

2

# REPORT DOCUMENTATION PAGE

Form Approved

OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and reviewing the collection of information. Send comments regarding this burden, to Washington Headquarters Service, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Information and Regulatory Affairs, Office of Management and Budget, Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE June 28, 1995	3. REPORT TYPE AND DATES COVERED Final
----------------------------------	---------------------------------	---

4. TITLE AND SUBTITLE: Ada Compiler Validation Summary Report, VC# 95062211.11395 Tartan, Inc. -- Tartan Ada SPARC/960MC/PMRT, Version 5.1	5. FUNDING NUMBERS
--	--------------------

6. AUTHOR(S)  IABG, Abt. ITE	
------------------------------------	--

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Ada Validation Facility IABG, Abt. ITE Einsteinstrasse 20 D-85521 Ottobrunn, GERMANY	8. PERFORMING ORGANIZATION REPORT NUMBER  IABG-VSR 120
---	--

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Ada Joint Program Office, Defense Information System Agency Code JEXEV, 701 S. Courthouse Rd., Arlington, VA 22204-2199	10. SPONSORING/MONITORING AGENCY REPORT NUMBER
--	--

11. SUPPLEMENTARY NOTES  <div style="text-align: right; font-size: 2em; font-weight: bold;">19950905 008</div>
--

12a. DISTRIBUTION/AVAILABILITY STATEMENT  Approved for public release; Distribution is unlimited.	
---	--

13. ABSTRACT (Maximum 200 words)  This Ada implementation was tested and determined to pass ACVC 1.11. Testing was completed on 22 June 1995. Host Computer System: Sun SPARCstation 5 under SunOS, Release 4.1.4 Target Computer System: Intel EXV80960MC board (bare machine)
---

14. SUBJECT TERMS Ada Programming Language, Ada Compiler Validation Summary Report, Ada Compiler Validation Capability, Validation Testing, Ada Validation Office, Ada Validation Facility, ANSI/MIL-STD-1815A, Ada Joint Program Office	15. NUMBER OF PAGES 68
	16. PRICE

17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UNCLASSIFIED
---	--	---	--

3

AVF Control Number: IABG-VSR 120  
28 June, 1995

**Ada COMPILER  
VALIDATION SUMMARY REPORT:  
Certificate Number: 95062211.11395  
Tartan, Inc.  
Tartan Ada SPARC/960MC/PMRT Version 5.1  
Sun SPARCstation 5 =>  
Intel EXV80960MC Board**

Accession For	
NTIS CRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification _____	
By _____	
Distribution / _____	
Availability Codes	
Dist	Avail and / or Special
A-1	

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

### Certificate Information

The following Ada implementation was tested and determined to pass ACVC 1.11. Testing was completed on 22 June, 1995.

**Compiler Name and Version:** Tartan Ada SPARC/960MC/PMRT Version 5.1

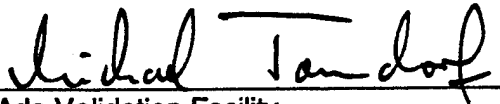
**Host Computer System:** Sun SPARCstation 5 under SunOS Release 4.1.4

**Target Computer System:** Intel EXV80960MC board (bare machine)

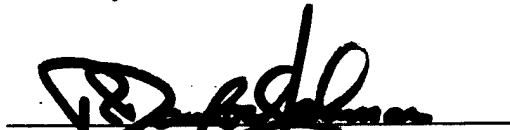
See section 3.1 for any additional information about the testing environment.

As a result of this validation effort, Validation Certificate 95062211.11395 is awarded to Tartan, Inc. This certificate expires on 31 March 1998.

This report has been reviewed and is approved.



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



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



Ada Joint Program Office  
Center for Information Management  
Defense Information Systems Agency  
Don Reifer, Director  
Arlington VA 22204, USA

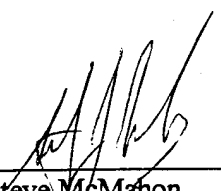
## Declaration of Conformance

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

Ada Implementation:  
Ada CompilerName and Version: Tartan Ada SPARC/960MC/PMRT Version 5.1  
Host Computer System: SPARCstation 5 under SunOS Release 4.1.4  
Target Computer System: Intel EXV80960MC board (bare machine)

### 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.

  
\_\_\_\_\_  
Steve McMahon  
Senior Vice President

6/28/95  
\_\_\_\_\_  
Date

## TABLE OF CONTENTS

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

## CHAPTER 1

### INTRODUCTION

The Ada implementation described above was tested according to the Ada Validation Procedures [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

## 1.2 REFERENCES

- [Ada83]     Reference Manual for the Ada Programming Language,  
ANSI/MIL-STD-1815A, February 1983 and ISO 8652-1987.
- [Pro92]     Ada Compiler Validation Procedures,  
Version 3.1, Ada Joint Program Office, August 1992.
- [UG90]     Ada Compiler Validation Capability User's Guide,  
3 April 1990.

## 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.

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

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

## INTRODUCTION

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 [UG90]).

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

### 1.4 DEFINITION OF TERMS

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

## INTRODUCTION

Conformity	Fulfillment by a product, process or service of all requirements specified.
Customer	An individual or corporate entity who enters into an agreement with an AVF which specifies the terms and conditions for AVF services (of any kind) to be performed.
Declaration of Conformance	A formal statement from a customer assuring that conformity is 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>."
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.

## INTRODUCTION

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	B83022B	B83022H	B83025B	B83025D
C83026A	B83026B	C83041A	B85001L	C86001F	C94021A
C97116A	C98003B	BA2011A	CB7001A	CB7001B	CB7004A
CC1223A	BC1226A	CC1226B	BC3009B	BD1B02B	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 AI-ddddd. For this implementation, the following tests were determined to be inapplicable for the reasons indicated; references to Ada Commentaries are included as appropriate.

## IMPLEMENTATION DEPENDENCIES

The following 159 tests have floating-point type declarations requiring more digits than SYSTEM.MAX\_DIGITS:

C24113O..Y (11 tests)	C35705O..Y (11 tests)
C35706O..Y (11 tests)	C35707O..Y (11 tests)
C35708O..Y (11 tests)	C35802O..Z (12 tests)
C45241O..Y (11 tests)	C45321O..Y (11 tests)
C45421O..Y (11 tests)	C45521O..Z (12 tests)
C45524O..Z (12 tests)	C45621O..Z (12 tests)
C45641O..Y (11 tests)	C46012O..Z (12 tests)

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

C45531M..P and C45532M..P (3 tests) check fixed-point operations for types that require a SYSTEM.MAX\_MANTISSA of 47 or greater; for this implementation, MAX\_MANTISSA is less than 47.

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

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

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

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.

## IMPLEMENTATION DEPENDENCIES

AE2101C uses instantiations of package SEQUENTIAL\_IO with unconstrained array types and with record types with discriminants without defaults; these instantiations are rejected by this compiler.

AE2101H uses instantiations of package DIRECT\_IO with unconstrained array types and record types with discriminants without defaults; these instantiations are rejected by this compiler.

The following 264 tests check operations on sequential, text, and direct access files; this implementation does not support external files:

CE2102A..C (3)	CE2102G..H (2)	CE2102K	CE2102N..Y (12)
CE2103C..D (2)	CE2104A..D (4)	CE2105A..B (2)	CE2106A..B (2)
CE2107A..H (8)	CE2107L	CE2108A..H (8)	CE2109A..C (3)
CE2110A..D (4)	CE2111A..I (9)	CE2115A..B (2)	CE2120A..B (2)
CE2201A..C (3)	EE2201D..E (2)	CE2201F..N (9)	CE2203A
CE2204A..D (4)	CE2205A	CE2206A	CE2208B
CE2401A..C (3)	EE2401D	CE2401E..F (2)	EE2401G
CE2401H..L (5)	CE2403A	CE2404A..B (2)	CE2405B
CE2406A	CE2407A..B (2)	CE2408A..B (2)	CE2409A..B (2)
CE2410A..B (2)	CE2411A	CE3102A..C (3)	CE3102F..H (3)
CE3102J..K (2)	CE3103A	CE3104A..C (3)	CE3106A..B (2)
CE3107B	CE3108A..B (2)	CE3109A	CE3110A
CE3111A..B (2)	CE3111D..E (2)	CE3112A..D (4)	CE3114A..B (2)
CE3115A	CE3119A	EE3203A	EE3204A
CE3207A	CE3208A	CE3301A	EE3301B
CE3302A	CE3304A	CE3305A	CE3401A
CE3402A	EE3402B	CE3402C..D (2)	CE3403A..C (3)
CE3403E..F (2)	CE3404B..D (3)	CE3405A	EE3405B
CE3405C..D (2)	CE3406A..D (4)	CE3407A..C (3)	CE3408A..C (3)
CE3409A	CE3409C..E (3)	EE3409F	CE3410A
CE3410C..E (3)	EE3410F	CE3411A	CE3411C
CE3412A	EE3412C	CE3413A..C (3)	CE3414A
CE3602A..D (4)	CE3603A	CE3604A..B (2)	CE3605A..E (5)
CE3606A..B (2)	CE3704A..F (6)	CE3704M..O (3)	CE3705A..E (5)
CE3706D	CE3706F..G (2)	CE3804A..P (16)	CE3805A..B (2)
CE3806A..B (2)	CE3806D..E (2)	CE3806G..H (2)	CE3904A..B (2)
CE3905A..C (3)	CE3905L	CE3906A..C (3)	CE3906E..F (2)

CE2103A, CE2103B, and CE3107A use an illegal file name in an attempt to create a file and expect NAME\_ERROR to be raised; this implementation does not support external files and so raises USE\_ERROR. (See section 2.3.)

