			(D)
RFPORT DOCUM	ENTATION F	PAGE	Form Approved . OPM No.
AD-A279 366	o average 1 hour per response, in mation. Send comments regardi ervice, Directorate for Information Office of Management and Budge EPORT	ncluding the time for review of this burden estimate or Operations and Reports, I, Washington, DC 20503 3. REPORT	ning instructions, searching existing data sources gath any other aspect of this collection of information, inclu 1215 Jefferson Davis Highway, Suite 1204, Arlington, TYPE AND DATES
Tartan, Inc. / Tart.	mWorks Ada	berxxyVersion	5. FUNDING
Authors: TABG, A	SunOS 4.3.1)/Targe 1.6+. ITE	k Motorola MU	mE167(E8040) (usim: VxWorks 5.1)
PERFORMING ORGANIZATION NAME(S) AND IFABG, Abt. ITE Einsteinstr. 20 D-85521 Ottobrunn)		8. PERFORMING ORGANIZATION
SPONSORING/MONITORING AGENCY NAME(da Joint Program Office he Pentagon, Rm 3E118 /ashington, DC 20301-3080	S) AND DT ELEC MAY 12	TE 1994	10. SPONSORING/MONITORING AGENCY
. SUPPLEMENTARY			94-14197
a DISTRIBUTIONAVAILABILITY Approved for public unlimited	release ; dist	ribution	120. DISTRIBUTION
. (Maximum 200 *:940221 I 1.11340	, AVF#: VSR	-116	
Tortan, Inc.			
		gre pro fre State of the State	1 - 5 8
		Ker Val. Sun	15. NUMBER OF
eport, Ada programming lar eport, Ada Compiler Vall Co frice, Ada Val. Facility, A	pacility Valit NSI/MIL-ST	esting, Ada TD - 1815A,	AJPO
4. SUBJECT Ada programming law eport, Ada Compiler Vali Co Hice, Ada Val. Facility, A 7. SECURITY LASSIFICATION INCLASSIFIED UNCLASS	NSI/MIL-S- NSI/MIL-S- SIFIED	esting, Ada FD - 1815 A, ECURITY SSIFICATION CLASSIFIED	AJPO 16. PRICE 20, LIMITATION OF UNCLASSIFIED Standard Form 200 (Part)

:

AVF Control Number: IABG-VSR 115 22 February, 1994

Ada COMPILER VALIDATION SUMMARY REPORT: Certificate Number: 94022111.11340 Tartan, Inc. TartanWorks Ada 68xxx Version 4.3.1 Sun SPARCstation ELC = > Motorola MVME167 (MC68040)

÷

Accesio	on For			
NTIS CRA&I DTIC TAB Unannounced Justification				
By Distribution /				
A	vailabilit	y Codes		
Dist	Avail a Spe	and / or cial		
A-1				

Prepared By: IABG, Abt. ITE Einsteinstr. 20 D-85521 Ottobrunn Germany

1 ...

The following Ada implementation was tested and determined to pass ACVC 1.11. Testing was completed on 21 February, 1994.

Compiler Name and Version:	TartanWorks Ada 68xxx Version 4.3.1
Host Computer System:	Sun SPARCstation ELC under SunOS Version 4.3.1
Target Computer System:	Motorola MVME167 (MC68040) running VxWorks Version 5.1

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

As a result of this validation effort, Validation Certificate 94022111.11340 is awarded to Tartan, Inc. This certificate expires 24 months after ANSI approval of MIL-STD 1815B.

This report has been reviewed and is approved.

Michael Tonndorf IABG, Abt. ITE Einsteinstr. 20 D-85521 Ottobrunn Germany

Ada Valdation Organization Director, Computer & Software Engineering Division Institute for Defense Analyses Alexandria VA 22311, USA

Donett J. Reifer David R. Basel

Deputy Director, Ada Joint Program Office Defense Information Systems Agency, **Center for Information Management**

The following Ada implementation was tested and determined to pass ACVC 1.11. Testing was completed on 21 February, 1994.

Compiler Name and Version:	TartanWorks Ada 68xxx Version 4.3.1
Host Computer System:	Sun SPARCstation ELC under SunOS Version 4.3.1
Target Computer System:	Motorola MVME167 (MC68040) running VxWorks Version 5.1

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

As a result of this validation effort, Validation Certificate 94022111.11340 is awarded to Tartan, Inc. This certificate expires 24 months after ANSI approval of MIL-STD 1815B.

This report has been reviewed and is approved.

10

Michael Tonndorf IABG, Abt. ITE Einsteinstr. 20 D-85521 Ottobrunn Germany

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

Deputy Director, Ada Joint Program Office Defense Information Systems Agency, Center for Information Management

Declaration of Conformance

Customer:	Tartan, Inc.
Certificate Awardee:	Tartan, Inc.
Ada Validation Facility:	IABG mbH
ACVC Version:	1.11

Ada Implementation: Ada Compiler Name and Version: Host Computer System: Target Computer System:

TartanWorks Ada 68xxx Version 4.3.1SPARC Station/ELC SunOS version 4.3.1MotorolaMVME167 (MC68040) runningVxWorks version 5.1

Declaration:

I, the undersigned, declare that I have no knowledge of deliberate deviations from the Ada Language Standard ANSI/MIL-STD-1815A, ISO 8652-1987, FIPS 119 as tested in this validation and documented in the Validation Summary Report.

Lee B. Ehrlichman Tartan, Inc. President and Chief Executive Officer

1/22/94

Date

CHAPTER 1	
1.1	USE OF THIS VALIDATION SUMMARY REPORT
1.2	REFERENCES
1.3	ACVC TEST CLASSES
1.4	DEFINITION OF TERMS
CHAPTER 2	
2.1	WITHDRAWN TESTS
2.2	INAPPLICABLE TESTS
2.3	TEST MODIFICATIONS
CHAPTER 3	
3.1	TESTING ENVIRONMENT
3.2	SUMMARY OF TEST RESULTS
3.3	TEST EXECUTION

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 [Pro92] 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 [Pro92]. 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, USA

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, USA

1-1

1.2 **REFERENCES**

- [Ada83] <u>Reference Manual for the Ada Programming Language</u>, ANSI/MIL-STD-1815A, February 1983 and ISO 8652-1987.
- [Pro92] <u>Ada Compiler Validation Procedures</u>, Version 3.1, Ada Joint Program Office, August 1992.
- [UG89] <u>Ada Compiler Validation Capability User's Guide</u>, 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 SPPRT13, 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 SPPRT13 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.

Ciass 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)

Implementation

Office (AJPO)

Facility (AVF)

Ada Joint

Validation

Validation

Organisation

Program

Ada

Ada

Ada

consisting of the test suite, the support programs, the ACVC user's guide and the template for the validation summary report.

The means for testing compliance of Ada implementations,

An Ada compiler with its host computer system and its target computer system.

The part of the certification body which provides policy and guidance for the Ada certification system.

The part of the certification body which carries out the procedures required to establish the compliance of an Ada implementation.

The part of the certification body that provides technical guidance for operations of the Ada certification system.

The ability of the implementation to pass an ACVC version.

Compliance of an Ada Implementation

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 userdesignated 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.

Conformity Fulfillment by a product, process or service of all requirements specified.

1-3

	INTRODUCTION
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 realised or is attainable on the Ada implementation for which validation status is realised.
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 Organisation for Standardisation.
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>.</paragraph></subsection></section>
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 [Pro92].
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 104 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 November 22, 1993.

B27005A	E28005C	B28006C	C32303A	C34006D	C35507K
C35507L	C35507N	C35507O	C35507P	C35508I	C35508J
C35508M	C35508N	C35702A	C35702B	C35310A	B41308B
C43004A	C45114A	C45346A	C45612A	C45612B	C45612C
C45651A	C46022A	B49008A	B49008B	A54B02A	C55B06A
A74006A	C74308A	383022B	B83022H	B83025B	B83025D
C83026A	B83026B	C83041A	B85001L	C86001F	C94021A
C97116A	C98003B	BA2011A	CB7001A	CB7001B	CB7004A
CC1223A	BC1226A	CC1226B	BC3009B	BD1802B	BD1B06A
BD1B08A	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 Al-ddddd. For this implementation, the following tests were determined to be inapplicable for the reasons indicated; references to Ada Commentaries are included as appropriate.

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

C24113LY (14 te	əsts) (C35705LY	(14	tests)
C35706LY (14 to	ests) (C35707LY	(14	tests)
C35708LY (14 te	əsts) (C35802LZ	(15	tests)
C45241LY (14 te	ests) (C45321LY	(14	tests)
C45421LY (14 te	ests) (C45521LZ	(15	tests)
C45524LY (15 te	ests) (C45621LZ	(15	tests)
C45641LY (14 te	ests) (C46012LZ	(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

C35713B, C45423B, B86001T, and C86006H check for the predefined type SHORT FLOAT; for this implementation, there is no such type.

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 hat 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_OVERFLOWS 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 OVERFLOWS is TRUE.

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

CA2009A, CA2009C..D (2 tests), CA2009F and BC3009C instantiate generic units before their bodies are compiled; this implementation creates a dependence on generic units as allowed by AI-00408 & AI-00506 such that the compilation of the generic unit bodies makes the instantiating units obsolete. (see 2.3.)

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.)

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.

CD2B15B checks that STORAGE_ERROR is raised when the storage size specified for a collection is too small to hold a single value of the designated type; this implementation allocates more space than was specified by the length clause, as allowed by AI-00558.

AE2101C and EE2201D..E (2 tests) use instantiations of package SEQUENTIAL_IO with unconstrained array types and record types with discriminants without defaults; these instantiations are rejected by this compiler.

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
CE2102D	CREATE	IN_FILE	SEQUENTIAL IO
CE2102E	CREATE	OUT_FILE	SEQUENTIALIO
CE2102F	CREATE	INOUT_FILE	DIRECT_IO
CE21021	CREATE	IN_FILE	DIRECTIO
CE2102J	CREATE	OUT_FILE	DIRECTIO
CE2102N	OPEN	IN_FILE	SEQUENTIAL_IO
CE21020	RESET	IN_FILE	SEQUENTIALIO
CE2102P	OPEN	OUT_FILE	SEQUENTIAL_IO
CE2102Q	RESET	OUT_FILE	SEQUENTIAL_IO
CE2102R	OPEN	INOUT_FILE	DIRECT_IO
CE2102S	RESET	INOUT_FILE	DIRECTIO
CE2102T	OPEN	IN_FILE	DIRECTIO
CE2102U	RESET	IN_FILE	DIRECT_IO
CE2102V	OPEN	OUT_FILE	DIRECTIO
CE2102W	RESET	OUT_FILE	DIRECTIO
CE3102E	CREATE	IN_FILE	TEXT_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

AE2101H, EE2401D, and EE2401G use instantiations of package DIRECT_IO with unconstrained array types and record types with discriminants without defaults; these instantiations are rejected by this compiler.

