its size is
-- 8 bits because N inherits the size specification of F.

The compiler fully implements size specifications. Nevertheless, as fixed point objects are represented using machine integers, the specified length cannot be greater than 32 bits.

Size of the objects of a fixed point subtype

Provided its size is not constrained by a record component clause or a pragma PACK, an object of a fixed point type has the same size as its subtype.

4.5 Access Types and Collections

Access Types and Objects of Access Types

The only size that can be specified for an access type using a size specification is its usual size (32 bits).

An object of an access subtype has the same size as its subtype, thus an object of an access subtype is always 32 bits long.

Collection Size

As described in RM 13.2, a specification of collection size can be provided in order to reserve storage space for the collection of an access type.

When no STORAGE_SIZE specification applies to an access type, no storage space is reserved for its collection, and the value of the attribute STORAGE_SIZE is then 0.

The maximum size is limited by the amount of memory available.

4.6 Task Types

Storage for a task activation

As described in RM 13.2, a length clause can be used to specify the storage space (that is, the stack size) for the activation of each of the tasks of a given type. Alsys also allows the task stack size, for all tasks, to be established using a Binder option. If a length clause is given for a task type, the value indicated at bind time is ignored for this task type, and the length clause is obeyed. When no length clause is used to specify the storage space to be

reserved for a task activation, the storage space indicated at bind time is used for this activation.

A length clause may not be applied to a derived task type. The same storage space is reserved for the activation of a task of a derived type as for the activation of a task of the parent type.

The minimum size of a task subtype is 32 bits.

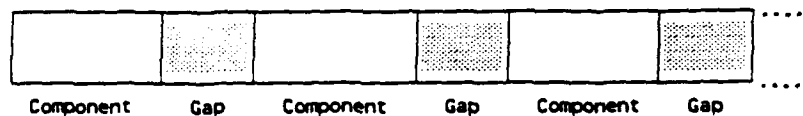
A size specification has no effect on a task type. The only size that can be specified using such a length clause is its usual size (32 bits).

An object of a task subtype has the same size as its subtype. Thus an object of a task subtype is always 32 bits long.

4.7 Array Types

Each array is allocated in a contiguous area of storage units. All the components have the same size. A gap may exist between two consecutive components (and after the last one). All the gaps have the same size.

4.7.1 Array Layout and Structure and Pragma PACK



If pragma PACK is not specified for an array, the size of the components is the size of the subtype of the components:

```
type A is array (1 .. 8) of BOOLEAN;  
-- The size of the components of A is the size of the type BOOLEAN: 8 bits.
```

```

type DECIMAL_DIGIT is range 0 .. 9;
for DECIMAL_DIGIT'SIZE use 4;
type BINARY_CODED_DECIMAL is
  array (INTEGER range <>) of DECIMAL_DIGIT;
-- The size of the type DECIMAL_DIGIT is 4 bits. Thus in an array of
-- type BINARY_CODED_DECIMAL each component will be represented on
-- 4 bits as in the usual BCD representation.

```

If pragma PACK is specified for an array and its components are neither records nor arrays, the size of the components is the minimum size of the subtype of the components:

```

type A is array (1 .. 8) of BOOLEAN;
pragma PACK(A);
-- The size of the components of A is the minimum size of the type BOOLEAN:
-- 1 bit.

```

```

type DECIMAL_DIGIT is range 0 .. 9;
for DECIMAL_DIGIT'SIZE use 32;
type BINARY_CODED_DECIMAL is
  array (INTEGER range <>) of DECIMAL_DIGIT;
pragma PACK(BINARY_CODED_DECIMAL);
-- The size of the type DECIMAL_DIGIT is 32 bits, but, as
-- BINARY_CODED_DECIMAL is packed, each component of an array of this
-- type will be represented on 4 bits as in the usual BCD representation.

```

Packing the array has no effect on the size of the components when the components are records or arrays, since records and arrays may be assigned addresses consistent with the alignment of their subtypes.

Gaps

If the components are records or arrays, no size specification applies to the subtype of the components and the array is not packed, then the compiler may choose a representation with a gap after each component; the aim of the insertion of such gaps is to optimize access to the array components and to their subcomponents. The size of the gap is chosen so that the relative displacement of consecutive components is a multiple of the alignment of the subtype of the components. This strategy allows each component and subcomponent to have an address consistent with the alignment of its subtype:

```

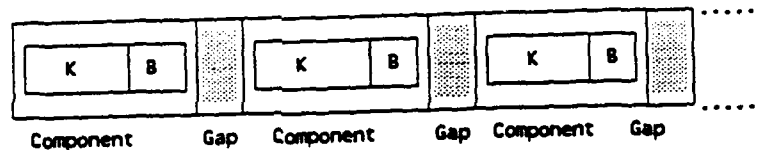
type R is
  record
    K : SHORT_INTEGER;
    B : BOOLEAN;
  end record;
for R use
  record
    K at 0 range 0 .. 31;
    B at 4 range 0 .. 0;
  end record;
-- Record type R is byte aligned. Its size is 33 bits.

```

```

type A is array (1 .. 10) of R;
-- A gap of 7 bits is inserted after each component in order to respect the
-- alignment of type R. The size of an array of type A will be 400 bits.

```



Array of type A: each subcomponent K has an even offset.

