



C User's Guide

Forte Developer 6 update 2
(Sun WorkShop 6 update 2)

Sun Microsystems, Inc.
901 San Antonio Road
Palo Alto, CA 94303
U.S.A. 650-960-1300

Part No. 806-7984-10
July 2001, Revision A

Send comments about this document to: docfeedback@sun.com

Copyright 2001 Sun Microsystems, Inc., 901 San Antonio Road • Palo Alto, CA 94303-4900 USA. All rights reserved.

This product or document is distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd. For Netscape™, Netscape Navigator™, and the Netscape Communications Corporation logo™, the following notice applies: Copyright 1995 Netscape Communications Corporation. All rights reserved.

Sun, Sun Microsystems, the Sun logo, docs.sun.com, AnswerBook2, Solaris, SunOS, JavaScript, SunExpress, Sun WorkShop, Sun WorkShop Professional, Sun Performance Library, Sun Performance WorkShop, Sun Visual WorkShop, and Forte are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

Sun f90/f95 is derived from Cray CF90™, a product of Cray Inc.

Federal Acquisitions: Commercial Software—Government Users Subject to Standard License Terms and Conditions.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 2001 Sun Microsystems, Inc., 901 San Antonio Road • Palo Alto, CA 94303-4900 Etats-Unis. Tous droits réservés.

Ce produit ou document est distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd. La notice suivante est applicable à Netscape™, Netscape Navigator™, et the Netscape Communications Corporation logo™: Copyright 1995 Netscape Communications Corporation. Tous droits réservés.

Sun, Sun Microsystems, the Sun logo, docs.sun.com, AnswerBook2, Solaris, SunOS, JavaScript, SunExpress, Sun WorkShop, Sun WorkShop Professional, Sun Performance Library, Sun Performance WorkShop, Sun Visual WorkShop, et Forte sont des marques de fabrique ou des marques déposées, ou marques de service, de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

Sun f90/f95 est dérivé de CRAY CF90™, un produit de Cray Inc.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REpondre A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.



Important Note on New Product Names

As part of Sun's new developer product strategy, we have changed the names of our development tools from Sun WorkShop™ to Forte™ Developer products. The products, as you can see, are the same high-quality products you have come to expect from Sun; the only thing that has changed is the name.

We believe that the Forte™ name blends the traditional quality and focus of Sun's core programming tools with the multi-platform, business application deployment focus of the Forte tools, such as Forte Fusion™ and Forte™ for Java™. The new Forte organization delivers a complete array of tools for end-to-end application development and deployment.

For users of the Sun WorkShop tools, the following is a simple mapping of the old product names in WorkShop 5.0 to the new names in Forte Developer 6.

Old Product Name	New Product Name
Sun Visual WorkShop™ C++	Forte™ C++ Enterprise Edition 6
Sun Visual WorkShop™ C++ Personal Edition	Forte™ C++ Personal Edition 6
Sun Performance WorkShop™ Fortran	Forte™ for High Performance Computing 6
Sun Performance WorkShop™ Fortran Personal Edition	Forte™ Fortran Desktop Edition 6
Sun WorkShop Professional™ C	Forte™ C 6
Sun WorkShop™ University Edition	Forte™ Developer University Edition 6

In addition to the name changes, there have been major changes to two of the products.

- Forte for High Performance Computing contains all the tools formerly found in Sun Performance WorkShop Fortran and now includes the C++ compiler, so High Performance Computing users need to purchase only one product for all their development needs.
- Forte Fortran Desktop Edition is identical to the former Sun Performance WorkShop Personal Edition, except that the Fortran compilers in that product no longer support the creation of automatically parallelized or explicit, directive-based parallel code. This capability is still supported in the Fortran compilers in Forte for High Performance Computing.

We appreciate your continued use of our development products and hope that we can continue to fulfill your needs into the future.

Contents

Before You Begin	1
How This Book Is Organized	1
Typographic Conventions	3
Shell Prompts	3
Supported Platforms	3
Accessing Sun WorkShop Development Tools and Man Pages	4
Accessing Sun WorkShop Documentation	6
Accessing Related Documentation	6
Ordering Sun Documentation	7
Sending Your Comments	7
1. Introduction to the C Compiler	9
Standards Conformance	9
Organization of the Compiler	10
C-Related Programming Tools	11
2. cc Compiler Options	13
Option Syntax	13
Options Summary	14
The cc Options	20

-# 21
-### 21
-Aname[(*tokens*)] 21
-B[static|dynamic] 21
-C 22
-c 22
-Dname[=*tokens*] 22
-d[y|n] 23
-dalign 23
-E 23
-erroff=*t* 23
-errtags=*a* 24
-errwarn=*t* 24
-fast 25
-fd 26
-flags 27
-fnonstd 27
-fns[={no,yes}] 27
-fprecision=*p* 28
-fround=*r* 28
-fsimple[=*n*] 28
-fsingle 29
-fstore 29
-ftrap=*t* 30
-G 30
-g 30
-H 31

-h *name* 31
-I[-|*dir*] 32
-i 32
-KPIC 32
-Kpic 32
-keeptmp 33
-L*dir* 33
-lname 33
-mc 33
-misalign 33
-misalign2 33
-mr[,*string*] 34
-mt 34
-native 34
-nofstore 34
-noqueue 34
-O 35
-o *filename* 35
-P 35
-p 35
-Q[y|n] 35
-qp 35
-R*dir*[:*dir*] 36
-S 36
-s 36
-U*name* 36
-V 37

-v 37
-Wc,*arg* 37
-w 38
-X[c|a|t|s] 38
-x386 39
-x486 39
-xa 39
-xalias_level[=*l*] 40
-xarch=*isa* 42
-xautopar 47
-xbuiltin[=(%all|%none)] 48
-xCC 48
-xc99=*o* 49
-xcache=*c* 49
-xcg[89|92] 50
-xchar_byte_order=*o* 51
-xchip=*c* 51
-xcode=*v* 52
-xcrossfile=[*n*] 54
-xcsi 54
-xdepend 55
-xe 55
-xexplicitpar 55
-xF 56
-xhelp=*f* 56
-xildoff 56
-xildon 57

-xinline=*list* 57
-xipo[=*a*] 57
-xlibmieee 59
-xlibmil 59
-xlic_lib=sunperf 59
-xlicinfo 60
-xloopinfo 60
-xM 60
-xM1 61
-xMerge 61
-xmaxopt[=*v*] 61
-xmemalign=*ab* 62
-xnolib 63
-xnolibmil 63
-xO[1|2|3|4|5] 64
-xopenmp[=*i*] 66
-xP 67
-xparallel 67
-xpentium 67
-xpg 68
-xprefetch=[*val*],*val* 68
-xprofile=*p* 68
-xreduction 70
-xregs=*r*[,*r*...] 70
-xrestrict[=*f*] 71
-xs 72
-xsafe=mem 72