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.

2.3 TEST MODIFICATIONS

Modifications (see Section 1.3) were required for 108 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.

B22003A	B24007A	B24009A	B25002B	B32201A	B33204A
B33205A	B35701A	B36171A	B36201A	B37101A	B37102A
B37201A	B37202A	B37203A	B37302A	B38003A	B38003B
B38008A	B38008B	B38009A	B38009B	B38103A	B38103B
B38103C	B38103D	B38103E	B43202C	B44002A	B48002A
B48002B	B48002D	B48002E	B48002G	B48003E	B49003A
B49005A	B49006A	B49006B	B49007A	B49007B	B49009A
B4A010C	B54A20A	B54A25A	B58002A	B58002B	B59001A
B59001C	B590011	B62006C	B67001A	B67001B	B67001C
B67001D	B74103E	B74104A	B74307B	B83E01A	B85007C
B85008G	B85008H	B91004A	B91005A	B95003A	B95007B
B95031A	B95074E	BA1001A	BC1002A	BC1109A	BC1109C
BC1206A	BC2001E	BC3005B	BD2A06A	BD2B03A	BD2D03A
BD4003A	BD4006A	BD8003A			

E28002B was graded inapplicable by Evaluation and Test Modification as directed by the AVO. This test checks that pragmas may have unresolvable arguments, and it includes a check that pragma LIST has the required effect; but, for this implementation, pragma LIST has no effect if the compilation results in errors or warnings, which is the case when the test is processed without modification. This test was also processed with the pragmas at lines 46, 58, 70 and 71 commented out so that pragma LIST had effect.

Tests C45524A..K (11 tests) were graded passed by Test Modification as directed by the AVO. These tests expect that a repeated division will result in zero; but the Ada standard only requires that the result lie in the smallest safe interval. Thus, the tests were modified to check that the result was within the smallest safe interval by adding the following code after line 141; the modified tests were passed:

ELSIF VAL < = F'SAFE_SMALL THEN COMMENT ("UNDERFLOW SEEMS GRADUAL");

C83030C and C86007A were graded passed by Test Modification as directed by the AVO. These tests were modified by inserting "PRAGMA ELABORATE (REPORT);" before the package declarations at lines 13 and 11, respectively. Without the pragma, the packages may be elaborated prior to package report's body, and thus the packages' calls to function Report.Ident_Int at lines 14 and 13, respectively, will raise PROGRAM ERROR.

.

B83E01B was graded passed by Evaluation Modification as directed by the AVO. This test checks that a generic subprogram's formal parameter names (i.e. both generic and subprogram formal parameter names) must be distinct; the duplicated names within the generic declarations are marked as errors, whereas their recurrences in the subprogram bodies are marked as "optional" errors--except for the case at line 122, which is marked as an error. This implementation does not additionally flag the errors in the bodies and thus the expected error at line 122 is not flagged. The AVO ruled that the implementation's behavior was acceptable and that the test need not be split (such a split would simply duplicate the case in B83E01A at line 15).

CA2009A, CA2009C..D (2 tests), CA2009F and BC3009C were graded inapplicable by Evaluation Modification as directed by the AVO. These tests instantiate generic units before those units' bodies are compiled; this implementation creates dependences as allowed by AI-00408 & AI-00506 such that the compilation of the generic unit bodies makes the instantiating units obsolete, and the objectives of these tests cannot be met.

- BC3204C and BC3205D were graded passed by Processing Modification as directed by the AVO. These tests check that instantiations of generic units with unconstrained types as generic actual parameters are illegal if the generic bodies contain uses of the types that require a constraint. However, the generic bodies are compiled after the units that contain the instantiations, and this implementation creates a dependence of the instantiating units on the generic units as allowed by Al-00408 & Al-00506 such that the compilation of the generic bodies makes the instantiating units obsolete--no errors are detected. The processing of these tests was modified by compiling the seperate files in the following order (to allow re-compilation of obsolete units), and all intended errors were then detected by the compiler:

BC3204C: C0, C1, C2, C3M, C4, C5, C6, C3M

BC3205D: D0, D1M, D2, D1M

BC3204D and BC3205C were graded passed by Test Modification as directed by the AVO. These tests are similar to BC3204C and BC3205D above, except that all compilation units are contained in a single compilation. For these two tests, a copy of the main procedure (which later units make obsolete) was appended to the tests; all expected errors were then detected.

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

AD9001B and AD9004A were graded passed by Processing Modification as directed by the AVO. These tests check that various subprograms may be interfaced to external routines (and hence have no Ada bodies). This implementation requires that a file specification exists for the foreign subprogram bodies. The following command was issued to the Librarian to inform it that the foreign bodies will be supplied at link time (as the bodies are not actually needed by the program, this command alone is sufficient):

interface -sys -L=library ad9001b & ad9004a

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 information about this Ada implementation, contact:

Mr Ken Butler Director of Ada Products Tartan Inc. 300 Oxford Drive Monroeville, PA 15146 USA Tel. (412) 856-3600

For sales information about this Ada implementation, contact:

Ms. Marlyse Bennett Director of Sales Tartan, Inc. 12110 Sunset Hills Road Suite 450 Reston, VA 22090 USA Tel. (703) 715-3044

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 [Pro92].

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

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	.3779
b)	Total Number of Withdrawn Tests	104
C)	Processed Inapplicable Tests	86
d)	Non-Processed I/O Tests	0
e)	Non-Processed Floating-Point Precision Tests	201
f)	Total Number of Inapplicable Tests	287 (c+d+e)
g)	Total Number of Tests for ACVC 1.11	4170 (a+b+f)

3.3 TEST EXECUTION

A magnetic data cartridge containing the customized test suite (see section 1.3) was taken on-site by the validation team for processing. The contents of the magnetic data cartridge were loaded directly onto the host computer.

The tests were compiled and linked on the host computer system, as appropriate. The executable images were transferred to the target computer system by the communications link, an ethernet interface, and run. The results were captured on the host 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 for compiling:

- -f forces the compiler to accept an attempt to compile a unit imported from another library which is normally prohibited.
- -c suppresses the creation of a registered copy of the source code in the library directory for use by the REMAKE and MAKE subcommands.

For this validation the default optimization level -Op2 was used. No explicit Linker options were set.

Test output, compiler and linker listings, and job logs were captured on magnetic data cartridge 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	240
\$BIG_ID1	(1V-1 => 'A', V => '1')
\$BIG_1D2	(1V-1 => 'A', V => '2')
\$BIG_ID3	(1V/2 => 'A') & '3' & (1V-1-V/2 => 'A')
\$BIG_ID4	(1V/2 => 'A') & '4' & (1V-1-V/2
\$BIG_INT_LIT	(1V-3 => '0') & "298"
\$BIG_REAL_LIT	(1V-5 => '0') & "690.0"
\$BIG_STRING1	'"' & (1V/2 => 'A') & '"'
\$BIG_STRING2	'"' & (1V-1-V/2 => 'A') & '1' & '"'
\$BLANKS	(1V-20 => ' ')
\$ILLEGAL_EXTERNAL_FILE_NAM	E1 GAL_EXTERNAL_FILE_NAME1" & (1V => '-')
\$ILLEGAL_EXTERNAL_FILE_NAM "ILLEG	E2 GAL_EXTERNAL_FILE_NAME2" & (1V = > '-')
\$MAX_LEN_INT_BASED_LITERAL	. "2:" & (1V-5 => '0') & "11:"
\$MAX_LEN_REAL_BASED_LITER	AL"16:" & (1V-7 => '0') & "F.E:"
\$MAX STRING LITERAL	'"' & (1V-2 => 'A') & '"'

Macro Parameter	Macro Value
\$ACC_SIZE	32
\$ALIGNMENT	4
\$COUNT_LAST	2147483646
\$DEFAULT_MEM_SIZE	1_000_000
\$DEFAULT_STOR_UNIT	8
\$DEFAULT_SYS_NAME	MC68020
\$DELTA_DOC	2#1.0#E-31
\$ENTRY_ADDRESS	SYSTEM.ADDRESS'(16#F0#)
\$ENTRY_ADDRESS1	SYSTEM.ADDRESS'(16#F1#)
\$ENTRY_ADDRESS2	SYSTEM.ADDRESS'(16#F2#)
\$FIELD_LAST	240
\$FILE_TERMINATOR	
\$FIXED_NAME	NO_SUCH_FIXED_TYPE
\$FLOAT_NAME	THERE_IS_NO_SUCH_FLOAT_NAME
\$FORM_STRING	11 FF
\$FORM_STRING2	"CANNOT_RESTRICT_FILE_CAPACITY"
\$GREATER_THAN_DURATION	100_000.0
\$GREATER_THAN_DURATION_BA	NSE_LAST 100_000_000.0
\$GREATER_THAN_FLOAT_BASE_	LAST 3.5E+38
\$GREATER_THAN_FLOAT_SAFE_	LARGE 1.0E + 38
\$GREATER_THAN_SHORT_FLOA	T_SAFE_LARGE 1.0E+38

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

÷

A-2

MACRO PARAMETERS

\$HIGH_PRIORITY	200
\$INAPPROPRIATE_LINE_LENGTH	-1
\$INAPPROPRIATE_PAGE_LENGTH	-1
\$INCLUDE_PRAGMA	PRAGMA INCLUDE ("A28006D1.TST")
\$INCLUDE_PRAGMA2	PRAGMA INCLUDE ("B28006F1.TST")
\$INTEGER_FIRST	-2147483648
\$INTEGER_LAST	2147483647
\$INTEGER_LAST_PLUS_1	2147483648
\$INTERFACE_LANGUAGE	С.
\$LESS_THAN_DURATION	-100_000.0
\$LESS_THAN_DURATION_BASE_FI	RST -100_000_000.0
\$LINE_TERMINATOR	••
\$LOW_PRIORITY	10
\$MACHINE_CODE_STATEMENT	Two_Opnds'(MOVE_L, (IMM,2),(DR, d7));
\$MACHINE_CODE_TYPE	Address_Mode
\$MANTISSA_DOC	31
\$MAX_DIGITS	15
\$MAX_INT	2147483647
\$MAX_INT_PLUS_1	2147483648
\$MIN_INT	-214748368
\$NAME	BYTE_INTEGER
\$NAME_LIST	MC68020
\$NEG_BASED_INT	8#77777777776#
\$NEW_MEM_SIZE	1_000_000
\$NEW_STOR_UNIT	8
\$NEW_SYS_NAME	MC68020

:'

-

A-3

\$PAGE_TERMINATOR

\$RECORD_DEFINITION

record Operation: Instruction_Mnemonic; Operand_1: Operand; end record;

\$RECORD_NAME \$TASK_SIZE

\$TASK_STORAGE_SIZE

\$TICK

\$VARIABLE_ADDRESS

\$VARIABLE_ADDRESS1

\$VARIABLE_ADDRESS2

One_Opnds

96

4096

. .