If a size specification applies to the subtype of the components or if the array is packed, no gaps are inserted:

```

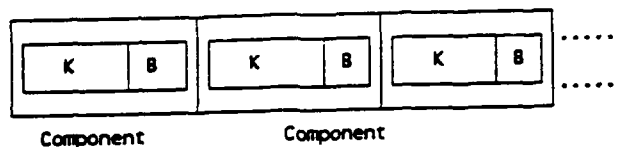
type R is
  record
    K : SHORT_INTEGER;
    B : BOOLEAN;
  end record;

type A is array (1 .. 10) of R;
pragma PACK(A);
-- There is no gap in an array of type A because A is packed.
-- The size of an object of type A will be 330 bits.

type NR is new R;
for NR'SIZE use 24;

```

type B is array (1 .. 10) of NR;
 -- There is no gap in an array of type B because
 -- NR has a size specification.
 -- The size of an object of type B will be 240 bits.



Array of type A or B

4.7.2 Array Subtype and Object Size

Size of an array subtype

The size of an array subtype is obtained by multiplying the number of its components by the sum of the size of the components and the size of the gaps (if any). If the subtype is unconstrained, the maximum number of components is considered.

The size of an array subtype cannot be computed at compile time

- if it has non-static constraints or is an unconstrained array type with non-static index subtypes (because the number of components can then only be determined at run time).
- if the components are records or arrays and their constraints or the constraints of their subcomponents (if any) are not static (because the size of the components and the size of the gaps can then only be determined at run time).

As has been indicated above, the effect of a pragma PACK on an array type is to suppress the gaps. The consequence of packing an array type is thus to reduce its size.

If the components of an array are records or arrays and their constraints or the constraints of their subcomponents (if any) are not static, the compiler ignores any pragma PACK applied to the array type but issues a warning message. Apart from this limitation, array packing is fully implemented by the compiler.

A size specification applied to an array type or first named subtype has no effect. The only size that can be specified using such a length clause is its usual size. Nevertheless, such a length clause can be useful to verify that the layout of an array is as expected by the application.

Size of the objects of an array subtype

The size of an object of an array subtype is always equal to the size of the subtype of the object.

4.8 Record Types

4.8.1 Basic Record Structure

Layout of a record

Each record is allocated in a contiguous area of storage units. The size of a record component depends on its type.

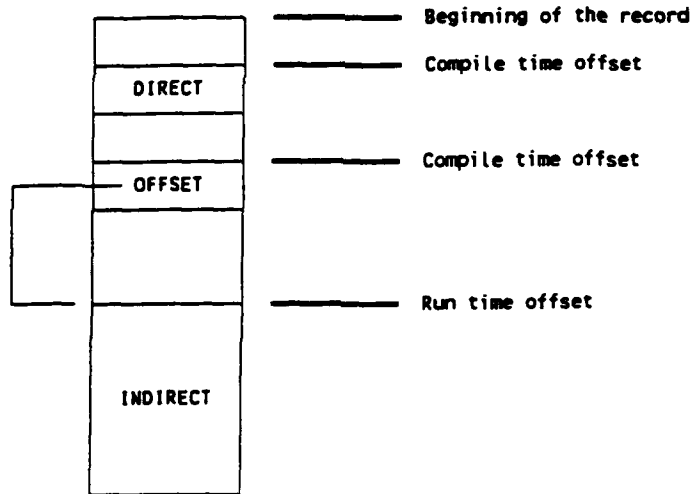
The positions and the sizes of the components of a record type object can be controlled using a record representation clause as described in RM 13.4. In the Alsys implementation for I80x86 machines there is no restriction on the position that can be specified for a component of a record. If a component is not a record or an array, its size can be any size from the minimum size to the size of its subtype. If a component is a record or an array, its size must be the size of its subtype.

Pragma PACK has no effect on records. It is unnecessary because record representation clauses provide full control over record layout.

A record representation clause need not specify the position and the size for every component. If no component clause applies to a component of a record, its size is the size of its subtype.

4.8.2 Indirect Components

If the offset of a component cannot be computed at compile time, this offset is stored in the record objects at run time and used to access the component. Such a component is said to be indirect while other components are said to be direct:



A direct and an indirect component

If a record component is a record or an array, the size of its subtype may be evaluated at run time and may even depend on the discriminants of the record. We will call these components dynamic components:

type DEVICE is (SCREEN, PRINTER);

type COLOR is (GREEN, RED, BLUE);

type SERIES is array (POSITIVE range <>) of INTEGER;

type GRAPH (L : NATURAL) is

record

X : SERIES(1 .. L); -- The size of X depends on L

Y : SERIES(1 .. L); -- The size of Y depends on L

end record;