### 2.3 TEST MODIFICATIONS

Modifications (see Section 1.3) were required for 114 tests.

The following 81 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.

## IMPLEMENTATION DEPENDENCIES

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	B59001I	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 passed 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..N (14 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).

## IMPLEMENTATION DEPENDENCIES

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 AI-00408 & AI-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 separate 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 as small 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
```

CE2103A, CE2103B, and CE3107A were graded inapplicable by Evaluation Modification as directed by the AVO. The tests abort with an unhandled exception when USE\_ERROR is raised on the attempt to create an external file. This is acceptable behavior because this implementation does not support external files (cf. AI-00332).

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 Wayne Lieberman  
Ada Business Unit Leader  
Tartan Inc.  
300 Oxford Drive  
Monroeville, PA 15146, USA  
Tel. (412) 856-3600

For sales information about this Ada implementation, contact:

Mr Keith Franz  
Vice President Sales  
Tartan Inc.  
300 Oxford Drive  
Monroeville, PA 15146, USA  
Tel. (412) 856-3600

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).

Summary of Test Counts		
a	Total Number of Applicable Tests	3605
b	Total Number of Withdrawn Tests	104
c	Processed Inapplicable Tests	38
d	Non-Processed I/O Tests	264
e	Non Processsd Floating Point Precision Tests	159
f	Total Number of Inapplicable Tests (c+d+e)	461
g	Total Number of Tests for ACVC 1.11 (a+b+f)	4170

### 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 on a computer with an attached tape drive and copied directly onto the host computer using networking facilities.

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:

## PROCESSING INFORMATION

- 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.
- La forces a listing to be produced, default is to only produce a listing when an error occurs

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 [UG90]. 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

Macro Parameter	Macro Value
\$MAX_IN_LEN	240
\$BIG_ID1	(1..V-1 => 'A', V => '1')
\$BIG_ID2	(1..V-1 => 'A', V => '2')
\$BIG_ID3	(1..V/2 => 'A') & '3' & (1..V-1-V/2 => 'A')
\$BIG_ID4	(1..V/2 => 'A') & '4' & (1..V-1-V/2
\$BIG_INT_LIT	(1..V-3 => '0') & "298"
\$BIG_REAL_LIT	(1..V-5 => '0') & "690.0"
\$BIG_STRING1	"" & (1..V/2 => 'A') & ""
\$BIG_STRING2	"" & (1..V-1-V/2 => 'A') & '1' & ""
\$BLANKS	(1..V-20 => '')
\$ILLEGAL_EXTERNAL_FILE_NAME1	"ILLEGAL_EXTERNAL_FILE_NAME1" & (1..V => '-')
\$ILLEGAL_EXTERNAL_FILE_NAME2	"ILLEGAL_EXTERNAL_FILE_NAME2" & (1..V => '-')
\$MAX_LEN_INT_BASED_LITERAL	"2:" & (1..V-5 => '0') & "11:"
\$MAX_LEN_REAL_BASED_LITERAL	"16:" & (1..V-7 => '0') & "F.E:"
\$MAX_STRING_LITERAL	"" & (1..V-2 => 'A') & ""

## MACRO PARAMETERS

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

Macro Parameter	Macro Value
\$ACC_SIZE	32
\$ALIGNMENT	4
\$COUNT_LAST	2147483646
\$DEFAULT_MEM_SIZE	2_097_152
\$DEFAULT_STOR_UNIT	8
\$DEFAULT_SYS_NAME	1960MC
\$DELTA_DOC	2#1.0#E-31
\$ENTRY_ADDRESS	SYSTEM.ADDRESS'(16#0000_00C8#)
\$ENTRY_ADDRESS1	SYSTEM.ADDRESS'(16#0000_00C9#)
\$ENTRY_ADDRESS2	SYSTEM.ADDRESS'(16#0000_00CA#)
\$FIELD_LAST	220
\$FILE_TERMINATOR	''
\$FIXED_NAME	NO_SUCH_FIXED_TYPE
\$FLOAT_NAME	EXTENDED_FLOAT
\$FORM_STRING	""
\$FORM_STRING2	"CANNOT_RESTRICT_FILE_CAPACITY"
\$GREATER_THAN_DURATION	100_000.0
\$GREATER_THAN_DURATION_BASE_LAST	100_000_000.0
\$GREATER_THAN_FLOAT_BASE_LAST	3.5E+38

## MACRO PARAMETERS

\$GREATER_THAN_FLOAT_SAFE_LARGE	1.0E+38
\$GREATER_THAN_SHORT_FLOAT_SAFE_LARGE	1.0E+38
\$HIGH_PRIORITY	17
\$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	Use_Call
\$LESS_THAN_DURATION	-100_000.0
\$LESS_THAN_DURATION_BASE_FIRST	-100_000_000.0
\$LINE_TERMINATOR	" "
\$LOW_PRIORITY	2
\$MACHINE_CODE_STATEMENT	Two_Format'(MOV, (Reg_' it,5),(Reg,R5));
\$MACHINE_CODE_TYPE	Mnemonic
\$MANTISSA_DOC	31
\$MAX_DIGITS	18
\$MAX_INT	9223372036854775807
\$MAX_INT_PLUS_1	9223372036854775808
\$MIN_INT	-9223372036854775808
\$NAME	BYTE_INTEGER
\$NAME_LIST	I960MC

MACRO PARAMETERS

\$NEG_BASED_INT	16#FFFFFFFFFFFFFFE#
\$NEW_MEM_SIZE	2097152
\$NEW_STOR_UNIT	8
\$NEW_SYS_NAME	I960MC
\$PAGE_TERMINATOR	"
\$RECORD_DEFINITION	record Operation: Instruction_Mnemonic; Operand_1: Operand; Operand_2: Operand; end record;
\$RECORD_NAME	Two_Format
\$TASK_SIZE	32
\$TASK_STORAGE_SIZE	4096
\$TICK	0.0156125
\$VARIABLE_ADDRESS	SYSTEM.ADDRESS'(16#0000_1000#)
\$VARIABLE_ADDRESS1	SYSTEM.ADDRESS'(16#0000_1004#)
\$VARIABLE_ADDRESS2	SYSTEM.ADDRESS'(16#0000_1008#)

## 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.

### 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:

- |    |   |
|----|---|
| 9x | Notifies the compiler that the source file to be compiled contains Ada 9X constructs. The option <i>no9x</i> (Ada 83) is the default.   |
| A  | Generates an assembly code file with interleaved source code. The assembly code file has an extension <i>.s</i> for a body or <i>.ss</i> for a specification. The assembly code file is not intended to be input to an assembler, but serves as documentation only.<br><br>To produce this type of file, debug information must be generated for the compilation unit. When this option is specified, debug information will be generated, even if the <i>debug</i> option <sup>2</sup> has not been specified.   |
| a  | Generates an assembly code file. The assembly code file has an extension <i>.s</i> for a body or <i>.ss</i> for a specification (sec. 3.5). The assembly code file is not intended to be input to an assembler, but serves as documentation only. In the default mode, no assembly code is generated.   |
| Ba | Specifies that the compiler will produce an optimization ( <i>.opt</i> ) 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 try 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.<br><br>By default, the compiler will <i>not</i> produce an optimization file for the current unit (effectively preventing the creation of dependencies on this body), but will read the <i>.opt</i> 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. |
| Bm | Specifies that the compiler will neither produce an optimization ( <i>.opt</i> ) 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.<br><br>By default, the compiler will <i>not</i> produce an optimization file for the current unit, but will read the <i>.opt</i> 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.  |
| Cl | Controls the type of calls generated through the option supplied. With this option, the compiler generates all long calls in the compiled code. With the default, the   |

---

<sup>2</sup>UNIX: -g

- compiler generates short calls within the application's code and long calls from applications to runtime routines.
- Cs Controls the type of calls generated through the option supplied. With this option, the compiler generates all short calls in the compiled code. Inappropriate use of this option will cause a failure at link-time. With the default, the compiler generates short calls within the application's code and long calls from applications to runtime routines.
- c Normally, the compiler creates a registered copy of the user's source code in the library directory for use with the librarian *remake*<sup>3</sup> subcommand. This option suppresses the creation of this copy. If this option is specified, it is very important to provide AdaScope with the directories in which the source files reside. See the section "Locating Source Files" in the *AdaScope Manual*.
- d 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* option<sup>4</sup> can be used in conjunction with this option to check for possible refinements without risking a change to the program library.
- e=*n* 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.
- f Forces the compiler to accept an attempt to compile a unit imported from another library, which is normally prohibited.
- g 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 execution-time penalty is incurred with this option.
- i Causes the compiler to omit data segments with the text of enumeration literals. This text is normally produced for exported enumeration types 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.
- K 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 overridden by compiler options specified on the *(re)make* command line (sec. 20.2.5).
- L= [*project*: ] *library* Selects the library and optionally the project for this compilation. This option takes effect after all commands from the librarian initialization file<sup>5</sup> have been executed, thereby possibly overriding its effects.

<sup>3</sup>UNIX: remakecu

<sup>4</sup>UNIX: -n

<sup>5</sup>UNIX: .adalibr

La	Always generates a listing. The default is to generate a listing only if a diagnostic message is issued.
ll= <i>n</i>	Specifies the number of lines printed on a page of the listing file, where <i>n</i> is an integer in the range 6..9999. The default value for <i>n</i> is 60.
Ln	Never generates a listing. The default is to generate a listing only if a diagnostic message is issued.
lw= <i>n</i>	Specifies the line width used in a listing, where <i>n</i> is an integer in the range 80..132. The default value for <i>n</i> is 80. <sup>6</sup>
m= <i>argument</i>	Controls the type of messages that will be generated by the compiler. The following <i>arguments</i> can be specified: <ul style="list-style-type: none"> <li>e Reports only errors.</li> <li>i Reports errors, warnings, and informational messages.</li> </ul> <p>The default condition, for which there is no option, is to report errors and warnings.</p>
Me	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.
Mw	The compiler attempts to generate extra instructions to fix incorrect address modes. A warning message is issued if such a <i>fixup</i> 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.
n	Specifies that the program library will not be updated with the result of this compilation.
nbl	Suppresses the CALLJ optimization. This optimization enables the linker to choose CALL or BAL calling instructions as appropriate. This option should always be used when compiling the runtime sources (sec. 5.5.3).
no9x	Notifies the compiler that the source file to be compiled contains Ada 83 constructs. This is the default.
nss	Causes the compiler to call runtime routines to perform 8 and 16-bit stores. This option accommodates users with customized 960 hardware that cannot perform 8 and 16-bit stores. The use of this option will produce less efficient code; therefore, it is recommended that this option be used only when absolutely necessary.
O <i>pn</i>	Controls the level of optimization performed by the compiler, where <i>n</i> is an integer in the range 0..4. However, when the code being compiled contains an OPTIMIZE pragma, the pragma takes precedence over the specification of this command line option.

<sup>6</sup>Note that 10 fewer 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.

The following optimization levels can be specified:

- $n = 0$      *minimum*—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.

- p            Loads a syntactically correct compilation unit(s) from the source file into the program 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*<sup>7</sup> is subsequently used to compile the units in the correct sequence (sec. 20.2.5.2).
- pr           Produces the additional code and data necessary to perform profile analysis (see *AdaTrak Manual*).
- REMOTE *host*     Specifies the name of the remote *host* machine on which the compiler will be invoked to perform the current compilation. Section 2.6.4.2 describes the use of this option to perform distributed compilation.
- RS           Causes the compiler to accept non-Ada input, necessary to replace package System. This option should not be used for compiling user-defined packages with illegal code. Any changes of package System must conform to the requirements stated in ARM 4-5, 13.7, and 13.7.1, and must not change the given definition of type Address, to preserve validatability of the Ada system.
- s            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]     Suppresses the given set of checks:

---

<sup>7</sup>UNIX: remakecu

A	ACCESS_CHECK
C	CONSTRAINT_CHECK
D	DISCRIMINANT_CHECK
E	ELABORATION_CHECK
I	INDEX_CHECK
L	LENGTH_CHECK
O	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; i.e., 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      OVERFLOW\_CHECK and "ZERO" DIVISION\_CHECK are suppressed.

S          Suppresses all checks. Invoking this option will not remove all checks if the resulting code without checks will be less efficient.

SC        Suppresses CONSTRAINT\_CHECK, equivalent to SADILR.

v                      Display compiler phase names to standard output. The compiler displays a short description as it progresses through each phase of compilation.

ws=*n*                   Specifies the number of wait states for the memory in which the program code will be executed, where *n* is an integer in the range 0..7. This option affects the code generated for delayed branches and loops, and to account for differences in memory access times. The default value for *n* is 6.

x                      Includes cross-reference information in the object code file (see *AdaRef Manual*).

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