0.01

SYSTEM.ADDRESS'(16#99000#)

SYSTEM.ADDRESS'(16#99004#)

SYSTEM.ADDRESS'(16#99008#)

APPENDIX B

COMPILATION AND LINKER SYSTEM OPTIONS

The compiler and 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 compiler documentation and not to this report.

•

B-1

3.2. UNIX COMMAND LINE OPTIONS

Command line options indicate special actions to be performed by the compiler or special output file properties.

The following UNIX command line options may be specified:

- a Generates an assembly code file. The assembly code file has an extension . s for a body or . s s for a specification (sec. 3.5). In the default mode, no assembly code is generated.
 - Generates an assembly code file with interleaved source code. The assembly code file has an extension . s for a body or . s s for a specification.

Ba

A

Bm

C

Cs

Specifies that the compiler will produce an optimization (.opt) file which contains special optimization information, when the unit being compiled is a body. When another unit is compiled which refers to this unit in its context clause, a dependency may be created on this unit's body (in addition to the specification) due to the utilization of this optimization information. Also, the compiler will attempt to utilize optimization information from the optimization files, of units named in the context clause of the current unit. Dependencies will be created on both the specification and the body (if any) of the units from which optimization information is utilized. This option will allow maximum optimization at the expense of increased recompilations when changes are made.

By default, the compiler will *not* produce an optimization file for the current unit (effectively preventing the creation of dependencies on this body), but will read the .opt files of units in the current unit's compilation closure to obtain information which may be used to improve the optimizations performed on the current unit.

Specifies that the compiler will neither produce an optimization (.opt) file when the unit being compiled is a body (effectively preventing the creation of dependencies on this body), nor will the compiler attempt to utilize optimization information from units named in the context clause of the current unit (preventing the possibility of creating a dependency on another body). When compiling an entire system, this strategy will lead to minimal dependencies between the compilation units in the system.

By default, the compiler will *not* produce an optimization file for the current unit, but will read the .opt files of units in the current unit's compilation closure to obtain information which may be used to improve the optimizations performed on the current unit.

Normally, the compiler creates a registered copy of the user's source code in the library directory for use with the librarian $remake^2$ subcommand. This option suppresses the creation of this copy.

Controls whether the compiler generates 16-bit instead of 32-bit PC relative address modes when addressing objects whose distance from the current location cannot be determined at compile time. By using this option, the user asserts that the program space for the final program will be small enough for all calls to use the 16-bit PC-relative address modes in call instructions. If this assertion is incorrect, erroneous code could result.

²UNIX: remakecu

When compiling a library unit, determines whether the unit is a refinement of its previous version and, if so, do not make dependent units obsolete. This check is not done by default. A warning message is given if the unit is not a refinement of its previous version. The *no update* option³ may be used in conjunction with this option to check for possible refinements without risking a change to the program library.

Stops compilation and produces a listing after n errors are encountered, where n is an integer in the range 0...255. The default value for n is 255.

Forces the compiler to accept an attempt to compile a unit imported from another library, which is normally prohibited.

Produces debugging information for AdaScope, the Tartan Ada symbolic debugger. It is not necessary for all object modules to include debugging information to obtain a linkable image, but use of this option is encouraged for all compilations. No significant execution-time penalty is incurred with this option.

Causes the compiler to omit data segments with the text of enumeration literals. This text is normally produced for exported enumeration types in order to support the text attributes ('IMAGE, 'VALUE and 'WIDTH). You should use this option only when you can guarantee that no unit that will import the enumeration type will use any of its text attributes. However, if you are compiling a unit with an enumeration type that is not visible to other compilation units, this option is not needed. The compiler can recognize when the text attributes are not used and will not generate the supporting strings.

Causes the compiler options specified for this compilation unit to be saved in the program library. These options will be used when a subsequent (re)make command is issued on this unit, unless overidden by compiler options specified on the (re)make command line (sec. 10.2.5).

L=[project:] library Selects the library and optionally the project for this compilation. This option takes effect after all commands from the librarian initialization file⁴ have been executed, thereby possibly overriding its effects.

Always generates a listing. The default is to generate a listing only if a diagnostic message is issued.

Never generates a listing. The default is to generate a listing only if a diagnostic message is issued.

Specifies the line width used in a listing, where n is an integer in the range 80 ... 132. The default value for n is 80.⁵

Controls the type of messages that will be generated by the compiler. The available options are:

e Reports only errors.

i.

Reports errors, warnings, and informational messages.

³UNIX: -n

d

e=/

f

a

i

ĸ

La

1.n

lw=n

m-option

UNIX: .adalibrc

⁵Note that 10 less characters than the user specifies on the command line will actually appear on a line in the listing file due to the left and right margins and the line numbers.

v4.3.1

٤

3-4

The default condition, for which there is no option, is to report errors and warnings.

When package MACHINE_CODE is used, controls whether the compiler attempts to alter operand address modes when those address modes are used incorrectly. With this option, the compiler does not attempt to fix any machine code insertion that has incorrect address modes. An error message is issued for any incorrect machine code insertion. By default, when neither Me or Mw is specified, the compiler attempts to generate extra instructions to fix incorrect address modes in the array aggregates operand field.

The compiler attempts to generate extra instructions to fix incorrect address modes. A warning message is issued if such a *fixup* is required. By default, when neither Me or Mw is specified, the compiler attempts to generate extra instructions to fix incorrect address modes in the array aggregates operand field.

Specifies that the program library will not be updated with the result of this compilation.

Controls the level of optimization performed by the compiler, where n is an integer in the range 0..4.

The following optimization levels may be specified:

- n = 0 reinimum Performs context determination, constant folding, algebraic manipulation, and short circuit analysis. Pragma INLINEs are not obeyed.
- n=1 low Performs minimum optimizations plus evaluation order determination as well as common subexpression elimination and equivalence propagation within basic blocks. Again, pragma INLINEs are not obeyed.
- n=2 standard (default) Best tradeoff for space/time. Performs low optimizations plus flow analysis for common subexpression elimination and equivalence propagation across basic blocks, invariant hoisting, dead code elimination, assignment killing, strength reduction, lifetime analysis for improved register allocation, tail recursion elimination, interprocedural side-effect analysis and some inline expansion.
- n=3 time Performs standard optimizations plus loop unrolling and aggressive inline expansion. This optimization level usually produces the fastest code, however, optimization level standard may produce faster code under certain circumstances.
- n = 4 space Performs standard optimizations minus any optimization that may increase code size. This optimization level usually produces the smallest code, however, optimization level standard may produce smaller code under certain circumstances.

Loads a syntactically correct compilation unit(s) in the source file into the library as a parsed unit(s). Parsed units are, by definition, inconsistent. This option allows you to load units into the library without regard to correct compilation order. The librarian *remake* subcommand⁶ is subsequently used to compile the units in the correct sequence (sec. 10.2.5.2).

v4.3.1

p

0p*n*

Mw

n

Me

UNIX: remakecu

Examines units for syntax errors, then stops compilation. No semantic checking is performed. Nothing is entered in the program library. No library need be specified when using this option.

S[ACDEILORSZ]

s

Suppresses the given set of checks:

A	ACCESS CHECK
С	CONSTRAINT CHECK
D	DISCRIMINANT CHECK
Ε	ELABORATION CHECK
I	INDEX CHECK
L	LENGTH CHECK
0	OVERFLOW CHECK
R	RANGE CHECK
S	STORAGE CHECK
Z	"ZERO" DIVISION CHECK

The S option has the same effect as an equivalent pragma SUPPRESS applied to the source file. If the source program also contains a pragma SUPPRESS, a given check is suppressed if either the pragma or the option specifies it; that is, the effect of a pragma SUPPRESS cannot be negated with the command line option. See LRM 11.7 for further details. Supplying the S option can significantly decrease the size and execution time of the compiled code. Examples are:

SOZ	Suppresses	OVERFLOW_CHECK	and
	"ZERO" DIVISION	CHECK.	
S	Suppresses all checks, checks if the resulting	Invoking this option will ne code without checks will be les	ot remove all s efficient.
	-		

tgt=processor

SC Suppresses CONSTRAINT_CHECK, equivalent to SADILR.

Specifies the Motorola target processor, where *processor* may be one of the following:

mc 68020 for all supported boards using the MC68020 processor

mc 68030 for all supported boards using the MC68030 processor

mc 68040 for all supported boards using the MC 68040 processor

cpu32 for all supported boards using the CPU32 processor

Compilations for a 68xxx processor take place in a specific development environment called a *universe*. Tartan supplies three universes for the 68xxx product:

- 68020/68030
- 68040
- CPU32

When using the librarian (re)make subcommand, the librarian sets the compiler target option to mc 68020 for the 68020/68030 universe. To override the default setting for 68030 specific compilations, you must specify mc 68030 as an argument to the (re)make command line option -q "-tgt=mc 68030" (sec. 10.3.25).

Prints out compiler phase names. The compiler prints out a short description of each compilation phase in progress.

x

ν

Includes cross reference information for the source in the object file (sec. 3.7).

Note: On UNIX, the output from the compiler may be redirected using the UNIX redirection facility including '6' for stderr; for example:

tadaprocessor tax spec.ada >6 tax_spec.txt

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, are outlined below for convenience.

package STANDARD is

type BYTE_INTEGER is range -128 .. 127; type SHORT_INTEGER is range -32768 .. 32767; type INTEGER is range -2_147 483 648 .. 2_147_483 647;

type DURATION is delta 0.0001 range -86400.0 .. 86400.0;

end STANDARD;

CHAPTER 4

APPENDIX F TO MIL-STD-1815A

This chapter contains the required Appendix F to the LRM, which is Military Standard, Ada Programming Language, ANSI/MIL-STD-1815A.