-xsb 72
-xsbfast 73
-xsfpconst 73
-xspace 73
-xstrconst 73
-xtarget=*t* 73
-xtemp=*dir* 79
-xtime 79
-xtransition 79
-xtrigraphs 80
-xunroll=*n* 80
-xvector[={yes|no}] 81
-xvpara 81
-Yc, *dir* 81
-YA, *dir* 81
-YI, *dir* 82
-YP, *dir* 82
-YS, *dir* 82
-Z11 82

Options Passed to the Linker 82

3. Sun ANSI/ISO C Compiler-Specific Information 83

Environment Variables 83

TMPDIR 83
PARALLEL 84
SUNPROF_DATA 84
SUNPROF_DATA_DIR 84
SUNPRO_SB_INIT_FILE_NAME 84

SUNW_MP_THR_IDLE	85
Global Behavior: Value Versus unsigned Preserving	85
Keywords	86
__asm Keyword	86
_restrict Keyword	86
long long Data Type	86
Printing long long Data Types	87
Usual Arithmetic Conversions	87
Constants	88
Integral Constants	88
Character Constants	89
Include Files	89
Nonstandard Floating Point	92
Preprocessing Directives and Names	93
Assertions	93
Pragmas	94
Predefined Names	103
4. Parallelizing Sun ANSI/ISO C Code	105
Overview	105
Parallelizing for OpenMP	106
Handling OpenMP Runtime Warnings	106
Environment Variables	107
Data Dependence and Interference	109
Parallel Execution Model	111
Private Scalars and Private Arrays	112
Storeback	114
Reduction Variables	115

Speedups	116
Amdahl's Law	117
Load Balance and Loop Scheduling	121
Static or Chunk Scheduling	121
Self Scheduling	121
Guided Self Scheduling	122
Loop Transformations	122
Loop Distribution	122
Loop Fusion	123
Loop Interchange	125
Aliasing and Parallelization	126
Array and Pointer References	126
Restricted Pointers	127
Explicit Parallelization and Pragmas	129
5. Incremental Link Editor (<code>ild</code>)	137
Introduction	137
Overview of Incremental Linking	138
How to Use <code>ild</code>	138
How <code>ild</code> Works	140
What <code>ild</code> Cannot Do	141
Reasons for Full Relinks	142
<code>ild</code> Deferred-Link Messages	142
<code>ild</code> Relink Messages	142
Example 1: Internal Free Space Exhausted	143
Example 2: Running <code>strip</code>	143
Example 3: <code>ild</code> Version	144
Example 4: Too Many Files Changed	144

Example 5: Full Relink 144

Example 6: New Working Directory 145

`ild` Options 145

`-a` 145

`-Bdynamic | static` 146

`-d y|n` 146

`-e epsym` 146

`-g` 146

`-I name` 146

`-i` 147

`-Lpath` 147

`-lx` 147

`-m` 147

`-o outfile` 147

`-Q y|n` 148

`-Rpath` 148

`-s` 148

`-t` 148

`-u symname` 148

`-V` 148

`-xildoff` 149

`-xildon` 149

`-YP,dirlist` 149

`-z allextact|defaultextract| weakextract` 149

`-z defs` 149

`-z i_dryrun` 150

`-z i_full` 150

- z *i_noincr* 150
- z *i_quiet* 150
- z *i_verbose* 150
- z *nodefs* 150

Options Passed to `ild` From the Compilation System 150

- a 151
- e *epsym* 151
- I *name* 151
- m 151
- t 151
- u *symname* 152

Environment 152

ld Options not Supported by `ild` 154

- B *symbolic* 154
- b 154
- G 154
- h *name* 155
- z *muldefs* 155
- z *text* 155

Additional Unsupported Commands 155

- D *token,token,...* 155
- F *name* 155
- M *mapfile* 156
- r 156

Files That `ild` Uses 156

6. `lint` Source Code Checker 157

Basic and Enhanced `lint` Modes 157

Using lint	158
The lint Options	160
-#	160
-###	161
-a	161
-b	161
-C <i>filename</i>	161
-c	161
-dirout= <i>dir</i>	161
-err=warn	161
-errchk= <i>l, l</i>	162
-errfmt= <i>f</i>	164
-errhdr= <i>h</i>	164
-erroff= <i>tag, tag</i>	165
-errtags= <i>a</i>	166
-errwarn= <i>t</i>	166
-F	167
-fd	167
-flagsrc= <i>file</i>	167
-h	167
-I <i>dir</i>	167
-k	167
-L <i>dir</i>	167
-lx	168
-m	168
-Ncheck= <i>c</i>	168
-Nlevel= <i>n</i>	169

-n	170
-ox	170
-p	170
-R <i>file</i>	170
-s	170
-u	171
-V	171
-v	171
-W <i>file</i>	171
-x	171
-XCC= <i>a</i>	171
-Xalias_level[= <i>l</i>]	171
-Xarch=v9	172
-Xexplicitpar= <i>a</i>	172
-Xkeeptmp= <i>a</i>	172
-Xtemp= <i>dir</i>	172
-Xtime= <i>a</i>	172
-Xtransition= <i>a</i>	173
-y	173
lint Messages	173
Options to Suppress Messages	173
lint Message Formats	174
lint Directives	177
Predefined Values	177
Directives	177
lint Reference and Examples	181
Diagnostics Performed by lint	181

lint Libraries 186

lint Filters 187

7. Type-Based Alias Analysis 189

Introduction to Type-Based Analysis 189

Using Pragmas for Finer Control 190

Checking With `lint` 193

Struct Pointer Cast of Scalar Pointer 194

Struct Pointer Cast of Void Pointer 194

Cast of Struct Field to Structure Pointer 195

Explicit Aliasing Required 195

Examples of Memory Reference Constraints 196

8. Transitioning to ANSI/ISO C 207

Basic Modes 207

-xa 207

-xc 208

-xs 208

-xt 208

A Mixture of Old- and New-Style Functions 208

Writing New Code 208

Updating Existing Code 209

Mixing Considerations 209

Functions With Varying Arguments 212

Promotions: Unsigned Versus Value Preserving 215

Background 215

Compilation Behavior 215

First Example: The Use of a Cast 216

Bit-fields 217

Second Example: Same Result	217
Integral Constants	218
Third Example: Integral Constants	218
Tokenization and Preprocessing	219
ANSI/ISO C Translation Phases	219
Old C Translation Phases	221
Logical Source Lines	221
Macro Replacement	221
Using Strings	222
Token Pasting	223
const and volatile	224
Types, Only for lvalue	224
Type Qualifiers in Derived Types	224
const Means readonly	226
Examples of const Usage	226
volatile Means Exact Semantics	227
Examples of volatile Usage	227
Multibyte Characters and Wide Characters	228
Asian Languages Require Multibyte Characters	228
Encoding Variations	228
Wide Characters	229
Conversion Functions	229
C Language Features	230
Standard Headers and Reserved Names	231
Balancing Process	231
Standard Headers	232
Names Reserved for Implementation Use	232