```
tada960runtime tax_spec.ada >& 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 LONG\_INTEGER is range -9\_223\_372\_036\_854\_775\_808 ..  
9\_223\_372\_036\_854\_775\_807;

type FLOAT is digits 6 range  
-16#0.FFF\_FFF#E+32 ..  
..16#0.FFF\_FFF#E+32;

type LONG\_FLOAT is digits 15 range  
-16#0.FFFF\_FFFF\_FFFF\_F8#E+256  
..16#0.FFFF\_FFFF\_FFFF\_F8#E+256 ;

type EXTENDED\_FLOAT is digits 18 range  
-16#0.FFFF\_FFFF\_FFFF\_FFF#E+4096  
..16#0.FFFF\_FFFF\_FFFF\_FFF#E+4096 ;

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

...

end STANDARD;

## APPENDIX F OF THE Ada STANDARD

### 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 LKM, 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 (sec. 4.1.1.6)
- 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 root library (sec. 19.3.15 and 19.3.19).

## APPENDIX F OF THE Ada STANDARD

- The optimization level, as set by the compiler command line option or an OPTIMIZE pragma, 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 an INLINE pragma 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 syntax of the pragma is:

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

The pragma is placed just after the subprogram declaration and will be applied to all subprograms declared thus far with the name subprogram\_name.

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. 19.3.22) must be used to identify the associated object code module.

The language\_name may be Use\_Bal, Use\_Call, or Intel\_C. Any other language\_name will be accepted, but ignored, and the default, Use\_Call, will be used.

While the BAL calling convention is faster than the standard calling convention, be aware that BAL must be used carefully. In particular, when a routine is called with BAL:

- No new stack frame is allocated. The called routine must not change the stack pointer, or must at least ensure that the stack pointer is restored before the routine returns.
- No new local registers are allocated.
- The called routine must return via a bx(reg) instruction. The BAL instruction will automatically store the return address in register g14.
- If a called routine has more than 12 words worth of parameters, the compiler will store the argument block pointer in G14. Since the BAL instruction will place the return address in G14, the called routine could find that its argument block pointer has been corrupted.

Please see Chapter 5 for a complete list of BAL calling convention restrictions.

Pragma INTERFACE(Intel\_C, subprogram\_name) causes the compiler to generate a calling sequence which will be compatible with the Intel C compiler ic960 v2.0. Please see the ic960 Compiler User's Manual, and Chapter 5 of this manual for details.

## APPENDIX F OF THE Ada STANDARD

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. Pragma `MEMORY_SIZE`, `STORAGE_UNIT`, and `SYSTEM_NAME`

This section details the procedure for compiling one of these pragmas. The compilation unit 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 the library `tartan:standard_packages`.
2. Compile the pragma into the library `tartan:standard_packages`. 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 the library `tartan:standard_packages`.