4.1. PRAGMAS

4.1.1. Predefined Pragmas

The Tartan Ada Compiler supports all of the predefined pragmas described in the LRM, annex B.

- pragma CONTROLLED (sec. 4.1.1.1)
- pragma ELABORATE
- pragma INLINE (sec. 4.1.1.2)
- pragma INTERFACE (sec. 4.1.1.3)
- pragma LIST
- pragma MEMORY_SIZE (sec. 4.1.1.4)
- pragma OPTIMIZE (sec. 4.1.1.5)
- pragma PACK (sec. 4.4.6)
- , pragma PAGE
- pragma PRIORITY
- pragma SHARED
- pragma STORAGE_UNIT (sec. 4.1.1.4)
- pragma SUPPRESS
- pragma SYSTEM NAME (sec. 4.1.1.4)

The following sections summarize the effects of and restrictions on certain predefined pragmas.

4.1.1.1. Pragma CONTROLLED

Access collections are not subject to automatic storage reclamation so pragma CONTROLLED has no effect. Space deallocated by means of UNCHECKED DEALLOCATION will be reused by the allocation of new objects.

4.1.1.2. Pragma INLINE

Pragma INLINE is supported as described in the LRM 6.3.2, with the following restrictions and clarifications:

- The body of the subprogram to be expanded inline must be compiled before the unit that calls the subprogram. If the call is compiled prior to the subprogram body, inline expansion of that call will not be performed. A warning is issued when a call is not inlined because the body has not been compiled.
- If a unit contains a call that results in inlined code, any subsequent recompilation of the body of the called subprogram will make the unit containing the inlined call obsolete.

- When inlining across libraries, the body of the subprogram to be inlined must be exported from a frozen specification library (sec. 10.3.13 and 10.3.17).
- The optimization level, as set by the compiler command line option, determines whether an attempt is made to obey a pragma INLINE (sec. 8.2).
 If the compilation containing a call to the subprogram named in a pragma INLINE is compiled at the minimum or low optimization level, (UNIX: -Op0 or -Op1; VMS: /optimize = minimum or /optimize = low) inlining will not be attempted for that call.
- Inlining may not be performed if the compiler determines that the subprogram to be inlined is too complex. Typical examples are subprograms that recursively call themselves, or whose objects are referenced by enclosing subprograms.

4.1.1.3. Pragma INTERFACE

Pragma INTERFACE is supported as described in LRM 13.9.

The pragma associates a particular calling sequence with a subprogram whose implementation is provided in the form of an object code module. The librarian interface subcommand < All hosts: interface > (sec. 10.3.20) must be used to identify the associated object code module.

The language_name may be Ada, Assembly, or C. Any other language_name will be accepted, but ignored, and the default language, Ada, will be used.

It is almost always necessary to use a pragma LINKAGE_NAME (sec. 4.1.2.1) for interfaced subprograms. Without the LINKAGE_NAME pragma, the user must determine the compressed name the compiler generates and use that name in the provided object module.

An interfaced subprogram cannot have a direct Ada implementation, i.e., a body is not allowed for such a subprogram. It is possible to compile an Ada subprogram with a different name and then use the librarian interface subcommand to reference that subprogram.

4.1.1.4. Pragmas MEMORY_SIZE, STORAGE_UNIT, and SYSTEM_NAME

This section details the procedure for compiling one of these pragmas. The compilation containing the pragma must be compiled into a library that contains package SYSTEM. For most users, the Tartan Ada Standard Packages Library will be the library that includes package SYSTEM. In that case, the procedure is as follows:[This procedure will not cause any of the units in the Tartan Ada Standard Packages Library to become obsolete.]

- 1. Thaw standard_packages.spec.
- 2. Compile the pragma into the library standard_packages.root. This step updates package SYSTEM. Any unit that depends on SYSTEM becomes obsolete and will require recompilation before it may be used in a program.

3. Freeze standard_packages.spec.

For pragma STORAGE_UNIT, no value other than that already specified by SYSTEM.STORAGE UNIT (sec. 4.3) is allowed.

For pragma SYSTEM_NAME, no value other than that already specified by SYSTEM.SYSTEM_NAME (sec. 4.3) is allowed.

4.1.1.5. Pragma OPTIMIZE

Pragma OPTIMIZE is supported as described in the LRM, annex B with the following exception. A pragma OPTIMIZE only has an effect when placed in a subprogram's declarative part. The pragma is applied only to that subprogram and not to any nested subprograms.

The argument applied to pragma OPTIMIZE (space or time) directly corresponds to the same argument supplied with the optimization option on the compiler command line. For example, specifying time as the argument to pragma OPTIMIZE has the same effect as compiling the subprogram and specifying optimization level time(UNIX: -Op3; VMS: /optimize = time) on the command line.

4.1.2. Implementation-Defined Pragmas

Implementation-defined pragmas provided by Tartan are described in the following sections.

4.1.2.1. Pragma LINKAGE NAME

The pragma LINKAGE_NAME associates an Ada entity with a string that is meaningful externally; for example, to a linkage editor. It takes the form

pragma LINKAGE_NAME (name, string-constant)

The pragma is only allowed in a package specification or in a declarative part, or after a library subprogram in a compilation before any subsequent compilation unit.

If the pragma appears in a library package specification the name must denote an entity declared earlier in the same package. If the pragma appears in any other package specification or in a declarative part, the name must denote an entity declared earlier in the same package or declarative part, and must denote either a subprogram or an exception. If the pragma appears after a given library subprogram, the only name allowed is the name of this subprogram.

The name must be the simple name or operator symbol of an Ada entity. The name refers only to the most recently declared entity with the given name, not to all of the overloadings of the name.

The name should denote an entity that has a runtime representation; for example, a subprogram, an exception, or an object. If the name denotes an entity that has no runtime representation the pragma has no effect; for example, named numbers, generic units, and most constants with values known at compile-time do not have runtime representations. The pragma also will have

no effect if the name is one declared by a renaming declaration.

The effect of the pragma is to cause the string-constant to be used in the generated object code as an external name for the associated Ada entity. It is the responsibility of the user to guarantee that this string constant is meaningful to the linkage editor and that no illegal linkname clashes arise. < Names given in the string-constant argument of a pragma LINKAGE_NAME are case sensitive. For example, aNy_Old_LINKname is not equivalent to ANY OLD_LINKNAME. Therefore, a misspelled linkname will cause the link to fail. >

When determining the maximum allowable length for the external linkage name, keep in mind that the compiler will generate names for elaboration flags for subprograms simply by appending a five-character suffix to the linkage name. Therefore, a linkage name for a subprogram may have five fewer characters than the lower limit of other tools that need to process the name (e.g., the Tartan Linker limits names to 40 characters; therefore, your external linkage name should not exceed 35 characters).

4.1.2.2. Pragma FOREIGN_BODY

In addition to pragma INTERFACE, Tartan Ada supplies pragma FOREIGN_BODY as a way to access entities defined in programs written in other languages. Use of the pragma FOREIGN_BODY dictates that all subprograms and objects in the package are provided by means of a foreign object module. Unlike pragma INTERFACE, pragma FOREIGN BODY allows access to objects as well as subprograms.

The pragma is of the form:

pragma FOREIGN_BODY (language_name [, elaboration_routine_name])

A single such pragma may appear in any non-generic library package, and must appear in the visible part of the package before any declarations. The pragma is only permitted when the declarations in the visible and private parts of the package consist of subprogram declarations, number declarations, and object declarations with no explicit initialization and with a subtype given by a simple type mark. Use clauses and other pragmas may also appear in the package specification. If any of these restrictions are violated, the pragma is ignored and a warning is generated. Note in particular that types, exceptions, packages, and generic units may not be declared in the package.

The language_name argument is a string intended to identify the language processor used to create the foreign module. It is treated as a comment by the compiler.

The optional elaboration_routine_name argument is a string giving the linkage name of a routine to initialize the package. The routine specified will be called for the elaboration of this package body. It must be a global routine in the object module provided by the user.

The programmer must ensure that the calling convention and data representation of the foreign body subprograms and elaboration routine are compatible with those used by the Tartan Ada Compiler (sec. 5.4).

In order to successfully link a program including a foreign body, the object

module for that body must be provided to the library using the librarian foreign subcommand < UNIX: foreign; VMS: foreign > (sec. 2.3.3 and 10.3.16).

All entities declared by the package must be supplied by the foreign object module. Pragma LINKAGE_NAME will usually have to be used to ensure agreement between the linkage names used by the Tartan Ada Compiler and the foreign language processor.

The foreign body is entirely responsible for initializing objects declared in a package utilizing pragma FOREIGN_BODY. In particular, the user should be aware that the implicit initializations described in LRM 3.2.1 are not done by the compiler. (These implicit initializations are associated with objects of access types, certain record types and composite types containing components of the preceding kinds of types.)

The user may choose to override the pragma FOREIGN_BODY and compile a corresponding package body written in Ada. In this case the pragma is ignored (in particular the specified elaboration routine is not called), and no librarian foreign subcommand is required or allowed. This capability is useful for rapid prototyping, where an Ada package may serve to provide a simulated response for the functionality that a foreign body may eventually produce. It also allows the user to replace a foreign body with an Ada body without recompiling the specification.

If only subprograms are declared in the package specification it is more portable to use pragma INTERFACE on each of the subprograms instead of pragma FOREIGN_BODY on the package.

In the following example, we want to call a function plmn which computes polynomials and is written in C.

```
package MATH_FUNCTIONS is
```

```
pragma FOREIGN_BODY("C");
function POLYNOMIAL(X:INTEGER) return INTEGER;
-- Ada spec matching the C routine
pragma LINKAGE_NAME(POLYNOMIAL, "plmn");
-- Force compiler to use name "plmn" when referring to this
```

-- function

-- Note: The linkage name "plmn" may need to be " plmn",

-- if the C compiler produces leading underscores

-- for external symbols.

end MATH_FUNCTIONS;

```
with MATH_FUNCTIONS; use MATH_FUNCTIONS;
procedure MAIN is
X:INTEGER := POLYNOMIAL(10);
-- Will generate a call to "plmn"
begin ...
end MAIN;
```

To compile, link and run the above program, you must:

1. Compile MATH_FUNCTIONS.

2. Compile MAIN.

3. Provide the object module (for example, math.tof) containing the compiled `C'' code for plmn, converted to Tartan Object File Format (TOFF); if the module is written in assembly code, for example, using the IEEE-to-TOFF utility (Object File Utilities Manual, ch. 4).