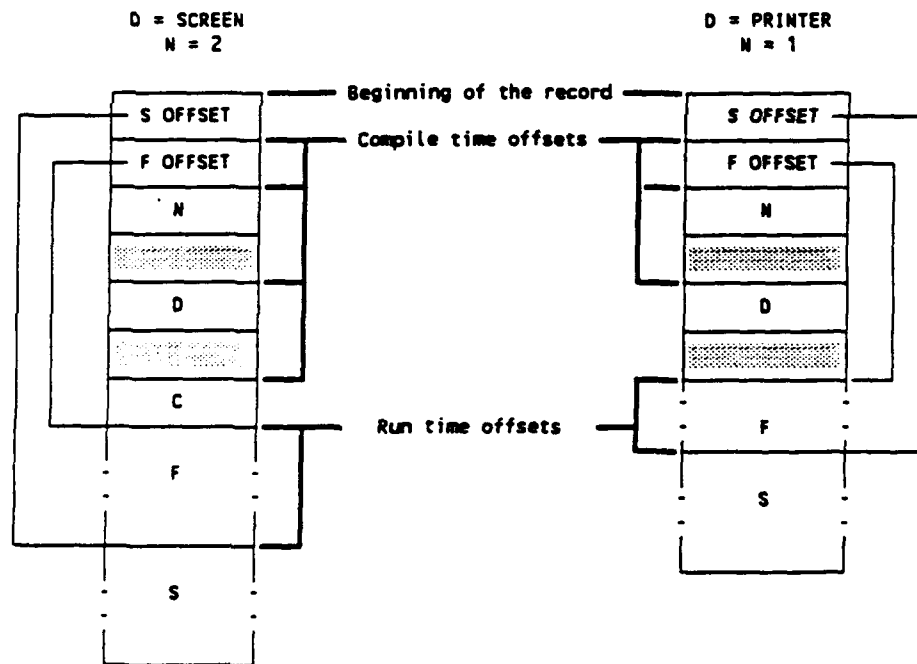
Q : POSITIVE;

```

type PICTURE (N : NATURAL; D : DEVICE) is
  record
    F : GRAPH(N); -- The size of F depends on N
    S : GRAPH(Q); -- The size of S depends on Q
  case D is
    when SCREEN =>
      C : COLOR;
    when PRINTER =>
      null;
  end case;
end record;

```

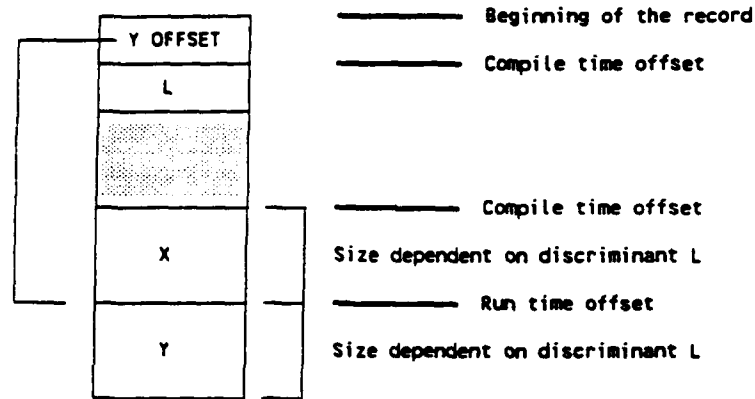
Any component placed after a dynamic component has an offset which cannot be evaluated at compile time and is thus indirect. In order to minimize the number of indirect components, the compiler groups the dynamic components together and places them at the end of the record:



The record type PICTURE: F and S are placed at the end of the record

Note that Ada does not allow representation clauses for record components with non-static bounds [RM 13.4.7], so the compiler's grouping of dynamic components does not conflict with the use of representation clauses.

Because of this approach, the only indirect components are dynamic components. But not all dynamic components are necessarily indirect: if there are dynamic components in a component list which is not followed by a variant part, then exactly one dynamic component of this list is a direct component because its offset can be computed at compilation time (the only dynamic components that are direct components are in this situation):



The record type GRAPH: the dynamic component X is a direct component.

The offset of an indirect component is always expressed in storage units.

The space reserved for the offset of an indirect component must be large enough to store the size of any value of the record type (the maximum potential offset). The compiler evaluates an upper bound *MS* of this size and treats an offset as a component having an anonymous integer type whose range is $0..MS$.

If *C* is the name of an indirect component, then the offset of this component can be denoted in a component clause by the implementation generated name *C'OFFSET*.

4.8.3 Implicit Components

In some circumstances, access to an object of a record type or to its components involves computing information which only depends on the discriminant values. To avoid recomputation (which would degrade performance) the compiler stores this information in the record objects, updates it when the values of the discriminants are modified and uses it when the objects or its components are accessed. This information is stored in special components called implicit components.

An implicit component may contain information which is used when the record object or several of its components are accessed. In this case the component will be included in any record object (the implicit component is considered to be declared before any variant part in the record type declaration). There can be two components of this kind; one is called `RECORD_SIZE` and the other `VARIANT_INDEX`.

On the other hand an implicit component may be used to access a given record component. In that case the implicit component exists whenever the record component exists (the implicit component is considered to be declared at the same place as the record component). Components of this kind are called `ARRAY_DESCRIPTORs` or `RECORD_DESCRIPTORs`.

RECORD_SIZE

This implicit component is created by the compiler when the record type has a variant part and its discriminants are defaulted. It contains the size of the storage space necessary to store the current value of the record object (note that the storage effectively allocated for the record object may be more than this).

The value of a `RECORD_SIZE` component may denote a number of bits or a number of storage units. In general it denotes a number of storage units, but if any component clause specifies that a component of the record type has an offset or a size which cannot be expressed using storage units, then the value designates a number of bits.

The implicit component `RECORD_SIZE` must be large enough to store the maximum size of any value of the record type. The compiler evaluates an upper bound `MS` of this size and then considers the implicit component as having an anonymous integer type whose range is `0 .. MS`.

If `R` is the name of the record type, this implicit component can be denoted in a component clause by the implementation generated name `R'RECORD_SIZE`. This allows user control over the position of the implicit component in the record.

VARIANT_INDEX

This implicit component is created by the compiler when the record type has a variant part. It indicates the set of components that are present in a record value. It is used when a discriminant check is to be done.