Names Reserved for Expansion	233
Names Safe to Use	234
Internationalization	234
Locales	234
The <code>setlocale()</code> Function	235
Changed Functions	236
New Functions	237
Grouping and Evaluation in Expressions	238
Definitions	238
The K&R C Rearrangement License	239
The ANSI/ISO C Rules	239
The Parentheses	240
The As If Rule	240
Incomplete Types	241
Types	241
Completing Incomplete Types	241
Declarations	242
Expressions	242
Justification	243
Examples	243
Compatible and Composite Types	244
Multiple Declarations	244
Separate Compilation Compatibility	244
Single Compilation Compatibility	245
Compatible Pointer Types	245
Compatible Array Types	245
Compatible Function Types	245

Special Cases	246
Composite Types	246
9. Converting Applications	247
Overview of the Data Model Differences	248
Implementing Single Source Code	249
Derived Types	249
Tools	252
Converting to the LP64 Data Type Model	253
Integer and Pointer Size Change	253
Integer and Long Size Change	254
Sign Extension	254
Pointer Arithmetic Instead of Address Arithmetic	256
Structures	256
Unions	257
Type Constants	258
Beware of Implicit Declarations	258
sizeof() is an Unsigned Long	259
Use Casts to Show Your Intentions	259
Check Format String Conversion Operation	260
Other Considerations	261
Derived Types That Have Grown in Size	261
Check for Side Effects of Changes	261
Check Whether Literal Uses of long Still Make Sense	261
Use #ifdef for Explicit 32-bit Versus 64-bit Prototypes	262
Calling Convention Changes	262
Algorithm Changes	262
Checklist for Getting Started	263

10.	cscope: Interactively Examining a C Program	265
	The cscope Process	265
	Basic Use	266
	Step 1: Set Up the Environment	266
	Step 2: Invoke the cscope Program	267
	Step 3: Locate the Code	268
	Step 4: Edit the Code	274
	Command-Line Options	275
	View Paths	278
	cscope and Editor Call Stacks	279
	Examples	279
	Command-Line Syntax for Editors	283
	Unknown Terminal Type Error	284
A.	ANSI/ISO C Data Representations	285
	Storage Allocation	285
	Data Representations	286
	Integer Representations	287
	Floating-Point Representations	288
	Exceptional Values	290
	Hexadecimal Representation of Selected Numbers	291
	Pointer Representation	291
	Array Storage	292
	Arithmetic Operations on Exceptional Values	292
	Argument-Passing Mechanism	294
B.	Implementation-Defined ISO/IEC C Behavior	297
	Implementation Compared to the ANSI/ISO Standard	298
	Translation (G.3.1)	298

Environment (G.3.2)	298
Identifiers (G.3.3)	299
Characters(G.3.4)	299
Integers(G.3.5)	301
Floating-Point(G.3.6)	303
Arrays and Pointers(G.3.7)	304
Registers(G.3.8)	305
Structures, Unions, Enumerations, and Bit-Fields(G.3.9)	305
Qualifiers(G.3.10)	307
Declarators(G.3.11)	307
Statements(G.3.12)	307
Preprocessing Directives(G.3.13)	308
Library Functions(G.3.14)	309
Locale-Specific Behavior(G.4)	316
C. Supported Features of C99	319
D. Performance Tuning (SPARC)	321
Limits	321
<code>libfast.a</code> Library	322
E. The Differences Between K&R Sun C and Sun ANSI/ISO C	323
K&R Sun C Incompatibilities with Sun ANSI/ISO C	324
The Difference Between Sun C and ANSI/ISO C as Set by <code>-xs</code>	332
Keywords	334
F. Implementation-Specific Information of OpenMP	335
Index	337

List of Tables

TABLE 1-1	Components of the C Compilation System	10
TABLE 2-1	Compiler Options Grouped by Functionality	14
TABLE 2-2	-erroff Arguments	23
TABLE 2-3	-errwarn Values	24
TABLE 2-4	The Levels of Alias-Disambiguation	41
TABLE 2-5	-xarch ISA Keywords	42
TABLE 2-6	-xarch Matrix	43
TABLE 2-7	-xarch Values for SPARC Platforms	45
TABLE 2-8	-xarch Values on x86	47
TABLE 2-9	The -xcache Values	50
TABLE 2-10	The -xchip Values	51
TABLE 2-11	-xmemalign Alignment and Behavior Values	62
TABLE 2-12	Examples of -xmemalign	63
TABLE 2-13	The -xregs Values	71
TABLE 2-14	The -xtarget Values	74
TABLE 2-15	-xtarget Expansions on SPARC	74
TABLE 2-16	-xtarget Expansions on Intel Architecture	78
TABLE 3-1	Data Type Suffixes	88
TABLE 3-2	Predefined Identifier	103

TABLE 6-1	The <code>-errfmt</code> Values	164
TABLE 6-2	The <code>-errhdr</code> Values	164
TABLE 6-3	The <code>-erroff</code> Values	165
TABLE 6-4	<code>-errwarn</code> Values	166
TABLE 6-5	The <code>-Ncheck</code> Values	168
TABLE 6-6	<code>lint</code> Options to Suppress Messages	174
TABLE 6-7	<code>lint</code> Directives	178
TABLE 8-1	Trigraph Sequences	220
TABLE 8-2	Multibyte Character Conversion Functions	229
TABLE 8-3	Standard Headers	232
TABLE 8-4	Names Reserved for Expansion	233
TABLE 9-1	Data Type Size for ILP32 and LP64	248
TABLE 10-1	<code>cscope</code> Menu Manipulation Commands	268
TABLE 10-2	Commands for Use After an Initial Search	270
TABLE 10-3	Commands for Selecting Lines to Be Changed	280
TABLE A-1	Storage Allocation for Data Types	285
TABLE A-2	Representation of <code>short</code>	287
TABLE A-3	Representation of <code>int</code>	287
TABLE A-4	Representation of <code>long</code> on Intel and SPARC v8 versus SPARC v9	287
TABLE A-5	Representation of <code>long long</code>	288
TABLE A-6	<code>float</code> Representation	289
TABLE A-7	<code>double</code> Representation	289
TABLE A-8	<code>long double</code> Representation (<i>SPARC</i>)	289
TABLE A-9	<code>long double</code> Representation (<i>Intel</i>)	289
TABLE A-10	<code>float</code> Representations	290
TABLE A-11	<code>double</code> Representations	290
TABLE A-12	<code>long double</code> Representations	290
TABLE A-13	Hexadecimal Representation of Selected Numbers (<i>SPARC</i>)	291
TABLE A-14	Hexadecimal Representation of Selected Numbers (<i>Intel</i>)	291