4. Issue the command:

UNIX:

adalib < processor > foreign math_functions math.tof

VMS:

al68 foreign math functions math.tof

5. Issue the command:

UNIX:

adalib < processor > link main

VMS:

al68 link main

Without step 4, an attempt to link will produce an error message informing you of a missing package body for MATH FUNCTIONS.

4.1.2.3. Pragma UNCHECKED_NO_STATE_WRITTEN and Pragma UNCHECKED_NO_STATE_WRITTEN_OR_READ

The pragmas UNCHECKED_NO_STATE_WRITTEN and UNCHECKED_NO_STATE_WRITTEN_OR_READ take the form:

pragma UNCHECKED_NO_STATE_WRITTEN(name [, name...]) pragma UNCHECKED_NO_STATE_WRITTEN_OR_READ(name [, name...])

Each name must be the simple name of an Ada subprogram declared in the declarative part or package specification where the pragma appears. The name refers only to the most recently declared subprogram with the given name, not to all of the overloadings of the name.

The pragma UNCHECKED_NO_STATE_WRITTEN notifies the compiler that the named subprogram has no side effects on any objects outside the subprogram. Assignment to in out or out parameters is not considered a side effect.Function results are also not considered to be side effects. Calling another subprogram is considered to be a side effect, unless the called subprogram is also named in either a pragma UNCHECKED_NO_STATE_WRITTEN or pragma UNCHECKED NO STATE WRITTEN OR READ.

This pragma permits the compiler to improve the optimization performed near calls to the named subprogram without introducing a dependency on the body of the subprogram. In effect, global side effect analysis is achieved without creating additional dependencies which may require recompilation.

Any function which writes only to its result, or any subprogram which writes only to its in out or out parameters is an excellent candidate for this pragma.

The pragma UNCHECKED_NO_STATE_WRITTEN_OR_READ indicates that the named subprogram behaves strictly as a mathematically pure function. In essence, this means that the subprogram will always return the same result when called with identical parameters. The named subprogram must follow all of the rules for an UNCHECKED_NO_STATE_WRITTEN subprogram. In addition, the named subprogram may not read the value of any variable not contained within is undefined and may be unpredictable.

4.2. IMPLEMENTATION-DEPENDENT ATTRIBUTES

4.2.1. 'EXCEPTION ADDRESS

The attribute 'EXCEPTION_ADDRESS used with a prefix that denotes an exception yields the storage address associated with the exception. The value of this attribute is of the type ADDRESS defined in the package SYSTEM.

4.3. SPECIFICATION OF THE PACKAGE SYSTEM

The parameter values specified for the 68xxx target in package SYSTEM (LRM 13.7.1 and Annex C) are:

package SYSTEM is type ADDRESS is new U_ADDRESS; type NAME is (MC68020); SYSTEM_NAME : constant NAME := MC68020; STORAGE_UNIT : constant := 8; MEMORY_SIZE : constant := 1_000_000; MAX_INT : constant := 2_147_483_647; MIN_INT : constant := 2_147_483_647; MIN_INT : constant := -MAX_INT - 1; MAX_DIGITS : constant := 15; MAX_MANTISSA : constant := 31; FINE_DELTA : constant := 2#1.0#e-31; TICK : constant := 0.01;

subtype PRIORITY is INTEGER range 10 .. 200; DEFAULT_PRIORITY : constant PRIORITY := PRIORITY'FIRST; RUNTIME_ERROR : exception; end SYSTEM;

4.4. RESTRICTIONS ON REPRESENTATION CLAUSES

The following sections explain the basic restrictions for representation specifications followed by additional restrictions applying to specific kinds of clauses.

4.4.1. Basic Restriction

The basic restriction on representation specifications (LRM 13.1) is that they may be given only for types declared in terms of a type definition, excluding

- a G type in a smaller size, even if possible. The following rules apply with regard to feasibility:
 - An object that is not a component of a composite object is allocated with a size and alignment that is referable on the target machine (i.e., no attempt is made to create objects of non-referable size on the stack). If such stack compression is desired, it can be achieved by the user by combining multiple stack size and restrictions.

4.4.2.1. Size Specifications for Types

The rules and restrictions for size specifications applied to types of various classes are described below.

The following principle rules apply:

- 1. The size is specified in bits and must be given by a static expression.
- 2. The specified size is taken as a mandate to store objects of the type in the given size wherever feasible. No attempt is made to store values of the type in a smaller size, even if possible. The following rules apply with regard to feasibility:
 - An object that is not a component of a composite object is allocated with a size and alignment that is referable on the target machine (i.e., no attempt is made to create objects of non-referable size on the stack). If such stack compression is desired, it can be achieved by the user by combining multiple stack variables in a composite object; for y other

components of the composite object (i.e., whenever possible, a component of non-referable size is made referable).

In all cases, the compiler generates correct code for all operations on objects of the type, even if they are stored with differing representational sizes in different contexts.

Note: A size specification cannot be used to force a certain size in value operations of the type; for example:

type MY_INT is range 0..65535; for MY_INT'SIZE use 16; -- o.k. A,B: MY_INT; ...A + B... -- this operation will generally be -- executed on 32-bit values

3. A size specification for a type specifies the size for objects of this type and of all its subtypes. For components of composite types, whose subtype would allow a shorter representation of the component, no attempt is made to take advantage of such shorter representations.

C-9

For example, consider the following:

type MY_INT is range 0..2**17-1; for MY_INT'SIZE use 17; -- (1) subtype SMALL_MY_INT is MY_INT range 0..255; type R is record

X: SMALL_MY_INT;

end record;

The component R.X will occupy 17 bits even though it can be represented in 8 bits. If a pragma PACK(R) is added, R.X will still be allocated in 17 bits.

In contrast, for types without a size specification, such components may be represented in a lesser number of bits than the number of bits required to represent all values of the type. In the example above, if the size specification at (1) is removed, R.X will be represented in 32 bits (the size of MY_INT). However, a pragma PACK(R) will now cause R.X to be allocated in 8 bits.

Size specifications for access types must coincide with the default size chosen by the compiler for the type.

Size specifications are not supported for floating-point types or task types.

No useful effect can be achieved by using size specifications for these types.

4.4.2.2. Size Specification for Scalar Types

The specified size must accommodate all possible values of the type including the value O, even if O is not in the range of the values of the type. For numeric types with negative values, the number of bits must acccunt for the sign bit. No skewing of the representation is attempted. Thus,

type MY_INT is range 100..101;

requires at least 7 bits, although it has only two values, while

type MY INT is range -101..-100;

requires 8 bits to account for the sign bit.

A size specification for a real type does not affect the accuracy of operations on the type. Such influence should be exerted via the ACCURACY_DEFINITION of the type (LRM 3.5.7, 3.5.9).

A size specification for a scalar type may not specify a size larger than the largest operation size supported by the target architecture for the respective class of values of the type.

4.4.2.3. Size Specification for Array Types

A size specification for an array type must be large enough to accommodate all components of the array under the densest packing strategy. Any alignment constraints on the component type (sec. 4.4.7) must be met.

The size of the component type cannot be influenced by a length clause for an array. Within the limits of representing all possible values of the component subtype (but not necessarily of its type), the representation of components may, however, be reduced to the minimum number of bits, unless the component type carries a size specification.

If there is a size specification for the component type, but not for the array type, the component size is rounded up to a referable size, unless pragma PACK is given. This rule applies even to boolean types or other types that require only a single bit for the representation of all values.

4.4.2.4. Size Specification for Record Types

A size specification for a record type does not influence the default type mapping of a record type. The size must be at least as large as the number of bits determined by type mapping. Influence over packing of components can be exerted by means of (partial) record representation clauses or by pragma PACK.

Neither the size of component types, nor the representation of component subtypes can be influenced by a length clause for a record.

The only implementation-dependent components allocated by Tartan Ada in records contain either dope information for arrays whose bounds depend on discriminants of the record or relative offsets of components within a record layout for record components of dynamic size. These implementation-dependent components cannot be named or sized by the user.

A size specification cannot be applied to a record type with components of dynamically determined size.

Note: Size specifications for records can be used only to widen the representation accomplished by padding at the beginning or end of the record. Any narrowing of the representation over default type mapping must be accomplished by representation clauses or pragma PACK.

4.4.2.5. Specification of Collection Sizes

The specification of a collection size causes the collection to be allocated with the specified size. It is expressed in storage units and need not be static; refer to package SYSTEM for the meaning of storage units.

Any attempt to allocate more objects than the collection can hold causes a STORAGE_ERROR exception to be raised. Dynamically sized records or arrays may carry hidden administrative storage requirements that must be accounted for as part of the collection size. Moreover, alignment constraints on the type of the allocated objects may make it impossible to use all memory locations of the allocated collection. No matter what the requested object size, the allocator must allocate a minimum of 2 words per object. This lower limit is

necessary for administrative overhead in the allocator. For example, a request of 5 words results in an allocation of 5 words; a request of one (1) word results in an allocation of 2 words.

In the absence of a specification of a collection size, the collection is extended automatically if more objects are allocated than possible in the collection originally allocated with the compiler-established default size. In this case, STORAGE_ERROR is raised only when the available target memory is exhausted. If a collection size of zero is specified, no access collection is allocated.

4.4.2.6. Specification of Task Activation Size

The specification of a task activation size causes the task activation to be allocated with the specified size. It is expressed in storage units; refer to package SYSTEM for the meaning of storage units.

Any attempt to exceed the activation size during execution causes a STORAGE_ERROR exception to be raised. Unlike collections, there is no extension of task activations.

4.4.2.7. Specification of 'SMALL

Only powers of 2 are allowed for 'SMALL.

The length of the representation may be affected by this specification. If a size specification is also given for the type, the size specification takes precedence; it must then be possible to accommodate the specification of 'SMALL within the specified size.

4.4.3. Enumeration Representation Clauses

For enumeration representation clauses (LRM 13.3), the following restrictions apply:

- The internal codes specified for the literals of the enumeration type may be any integer value between LONG_INTEGER'FIRST and INTEGER'LAST. It is strongly advised that you do not provide a representation clause that merely duplicates the default mapping of enumeration types which assigns consecutive numbers in ascending order starting with zero (0). Unnecessary runtime cost is incurred by such duplication. It should be noted that the use of attributes on enumeration types with user-specified encodings is costly at runtime.
- Array types, whose index type is an enumeration type with non-contiguous value encodings, consist of a contiguous sequence of components. Indexing into the array involves a runtime translation of the index value into the corresponding position value of the enumeration type.