Component lists in variant parts that themselves do not contain a variant part are numbered. These numbers are the possible values of the implicit component **VARIANT_INDEX**.

```
type VEHICLE is (AIRCRAFT, ROCKET, BOAT, CAR);
```

```
type DESCRIPTION (KIND : VEHICLE := CAR) is
```

```
  record
```

```
    SPEED : INTEGER;
```

```
    case KIND is
```

```
      when AIRCRAFT | CAR =>
```

```
        WHEELS : INTEGER;
```

```
      case KIND is
```

```
        when AIRCRAFT =>          -- 1
```

```
          WINGSPAN : INTEGER;
```

```
        when others =>          -- 2
```

```
          null;
```

```
      end case;
```

```
      when BOAT => -- 3
```

```
        STEAM : BOOLEAN;
```

```
      when ROCKET =>          -- 4
```

```
        STAGES : INTEGER;
```

```
      end case;
```

```
  end record;
```

The value of the variant index indicates the set of components that are present in a record value:

Variant Index	Set
1	(KIND, SPEED, WHEELS, WINGSPAN)
2	(KIND, SPEED, WHEELS)
3	(KIND, SPEED, STEAM)
4	(KIND, SPEED, STAGES)

A comparison between the variant index of a record value and the bounds of an interval is enough to check that a given component is present in the value:

Component	Interval
KIND	--
SPEED	--
WHEELS	1 .. 2
WINGSPAN	1 .. 1
STEAM	3 .. 3
STAGES	4 .. 4

The implicit component `VARIANT_INDEX` must be large enough to store the number `V` of component lists that don't contain variant parts. The compiler treats this implicit component as having an anonymous integer type whose range is `1 .. V`.

If `R` is the name of the record type, this implicit component can be denoted in a component clause by the implementation generated name `R'VARIANT_INDEX`. This allows user control over the position of the implicit component in the record.

ARRAY_DESCRIPTOR

An implicit component of this kind is associated by the compiler with each record component whose subtype is an anonymous array subtype that depends on a discriminant of the record. It contains information about the component subtype.

The structure of an implicit component of kind `ARRAY_DESCRIPTOR` is not described in this documentation. Nevertheless, if a programmer is interested in specifying the location of a component of this kind using a component clause, size of the component may be obtained using the `ASSEMBLY` parameter in the `COMPILE` command.

The compiler treats an implicit component of the kind `ARRAY_DESCRIPTOR` as having an anonymous array type. If `C` is the name of the record component whose subtype is described by the array descriptor, then this implicit component can be denoted in a component clause by the implementation generated name `C'ARRAY_DESCRIPTOR`. This allows user control over the position of the implicit component in the record.

RECORD_DESCRIPTOR

An implicit component of this kind is associated by the compiler with each record component whose subtype is an anonymous record subtype that depends on a discriminant of the record. It contains information about the component subtype.

The structure of an implicit component of kind `RECORD_DESCRIPTOR` is not described in this documentation. Nevertheless, if a programmer is interested in specifying the location of a component of this kind using a component clause, the size of the component may be obtained using the `ASSEMBLY` parameter in the `COMPILE` command.

The compiler treats an implicit component of the kind `RECORD_DESCRIPTOR` as having an anonymous array type. If `C` is the name of the record component whose subtype is described by the record descriptor, then this implicit component can be denoted in a component clause by the implementation generated name `C'RECORD_DESCRIPTOR`. This allows user control over the position of the implicit component in the record.

Suppression of Implicit Components

The Alsys implementation provides the capability of suppressing the implicit components `RECORD_SIZE` and/or `VARIANT_INDEX` from a record type. This can be done using an implementation defined pragma called `IMPROVE`. The syntax of this pragma is as follows:

```
pragma IMPROVE ( TIME | SPACE [,ON =>] simple_name );
```

The first argument specifies whether `TIME` or `SPACE` is the primary criterion for the choice of the representation of the record type that is denoted by the second argument.

If `TIME` is specified, the compiler inserts implicit components as described above. If on the other hand `SPACE` is specified, the compiler only inserts a `VARIANT_INDEX` or a `RECORD_SIZE` component if this component appears in a record representation clause that applies to the record type. A record representation clause can thus be used to keep one implicit component while suppressing the other.

A pragma `IMPROVE` that applies to a given record type can occur anywhere that a representation clause is allowed for this type.

4.8.4 Size of Record Types and Objects

Size of a record subtype

Unless a component clause specifies that a component of a record type has an offset or a size which cannot be expressed using storage units, the size of a record subtype is rounded up to a whole number of storage units.

The size of a constrained record subtype is obtained by adding the sizes of its components and the sizes of its gaps (if any). This size is not computed at compile time

- when the record subtype has non-static constraints,
- when a component is an array or a record and its size is not computed at compile time.

The size of an unconstrained record subtype is obtained by adding the sizes of the components and the sizes of the gaps (if any) of its largest variant. If the size of a component or of a gap cannot be evaluated exactly at compile time an upper bound of this size is used by the compiler to compute the subtype size.

A size specification applied to a record type or first named subtype has no effect. The only size that can be specified using such a length clause is its usual size. Nevertheless, such a length clause can be useful to verify that the layout of a record is as expected by the application.

Size of an object of a record subtype

An object of a constrained record subtype has the same size as its subtype.