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 exceptions:

- pragma `OPTIMIZE` is not implemented for the declarative part of a block
- in the declarative part of bodies, the effect of pragma `OPTIMIZE` on nested subprograms is dependent on its position in the declarative part, i.e., the pragma applies to subprograms declared after it.

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

```
pragma OPTIMIZE(TIME)
```

has the same effect as compiling the subprogram and specifying optimization level `time` (UNIX: `-Op3`; VMS: `/optimize=time`) on the command line.

When the code being compiled contains an `OPTIMIZE` pragma and the command line option to specify an optimization level is supplied, the pragma takes precedence over the command line option.

## APPENDIX F OF THE Ada STANDARD

### 4.1.1.6. Pragma SHARED

Pragma SHARED is supported as described in the LRM, Annex B.

Users should be aware that one consequence of applying pragma SHARED to a variable is to disable compiler optimizations that remove redundant reads and/or writes to that variable. Thus pragma SHARED allows the user to write loops that poll hardware devices until some change is seen.

```
with System;
package body Device_X is
  Control_Register_For_X : Integer;
  for Control_Register_For_X use at 16#123#;
  pragma SHARED(Control_Register_For_X);
  ...
  function Wait_Until_X_Signals return Integer is
    -- return the value of the control register as soon as it becomes
    -- non-zero
    T : Integer;
  begin
    loop
      T := Control_Register_For_X;
      exit when T /= 0;
    end loop;
    return T;
  end Wait_Until_X_Signals;
  ...
end Device_X;
```

Figure 4-1: Using pragma SHARED to Keep Redundant Reads

Note that pragma SHARED can be applied only to scalar variables. If the object is best described using a structure, one of the following alternatives must be used:

- Compile the code at low optimization levels, causing the compiler not to eliminate the redundant reads and writes. Since most applications cannot tolerate this kind of loss in performance across the board, it is suggested that the critical code be isolated into its own compilation unit.
- Declare the variable as a scalar (or adjacent scalars) and use unchecked conversion to copy to/from local variables of the proper type just after/before read/write operations.
- Use non-optimizable procedures that implement a read/write of the object. Package Low\_Level\_IO (LRM 14.6) provides one such set of routines. Other high-efficiency routines may be built using package Machine\_Code
- Another method is to create the obvious simple Ada routines, but isolate them in a separate package and do not use cross-compilation optimizations, such as pragma INLINE.

Figure 4-2 shows how package Low\_Level\_IO is used to force the compiler to keep redundant reads of a non-scalar object.

## APPENDIX F OF THE Ada STANDARD

```
package Device_Y is
end Device_Y;

with Unchecked_Conversion;
with Low_Level_IO;
with System;
package body Device_Y is
...
type Y_Regs_Type is
  record
    Control : Integer;
    Input   : Integer;
    Output  : Integer;
  end record;
Y : Y_Regs_Type;
for Y use at 16#246#;
function Address_to_DeviceClass is
  new Unchecked_Conversion(System.Address,
                           Low_Level_IO.Device_Class);
...
function Wait_Until_Y_Signals return Integer is
  -- return the value of the control register as soon as it becomes
  -- non-zero
  T : Low_Level_IO.Word_Data;
begin
  loop
    Low_Level_IO.Receive_Control(
      Address_to_DeviceClass(Y.Control'address), T);
    exit when Integer(T) /= 0;
  end loop;
  return Integer(T);
end Wait_Until_Y_Signals;
...
end Device_Y;
```

Figure 4-2: Using Low\_Level\_IO to Force Compiler to Keep Redundant Reads

### 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

## APPENDIX F OF THE Ada STANDARD

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).

### 4.1.2.1.1. Calling Ada Subprograms from non-Ada Code

Pragma LINKAGE\_NAME can be used to allow non-Ada code to call an Ada subprogram. Calling Ada from non-Ada code is highly dependent on the language the call is being made from as well as the compiler for that language.

First, the Ada subprogram must be given a linkage name so that the non-Ada code will be able to call the Ada subprogram. Pragma LINKAGE\_NAME is used to perform this task. Next, an unimplemented subprogram must be defined in the non-Ada code. This subprogram must have the same linkage name as specified by the pragma LINKAGE\_NAME on the Ada subprogram. Calls to this unimplemented subprogram will then go to the Ada subprogram if the object file containing the Ada subprogram was linked into the application. Note that a conversion between the object file (file.tof) for the Ada subprogram and the non-Ada object file will probably be needed. See the Object File Utilities Manual for more on object file conversion.

For the call to be made correctly, you must ensure that the calling convention and parameter passing mechanisms of the non-Ada code are compatible with the Tartan Ada Compiler (sec. 5.4 and 5.5).

Pragma LINKAGE\_NAME can also be applied to an Ada object or exception. This allows non-Ada code to refer to the Ada entity. In doing so, you should be aware of the Ada rules governing implicit initialization of objects (Ada LRM 3.2.1 and 3.3) and the Tartan Ada Compiler's data representation conventions

## APPENDIX F OF THE Ada STANDARD

(sec. 5.1).

### 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).

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 19.3.18).

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

## APPENDIX F OF THE Ada STANDARD

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
4. Issue the command:

UNIX:

```
adalib960runtime foreign math_functions math.tof
```

VMS:

```
al960 foreign math_functions math.tof
```

5. Issue the command:

UNIX:

```
adalib960runtime link main
```

## APPENDIX F OF THE Ada STANDARD

VMS:

al960 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 SYSTEM\_TICK

This section details the procedure for compilation of a new unit, such as pragma SYSTEM\_TICK, with a system pragma. The new unit must be compiled into a library that contains package SYSTEM. For most users, the Tartan Ada Standard Packages Library will be the library that also includes package SYSTEM.

1. Thaw the library tartan:standard\_packages.
2. Compile this unit into the library tartan:standard\_packages This step updates package SYSTEM.
3. Freeze the library tartan:standard\_packages.

Following these steps will allow you to modify the value of SYSTEM\_TICK at compilation time.

### 4.1.2.4. 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 overloads 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.

## APPENDIX F OF THE Ada STANDARD

The pragma `UNCHECKED_NO_STATE_WRITTEN_OR_READ` indicates that the named subprogram behaves strictly as a mathematically pure function. In essence, this statement 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 its own scope, in and in out parameters, and objects declared as constants may be read freely. Called subprograms must themselves be `NO_STATE_WRITTEN_OR_READ`.

The compiler may choose to make common subexpressions of the results of calls to the named subprogram. It may also remove such calls entirely when the result of the subprogram is not used; calls may also be loop-invariant hoisted.

Subprograms which are likely candidates for this pragma include math package subprograms, matrix math subprograms, trigonometric functions, etc.

### Caution:

Pragmas `UNCHECKED_NO_STATE_WRITTEN` and `UNCHECKED_NO_STATE_WRITTEN_OR_READ` are strictly assertive in nature and are entirely unchecked. The compiler will not notify you if the body of the subprogram does not meet the requirements of the pragma. When one of these pragmas is incorrectly applied to a subprogram which does not meet its requirements, the behavior of your program 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 960 in package `System` (LRM 13.7.1 and Appendix C) are:

```
package System is
  type Address is new Integer;
  type Name is (I960MC);
  System_Name : constant Name := I960MC;
  Storage_Unit : constant := 8;
  Memory_Size : constant := 2_097_152;
  Max_Int : constant := 9_223_372_036_854_775_807;
  Min_Int : constant := -Max_Int - 1;
  Max_Digits : constant := 18;

  Max_Mantissa : constant := 31;
  Fine_Delta : constant := 2#1.0#e-31;
  Tick : constant := 0.015625; [For the CVME962
target board the value of System.Tick is 0.000061.]
  subtype Priority is Integer range 2 .. 17;
  Default_Priority : constant Priority := Priority'First;
  Runtime_Error : exception;
end System;
```

## APPENDIX F OF THE Ada STANDARD

### 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 `Generic_Type_Definition` (LRM 12.1) and a `Private_Type_Definition` (LRM 7.4). Any representation clause in violation of these rules is not obeyed by the compiler; an error message is issued.

Further restrictions are explained in the following sections. Any representation clauses violating those restrictions cause compilation to stop and a diagnostic message to be issued.