4.4.4. Record Representation Clauses

The alignment clause of record representation clauses (LRM 13.4) is observed.

Static objects may be aligned at powers of 2 up to a page boundary. The

specified alignment becomes the minimum alignment of the record type, unless the minimum alignment of the record forced by the component allocation and the minimum alignment requirements of the components is already more stringent than the specified alignment.

The component clauses of record representation clauses are allowed only for components and discriminants of statically determinable size. Not all components need to be present. Component clauses for components of variant parts are allowed only if the size of the record type is statically determinable for every variant.

The size specified for each component must be sufficient to allocate all possible values of the component subtype, but not necessarily the component type. The location specified must be compatible with any alignment constraints of the component type; an alignment constraint on a component type may cause an implicit alignment constraint on the record type itself.

If some, but not all, discriminants and components of a record type are described by a component clause, the discriminants and components without component clauses are allocated after those with component clauses; no attempt is made to utilize gaps left by the user-provided allocation.

4.4.5. Address clauses

Address clauses (LRM 13.5) are only supported for variables. If an address clause is applied to an Ada Routine, the address clause will not be flagged by the compiler; however, the linker will not place the code for the routine at the specified address.

If an address clause is applied to an Ada variable, the linker will not reserve storage at the specified address. It is up to the user to ensure that use of the specified address will not result in a loss of significant data. This limitation will not affect the typical uses of address clauses for mapping memory-resident registers and device control words onto Ada variables.

4.4.6. Pragma PACK

Pragma PACK (LRM 13.1) is supported. For details, refer to the following sections.

4.4.6.1. Pragma PACK for Arrays

If pragma PACK is applied to an array, the densest possible representation is chosen. For details of packing, refer to the explanation of size specifications for arrays (sec. 4.4.2.3).

If, in addition, a length clause is applied to the array type, the pragma has no effect, since such a length clause already uniquely determines the array packing method.

If a length clause is applied to the component type, the array is packed densely, observing the component's length clause. Note that the component length clause may have the effect of preventing the compiler from packing as densely as would be the default if pragma PACK is applied where there was no length clause given for the component type.

4.4.6.2. The Predefined Type STRING

Package STANDARD applies pragma PACK to the type STRING. However, when applied to character arrays, this pragma cannot be used to achieve denser packing than is the default for the target: 1 character per 8-bit word.

4.4.6.3. Pragma PACK for Records

If pragma PACK is applied to a record, the densest possible representation is chosen that is compatible with the sizes and alignment constraints of the individual component types. Pragma PACK has an effect only if the sizes of some component types are specified explicitly by size specifications and are of non-referable nature. In the absence of pragma PACK, such components generally consume a referable amount of space.

It should be noted that the default type mapping for records maps components of boclean or other types that require only a single bit to a single bit in the record layout, if there are multiple such components in a record. Otherwise, it allocates a referable amount of storage to the component.

If pragma PACK is applied to a record for which a record representation clause has been given detailing the allocation of some but not all components, the pragma PACK affects only the components whose allocation has not been detailed. Moreover, the strategy of not utilizing gaps between explicitly allocated components still applies.

4.4.7. Minimal Alignment for Types

Certain alignment properties of values of certain types are enforced by the type mapping rules. Any representation specification that cannot be satisfied within these constraints is not obeyed by the compiler and is appropriately diagnosed.

Alignment constraints are caused by properties of the target architecture, most notably by the capability to extract non-aligned component values from composite values in a reasonably efficient manner. Typically, restrictions exist that make extraction of values that cross certain address boundaries very expensive, especially in contexts involving array indexing. Permitting data layouts that require such complicated extractions may impact code quality on a broader scale than merely in the local context of such extractions.

Instead of describing the precise algorithm of establishing the minimal alignment of types, we provide the general rule that is being enforced by the alignment rules:

- No object of scalar type (including components or subcomponents of a composite type) may span a target-dependent address bour cary that would mandate an extraction of the object's value to be performed by two or more extractions.

4.5. IMPLEMENTATION-GENERATED COMPONENTS IN RECORDS

The only implementation-dependent components allocated by Tartan Ada in records are fields containing either dope information for arrays whose bounds depend on discriminants of the record or relative offsets of components within a record layout for record components of dynamic size. These components cannot be named by the user.

4.6. INTERPRETATION OF EXPRESSIONS APPEARING IN ADDRESS CLAUSES

Section 13.5.1 of the LRM describes a syntax for associating interrupts with task entries. Tartan Ada implements the address clause

for TOENTRY use at intID;

by associating the interrupt specified by intID with the TOENTRY entry of the task containing this address clause. The interpretation of intID is both machine and compiler dependent.

The Motorola 68xxx specification provides 256 interrupts that may be associated with task entries. These interrupts are identified by an integer in the range 0..255, corresponding to the interrupt vector numbers in section 9.2 of the MC68040 32-Bit Microprocessor User's Manual. When you specify an interrupt address clause, the intID argument is interpreted as follows:

- If the argument is in the range 0..255, a full support interrupt association is made between the interrupt specified by the argument and the.task entry. That is, the runtimes make no assumptions about the task in question. This method is the slower.
- If the argument is in the range 256..511, a fast interrupt association is made between the interrupt number (argument-256) and the task entry. This method provides faster execution because the runtimes can depend upon the assumptions previously described.

For the difference between full support and fast interrupt handling, refer to section 9.5.12.

4.7. RESTRICTIONS ON UNCHECKED CONVERSIONS

Tartan supports UNCHECKED_CONVERSION as documented in Section 13.10 of the LRM. The sizes need not be the same, nor need they be known at compile time. The only exception is unconstrained array types which may not be used as the target of an UNCHECKED_CONVERSION.

If the value in the source is wider than that in the target, the source value will be truncated. If narrower, it will be zero-extended. Calls on instantiations of UNCHECKED_CONVERSION are made inline automatically.

4.8. IMPLEMENTATION-DEPENDENT ASPECTS OF INPUT-OUTPUT PACKAGES

Tartan Ada supplies the predefined input/output packages DIRECT IO, SEQUENTIAL_IO, TEXT_IO, and LOW_LEVEL_IO as required by LRM Chapter 14. The functionality of DIRECT_IO, SEQUENTIAL_IO, and TEXT_IO is fully supported.

4.9. OTHER IMPLEMENTATION CHARACTERISTICS

The following information is supplied in addition to that required by Appendix F to MIL-STD-1815A.

4.9.1. Definition of a Main Program

Any Ada library subprogram unit may be designated the main program for purposes of linking (using the Ada librarian's link subcommand) provided that the subprogram has no parameters.

Tasks initiated in imported library units follow the same rules for termination as other tasks (described in LRM 9.4 (6-10)). Specifically, these tasks are not terminated simply because the main program has terminated. Terminate alternatives in selective wait statements in library tasks are therefore strongly recommended.

4.9.2. Implementation of Generic Units

All instantiations of generic units, except the predefined generic UNCHECKED_CONVERSION and UNCHECKED_DEALLOCATION subprograms, are implemented by code duplications. No attempt at sharing code by multiple instantiations is made in this release of Tartan Ada.

Tartan Ada enforces the restriction that the body of a generic unit must be compiled before the unit can be instantiated. It does not impose the restriction that the specification and body of a generic unit must be provided as part of the same compilation. A recompilation of the body of a generic unit will cause any units that instantiated this generic unit to become obsolete.

4.9.3. Implementation-Defined Characteristics in Package STANDARD

The implementation-dependent characteristics in package STANDARD (Annex C) are:

package STANDARD is

type BYTE_INTEGER is range -128 .. 127; type SHORT_INTEGER is range -32768 .. 32767; type INTEGER is range -2_147_483_648 .. 2_147_483_647; type FLOAT is digits 6 range -16#0.FFFFFF#E+32 .. 16#0.FFFFFF#E+32

type LONG_FLOAT is digits 9 range -16#0.FFFFFFFFFFFF8#E + 256 .. 16#0.FFFFFFFFFFF8#E + 256 ; type DURATION is delta 0.0001 range -86400.0 .. 86400.0;

end STANDARD;

4.9.4. Attributes of Type DURATION

The type DURATION is defined with the following characteristics:

Attribute	Value
DURATION'DELTA	0.0001 sec
DURATION' SMALL	-5 6.103516E sec
DURATION' FIRST	-86400.0 sec
DURATION'LAST	86400.0 sec

4.9.5. Values of Integer Attributes

Tartan Ada supports the predefined integer types INTEGER, SHORT_INTEGER and BYTE_INTEGER. The range bounds of these predefined types are:



The range bounds for subtypes declared in package TEXT_IO are:

Attribute	Value
COUNT'FIRST	0
COUNT ' LAST	INTEGER'LAST - 1
POSITIVE_COUNT'FIRST	1
POSITIVE_COUNT'LAST	INTEGER'LAST - 1
FIELD'FIRST	0
FIELD'LAST	240

The range bounds for subtypes declared in packages DIRECT_IO are:

Attribute	Value
COUNT'FIRST	0
COUNT'LAST	INTEGER'LAST
POSITIVE_COUNT'FIRST	1
POSITIVE_COUNT'LAST	COUNT'LAST

4.9.6 Values of Floating-Point Attributes

Tartan Ada supports the predefined floating-point types FLOAT and LONG_FLOAT.

Attribute	Value for FLOAT
DIGITS	6
MANTISSA	21
EMAX	84
EPSILON	16#0.1000_00#E-4 (approximately 9.53674E-07)
SMALL	16#0.8000_00#E-21 (approximately 2.58494E-26)
LARGE	 16#0.FFFF_F8#E+21 (approximately 1.93428E+25)
SAFE_EMAX	126
SAFE_SMALL	16#0.2000_000#E-31 (approximately 5.87747E-39)
SAFE_LARGE	i6#0.3FFF_FE0#E+32 (approximately 8.50706E+37)
FIRST	-16#0.FFFFFF#E+32 (approximately -3.40282E+38)
LAST	16#0.FFFFFF#E+32 (approximately 3.40282E+38)
MACHINE_RADIX	2
MACHINE_MANTISSA	24
MACHINE_EMAX	128
MACHINE_EMIN	-125
MACHINE_ROUNDS	TRUE
MACHINE_OVERFLOWS	TRUE

,

C-19

Attribute	Value for LONG_FLOAT
DIGITS	15
MANTISSA	51
EMAX	204
EPSILON	16#0.4000_0000_0000_000#E-12 (approximately 8.8817841970013E-16)
SMALL	 16#0.8000_0000_0000_000#E-51 (approximately 1.9446922743316E-62)
LARGE	16#0.FFFF_FFFF_FFF_E00#E+51 (approximately 2.5711008708143E+61)
SAFE_EMAX	1022
SAFE_SMALL	16#0.2000_0000_0000_000#E-255 (approximately 1.1125369292536-308)
SAFE_LARGE	16#0.3FFF_FFFF_FFF_F80#E+256 (approximately 4.4942328371557E+307)
FIRST	-16#0.FFFFFFFFFFFFF8#E+256 (approximately -1.79769313486232E+308)
LAST	16#0.FFFFFFFFFFFF8#E+256 (approximately ~1.79769313486232E+308)

.