An object of an unconstrained record subtype has the same size as its subtype if this size is less than or equal to 8 kb. If the size of the subtype is greater than this, the object has the size necessary to store its current value; storage space is allocated and released as the discriminants of the record change.

Section 5

Conventions for Implementation-Generated Names

The Alsys Ada Software Development Environment for RS/6000 may add fields to record objects and have descriptors in memory for record or array objects. These fields are accessible to the user through implementation-generated attributes (See Section 2.3).

Section 6

Address Clauses

6.1 Address Clauses for Objects

An address clause can be used to specify an address for an object as described in RM 13.5. When such a clause applies to an object the compiler does not cause storage to be allocated for the object. The program accesses the object using the address specified in the clause. It is the responsibility of the user therefore to make sure that a valid allocation of storage has been done at the specified address.

An address clause is not allowed for task objects, for unconstrained records whose size is greater than 8k bytes or for a constant.

There are a number of ways to compose a legal address expression for use in an address clause. The most direct ways are:

- For the case where the memory is defined in Ada as another object, use the 'ADDRESS attribute to obtain the argument for the address clause for the second object.
- For the case where an absolute address is known to the programmer, instantiate the generic function SYSTEM.REFERENCE on a 16 bit unsigned integer type (either from package UNSIGNED, or by use of a length clause on a derived integer type or subtype) and on type SYSTEM.ADDRESS. Then the values of the desired segment and offset can be passed as the actual parameters to the instantiated function in the simple expression part of the address clause. See Section 3 for the specification of package SYSTEM.
- For the case where the desired location is memory defined in assembly or another non-Ada language (is relocatable), an interfaced routine may be used to obtain the appropriate address from referencing information known to the other language.

6.2 Address Clauses for Program Units

Address clauses for program units are not implemented in the current version of the compiler.

6.3 Address Clauses for Interrupt Entries

Address clauses for entries are not implemented in the current version of the compiler.

Section 7

Unchecked Conversions

Unchecked conversions are allowed between any types provided the instantiation of `UNCHECKED_CONVERSION` is legal Ada. It is the programmer's responsibility to determine if the desired effect is achieved.

If the target type has a smaller size than the source type then the target is made of the least significant bits of the source.

Section 8

Input-Output Packages

In this part of the Appendix the implementation-specific aspects of the input-output system are described.

8.1 Introduction

In Ada, input-output operations (IO) are considered to be performed on *objects* of a certain file type rather than being performed directly on external files. An external file is anything external to the program that can produce a value to be read or receive a value to be written. Values transferred for a given file must be all of one type.

Generally, in Ada documentation, the term *file* refers to an object of a certain file type, whereas a physical manifestation is known as an *external file*. An external file is characterized by

- Its name, which is a string defining a legal path name under the current version of the operating system.
- Its form, which gives implementation-dependent information on file characteristics.

Both the name and the form appear explicitly as parameters of the Ada CREATE and OPEN procedures. Though a file is an object of a certain file type, ultimately the object has to correspond to an external file. Both CREATE and OPEN associate a NAME of an external file (of a certain FORM) with a program file object.

Ada IO operations are provided by means of standard packages [14].

SEQUENTIAL_IO	A generic package for sequential files of a single element type.
DIRECT_IO	A generic package for direct (random) access files.
TEXT_IO	A generic package for human readable (text, ASCII) files.

IO_EXCEPTIONS A package which defines the exceptions needed by the above three packages.

The generic package `LOW_LEVEL_IO` is not implemented in this version.

The upper bound for index values in `DIRECT_IO` and for line, column and page numbers in `TEXT_IO` is given by

`COUNTLAST = 2**31 - 1`

The upper bound for field widths in `TEXT_IO` is given by

`FIELDLAST = 255`

8.2 The FORM Parameter

The `FORM` parameter of both the `CREATE` and `OPEN` procedures in Ada specifies the characteristics of the external file involved.

The `CREATE` procedure establishes a new external file, of a given `NAME` and `FORM`, and associates it with a specified program file object. The external file is created (and the file object set) with a specified (or default) file mode. If the external file already exists, the file will be erased. The exception `USE_ERROR` is raised if the file mode is `IN_FILE`.

Example:

```
CREATE (F, OUT_FILE, "MY_FILE",  
       FORM =>  
       "WORLD => READ, OWNER => READ_WRITE");
```

The `OPEN` procedure associates an existing external file, of a given `NAME` and `FORM`, with a specified program file object. The procedure also sets the current file mode. If there is an inadmissible change of mode, then the Ada exception `USE_ERROR` is raised.

The `FORM` parameter is a string, formed from a list of attributes, with attributes separated by commas (,). The string is not case sensitive (so that, for example, `HERE` and

here are treated alike). FORM attributes are distinct from Ada attributes. The attributes specify:

- File protection
- File sharing
- File structure
- Buffering
- Appending
- Blocking
- Terminal input

The general form of each attribute is a keyword followed by => and then a qualifier. The arrow and qualifier may sometimes be omitted. The format for an attribute specifier is thus either of

KEYWORD

KEYWORD => QUALIFIER

We will discuss each attribute in turn.

8.2.1 File Protection

These attributes are only meaningful for a call to the CREATE procedure.

File protection involves two independent classifications. The first classification is related to *who* may access the file and is specified by the keywords:

OWNER	Only the owner of the directory may access this file.
GROUP	Only the members of a predefined group of users may access this file.
WORLD	Any user may access this file.