#### 4.4.2. Length Clauses

Length clauses (LRM 13.2) are, in general, supported. The following sections detail use 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 example:

```
type My_Enum is (A, B);
for My_Enum'Size use 1;
V, W : My_Enum; -- will occupy two storage
                -- units on the stack
                -- (if allocated at all)

type Rec is record
  V, W : My_Enum;
end record;
pragma PACK(Rec);
O : Rec;        -- will occupy one storage unit
```

## APPENDIX F OF THE Ada STANDARD

- A formal parameter of the type is sized according to calling conventions rather than size specifications of the type. Appropriate size conversions upon parameter passing take place automatically and are transparent to the user.
- Adjacent bits to an object that is a component of a composite object, but whose size is non-referable, may be affected by assignments to the object, unless these bits are occupied by 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.

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.

## APPENDIX F OF THE Ada STANDARD

### 4.4.2.2. Size Specification for Scalar Types

The specified size must accommodate all possible values of the type including the value 0 (zero), even if 0 is not in the range of the values of the type. For numeric types with negative values, the number of bits must account for the sign bit. Biased representation is not 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 fixed-point type does not affect the accuracy of operations on the type. Such influence should be exerted via the Accuracy\_Definition of the type (LRM 3.5.9).

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.

## APPENDIX F OF THE Ada STANDARD

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 1 (one) word results in an allocation of 2 words.

Furthermore, the allocator must round non-word sized requests up to the nearest word. For example, a request of 11 bytes is rounded up to 12 bytes (3 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.

Collection sizes may not be specified for an access type whose designated type is a task type.

### 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. When using the Physical Mode runtime system, if the storage specified for a task activation (T.Storage\_Size) is not a multiple of 64 (one page), the compiler allocates the next higher multiple of 64, as permitted by the language. With the Small Virtual Mode runtime system, if the storage specified for a task activation (T.Storage\_Size) is not a multiple of 4,096 (one page), the compiler allocates the next higher multiple of 4,096.

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.

## APPENDIX F OF THE Ada STANDARD

### 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 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 0 (zero). 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 for library-level record objects.

The alignment clause has no effect on objects not defined on a library level (e.g., subprogram locals). The value given in the alignment clause is in addressable units and is restricted to the powers of two in the range of  $2^0 \dots 2^{15}$ . The specified alignment becomes the minimum alignment of the record type, unless the minimum alignment of the record forced by the component allocation is already more stringent than the value specified by the alignment clause.

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.

## APPENDIX F OF THE Ada STANDARD

### 4.4.5. Address clauses

Address clauses (LRM 13.5) are supported with the following restrictions:

- When applied to an object, an address clause becomes a linker directive to allocate the object at the given address. For any object not declared immediately within a top-level library package, the address clause is accepted but meaningless. Please refer to section 14.5 for details on how address clauses relate to linking with the Small Virtual Mode runtime system; refer to section 16.3.2 for an example.
- Address clauses applied to local packages are not supported by Tartan Ada. Address clauses applied to library packages are prohibited by the syntax; therefore, an address clause can be applied to a package only if it is a body stub.
- Address clauses applied to subprograms and tasks are implemented according to the LRM rules. When applied to an entry, the specified value identifies an interrupt in a manner customary for the target. Immediately after a task is created, a runtime call is made for each of its entries having an address clause, establishing the proper binding between the entry and the interrupt. Refer to section 10.3.4 for more details. A specified address must be an Ada static expression, as defined in LRM 4.9.
- Address clauses which are applied to objects, subprograms, packages, or task units specify virtual addresses. When using the 960 Physical Mode runtime system, virtual and physical addresses are identical unless the PHYSICAL command is used in the linker control file. The range of System.Address is -16#8000\_0000#..16#7fff\_ffff#. To represent a machine virtual address in the range 16#0000\_0000#..16#7fff\_ffff#, use the corresponding System.Address. To represent a machine virtual address greater than System.Address 16#7fff\_ffff#, use the negated radix-complement of the desired machine virtual address. For example, to express machine virtual address 16#C000\_0000#, use 16#C000\_0000#-2\*\*32.

Note: Creating an overlay of two objects by means of address clauses is possible with Tartan Ada. However, such overlays (which are considered erroneous by the Ada LRM 13.5(8)) will not be recognized by the compiler as an aliasing that prevents certain optimizations. Therefore, problems may arise if reading and writing of the two overlaid objects are intermingled. For example, if variables A and B are overlaid by means of address clauses, the Ada code sequence:

```
A := 5;
B := 7;
if A = 5 then raise Surprise; end if;
```

may well raise the exception Surprise, since the compiler believes the value of A to be 5 even after the assignment to B.

### 4.4.6. Pragma PACK

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

## APPENDIX F OF THE Ada STANDARD

### 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 were applied with 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: 4 characters per 32-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 non-referable. 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 boolean 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.

## APPENDIX F OF THE Ada STANDARD

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 boundary 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 implementation dependent.

The Ada runtimes provide interrupts that may be associated with task entries. These interrupts are of type System.Address in the ranges 8..243, 252..255, 264..499, and 508..511. Refer to section 10.3.4 for further explanation. If the argument is outside those ranges, the compiler will report an error.

### 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 or access to 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. However, since the i960 processors are used in embedded applications lacking both standard I/O devices and file systems, the functionality of packages Direct\_IO, Sequential\_IO, and Text\_IO is limited.

Packages Direct\_IO and Sequential\_IO raise Use\_Error if a file open or file access is attempted. Package Text\_IO is supported to Current\_Output and from Current\_Input. A routine that takes explicit file names raises Use\_Error. Package Low\_Level\_IO for the i960 processors provides an interface by which the user may read and write from memory mapped devices. In both the

## APPENDIX F OF THE Ada STANDARD

Send\_Control and Receive\_Control procedures, the device parameter specifies a device address while the data parameter is a byte, halfword, word, or doubleword of data transferred.

### 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. No Use of Numeric\_Error

No predefined operations will raise the exception Numeric\_Error. The Tartan Ada Compiler raises the predefined exception Constraint\_Error in situations where, according to the Ada LRM, the predefined exception Numeric\_Error should be raised.

This change in the compiler has been made in accordance with the approved Ada Interpretation AI-00387. ``An Ada program is portable if any handler intended to process Numeric\_Error provides a choice for Constraint\_Error and processes it in a similar manner.''

If this procedure has been followed, you will not experience any change in program behavior because of this change in the compilation system.

#### 4.9.3. 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.

#### 4.9.4. Attributes of Type Duration

The type Duration is defined with the following characteristics:

APPENDIX F OF THE Ada STANDARD

Attribute	Value
Duration'Delta	0.0001 sec
Duration'Small	0.000061 sec
Duration'First	-86400.0 sec
Duration'Last	86400.0 sec

4.9.5. Values of Integer Attributes

Tartan Ada supports the predefined integer type Integer. The range bounds of the predefined type Integer are:

Attribute	Value
Integer'First	$-2^{**31}$
Integer'Last	$2^{**31}-1$
Long_Integer'First	$-2^{**63}$
Long_Integer'Last	$2^{**63}-1$
Short_Integer'First	$-2^{**15}$
Short_Integer'Last	$2^{**15}-1$

APPENDIX F OF THE Ada STANDARD

Byte_Integer'First	-128
Byte_Integer'Last	127

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 package Direct\_IO are:

Attribute	Value
Count'First	0
Count'Last	Integer'Last

## APPENDIX F OF THE Ada STANDARD

Positive_Count'First	1
Positive_Count'Last	Count'Last

### 4.9.6. Ordinal Types

Ordinal types are supported via two separate packages which are included with the standard packages. The Ada specifications for the Tartan Ada standard packages are located online in the directory:

UNIX: /install-dir/960/version/std\_packages/src

VMS: [install-dir.960.version.std\_packages.src]      Package Ordinal\_Support provides support for unsigned arithmetic, including functions which convert between Integer and Ordinal types, and a complete set of Ordinal arithmetic operations. Package Intrinsic\_Ordinal\_Support provides ordinal wraparound arithmetic. Further information on Intrinsic\_Ordinal\_Support may be found in section 6.2.

### 4.9.7. Values of Floating-Point Attributes

Tartan Ada supports the predefined floating-point types Float, Long\_Float, and Extended\_Float.