TABLE A-15	Array Types and Storage	292
TABLE A-16	Abbreviation Usage	293
TABLE A-17	Addition and Subtraction Results	293
TABLE A-18	Multiplication Results	293
TABLE A-19	Division Results	294
TABLE A-20	Comparison Results	294
TABLE B-1	Representations and Sets of Values of Integers	301
TABLE B-2	Values for a float	303
TABLE B-3	Values for a double	303
TABLE B-4	Values for long double	303
TABLE B-5	Padding and Alignment of Structure Members	306
TABLE B-6	Character Sets Tested by <code>isalpha</code> , <code>islower</code> , Etc.	310
TABLE B-7	Values Returned on Domain Errors	311
TABLE B-8	Semantics for <code>signal</code> Signals	312
TABLE B-9	Names of Months	317
TABLE B-10	Days and Abbreviated Days of the Week	317
TABLE E-1	K&R Sun C Incompatibilities With Sun ANSI/ISO C	324
TABLE E-2	<code>-xs</code> Behavior	332
TABLE E-3	ANSI/ISO C Standard Keywords	334
TABLE E-4	Sun C (K&R) Keywords	334

Before You Begin

This manual describes the Sun WorkShop™ 6 C programming language compiler along with ANSI C compiler-specific information. It is intended for application developers who have a working knowledge of C, and UNIX®.

This manual describes the `lint` program that you can use to examine your code, provides instructions for parallelizing your code, explains how to transition to ANSI/ISO compliant code, describes the incremental linker, and the interactive program `cscope`. In the back of this manual, there are several appendices with reference material such as ANSI C data representations, implementation defined behavior, the differences between Sun C (K & R) and Sun ANSI C, performance tuning, and converting applications to compile for the 64-bit environment.

How This Book Is Organized

Chapter 1, “Introduction to the C Compiler,” provides information about the C compiler, including standards conformance, organization of the compiler, and C-related programming tools.

Chapter 2, “cc Compiler Options,” describes the C compiler options. It includes sections on option syntax, the `cc` options, and options passed to the linker.

Chapter 3, “Sun ANSI/ISO C Compiler-Specific Information,” documents those areas specific to the Sun ANSI C compiler.

Chapter 4, “Parallelizing Sun ANSI/ISO C Code,” explains how the Sun ANSI/ISO C compiler can optimize code to run on SPARC shared-memory multiprocessor machines.

Chapter 5, “Incremental Link Editor (ild),” describes `ild`, `ild`-specific features, example messages, and `ild` options.

Chapter 6, “lint Source Code Checker,” describes the `lint` program, its modes, options, messages, directives, and other helpful information.

Chapter 7, “Type-Based Alias Analysis,” explains how to use the `-xalias_level` option and several new pragmas to enable the compiler to perform type-based alias analysis and optimizations.

Chapter 8, “Transitioning to ANSI/ISO C,” provides tips and strategies for writing ANSI C compliant code.

Chapter 9, “Converting Applications,” provides the information you need to write code for the 32 bit or the 64-bit compilation environment.

Chapter 10, “`cscope`: Interactively Examining a C Program,” is a tutorial for the `cscope` browser which is provided with this release.

Appendix A, “ANSI/ISO C Data Representations,” describes how ANSI C represents data in storage and the mechanisms for passing arguments to functions.

Appendix B, “Implementation-Defined ISO/IEC C Behavior,” describes the implementation-defined features of the Sun WorkShop C compiler.

Appendix C, “Supported Features of C99,” describes which features of the C99 standard this release of Sun WorkShop supports.

Appendix D, “Performance Tuning (SPARC),” describes performance tuning on SPARC platforms.

Appendix E, “The Differences Between K&R Sun C and Sun ANSI/ISO C,” describes the differences between the previous K&R Sun C and Sun ANSI C.

Appendix F, “Implementation-Specific Information of OpenMP,” provides implementation-specific information for the OpenMP runtime system.

Typographic Conventions

Typeface	Meaning	Examples
AaBbCc123	The names of commands, files, and directories; on-screen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. % You have mail.
AaBbCc123	What you type, when contrasted with on-screen computer output	% su Password:
<i>AaBbCc123</i>	Book titles, new words or terms, words to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . These are called <i>class</i> options. You <i>must</i> be superuser to do this.
<code>AaBbCc123</code>	Command-line placeholder text; replace with a real name or value	To delete a file, type <code>rm filename</code> .

Shell Prompts

Shell	Prompt
C shell	%
Bourne shell and Korn shell	\$
C shell, Bourne shell, and Korn shell superuser	#

Supported Platforms

This Sun WorkShop™ release supports versions 2.6, 7, and 8 of the Solaris™ SPARC™ Platform Edition and Solaris™ Intel Platform Edition operating environments.

Accessing Sun WorkShop Development Tools and Man Pages

The Sun WorkShop product components and man pages are not installed into the standard `/usr/bin/` and `/usr/share/man` directories. To access the Sun WorkShop compilers and tools, you must have the Sun WorkShop component directory in your `PATH` environment variable. To access the Sun WorkShop man pages, you must have the Sun WorkShop man page directory in your `MANPATH` environment variable.

For more information about the `PATH` variable, see the `cs(1)`, `sh(1)`, and `ksh(1)` man pages. For more information about the `MANPATH` variable, see the `man(1)` man page. For more information about setting your `PATH` and `MANPATH` variables to access this release, see the *Sun WorkShop 6 update 2 Installation Guide* or your system administrator.

Note – The information in this section assumes that your Sun WorkShop 6 update 2 products are installed in the `/opt` directory. If your product software is not installed in the `/opt` directory, ask your system administrator for the equivalent path on your system.

Accessing Sun WorkShop Compilers and Tools

Use the steps below to determine whether you need to change your `PATH` variable to access the Sun WorkShop compilers and tools.