For each type of user who may access a file there are various access *rights*, and this forms the basis for the second classification. In general, there are four types of access right, specified by the qualifiers:

READ	The user may read from the external file.
WRITE	The user may write to the external file.
EXECUTE	The user may execute programs stored in the external file.
NONE	The user has no access rights to the external file. (This access right negates any prior privileges.)

More than one access right may be relevant for a particular file, in which case the qualifiers are linked with underscores (_).

For example, suppose that the WORLD may execute a program in an external file, but only the OWNER may modify the file.

```
WORLD =>  
    EXECUTE ,  
OWNER =>  
    READ_WRITE_EXECUTE,
```

Repetition of the same qualifier within the attributes is illegal:

```
WORLD =>  
    EXECUTE_EXECUTE, -- NOT legal
```

but repetition of the entire attribute is allowed:

```
WORLD =>  
    EXECUTE,  
WORLD =>  
    EXECUTE, -- Legal
```

8.2.2 File Sharing

An external file can be shared, which means associated simultaneously with several logical file objects created by the `OPEN` and `CREATE` procedures.

The file sharing attribute may restrict or suppress this capability by specifying one of the following access modes:

`NOT_SHARED`

Exclusive access - no other logical file may be associated with the external file

`SHARED => READERS`

Only logical files opened with mode `IN` are allowed

`SHARED => SINGLE_WRITER`

Only logical files opened with mode `IN` and at most one with mode `INOUT` or `OUT` are allowed

`SHARED => ANY`

No restriction

The exception `USE_ERROR` is raised if, for an external file already associated with an Ada file object:

- a further `OPEN` or `CREATE` specifies a file sharing attribute different from the current one
- a further `OPEN`, `CREATE` or `RESET` violates the conditions imposed by the current file sharing attribute.

The restrictions imposed by the file sharing attribute disappear when the last logical file object linked to the external file is closed.

The file sharing attribute provides control over multiple accesses within the program to a given external file.

This control does not extend to the whole system.

The default value for the file sharing attribute is `SHARED => ANY`

8.2.3 File Structure

Text Files

There is no FORM parameter to define the structure of text files.

A text file consists of a sequence of bytes holding the ASCII codes of characters.

The representation of Ada-terminators depends on the file's mode (IN or OUT) and whether it is associated with a terminal device or a mass-storage file:

- Mass-storage files
 - end of line: ASCII.LF
 - end of page: ASCII.LF ASCII.FF
 - end of file: ASCII.LF ASCII.EOT
- Terminal device with mode IN
 - end of line: ASCII.LF
 - end of page: ASCII.LF ASCII.FF
 - end of file: ASCII.LF ASCII.FF
- Terminal device with mode OUT
 - end of lin.: ASCII.LF
 - end of page: ASCII.FF
 - end of file: ASCII.EOT

Binary Files

Two FORM attributes, RECORD_SIZE and RECORD_UNIT, control the structure of binary files.

A binary file can be viewed as a sequence (sequential access) or a set (direct access) of consecutive RECORDS.

The structure of such a record is:

[HEADER] OBJECT [UNUSED_PART]

and it is formed from up to three items:

- an OBJECT with the exact binary representation of the Ada object in the executable program, possibly including an object descriptor
- a HEADER consisting of two fields (each of 32 bits):
 - the length of the object in bytes
 - the length of the descriptor in bytes
 - an UNUSED_PART of variable size to permit full control of the record's size

The HEADER is implemented only if the actual parameter of the instantiation of the IO package is unconstrained.

The file structure attributes take the form:

RECORD_SIZE => size_in_bytes

RECORD_UNIT => size_in_bytes

Their meaning depends on the object's type (constrained or not) and the file access mode (sequential or direct access):

- a) If the object's type is constrained:
 - The RECORD_UNIT attribute is illegal
 - If the RECORD_SIZE attribute is omitted, no UNUSED_PART will be implemented: the default RECORD_SIZE is the object's size
 - If present, the RECORD_SIZE attribute must specify a record size greater than or equal to the object's size, otherwise the exception USE_ERROR will be raised
- b) If the object's type is unconstrained and the file access mode is direct:
 - The RECORD_UNIT attribute is illegal
 - The RECORD_SIZE attribute has no default value, and if it is not specified, a USE_ERROR will be raised

- An attempt to input or output an object larger than the given RECORD_SIZE will raise the exception DATA_ERROR
- c) If the object's type is unconstrained and the file access mode is sequential:
- The RECORD_SIZE attribute is illegal
 - The default value of the RECORD_UNIT attribute is 1 (byte)
 - The record size will be the smallest multiple of the specified (or default) RECORD_UNIT that holds the object and its length. This is the only case where records of a file may have different sizes.

8.2.4 Buffering

The buffer size can be specified by the attribute

BUFFER_SIZE => size_in_bytes

The default value for BUFFER_SIZE is 0 (which means no buffering) for terminal devices; it is 1 block for disk files.

Only to be used with the procedure OPEN, the format of this attribute is simply

APPEND

and it means that any output will be placed at the end of the named external file.

In normal circumstances, when an external file is opened, an index is set which points to the beginning of the file. If the APPEND attribute is present for a sequential or for a text file, then data transfer will commence at the end of the file. For a direct access file, the value of the index is set to one more than the number of records in the external file.

This attribute is not applicable to terminal devices.

8.2.6 Blocking

This attribute has two alternative forms:

BLOCKING ,

or

`NON_BLOCKING`,

This attribute specifies the IO system behavior desired at any moment that a request for data transfer cannot be fulfilled. The stoppage may be due, for example, to the unavailability of data, or to the unavailability of the external file device.