Attribute	Value for LONG_FLOAT (cont'd)
MACHINE_RADIX	2
MACHINE_MANTISSA	53
MACHINE_EMAX	1024
MACHINE_EMIN	-1021
CHINE_ROUNDS	TRUE
MACHINE_OVERFLOWS	TRUE

4.10. SUPPORT FOR PACKAGE MACHINE_CODE

Package MACHINE_CODE provides the programmer with an interface to request the generation of any instruction that is available on the MC68020, MC68881/MC68882, MC68030, MC68040 or CPU32 processors. The implementation of package MACHINE_CODE is similar to that described in section 13.8 of the Ada LRM, with several added features. Please refer to Appendix B for the package MACHINE CODE specification.

4.10.1. Basic Information

As required by LRM, section 13.8, a routine which contains machine code inserts may not have any other kind of statement, and may not contain an exception handler. The only allowed declarative item is a use clause. Comments and pragmas are allowed as usual.

4.10.2. Instructions

A machine code insert has the form TYPE_MARK'RECORD_AGGREGATE, where the type must be one of the records defined in package MACHINE_CODE. Package MACHINE_CODE defines seven types of records. Each has an opcode and zero to 6 operands. These records are adequate for the expression of all instructions provided by the 68xxx.

4.10.3. Operands and Address Modes

An operand consists of a record aggregate which holds all the information to specify it to the compiler. All operands have an address mode and one or more other pieces of information. The operands correspond exactly to the operands of the instruction being generated.

Each operand in a machine code insert must have an Address_Mode. The address modes provided in package MACHINE_CODE provide access to all address modes supported by the 68xxx.

In addition, package MACHINE_CODE supplies the address modes SYMBOLIC_ADDRESS and SYMBOLIC_VALUE which allow the user to refer to Ada objects by specifying object'ADDRESS as the value for the operand. Any Ada object which has the 'ADDRESS attribute may be used in a symbolic operand. SYMBOLIC_ADDRESS should be used when the operand is a true address (for example, a branch target). SYMBOLIC_VALUE should be used when the operand is actually a value (for example, one of the source operands of an ADD instruction).

When an Ada object is used as a source operand in an instruction (that is, one from which a value is read), the compiler will generate code which fetches the value of the Ada object. When an Ada object is used as the destination operand of an instruction, the compiler will generate code which uses the address of the Ada object as the destination of the instruction.

4.10.4. Examples

The implementation of package MACHINE_CODE makes it possible to specify both simple machine code inserts such as:

TWO_OPNDS'(MOVEQ, (IMM, 3), (DR, DO))

and more complex inserts such as

TWO_OPNDS'(ADDI_L, (IMM, 10), (SYMBOLIC_VALUE, ARRAY_VAR(X, Y, 27)'ADDRESS))

In the first example, the compiler will emit the instruction MOVEQ 3, DO. In the second example, the compiler will first emit whatever instructions are needed to form the address of ARRAY_VAR(X, Y, 27) and then emit the ADDI_L instruction. The various error checks specified in the LRM will be performed on all compiler-generated code unless they are suppressed by the programmer (either through pragma SUPPRESS, or through command qualifiers).

4.10.5. Incorrect Operands

Under some circumstances, the compiler attempts to correct incorrect operands. Three modes of operation are supplied for package MACHINE_CODE to determine whether corrections are attempted and how much information about the necessary corrections is provided to the user.

The compiler command line options for the three modes of operation are:

UNIX VMS Me fixup=none Mw fixup=warn no option [default] fixup=quiet [default]

In the Me or fixup = none mode, the specification of incorrect operands for an instruction is considered to be a fatal error. In this mode, the compiler will not generate any extra instructions to help you to make a machine code insertion. Note that it is still legal to use 'ADDRESS constructs as long as the object which is used meets the requirements of the instruction.

In the default or fixup = quiet mode, if you specify incorrect operands for an instruction, the compiler will do its best to correct the machine code to provide the desired effect. For example, although it is illegal to use an address register as the destination of an ADDI instruction, the compiler will accept \cdot it and try to generate correct code. In this case, the compiler will load the value found in the address register indicated into a data register, use the data register in the ADDI instruction, and then store from that data register back to the desired address register.

TWO OPNDS'(ADDI_L, (IMM, 3), (AR, A1))

will produce a code sequence such as:

mov.l a1,d0 addi.l #3,d0 mov.l d0,a1

In the Mw or fixup = warn mode, the compiler will perform the same level of correction as in the default or fixup = quiet mode. However, a warning message is issued stating that the machine code insert required additional machine instructions to make its operands legal.

4.10.6. Assumptions Made in Correcting Operands

When compiling in the UNIX default or Mw modes, or in the VMS fixup = quiet or fixup = warn modes, the compiler attempts to emit additional code to move ``the right bits'' from an incorrect operand to a place which is a legal operand for the requested instruction. The compiler makes certain basic assumptions when performing these corrections. This section explains the assumptions the compiler makes and their implications for the generated code. Note that if

you want a correction which is different from that performed by the compiler, you must make explicit machine code insertions to perform it.

C-23

For source and source/destination operands:

- SYMBOLIC_ADDRESS means that the address specified by the 'ADDRESS expression is used as the source bits. When the Ada object specified by the 'ADDRESS instruction is bound to a register, this will cause a compile-time error message because it is not possible to ``take the address'' of a register.
- SYMBOLIC_VALUE means that the value found at the address specified bythe 'ADDRESS expression will be used as the source bits. An Ada object which is bound to a register is correct here, because the contents of a register can be expressed on the 68xxx. Any other non-register means that the value found at the address specified by the operand will be used as the source bits.

For destination operands:

- SYMBOLIC_ADDRESS means that the desired destination for the operation is the address specified by the 'ADDRESS expression. An Ada object which is bound to a register is correct here; a register is a legal destination on the 68xxx.
- SYMBOLIC_VALUE means that the desired destination for the operations is found by fetching 32 bits from the address specified by the 'ADDRESS expression, and storing the result to the address represented by the fetched bits. This is equivalent to applying one extra indirection to the address used in the SYMBOLIC_ADDRESS case.
- All other operands are interpreted as directly specifying the destination for the operation.

4.10.7. Register Usage

The compiler may need several registers to generate code for operand corrections in machine code inserts. If you use all the registers, corrections will not be possible. In general, when more registers are available to the compiler it is able to generate better code.

Since the compiler may need to allocate registers as temporary storage in machine code routines, there are some restrictions placed on your register usage. The compiler will automatically free all registers which are volatile across a call for your use (that is D0, D1, A0, A1, Fp0, Fp1).

If you reference any other register, the compiler will reserve it for your use until the end of the machine code routine. The compiler will not save the register automatically if this routine is inline expanded. This means that the first reference to a register which is not volatile across calls should be an instruction which saves its value in a safe place.

The value of the register should be restored at the end of the machine code routine. This rule will help ensure correct operation of your machine code insert even if it is inline explained in another routine. However, the compiler will save the register automatically in the prolog code for the routine and restore it in the epilog code for the routine if the routine is not inline expanded.

4.10.8. Data Directives

Four special instructions are included in package MACHINE_CODE to allow the user to place data into the code stream. These four instructions are DATA8, DATA16, DATA32 and DATA64. Each of these instructions can have from 1 to 6 operands.

DATA8 and DATA16 are used to place 8-bit and 16-bit integer data items into the code stream.

DATA32 is used to place 32-bit data into the code stream. The value of an integer, a floating-point literal, or the address of a label or a routine are the legal operands (i.e. operands whose address mode is either IMM, FLOAT LIT SINGLE, or SYMBOLIC ADDRESS of an Ada object).

<< L1 >> THREE_OPNDS'(DATA32, (SYMBOLIC_ADDRESS, L1'ADDRESS), (FLOAT_LIT_SINGLE, 2.0), (IMM, 99));

will produce a code sequence such as:

L1: .long L1 .long 1073741824 | 0.2e1 .long 99

DATA64 is used to place a 64-bit data into the code stream. The only legal operand is a floating literal (i.e operand whose address mode is FLOAT_LIT_SINGLE or FLOAT_LIT_DOUBLE).

4.10.9. Inline Expansion

Routines which contain machine code inserts may be inline expanded into the bodies of other routines. This may happen under programmer control through the use of pragma INLINE, or with the optimization levels standard and time(UNIX: -Op2 and -Op3; VMS: /optimize = standard and /optimize = time) when the compiler selects the inline optimization as an appropriate action for the given situation. The compiler will treat the machine code insert as if it were a call. Volatile registers will be saved and restored around it and similar optimizing steps will be taken.

4.10.10. Unsafe Assumptions

There are a variety of assumptions which should not be made when writing machine code inserts. Violation of these assumptions may result in the generation of code which does not assemble or which may not function correctly.

- Do not assume that a machine code insert routine has its own set of local registers. This may not be true if the routine is inline expanded into another routine. Explicitly save and restore any registers which are not volatile across calls. If you wish to guarantee that a routine will

never be inline expanded, you should use an Ada separate body for the routine or compile at an optimization level lower than the default. Then make sure there is no pragma INLINE for the routine.

Values should not be assigned to the frame pointer register in the middle of a machine code insert routine, even if your code saves and restores the contents of the register. A dangerous situation can arise if an exception is propagated through the procedure frame, or if a machine code insert references a variable that uses a frame pointer in the address formula.