To Determine If You Need to Set Your `PATH` Environment Variable

1. **Display the current value of the `PATH` variable by typing:**

```
% echo $PATH
```

2. **Review the output for a string of paths containing `/opt/SUNWspro/bin/`.**

If you find the path, your `PATH` variable is already set to access Sun WorkShop development tools. If you do not find the path, set your `PATH` environment variable by following the instructions in the next section.

To Set Your PATH Environment Variable to Enable Access to Sun WorkShop Compilers and Tools

1. If you are using the C shell, edit your home `.cshrc` file. If you are using the Bourne shell or Korn shell, edit your home `.profile` file.
2. Add the following to your `PATH` environment variable.

```
/opt/SUNWspro/bin
```

Accessing Sun WorkShop Man Pages

Use the following steps to determine whether you need to change your `MANPATH` variable to access the Sun WorkShop man pages.

To Determine If You Need to Set Your `MANPATH` Environment Variable

1. Request the `workshop` man page by typing:

```
% man workshop
```

2. Review the output, if any.

If the `workshop(1)` man page cannot be found or if the man page displayed is not for the current version of the software installed, follow the instructions in the next section for setting your `MANPATH` environment variable.

To Set Your `MANPATH` Environment Variable to Enable Access to Sun WorkShop Man Pages

1. If you are using the C shell, edit your home `.cshrc` file. If you are using the Bourne shell or Korn shell, edit your home `.profile` file.
2. Add the following to your `MANPATH` environment variable.

```
/opt/SUNWspro/man
```

Accessing Sun WorkShop Documentation

You can access Sun WorkShop product documentation at the following locations:

- **The product documentation is available from the documentation index installed with the product on your local system or network.**

Point your Netscape™ Communicator 4.0 or compatible Netscape version browser to the following file:

```
/opt/SUNWspr/docs/index.html
```

If your product software is not installed in the /opt directory, ask your system administrator for the equivalent path on your system.

- **Manuals are available from the docs.sun.comsm Web site.**

The docs.sun.com Web site (<http://docs.sun.com>) enables you to read, print, and buy Sun Microsystems manuals through the Internet. If you cannot find a manual, see the documentation index installed with the product on your local system or network.

Accessing Related Documentation

The following table describes related documentation that is available through the docs.sun.com Web site.

Document Collection	Document Title	Description
Numerical Computation Guide Collection	<i>Numerical Computation Guide</i>	Describes issues regarding the numerical accuracy of floating-point computations.

Document Collection	Document Title	Description
Solaris 8 Reference Manual Collection	See the titles of man page sections.	Provides information about the Solaris operating environment.
Solaris 8 Software Developer Collection	<i>Linker and Libraries Guide</i>	Describes the operations of the Solaris link-editor and runtime linker.
Solaris 8 Software Developer Collection	<i>Multithreaded Programming Guide</i>	Covers the POSIX and Solaris threads APIs, programming with synchronization objects, compiling multithreaded programs, and finding tools for multithreaded programs.

Ordering Sun Documentation

You can order product documentation directly from Sun through the `docs.sun.com` Web site or from Fatbrain.com, an Internet bookstore. You can find the Sun Documentation Center on Fatbrain.com at the following URL:

<http://www.fatbrain.com/documentation/sun>

Sending Your Comments

Sun is interested in improving its documentation and welcomes your comments and suggestions. Email your comments to Sun at this address:

`docfeedback@sun.com`

Introduction to the C Compiler

This chapter provides information about the C compiler, including operating environments, standards conformance, organization of the compiler, and C-related programming tools.

Standards Conformance

The compiler conforms to the following standards:

- The American National Standard for Programming Language - C, ANSI/ISO 9899-1990 standard.
- The ISO/IEC 9899:1990, Programming Languages - C standard. For information on implementation-specific behavior, see “Implementation-Defined ISO/IEC C Behavior” on page 297.
- The FIPS 160 standard.

This release also supports some of the features specified in the following standard:

- The ISO/IEC 9899:1999, Programming Language - C standard. For more information on supported features, see “Supported Features of C99” on page 319.

Because the compiler also supports traditional K&R C (Kernighan and Ritchie, or pre-ANSI C), it can ease your migration to ANSI/ISO C.

Organization of the Compiler

The C compilation system consists of a compiler, an assembler, and a link editor. The `cc` command invokes each of these components automatically unless you use command-line options to specify otherwise.

“`cc` Compiler Options” on page 13 discusses all the options available with `cc`.

The following figure shows the organization of the C compilation system.




FIGURE 1-1 Organization of the C Compilation System

The following table summarizes the components of the compilation system.

TABLE 1-1 Components of the C Compilation System

Component	Description	Notes on Use
<code>cpp</code>	Preprocessor	<code>-Xs</code>
<code>acomp</code>	Compiler (preprocessor built in for non-Xs modes)	
<code>iropt</code>	Code optimizer	(SPARC) <code>-o</code> , <code>-x0 [2-5]</code> , <code>-fast</code>
<code>cg386</code>	Intermediate language translator	(Intel) Always invoked
<code>inline</code>	Inline expansion of assembly language templates	<code>.il</code> file specified
<code>mwinline</code>	Automatic inline expansion of functions	(Intel) <code>-x04</code> , <code>-xinline</code>

TABLE 1-1 Components of the C Compilation System (*Continued*)

Component	Description	Notes on Use
<code>fbe</code>	Assembler	
<code>cg</code>	Code generator, inliner, assembler	(<i>SPARC</i>)
<code>codegen</code>	Code generator	(<i>Intel</i>)
<code>ld</code>	Linker	
<code>ild</code>	Incremental linker	(<i>SPARC</i>) <code>-g</code> , <code>-xildon</code>

The C compiler optimizer removes redundancies, optimally allocates registers, schedules instructions, and reorganizes code. Select from multiple levels of optimization to obtain the best balance between application speed and use of memory.

C-Related Programming Tools

There are a number of tools available to aid in developing, maintaining, and improving your C programs. The two most closely tied to C, `cscope` and `lint`, are described in this book. In addition, a man page exists for each of these tools. Refer to the preface of this book for a list of all the associated man pages.

Sun WorkShop also provides tools for source browsing, debugging and performance analysis. See “Accessing Related Documentation” on page 6 for more information.

cc Compiler Options

This chapter describes the C compiler options and is organized into the following sections:

- “Option Syntax” on page 13
- “Options Summary” on page 14
- “The cc Options” on page 20
- “Options Passed to the Linker” on page 82

If you are porting a K&R C program to ANSI/ISO C, make special note of the section on compatibility flags, “-x[c | a | t | s]” on page 38. Using them makes the transition to ANSI/ISO C easier. Also refer to the discussion on the transition in “Transitioning to ANSI/ISO C” on page 207.