Attribute	Value for Float
Digits	6
Mantissa	21
Emax	84
Epsilon	16#0.1000_000#E-4 (approx. 9.536743E-07)
Small	16#0.8000_000#E-21 (approx. 2.58494E-26)

APPENDIX F OF THE Ada STANDARD

Large	16#0.FFFF_F80#E+21 (approx. 1.93428E+25)
Safe_Emax	126
Safe_Small	16#0.2000_000#E-31 (approx. 5.87747E-39)
Safe_Large	16#0.3FFF_FE0#E+32 (approx. 8.50706E+37)
First	-16#0.FFFFFFF#E+32 (approx. -3.40282E+38)
Last	16#0.FFFFFFF#E+32 (approx. 3.40282E+38)
Machine_Radix	2
Machine_Mantissa	24
Machine_Emax	128
Machine_Emin	-125
Machine_Rounds	true
Machine_Overflows	true

APPENDIX F OF THE Ada STANDARD

Attribute	Value for Long_Float
Digits	15
Mantissa	51
E <sub>max</sub>	204
Epsilon	16#0.4000_0000_0000_000#E-12 (approx. 8.8817841970013E-16)
Small	16#0.8000_0000_0000_000#E-51 (approx. 1.9446922743316E-62)
Large	16#0.FFFF_FFFF_FFFF_E00#E+51 (approx. 2.5711008708143E+61)
Safe_E <sub>max</sub>	1022
Safe_Small	16#0.2000_0000_0000_000#E-255 (approx. 1.1125369292536E-308)
Safe_Large	16#0.3FFF_FFFF_FFFF_F80#E+256 (approx. 4.4942328371557E+307)
First	-16#0.FFFFFFFFFFFFF8#E+256 (approx. -1.7976931348623E+308)
Last	16#0.FFFFFFFFFFFFF8#E+256 (approx. 1.7976931348623E+308)

APPENDIX F OF THE Ada STANDARD

Machine_Radix	2
Machine_Mantissa	53
Machine_Emax	1024
Machine_Emin	-1021
Machine_Rounds	true
Machine_Overflows	true

Attribute	Value for Extended_Float
Digits	18
Mantissa	61
Emax	244
Epsilon	16#0.1000_0000_0000_0000_0#E-14 (approx. 8.67361737988403547E-19)
Small	16#0.8000_0000_0000_0000_0#E-61 (approx. 1.76868732008334226E-74)
Large	16#0.FFFF_FFFF_FFFF_FFF8_0#E+61 (approx. 2.82695530364541493E+73)

APPENDIX F OF THE Ada STANDARD

Safe_Emax	16382
Safe_Small	16#0.2000_0000_0000_0000_0#E-4096 (approx. 1.68105157155604675E-4932)
Safe_Large	16#0.3FFF_FFFF_FFFF_FFFF_0#E+4096 (approx. 2.97432873839307941E+4931)
First	-16#0.FFFFFFFFFFFFFFFF#E+4096 (approx. -1.189731495357231766E+4932)
Last	16#0.FFFFFFFFFFFFFFFF#E+4096 (approx. 1.189731495357231766E+4932)
Machine_Radix	2
Machine_Mantissa	64
Machine_Emax	16384
Machine_Emin	-16381
Machine_Rounds	true
Machine_Overflows	true

4.10. SUPPORT FOR PACKAGE Machine\_Code

Package Machine\_Code provides the programmer with an interface through which to request the generation of any instruction that is available on the i960 processors. The Tartan Ada 960 implementation of package Machine\_Code is similar to that described in section 13.8 of the Ada LRM, with several added features. The Ada specifications for the Tartan Ada standard packages are located online in the directory:

## APPENDIX F OF THE Ada STANDARD

UNIX: /install-dir/960/version/std\_packages/src  
VMS: [install-dir.960.version.std\_packages.src]

### 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 four types of records. Each has an opcode and zero to three operands. These records are adequate for the expression of all instructions provided by the i960 processors.

### 4.10.3. Operands

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.

#### 4.10.3.1. Address Modes

Each operand in a machine code insert must have an `Address_Mode_Name`. The address modes provided in package `Machine_Code` provide access to all address modes supported by the i960 processors.

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 (i.e., a branch target or the source of an LDA instruction). `Symbolic_Value` should be used when the operand is actually a value (that is, 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 (sec. 4.10.10).

### 4.10.4. Examples

The Tartan Ada 960 implementation of package `Machine_Code` makes it possible to specify both simple machine code inserts such as

```
Two_Format'(MOV, (REG_LIT, 5), (REG, R5))
```

and more complex inserts such as

## APPENDIX F OF THE Ada STANDARD

```

Three_Format' (MULI,
               (Symbolic_Value, Array_Var(X, Y, 27)'Address),
               (LIT, 123456),
               (Symbolic_Address, Parameter_1'Address))
    
```

In the first example, the compiler will emit the instruction `Mov 5, r5`. In the second example, the compiler will first emit whatever instructions are needed to form the address of `Array_Var(X, Y, 27)`, load the value found at that address into a register, load 123456 into a register, and then emit the `MULI` instruction. If `Parameter_1` is not found in a register, the compiler will put the result of the multiplication in a temporary register and then store it to `Parameter_1'Address`. Note that the destination operand of the `MULI` instruction is given as a `Symbolic_Address`. This behavior holds true for all destination operands. 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 a memory address as the destination of an `ADD` instruction, the compiler will accept it and try to generate correct code. In this case, the compiler will allocate a temporary register to use as the destination of the `ADD`, and then store from that register to the desired location in memory.

## APPENDIX F OF THE Ada STANDARD

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.

The compiler will always emit the instruction named in the machine code insert, even if it was necessary to correct all of its operands. In extreme cases, this action can lead to surprising code sequences. Consider, for example, the machine code insert

```
TWO_FORMAT' (MOV, (REG_IND, G0), (REG_IND_DISP, G1, 128))
```

The MOV instruction requires two registers, but both operands are memory addresses. The compiler will generate a code sequence like

```
LD      (G0), G12
MOV     G12, G13
ST      G13, 128(G1)
```

Note that the MOV instruction is generated even though a LD ST combination would have been sufficient. As a result of always emitting the instruction specified by the programmer, the compiler will never optimize away instructions which it does not understand (such as SENDSERV), unless they are unreachable by ordinary control flow.

### 4.10.6. Assumptions Made in Correcting Operands

When compiling in -Mw or fixup=warn modes, or the default or fixup=quiet 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.

For source 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 binding 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 by the '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 960.
- Label indicates that the address of the label will be used as the source bits.
- Any other non-register means that the value found at the address specified by the operand will be used as the source bits.

## APPENDIX F OF THE Ada STANDARD

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 960.
- `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 action 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.

Table 4-1 below describes the fixup operation attempted for each possible instruction-operation combination. The fixup actions shown in the table have the following meanings:

### Load to Register 1

The operand given represents a memory location, but the instruction requires a register. The operand is used as a source. The compiler will load from the operand to a temporary register.

### Load to Register 2

The operand given represents a register, but the instruction requires a memory location. The operand is a destination. The compiler will store the result value to a scratch memory location, and then load it into the specified register.

### Store to Memory 1

The operand given represents a register, but the instruction requires a memory location. The operand is a source. The compiler will store the value to a scratch memory location so that it will be in the proper place for the instruction.

### Store to Memory 2

The operand given represents a memory location, but the instruction requires a register. The operand is a destination. The compiler will allocate a scratch register, use that register as the destination for the instruction, and then store the result value to the specified memory address.

### Store to Memory 3

The operand given is not the address of a label. The operand will be stored to a scratch memory location, and then used as the indirect branch target.

### Error 1

The only incorrect operand for the source of an LDA is a register. It is not possible to take the address of a register on the 960.