 Do not assume that the 'ADDRESS on SYMBOLIC_ADDRESS or SYMBOLIC_VALUE operands means that you are getting an ADDRESS to operate on. The Address- or Value-ness of an operand is determined by your choice of SYMBOLIC_ADDRESS or SYMBOLIC_VALUE. This means that to add the contents of X to D0, you should write

TWO_OPNDS'(ADD_L, (SYMBOLIC_VALUE, X'ADDRESS), (DR, D0));

but to add the address of X to DO, you should write

TWO_OPNDS'(ADD_L, (SYMBOLIC_ADDRESS, X'ADDRESS), (DR, DO));

- The compiler will not generate call site code for you if you emit a call instruction. You must save and restore any volatile registers (DO, D1, AO, A1, FpO, Fp1) which currently have values in them. If the routine you call has out parameters, a large function return result, or an unconstrained result, it is your responsibility to emit the necessary instructions to deal with these constructs as the compiler expects. In other words, when you emit a call, you must follow the linkage conventions of the routine you are calling. For further details on call site code, see sections 5.4, 5.5, and 5.6.
- Do not attempt to move multiple Ada objects with a single long instruction such as MOVE. Although the objects may be contiguous under the current circumstances, there is no guarantee that later changes will permit them to remain contiguous. If the objects are parameters, it is virtually certain that they will not be contiguous if the routine is inline expanded into the body of another routine. In the case of locals, globals, and own variables, the compiler does not guarantee that objects which are declared textually ``next'' to each other will be contiguous in memory. If the source code is changed such that it declares additional objects, this may change the storage allocation such that objects which were previously adjacent are no longer adjacent.

4.10.11. Limitations

- The current implementation of the compiler is unable to fully support automatic correction of certain kinds of operands. In particular, the compiler assumes that the size of a data object is the same as the number of bits which is operated on by the instruction chosen in the machine code insert. This means that in the insert:

TWO_OPNDS'(ADD_B, (SYMBOLIC_VALUE, INTEGER_VARIABLE'ADDRESS), (DR, DO))

the compiler will assume that INTEGER_VARIABLE is 8 bits wide, when in fact it is stored in 32 bits of memory.

Note that the use of X'ADDRESS in a machine code insert does not guarantee that X will be bound to memory. This is a result of the use of 'ADDRESS to provide a ``typeless'' method for naming Ada objects in machine code inserts. For example, it is legal to say (SYMBOLC_VALUE, X'ADDRESS) in an insert even when X is found in a register.

- Absolute Short Address Mode with a symbolic operand is not supported. For example, the following operand is illegal:

(ABS_SHORT, SOME_VARIABLE'ADDRESS)

- In Address Modes in which two displacements are allowed only base displacement can be represented by a symbolic address. Outer displacement must be an integer. For example, this operand is legal:

(MEMPOST2, BD_MEMPOST2 => SOME_ROUTINE'ADDRESS, -- base AN_MEMPOST2 => A0, -- displacement XN_MEMPOST2 => D0, XN_SIZE_MEMPOST2 => LONG, SCALE_MEMPOST2 => ONE, OD_MEMPOST2 => 16) -- outer_Displacement

while the following operand is illegal:

(MEMPOST2, BD_MEMPOST2 => ROUTINE_1'ADDRESS, --base displacement AN_MEMPOST2 => A0, XN_MEMPOST2 => D0, XN_SIZE_MEMPOST2 => LONG, SCALE_MEMPOST2 => ONE, OD_MEMPOST2 => ROUTINE 2'ADDRESS) --outer Displacement

- PC-relative Address Modes with a suppressed base register field can sometimes be handled incorrectly by the current implementation of the compiler.

- Extended-precision floating-point literals are not supported.

4.10.12. ADDRESS_MODE Usage

- Addressing modes that accept 16 or 32-bit displacements are represented by two entries in package MACHINE_CODE's ADDRESS_MODE enumeration: one that accepts an integer, and one that accepts a symbolic address. For example, Memory Indirect Pre-Indexed addressing mode is represented by MEMPRE and MEMPRE2 Address Modes.
- DARI (Data or Address Register Indirect) ADDRESS_MODE is provided exclusively for use with operands five and six of the CAS2 instruction.
- ARIDX (Address Register Indirect with Index and Displacement) ADDRESS_MODE represents both the 8-bit displacement and the base displacement sub-modes of the Address Register Indirect with Index

C-27

addressing mode. The compiler will pick the most economical form.

- PCIDX (Program Counter Indirect with Index and Displacement) ADDRESS_MODE represents both the 8-bit displacement and the base displacement sub-modes of the Program Counter Indirect with Index addressing mode. The compiler will pick the most economical form.

4.10.13. INSTRUCTION_MNEMONIC Usage

- INSTRUCTION_MNEMONIC names in package MACHINE_CODE are formed by concatenating the base instruction name with a suffix representing the size of the instruction. For example, CMP_B, CMP_W, and CMP_L are package MACHINE_CODE entries for the 68xxx CMP instruction. If the instruction exists in a single size only, it is represented by two entries in package MACHINE_CODE: one with and one without a suffix. For example, the 68xxx LEA instruction is represented by LEA and LEA_L. Unsized instructions are represented by their base names with no suffix.
- For instructions that operate on control registers the control register operand needs to be explicitly supplied in the machine code insert:

TWO_OPNDS'(ANDITOCCR, (IMM, 3), (CR, CCR));

- For Conditional Branch, Branch Always, and Branch to Subroutine instructions an unsized entry (for example, BEQ) lets the compiler pick the instruction of the optimal size.
- BSRNORET and JSRNORET mnemonics are aliases for BSR and JSR respectively. Use them when a called routine is known to never return.
- The digit in the mnemonics of the Co-Processor instructions (cp*) indicates the number of optional co-processor defined extension words.
- When using MC68881/MC68882 unary instructions which operate on a single floating-point register, the register operand needs to be supplied as both the source and the destination operand:

TWO OPNDS'(FCOSH X, (FPR, FP1), (FPR, FP1));

The FTST instruction is the exception to this rule:

ONE_OPNDS'(FTST_X, (FPR, FP1));

- When using CPU32 Table Lookup and Interpolation instructions, the Dn register-pair operand needs to be supplied as two separate operands of the machine-code insert. For example,

THREE OPNDS' (TBLS L, (DR, D0), (DR, D1), (DR, D7));

will emit the TBLS.L D0:D1,D7 instruction.

- MC68881/MC68882 instruction Move System Control Register, FMOVE is represented by several individual instructions, each of which requires an explicit control register operand:

TWO OPNDS'(FMOVETOFPCR, (DR, D0), (CR, FPCR));

- FSINCOS instruction returns the sine in its second operand and the cosine in its third operand.
- MC68881/MC68882 Move Multiple Data Registers, FMOVEM_X and Move Multiple Control Registers, FMOVEM_L instructions, expect the register mask operand to be represented by an integer literal:

TWO OPNDS'(FMOVEM_X, (IMM, 3), (ARI, A0));

4.10.14. Example

with MACHINE_CODE; use MACHINE_CODE; with SYSTEM; use SYSTEM;

procedure SINCOS(SOURCE : in LONG_FLOAT;

SIN : out LONG_FLOAT;

COS : out LONG_FLOAT) is

begin

2 .

-- Compute sine and cosine of Source and return them in

-- parameters Sin and Cos, respectively

THREE_OPNDS'(FSINCOS_D, (SYMBOLIC_VALUE, SOURCE'ADDRESS),

(FPR, FPO), (FPR, FP1));

TWO_OPNDS'(FMOVE_D, (FPR, FP0), (SYMBOLIC_ADDRESS, SIN'ADDRESS)); TWO_OPNDS'(FMOVE_D, (FPR, FP1), (SYMBOLIC_ADDRESS, COS'ADDRESS)); end SINCOS;

Assembly code output:

.data

.globl _aOfsincos

.text

_aOfsincos: link

link a6,#0 clrl a7@-

fsincosd:a6@(8:w),fp1:fp0 fmoved fp0,a6@(16:w) fmoved fp1,a6@(24:w)

unik a6 rts

| Total bytes of code in the above routine = 28

.text

.even | Total bytes of code = 28 | Total bytes of data = 0

4.11. INLINE GUIDELINES

The following discussion on inlining is based on the next two examples. From these sample programs, general rules, procedures, and cautions are illustrated.

Consider a package with a subprogram that is to be inlined.

package IN_PACK is procedure I_WILL_BE_INLINED; pragma INLINE (I_WILL_BE_INLINED); end IN_PACK;

Consider a procedure that makes a call to an inlined subprogram in the package.

with IN_PACK; use IN_PACK; procedure USES_INLINED_SUBP is begin I_WILL_BE_INLINED; end USES INLINED SUBP;

After the package specification for IN_PACK has been compiled, it is possible to compile the unit USES_INLINED_SUBP that makes a call to the subprogram I_WILL_BE_INLINED. However, because the body of the subprogram is not yet available, the generated code will not have an inlined version of the

subprogram. The generated code will use an out of line call for I_WILL_BE_INLINED. The compiler will issue warning message #2429 that the call was not inlined when USES_INLINED_SUBP was compiled.

1

If IN_PACK is used across libraries, it can be exported as part of a specification library after having compiled the package specification. Note that if only the specification is exported, there will be no inlined calls to IN_PACK in all units within libraries that import IN_PACK. If only the specification is exported, all calls that appear in other libraries will be out of line calls. The compiler will issue warning message #2429 to indicate the call was not inlined.

There is no warning at link-time that subprograms have not been inlined.

If the body for package IN_PACK has been compiled before the call to I_WILL_BE_INLINED is compiled, the compiler will inline the subprogram. In the example above, if the body of IN_PACK has been compiled before USES_INLINED_SUBP, the call will be inlined when USES_INLINED_SUBP is compiled.

Having an inlined call to a subprogram makes a unit dependent on the unit that contains the body of the subprogram. In the example, once USES_INLINED_SUB has been compiled with an inlined call to I_WILL_BE_INLINED, the unit USES_INLINED_SUBP will have a dependency on the package body IN_PACK. Thus, if the body for package body IN_PACK is recompiled, USES_INLINED_SUBP will become obsolete, and must be recompiled before it can be linked.

It is possible to export the body for a library unit. If the body for package IN_PACK is added to the specification library using the Ada librarian subcommand export compilation unit, other libraries that import package IN_PACK will be able to compile inlined calls across library units.

At optimization levels lower than the default, the compiler will not inline calls, even when pragma INLINE has been used and the body of the subprogram is in the library prior to the unit that makes the call. Lower optimization levels avoid any changes in flow of the code that causes movement of code sequences, as happens in a pragma INLINE. If the compiler is running at a low optimization level, the user will not be warned that inlining is not happening.