Option Syntax

The syntax of the cc command is:

```
% cc [options] filenames [libraries] . . .
```

where:

- *options* represents one or more of the options described in “The cc Options” on page 20
- *filenames* represents one or more files used in building the executable program
cc accepts a list of C source files and object files contained in the list of files specified by *filenames*. The resulting executable code is placed in a .out, unless the -o option is used. In this case, the code is placed in the file named by the -o option.

Use `cc` to compile and link any combination of the following:

- C source files, with a `.c` suffix
- Inline template files, with a `.i1` suffix (only when specified with `.c` files)
- C preprocessed source files, with a `.i` suffix
- Object-code files, with `.o` suffixes
- Assembler source files, with `.s` suffixes

After linking, `cc` places the linked files, now in executable code, into a file named `a.out`, or into the file specified by the `-o` option.

- *libraries* represents any of a number of standard or user-provided libraries containing functions, macros, and definitions of constants.

See option `-YP, dir` to change the default directories used for finding libraries. *dir* is a colon-separated path list. The default library search order for `cc` is:

```
/opt/SUNWsprow/WS6/lib
```

```
/usr/ccs/lib
```

```
/usr/lib
```

`cc` uses `getopt` to parse command-line options. Options are treated as a single letter or a single letter followed by an argument. See `getopt(3c)`.

Options Summary

In this section, the compiler options are grouped by function to provide an easy reference. The details are in the sections of the following pages. The following table summarizes the `cc` compiler options by functionality. Some flags serve more than one purpose and appear more than once.

TABLE 2-1 Compiler Options Grouped by Functionality

Licensing	Option Flag
Instructs the compiler not to queue this compile request if a license is not available.	<code>-noqueue</code>
Returns information about the licensing system.	<code>-xlicinfo</code>
Optimization and Performance	Option Flag
Selects the optimum combination of compilation options for speed.	<code>-fast</code>
Prepares the object code to collect data for profiling	<code>-p</code>

TABLE 2-1 Compiler Options Grouped by Functionality (*Continued*)

Optimizes for the 80386 processor.	-x386
Optimizes for the 80486 processor.	-x486
Inserts code to count how many times each basic block is executed.	-xa
Enables the compiler to perform type-based alias analysis and optimizations.	-xalias_level
Improve the optimization of code that calls standard library functions.	-xbuiltin
Enables optimization and inlining across source files.	-xcrossfile
Analyzes loops for inter-iteration data dependencies and does loop restructuring.	-xdepend
Enables performance analysis of the executable using the Analyzer.	-xF
Tries to inline only those functions specified.	-xinline
Performs whole-program optimizations by invoking an interprocedural analysis pass.	-xipo
Inlines some library routines for faster execution.	-xlibmil
Links in the Sun-supplied performance libraries.	-xlic_lib= sunperf
This command limits the level of <code>pragma opt</code> to the level specified.	-xmaxopt
Optimizes the object code.	-x0
Optimizes for the Pentium™ processor.	-xpentium
Enable prefetch instructions.	-xprefetch
Collects data for a profile or uses a profile to optimize.	-xprofile
Treats pointer-valued function parameters as restricted pointers.	-xrestrict
Allows the compiler to assume no memory-based traps occur.	-xsafe
Does no optimizations or parallelization of loops that increase code size.	-xspace
Suggests to the optimizer to unroll loops <i>n</i> times.	-xunroll
Data Alignment	Option Flag
Produce an integer constant by placing the characters of a multi-character character-constant in the specified byte order.	-xchar_byte_order
Specify maximum assumed memory alignment and behavior of misaligned data accesses.	-xmemalign

TABLE 2-1 Compiler Options Grouped by Functionality (*Continued*)

Does not inline math library routines.	-xnolibmil
Numerics and Floating-Point	Option Flag
Causes nonstandard initialization of floating-point arithmetic hardware.	-fnonstd
Turns on the SPARC nonstandard floating-point mode.	-fns
Initializes the rounding-precision mode bits in the Floating-point Control Word	-fprecision
Sets the IEEE 754 rounding mode that is established at runtime during the program initialization.	-fround
Allows the optimizer to make simplifying assumptions concerning floating-point arithmetic.	-fsimple
Causes the compiler to evaluate <code>float</code> expressions as single precision rather than double precision.	-fsingle
Causes the compiler to convert the value of a floating-point expression or function to the type on the left-hand side of an assignment	-fstore
Sets the IEEE 754 trapping mode in effect at startup.	-ftrap
Does not convert the value of a floating-point expression or function to the type on the left-hand side of an assignment	-nofstore
Forces IEEE 754 style return values for math routines in exceptional cases.	-xlibmieee
Represents unsuffixed floating-point constants as single precision	-xsfpconst
Enable automatic generation of calls to the vector library functions.	-xvector
Parallelization	Option Flag
Macro option that expands to <code>-D_REENTRANT -lthread</code> .	-mt
Turns on automatic parallelization for multiple processors.	-xautopar
Analyzes loops for inter-iteration data dependencies and does loop restructuring.	-xdepend
Generates parallelized code based on specification of <code>#pragma MP</code> directives.	-xexplicitpar
Shows which loops are parallelized and which are not.	-xloopinfo
Supports the OpenMP interface for explicit parallelization including a set of source code directives, run-time library routines, and environment variables	-xopenmp

TABLE 2-1 Compiler Options Grouped by Functionality (*Continued*)

Parallelizes loops both automatically by the compiler and explicitly specified by the programmer.	-xparallel
Turns on reduction recognition during automatic parallelization.	-xreduction
Treats pointer-valued function parameters as restricted pointers.	-xrestrict
Warns about loops that have <code>#pragma MP</code> directives specified but may not be properly specified for parallelization.	-xvpara
Creates the program database for <code>lock_lint</code> , but does not actually compile.	-Zll
Source Code	Option Flag
Associates <i>name</i> as a predicate with the specified <i>tokens</i> as if by a <code>#assert</code> preprocessing directive.	-A
Prevents the preprocessor from removing comments, except those on the preprocessing directive lines.	-C
Associates <i>name</i> with the specified <i>tokens</i> as if by a <code>#define</code> preprocessing directive.	-D
Runs the source file through the preprocessor only and sends the output to <code>stdout</code> .	-E
Reports K&R-style function definitions and declarations.	-fd
Prints to standard error, one per line, the path name of each file included during the current compilation.	-H
Adds directories to the list that is searched for <code>#include</code> files with relative file names.	-I
Runs the source file through the C preprocessor only.	-P
Removes any initial definition of the preprocessor symbol <i>name</i> .	-U
Accepts the C++-style comments.	-xCC
Controls compiler recognition of supported C99 features.	-xc99
Allows the C compiler to accept source code written in locales that do not conform to the ISO C source character code requirements	-xcsi
Runs only the preprocessor on the named C programs, requesting that it generate makefile dependencies and send the result to the standard output	-xM
Collects dependencies like -xM, but excludes <code>/usr/include</code> files.	-xMl
Prints prototypes for all K&R C functions defined in this module	-xP