NON_BLOCKING

If this attribute is set, then the task that ordered the data transfer is suspended - meaning that other tasks can execute. The suspended task is kept in a 'ready' state, together with other tasks in a ready state at the same priority level (that is, it is rescheduled).

When the suspended task is next scheduled, the data transfer request is reactivated. If ready, the transfer is activated, otherwise the rescheduling is repeated. Control returns to the user program after completion of the data transfer.

BLOCKING

In this case the task waits until the data transfer is complete, and all other tasks are suspended (or 'blocked'). The system is busy waiting.

The default for this attribute depends on the actual program: it is `BLOCKING` for programs without task declarations and `NON_BLOCKING` for a program containing tasks.

8.2.7 Terminal Input

This attribute takes one of two alternative forms:

`TERMINAL_INPUT => LINES`,

`TERMINAL_INPUT => CHARACTERS`,

Terminal input is normally processed in units of a line at a time, where a line is delimited by a special character. A process attempting to read from the terminal as an external file will be suspended until a complete line has been typed. At that time, the outstanding read call (and possibly also later calls) will be satisfied.

The first option specifies line-at-a-time data transfer, which is the default case.

The second option means that data transfer is character by character, and so a complete line does not have to be entered before the read request can be satisfied. For this option the `BUFFER_SIZE` must be zero.

The `TERMINAL_INPUT` attribute is only applicable to terminal devices.

Section 9

Characteristics of Numeric Types

9.1 Integer Types

The ranges of values for integer types declared in package STANDARD are as follows:

SHORT_SHORT_INTEGER	-128 .. 127	-- $2^{**7} - 1$
SHORT_INTEGER	-32768 .. 32767	-- $2^{**15} - 1$
INTEGER	-2147483648 .. 2147483647	-- $2^{**31} - 1$

For the packages DIRECT_IO and TEXT_IO, the range of values for types COUNT and POSITIVE_COUNT are as follows:

COUNT	0 .. 2147483647	-- $2^{**31} - 1$
POSITIVE_COUNT	1 .. 2147483647	-- $2^{**31} - 1$

For the package TEXT_IO, the range of values for the type FIELD is as follows:

FIELD	0 .. 255	-- $2^{**8} - 1$
-------	----------	------------------

9.2 Floating Point Type Attributes

	SHORT_FLOAT and FLOAT	LONG_FLOAT
DIGITS	6	15
MANTISSA	21	51
EMAX	84	204

EPSILON	9.53674E-07	8.88178E-16
LARGE	1.93428E+25	2.57110E+61
SAFE_EMAX	125	1021
SAFE_SMALL	1.17549E-38	2.22507E-308
SAFE_LARGE	4.25353E+37	2.24712E+307
FIRST	-3.40282E+38	-1.79769E+308
LAST	3.40282E+38	1.79769E+308
MACHINE_RADIX	2	2
MACHINE_EMAX	128	1024
MACHINE_EMIN	-125	-1021
MACHINE_ROUNDS	true	true
MACHINE_OVERFLOWS	false	false
SIZE	32	64

9.3 Attributes of Type DURATION

DURATION'DELTA	2.0 ** (-14)
DURATION'SMALL	2.0 ** (-14)
DURATION'FIRST	-131_072.0
DURATION'LAST	131_072.0
DURATION'LARGE	same as DURATION'LAST

Section 10

Other Implementation-Dependent Characteristics

10.1 Characteristics of the Heap

All objects created by allocators go into the heap. Also, portions of the *Runtime Executive* representation of task objects, including the task stacks, are allocated in the heap.

UNCHECKED_DEALLOCATION is implemented for all Ada access objects except access objects to tasks. Use of UNCHECKED_DEALLOCATION on a task object will lead to unpredictable results.

All objects whose visibility is linked to a subprogram, task body, or block have their storage reclaimed at exit, whether the exit is normal or due to an exception. Effectively *pragma CONTROLLED* is automatically applied to all access types. Moreover, all compiler temporaries on the heap (generated by such operations as function calls returning unconstrained arrays, or many concatenations) allocated in a scope are deallocated upon leaving the scope.

Note that the programmer may force heap reclamation of temporaries associated with any statements by enclosing the statement in a *begin .. end* block. This is especially useful when complex concatenations or other heap-intensive operations are performed in loops, and can reduce or eliminate STORAGE_ERRORS that might otherwise occur.

The maximum size of the heap is limited only by available memory. This includes the amount of physical memory (RAM) and the amount of virtual memory (hard disk swap space).

10.2 Characteristics of Tasks

The default task stack size is 1K bytes (32K bytes for the environment task), but by using the Binder option STACK.TASK the size for all task stacks in a program may be set to a size from 1K bytes to 32767 bytes.

Normal priority rules are followed for preemption, where PRIORITY values are in the range 1 .. 10. A task with *undefined* priority (no pragma PRIORITY) is considered to be lower than priority 1.

The minimum timeable delay is 0.01562 seconds.

The maximum number of active tasks is restricted only by memory usage.

The acceptor of a rendezvous executes the accept body code in its own stack. Rendezvous with an empty accept body (for synchronization) does not cause a context switch.

The main program waits for completion of all tasks dependent upon library packages before terminating.