## APPENDIX F OF THE Ada STANDARD

Error 2	The operand must be a Label Address.		
Inst	Opnd1	Opnd2	Opnd3
addo, addi, addc, addr, addr..	Load to Register 1	Load to Register 1	Store to Memory 2
alterbit	Load to Register 1	Load to Register 1	Store to Memory 2
and, andnot	Load to Register 1	Load to Register 1	Store to Memory 2
atadd	Load to Register 1	Load to Register 1	Store to Memory 2
atanr, atanrl	Load to Register 1	Load to Register 1	Store to Memory 2
atmod	Load to Register 1	Load to Register 1	Store to Memory 2
b	Error 2		
bx	Store to Memory 3		
bal	Error 2		
balx	Store to Memory 3		
bbc, bbs	Load to Register 1	Load to Register 1	Error 2
BRANCH IF	Error 2		
call	Error 2		
calls	Load to Register 1		
callx	Store to Memory 3		
chkbit	Load to Register 1	Load to Register 1	
classr, classrl	Load to Register 1		
clrbit	Load to Register 1	Load to Register 1	Store to Memory 2
cmpi, cmpo	Load to Register 1	Load to Register 1	
cmpdeci, cmpdeco	Load to Register 1	Load to Register 1	Store to Memory 2
cmpinci, cmpinco	Load to Register 1	Load to Register 1	Store to Memory 2
cmpor, cmporl	Load to Register 1	Load to Register 1	
cmprr, cmprrl	Load to Register 1	Load to Register 1	
cmpstr	Load to Register 1	Load to Register 1	Load to Register 1
COMPARE AND BRANCH	Load to Register 1	Load to Register 1	Error 2
concmpi, concmpo	Load to Register 1	Load to Register 1	
condrec	Load to Register 1	Load to Register 1	
condwait	Load to Register 1		
cosr, cosrl	Load to Register 1	Store to Memory 2	
cpysre, cpysre	Load to Register 1	Load to Register 1	Store to Memory 2
cvtilr	Load to Register 1	Store to Memory 2	
	(64 bits)		
cvtir	Load to Register 1	Store to Memory 2	
cvtri	Load to Register 1	Store to Memory 2	
cvtril	Load to Register 1	Store to Memory 2 (64 bits)	
cvtzri	Load to Register 1	Store to Memory 2	
cvtzril	Load to Register 1	Store to Memory 2 (64 bits)	
daddc	Load to Register 1	Load to Register 1	Store to Memory 2
divo, divi, divr, divrl	Load to Register 1	Load to Register 1	Store to Memory 2
dmovt	Load to Register 1	Store to Memory 2	
dsubc	Load to Register 1	Load to Register 1	Store to Memory 2
ediv	Load to Register 1	Load to Register 1	Store to Memory 2
		(64 bits)	(64 bits)
emul	Load to Register 1	Load to Register 1	Store to Memory 2
			(64 bits)
expr, exprl	Load to Register 1	Store to Memory 2	
extract	Load to Register 1	Load to Register 1	Store to Memory 2
FAULT IF			
fill	Load to Register 1	Load to Register 1	Load to Register 1
flushreg			

## APPENDIX F OF THE Ada STANDARD

fmark			
inspace	Load to Register 1	Store to Memory 2	
LOAD	Store to Memory 1	Store to Memory 2	
lda	Error 1	Store to Memory 2	
ldphy	Load to Register 1	Store to Memory 2	
ldtime	Store to Memory 2		
logbnr, logbnrl	Load to Register 1	Store to Memory 2	
logepr, logeprl	Load to Register 1	Store to Memory 2	
logr, logrl	Load to Register 1	Store to Memory 2	
mark			
modac	Load to Register 1	Load to Register 1	Store to Memory 2
modi	Load to Register 1	Load to Register 1	Store to Memory 2
modify	Load to Register 1	Load to Register 1	Store to Memory 2
modpc	Load to Register 1	Load to Register 1	Store to Memory 2
modtc	Load to Register 1	Load to Register 1	Store to Memory 2
MOVE	Load to Register 1	Store to Memory 2	
movqstr, movstr	Load to Register 1	Load to Register 1	Load to Register 1
mulo, muli, mulr,	Load to Register 1	Load to Register 1	Store to Memory 2
mulrlnand	Load to Register 1	Load to Register 1	Store to Memory 2
nor	Load to Register 1	Load to Register 1	Store to Memory 2
not	Load to Register 1	Store to Memory 2	
notand	Load to Register 1	Load to Register 1	Store to Memory 2
notbit	Load to Register 1	Load to Register 1	Store to Memory 2
notor	Load to Register 1	Load to Register 1	Store to Memory 2
or, ornot	Load to Register 1	Load to Register 1	Store to Memory 2
receive	Load to Register 1	Load to Register 1	
remo, remi, remr,	Load to Register 1	Load to Register 1	Store to Memory 2
remrlresumprcs	Load to Register 1		
ret			
rotate	Load to Register 1	Load to Register 1	Store to Memory 2
roundr, roundrl	Load to Register 1	Store to Memory 2	
saveprcs			
scaler, scalerl	Load to Register 1	Load to Register 1	Store to Memory 2
scanbit	Load to Register 1	Store to Memory 2	
scanbyte	Load to Register 1	Load to Register 1	
schedprcs	Load to Register 1		
send	Load to Register 1	Load to Register 1	Load to Register 1
sendserv	Load to Register 1		
setbit	Load to Register 1	Load to Register 1	Store to Memory 2
SHIFT	Load to Register 1	Load to Register 1	Store to Memory 2
signal	Load to Register 1		
sinr, sinrl	Load to Register 1	Store to Memory 2	
spanbit	Load to Register 1	Store to Memory 2	
sqrtr, sqrtrl	Load to Register 1	Store to Memory 2	
STORE	Load to Register 1	Load To Register 2	
subo, subi, subc,	Load to Register 1	Load to Register 1	Store to Memory 2
subr, subrlsyncf			
synld	Load to Register 1	Store to Memory 2	
synmov, synmovl,	Load To Register 1	Load to Register 1	
synmovq			
tanr, tanrl	Load to Register 1	Store to Memory 2	
TEST	Store to Memory 2		
wait	Load to Reg 1		
xnor, xor	Load to Register 1	Load to Register 1	Store to Memory 2

Table 4-1: MACHINE\_CODE Fixup Operations

## APPENDIX F OF THE Ada STANDARD

### 4.10.7. Register Usage

Since the compiler may need to allocate registers as temporary storage in machine code routines, there are some restrictions placed on register usage. The compiler will automatically free all the registers which would be volatile across a call. Therefore you may use G0..G7, G13, and G14. However, it is advisable to leave at least three registers available for corrections.

Note: The compiler may need several registers to generate code for operand corrections in machine code inserts. If all registers are used, corrections will not be possible. If a fixup is needed, the compiler may require up to three registers to guarantee success. In general, when more registers are available to the compiler, it is able to generate better code.

If any other register or registers are referenced, the compiler will reserve it (them) for use by the user until the end of the machine code routine. The compiler will not save the register(s) automatically.

With the mode of operation as described above, Tartan recommends 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 a machine code insert even if it is inline expanded in another routine.

### 4.10.8. Inline Expansion

Routines which contain machine code inserts may be inline expanded into the bodies of other routines. This expansion 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 as well as the other call assumptions.

### 4.10.9. 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 assumption 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 and make sure that there is no pragma `INLINE` for it.

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.

## APPENDIX F OF THE Ada STANDARD

- Do not attempt to move multiple Ada objects with a single long instruction such as MOVL or STT. 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, storage allocation may be rearranged such that objects which were previously adjacent are no longer adjacent.
- The compiler will not generate call site code for you if you emit a call instruction. You must save and restore any volatile registers 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 assume that the 'Address on Symbolic\_Address or Symbolic\_Value operands provide an Address on which to operate. The ...\_Address or ...\_Value of an operand is determined by choice of Symbolic\_Address or Symbolic\_Value. For example, to add the contents of X to R3:

```
Three_Format'(ADDI, (Symbolic_Value, X'Address),
               (REG, R3), (REG, R3));
```

but to add the address of X to r3:

```
Three_Format'(ADDI, (Symbolic_Address, X'Address),
               (REG, R3), (REG, R3));
```

- The compiler will not prevent writing to register R3, which is used to hold the address of the current exception handler. This feature provides the opportunity to make a custom exception handler. However, there is considerable danger in creating it. Knowledge of the details on the structure of exception handlers will help; see the Tartan Ada Runtime Modification Manual, available only if you have purchased the Runtime Modification Kit.

### 4.10.10. Limitations

- When specifying absolute addresses in machine-code inserts, please note that the compiler will treat addresses as an Integer type and specifications of addresses may raise arithmetic overflow errors. Therefore, addresses must be in the range Integer'First..Integer'Last. To represent an address greater than Integer'Last, use the negated radix-complement of the desired address. For example, to express address 16#C000\_0000#, specify instead -16#4000\_0000#.
- 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 operated on by the instruction chosen in the machine-code insert.

## APPENDIX F OF THE Ada STANDARD

For example:

```
Three_Format'(ADDO, (Symbolic_Value, Byte_Variable'Address),
              (REG, R0), (REG, R1))
```

will not generate correct code when `Byte_Variable` is bound to memory. The compiler will assume that `Byte_Variable` is 32 bits, when in fact it is only 8, and will emit an LD instruction to load the value of `Byte_Variable` into a register. If, on the other hand, `Byte_Variable` was bound to a register, the insertion will function properly as no correction is needed.