TABLE 2-1 Compiler Options Grouped by Functionality (*Continued*)

Prepares the object code to collect data for profiling with <i>gprof</i> (1).	-xpg
Generates extra symbol table information for the Source Browser.	-xsb
Creates the database for the Source Browser.	-xsbfast
Determines recognition of trigraph sequences.	-xtrigraphs
Compiled Code	Option Flag
Directs the compiler to suppress linking with <i>ld</i> (1) and to produce a .o file for each source file	-c
Names the output file	-o
Directs the compiler to produce an assembly source file but not to assemble the program.	-S
Compilation Mode	Option Flag
Turns on verbose mode, showing each component as it is invoked.	-#
Shows each component as it would be invoked, but does not actually execute it.	-###
Retains temporary files created during compilation instead of deleting them automatically.	-keeptmp
Directs <i>cc</i> to print the name and version ID of each component as the compiler executes.	-V
Passes arguments to C compilation-system components.	-W
The -x options specify varying degrees of compliance to the ANSI/ISO C standard.	-X
Displays on-line help information.	-xhelp
Sets the directory for temporary files used by <i>cc</i> to <i>dir</i> .	-xtemp
Reports the time and resources used by each compilation component.	-xtime
Specifies a new directory for the location of a C compilation-system component.	-Y
Changes the default directory searched for components.	-YA
Changes the default directory searched for include files.	-YI
Changes the default directory for finding library files.	-YP
Changes the default directory for startup object files.	-YS

TABLE 2-1 Compiler Options Grouped by Functionality (*Continued*)

Diagnostics	Option Flag
Suppresses compiler warning messages.	-eroff
Displays the message tag for each warning message.	-errtags
If the indicated warning message is issued, cc exits with a failure status.	-errwarn
Directs the compiler to perform stricter semantic checks and to enable other lint-like checks.	-v
Suppresses compiler warning messages.	-w
Performs only syntax and semantic checking on the source file, but does not produce any object or executable code.	-xe
Issues warnings for the differences between K&R C and Sun ANSI/ISO C.	-xtransition
Warns about loops that have #pragma MP directives specified but may not be properly specified for parallelization.	-xvpara
Debugging	Option Flag
Produces additional symbol table information for the debugger.	-g
Removes all symbolic debugging information from the output object file.	-s
Disables Auto-Read for dbx.	-xs
Linking and Libraries	Option Flag
Specifies whether bindings of libraries for linking are static or dynamic.	-B
Specifies dynamic or static linking in the link editor.	-d
Passes the option to the link editor to produce a shared object rather than a dynamically linked executable.	-G
Assigns a name to a shared dynamic library as a way to have different versions of a library.	-h
Passes the option to the linker to ignore any LD_LIBRARY_PATH setting.	-i
Adds directories to the list that the linker searches for libraries.	-L
Links with object library <code>libname.so</code> , or <code>libname.a</code> .	-l
Removes duplicate strings from the <code>.comment</code> section of the object file.	-mc
Removes all strings from the <code>.comment</code> section. Can also insert a <i>string</i> in that section of the object file.	-mr

TABLE 2-1 Compiler Options Grouped by Functionality (*Continued*)

Emits or does not emit identification information to the output file.	-Q
Passes a colon-separated list of directories used to specify library search directories to the runtime linker.	-R
Merges data segments into text segments.	-xMerge
Specify code address space.	-xcode
Inserts string literals into the read-only data section of the text segment instead of the default data segment.	-xstrconst
Turns off the incremental linker and forces the use of ld.	-xildoff
Turns on the incremental linker and forces the use of ild in incremental mode.	-xildon
Does not link any libraries by default	-xnolib
Does not inline math library routines.	-xnolibmil
Target Platform	Option Flag
Specify instruction set architecture.	-xarch
Defines the cache properties for use by the optimizer.	-xcache
Specifies values for -xarch, -xchip, and -xcache.	-xcg
Specifies the target processor for use by the optimizer.	-xchip
Specifies the usage of registers for the generated code.	-xregs
Specifies the target system for instruction set and optimization.	-xtarget

The cc Options

This section describes the `cc` options, arranged alphabetically. These descriptions are also available in the man page, `cc(1)`. Use the `cc -flags` option for a one-line summary of these descriptions.

Options noted as being unique to one or more platforms are accepted without error and ignored on all other platforms. For an explanation of the typographic notations used with the options and arguments, refer to “Typographic Conventions” on page 3.

-#

Turns on verbose mode, showing each component as it is invoked.

-###

Shows each component as it would be invoked, but does not actually execute it.

-*Aname* [(*tokens*)]

Associates *name* as a predicate with the specified *tokens* as if by a #assert preprocessing directive. Preassertions:

- system(unix)
- machine(sparc) (SPARC)
- machine(i386) (x86)
- cpu(sparc) (SPARC)
- cpu(i386) (x86)

These preassertions are not valid in -xc mode.

-B[static|dynamic]

Specifies whether bindings of libraries for linking are *static* or *dynamic*, indicating whether libraries are non-shared or shared, respectively.

-Bdynamic causes the link editor to look for files named *libx.so* and then for files named *libx.a* when given the -lx option.

-Bstatic causes the link editor to look only for files named *libx.a*. This option may be specified multiple times on the command line as a toggle. This option and its argument are passed to ld(1).

Note – Many system libraries, such as *libc*, are only available as dynamic libraries in the Solaris 64-bit compilation environment. Therefore, do not use -Bstatic as the last toggle on the command line.

-C

Prevents the C preprocessor from removing comments, except those on the preprocessing directive lines.

-C

Directs `cc` to suppress linking with `ld(1)` and to produce a `.o` file for each source file. You can explicitly name a single object file using the `-o` option. When the compiler produces object code for each `.i` or `.c` input file, it always creates an object (`.o`) file in the current working directory. If you suppress the linking step, you also suppress the removal of the object files.