Abnormal completion of an aborted task takes place immediately, except when the abnormal task is the caller of an entry that is engaged in a rendezvous, or if it is in the process of activating some tasks. Any such task becomes abnormally completed as soon as the state in question is exited.

The message

GLOBAL BLOCKING SITUATION DETECTED

is printed to STANDARD_OUTPUT when the *Runtime Executive* detects that no further progress is possible for any task in the program. The execution of the program is then abandoned.

10.3 Definition of a Main Subprogram

A library unit can be used as a main subprogram if and only if it is a procedure that is not generic and that has no formal parameters.

10.4 Ordering of Compilation Units

The Alsys Ada Software Development Environment for RS/6000 imposes no additional ordering constraints on compilations beyond those required by the language.

Section 11

Limitations

11.1 Compiler Limitations

- The maximum identifier length is 255 characters.
- The maximum line length is 255 characters.
- The maximum number of unique identifiers per compilation unit is 2500.
- The maximum number of compilation units in a library is 4096.
- The maximum number of Ada libraries in a family is 2047.

11.2 Hardware Related Limitations

- The maximum amount of data in the heap is limited only by available memory.
- If an unconstrained record type can exceed 4096 bytes, the type is not permitted (unless constrained) as the element type in the definition of an array or record type.
- A dynamic object bigger than 4096 bytes will be indirectly allocated. Refer to ALLOCATION parameter in the COMPILE command. (Section 4.2 of the *User's Guide*.)

INDEX

- Abnormal completion 58
- Aborted task 58
- Access types 24
- Allocators 57
- Application Developer's Guide 3
- Array gaps 26
- Array objects 39
- Array subtype 7
- Array subtype and object size 28
- Array type 7
- ARRAY_DESCRIPTOR 35
 - Attribute 7
- ASSEMBLER 3
- Attributes of type DURATION 56

- Basic record structure 29
- Binder 57

- C 3
- Characteristics of tasks 57
- Collection size 24
- Collections 24
- Compiler limitations 61
 - maximum identifier length 61
 - maximum line length 61
 - maximum number of Ada libraries 61
 - maximum number of compilation units 61
 - maximum number of unique identifiers 61
- COUNT 55

- DIGITS 55
- DIRECT_IO 55
- DURATION'DELTA 56
- DURATION'FIRST 56
- DURATION'LARGE 56
- DURATION'LAST 56
- DURATION'SMALL 56

- E'EXCEPTION_CODE 7
- EMAX 55
- Empty accept body 58
- Enumeration literal encoding 16
- Enumeration subtype size 17
- Enumeration types 16
- EPSILON 56
- EXCEPTION_CODE
 - Attribute 7

- FIELD 55
- FIRST 56
- Fixed point type representation 21
- Fixed point type size 22
- Floating point type attributes 55
- Floating point type representation 20
- Floating point type size 21

- GLOBAL BLOCKING SITUATION DETECTED 58

- Hardware limitations
 - maximum data in the heap 61
 - maximum size of a single array or record object 61
- Heap 57

- Implementation generated names 39
- Implicit component 35, 36
- Implicit components 33
- INDENT 5
- Indirect record components 29
- Integer type and object size 18

Integer type representation 18
 Integer types 55
 INTERFACE 3, 4
 INTERFACE_NAME 3, 4
 IS_ARRAY
 Attribute 7

 LARGE 56
 LAST 56
 Layout of a record 29
 Library unit 59
 Limitations 61

 MACHINE_EMAX 56
 MACHINE_EMIN 56
 MACHINE_MANTISSA 56
 MACHINE_OVERFLOWS 56
 MACHINE_RADIX 56
 MACHINE_ROUNDS 56
 Main program 58
 Main subprogram 59
 MANTISSA 55
 Maximum data in the heap 61
 Maximum identifier length 61
 Maximum line length 61
 Maximum number of Ada libraries 61
 Maximum number of compilation units
 61
 Maximum number of unique identifiers
 61
 Maximum size of a single array or
 record object 61
 Minimum timeable delay 58

 Number of active tasks 58

 Ordering of compilation units 59

 P'ARRAY_DESCRIPTOR 7
 P'IS_ARRAY 7
 P'RECORD_DESCRIPTOR 7
 PACK 5
 Parameter passing 2
 POSITIVE_COUNT 55
 Pragma IMPROVE 5, 36
 Pragma INDENT 5
 Pragma INTERFACE 3, 4
 Pragma INTERFACE_NAME 4
 Pragma PACK 5, 25, 26, 29
 Pragma PRIORITY 5, 58
 Pragma SUPPRESS 6
 Predefined packages 39
 PRIORITY 5, 58

 Record objects 39
 RECORD_DESCRIPTOR 36
 Attribute 7
 RECORD_SIZE 33, 36
 Rendezvous 58
 Representation clauses 15
 Runtime Executive 2, 4, 57, 58

 SAFE_EMAX 56
 SAFE_LARGE 56
 SAFE_SMALL 56
 SIZE 56
 Size of record types 37
 SPACE 36
 STANDARD_OUTPUT 58
 Storage reclamation at exit 57
 STORAGE_SIZE 24
 SUPPRESS 6
 SYSTEM 5

Task activation 24
Task stack size 24, 57
Task stacks 57
Task types 24
Tasks
 characteristics of 57
TEXT_IO 55
TIME 36

Unchecked conversions 43
UNCHECKED_DEALLOCATION
 57
UNIX Linker 4

Variant part 34
VARIANT_INDEX 34, 35, 36