- The compiler generates incorrect code when the BAL and BALX instructions are used with symbolic operands which are not of the form `routine'Address`. To get the effect of an unconditional branch, use the B or BX instructions instead.
- Note that the use of `X'Address` in a machine code insert does not guarantee that X will be bound to memory because `'Address` provides a `''typeless''` method for naming Ada objects in machine code inserts. For example, it is legal to say to `(Symbolic_Value, X'Address)` in an insert even when X is a formal parameter of the machine code routine, and is thus found in a register.

### 4.10.11. Example

```
package Mtest is
  type Ary_TYPE is array(1..4) of Integer;

  procedure Inline_Into_Me;
end Mtest;

with Machine_Code;
use Machine_Code;
package body Mtest is
  Own_Var : Integer := -1;

  procedure Mach_Test(X, Y, Z: in Integer; Ary: in out Ary_Type) is
  begin
    -- The next instruction is only OK if this routine is not INLINED.
    -- If the routine is inlined, there is no guarantee that parm X will
    -- be either A) in an even numbered register, or B) "next to" parm
    -- Y. If the programmer uses an instruction like MOVL here, there
    -- may be false assumptions about the generated code, the program is
    -- erroneous. On the other hand, the use of X'Address does guarantee
    -- that the instruction will use X even when this routine is inline
    -- expanded into a caller.
    Two_Format'(MOVL, (Symbolic_Value, X'Address), (REG, G6));
    Two_Format'(MOV, (Symbolic_Value, X'Address), (REG, G6));
    Two_Format'(MOV, (Symbolic_Value, Y'Address), (REG, G7));
    Two_Format'(MOV, (Symbolic_Value, Z'Address), (REG, G8));
    Three_Format'(ADDI, (Symbolic_Value, X'Address), (REG, G8), (REG, G1));
    Three_Format'(MULI,
                  (REG, G7),
                  (Symbolic_Value, Y'Address),
                  (REG, G12));
```

## APPENDIX F OF THE Ada STANDARD

```

-- Note the use of a complicated Ada object in this instruction.
Two_Format'(ST,
             (REG, G12),
             (Symbolic_Address, Ary(1)'Address));

-- In this instruction, note that Ary(1)'Address is NOT kept in a
-- register and is thus NOT a legal source for XORi. It is OK,
-- because the compiler can fix it up for the user.
Three_Format'(XORi, (Symbolic_Value, Ary(1)'Address),
              (Symbolic_Value, Ary(2)'Address), (REG, G12));
Two_Format'(ST,
             (REG, G12),
             (Symbolic_Address, Ary(3)'Address));
Two_Format'(ST,
             (REG, G12),
             (Symbolic_Address, Ary(X)'Address));
Two_Format'(ST,
             (REG, G0),
             (Symbolic_Address, Own_Var'Address));
Two_Format'(LDA,
             (Symbolic_Value, Own_Var'Address),
             (REG, G14));
One_Format'(CALLX, (Symbolic_Address, Inline_Into_Me'Address));
end Mach_Test;

pragma INLINE(Mach_Test);
procedure Mtest1(First, Second, Third: in Integer; Fourth: out Ary_Type) is
begin
  -- Note the use of Fourth(1)'Address as the destination of the MOV
  -- instruction. The compiler will understand that the user "really
  -- wanted" something moved to Fourth(1)'Address, and will make sure
  -- that the bits get there. The compiler does NOT assume that it
  -- knows enough to second guess the user's choice of instructions.
  -- We generate the MOV, followed by a store to memory.
  Two_Format'(MOV,
              (Symbolic_Value, First'Address),
              (Symbolic_Address, Fourth(1)'Address));
end Mtest1;

procedure Inline_Into_Me is
  Array1 : Ary_Type := (1, 2, 3, 4);
begin
  if Array1(3) >= 0 then
    -- note that Mach_Test is inline expanded
    Mach_Test(22, 41, Array1(4), Array1);
  else
    Mtest1(1, 2, 3, Array1);
  end if;
end Inline_Into_Me;
end Mtest;

```

## APPENDIX F OF THE Ada STANDARD

Assembly code output:

```

.data
    .align 2
ADA.OWN: .space 4
    .align 2
ADA.GLOBAL: .space 1
    .globl xxmtest$inline_into_me$00
    .globl xxmtest$declare
    .globl xxmtest$body

.seto own_var$00,ADA.OWN,0
    .globl xxmtest$inline_into_me$goto$00
.seto xxmtest$inline_into_me$goto$00,ADA.GLOBAL,0

.text

    .align 4

xxmtest$inline_into_me$00:
    mov    0,r3
    st     sp,40(sp)
    lda   40(sp),sp
    st     g12,100(fp)
    lda   .L19,r3

    mov    1,g13    #           line 74

    st     g13,80(fp)
    mov    2,g13
    st     g13,84(fp)
    mov    3,g13
    st     g13,88(fp)
    mov    4,g13
    st     g13,92(fp)
    ldq   80(fp),g4
    stq   g4,64(fp)
    ld     72(fp),g13    #           line 76

    cmpibg 0,g13,.L17
    ld     76(fp),g13    #           line 78

    mov    22,r4
    addo   31,10,r5
    mov    g13,r6
    lda   64(fp),r8
    movl   r4,g6
    mov    r4,g6
    mov    r5,g7
    mov    r6,g8
    addi   r4,g8,g1
    muli   g7,r5,g12
    st     g12,96(fp)
    ld     96(fp),g13
    st     g13,(r8)
    ld     4(r8),r7

```

APPENDIX F OF THE Ada STANDARD

```

xor      g13,r7,g12
st       g12,8(r8)
subi    1,r4,g13      #           line 46

cmpto   g13,3
faultg
st       g12,-4(r8)[r4*4]
st       g0,ADA.OWN
lda     ADA.OWN,g14
callx   xxmtest$inline_into_me$00
b       .L19      #           line 76

.L17:   mov      1,g0      #           line 81

mov      2,g1
mov      3,g2
lda     64(fp),g3
bal     mtest1$00

.L19:

ld       100(fp),g12
ret

```

# Total bytes of code in the above routine = 216

.align 4

```

mach_test$00:
mov      0,r3
st       sp,8(sp)
addo    8,sp,sp
st       g12,68(fp)
lda     .L21,r3

movl    g0,g6
mov      g0,g6
mov      g1,g7
mov      g2,g8
addi    g0,g8,g1
muli    g7,g1,g12
st       g12,64(fp)
ld       64(fp),g13
st       g13,(g3)
ld       4(g3),g5
xor     g13,g5,g12
st       g12,8(g3)
subi    1,g0,g13      #           line 46

cmpto   g13,3
faultg
st       g12,-4(g3)[g0*4]
st       g0,ADA.OWN
lda     ADA.OWN,g14
callx   xxmtest$inline_into_me$00

.L21:   ld       68(fp),g12
ret

```

## APPENDIX F OF THE Ada STANDARD

```
# Total bytes of code in the above routine = 124
```

```
.align 4
```

```
mtest1$00:
```

```
mov    g14,g4  
mov    0,g14
```

```
mov    g0,g13  
st     g13,(g3)
```

```
bx     (g4)
```

```
# Total bytes of code in the above routine = 20
```

```
.align 4
```

```
xxmtest$declare:
```

```
stob   g14,ADA.GLOBAL  
ret
```

```
# Total bytes of code in the above routine = 12
```

```
.align 4
```

```
xxmtest$body:
```

```
not    0,g13    #                line 5
```

```
st     g13,ADA.OWN  
mov    1,g13  
stob   g13,ADA.GLOBAL
```

```
ret
```

```
# Total bytes of code in the above routine = 28
```

```
.text
```

```
.align 2
```

```
.align 2
```

```
# Total bytes of code = 400
```

```
# Total bytes of data = 5
```

### 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 that contains 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;
```

## APPENDIX F OF THE Ada STANDARD

Consider a procedure that calls the inlined subprogram in the previous 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 calls the subprogram `I_Will_Be_Inlined`. However, because the body of the subprogram is not yet available, the generated code will not contain an inlined version of the subprogram. The generated code will use an out of line call to `I_Will_Be_Inlined`. The compiler will issue warning message number 2429 to inform you that the call was not inlined when `Uses_Inlined_Subp` was compiled.

If `In_Pack` is used across libraries, it can be exported from the root library after you have 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 number 2429 to indicate that 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_Subp` 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 marked as exported in the root 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 the flow of the code that cause movement of code sequences, as happens in a `pragma INLINE`. When compiling at a low optimization level, you will not be warned that inlining is not being performed.