-Dname [=tokens]

Associates *name* with the specified *tokens* as if by a `#define` preprocessing directive. If no *=tokens* is specified, the token `1` is supplied.

Predefinitions (not valid in `-xc` mode):

- `sun`
- `unix`
- `sparc (SPARC)`
- `i386 (x86)`

The following predefinitions are valid in all modes.

```
__sparcv9 (-Xarch=v9, v9a)
__sun
__unix
__SUNPRO_C=0x530
__`uname -s`_`uname -r` (example: __SunOS_5_7)
__sparc (SPARC)
__i386 (x86)
__BUILTIN_VA_ARG_INCR
__SVR4
```

The following is predefined in `-xa` and `-xt` modes only:

```
__RESTRICT
```

The compiler also predefines the object-like macro

```
__PRAGMA_REDEFINE_EXTNAME, to indicate the pragma will be recognized.
```


-d[y|n]

-dy specifies dynamic linking, which is the default, in the link editor.

-dn specifies static linking in the link editor.

This option and its arguments are passed to ld(1).

Note – Many system libraries are only available as dynamic libraries in the Solaris 7 64-bit compilation environment. As a result, this option causes fatal errors if you use it in combination with -xarch=v9.

-dalign

-dalign is equivalent to -xmemalign=8s. See “-xmemalign=ab” on page 62.

-E

Runs the source file through the preprocessor only and sends the output to stdout. The preprocessor is built directly into the compiler, except in -Xs mode, where /usr/ccs/lib/cpp is invoked. Includes the preprocessor line numbering information. See also the -P option.

-erroff=t

Suppresses cc warning messages. Has no effect on error messages.

t is a comma-separated list that consists of one or more of the following: *tag*, *no%tag*, %all, %none. Order is important; for example, %all, *no%tag* suppresses all warning messages except *tag*. The following table lists the -erroff values:

TABLE 2-2 -erroff Arguments

Value	Meaning
<i>tag</i>	Suppresses the warning message specified by this <i>tag</i> . You can display the tag for a message by using the -errtags=yes option.

TABLE 2-2 -erroff Arguments (Continued)

Value	Meaning
no%tag	Enables the warning message specified by this <i>tag</i>
%all	Suppresses all warning messages
%none%none	Enables all warning messages (default)

The default is `-erroff=%none`. Specifying `-erroff` is equivalent to specifying `-erroff=%all`.

You can achieve finer control over error message suppression. See “`#pragma error_messages (on|off|default, tag... tag)`” on page 96.

`-errtags=a`

Displays the message tag for each warning message.

a can be either `yes` or `no`. The default is `-errtags=no`. Specifying `-errtags` is equivalent to specifying `-errtags=yes`.

`-errwarn=t`

If the indicated warning message is issued, `cc` exits with a failure status. *t* is a comma-separated list that consists of one or more of the following: *tag*, `no%tag`, `%all`, `%none`. Order is important; for example `%all, no%tag` causes `cc` to exit with a fatal status if any warning except *tag* is issued. The following table lists the `-errwarn` values:

TABLE 2-3 -errwarn Values

<i>tag</i>	Cause <code>cc</code> to exit with a fatal status if the message specified by this <i>tag</i> is issued as a warning message. Has no effect if <i>tag</i> is not issued.
no%tag	Prevent <code>cc</code> from exiting with a fatal status if the message specified by <i>tag</i> is issued only as a warning message. Has no effect if the message specified by <i>tag</i> is not issued. Use this option to revert a warning message that was previously specified by this option with <i>tag</i> or <code>%all</code> from causing <code>cc</code> to exit with a fatal status when issued as a warning message.

TABLE 2-3 `-errwarn` Values (*Continued*)

<code>%all</code>	Cause <code>cc</code> to exit with a fatal status if any warning messages are issued. <code>%all</code> can be followed by <code>no%tag</code> to exempt specific warning messages from this behavior.
<code>%none</code>	Prevents any warning message from causing <code>cc</code> to exit with a fatal status should any warning message be issued.

The default is `-errwarn=%none`. If you specify `-errwarn` alone, it is equivalent to `-errwarn=%all`.

`-fast`

Selects the optimum combination of compilation options for speed. This should provide close to the maximum performance for most realistic applications. Modules compiled with `-fast` must also be linked with `-fast`.

The `-fast` option is unsuitable for programs intended to run on a different target than the compilation machine. In such cases, follow `-fast` with the appropriate `-xtarget` option. For example:

```
cc -fast -xtarget=ultra ...
```

For C modules that depend on exception handling specified by SVID, follow `-fast` by `-xnolibmil`:

```
% cc -fast -xnolibmil
```

With `-xlibmil`, exceptions cannot be noted by setting `errno` or calling `matherr(3m)`.

The `-fast` option is unsuitable for programs that require strict conformance to the IEEE 754 Standard.

The following table lists the set of options selected by `-fast` across platforms.

Option	SPARC	x86
<code>-dalign</code>	X	-
<code>-fns</code>	X	X
<code>-fsimple=2</code>	X	-

Option	SPARC	x86
-fsingle	X	X
-ftrap=%none	X	X
-nofstore	-	X
-xarch	X	X
-xbuiltin=%all	X	X
-xlibmil	X	X
-xtarget=native	X	X
-x05	X	X

`-fast` acts like a macro expansion on the command line. Therefore, you can override the optimization level and code generation option aspects by following `-fast` with the desired optimization level or code generation option. Compiling with the `-fast -x04` pair is like compiling with the `-x02 -x04` pair. The latter specification takes precedence.

In previous releases, the `-fast` macro option included `-fnonstd`; now it includes `-fns` instead.

`-fast` also defines the macro `__MATHERR_ERRNO_DONTCARE`. This macro causes `math.h` in the Solaris 8 operating environment to assert performance-related pragmas such as the following for some math routines prototyped in `<math.h>`:

- `#pragma does_not_read_global_data`
- `#pragma does_not_write_global_data`
- `#pragma no_side_effect`

If your code relies on the return value of `errno` in exceptional cases as documented in the `matherr(3M)` man page, you must turn off the macro by issuing the `-U__MATHERR_ERRNO_DONTCARE` macro after the `-fast` option.

You can usually improve performance for most programs with this option.

Do not use this option for programs that depend on IEEE standard exception handling; you can get different numerical results, premature program termination, or unexpected SIGFPE signals.

`-fd`

Reports K&R-style function definitions and declarations.

-flags

Prints a brief summary of each available compiler option.

-fnonstd

Causes nonstandard initialization of floating-point arithmetic hardware. In addition, the `-fnonstd` option causes hardware traps to be enabled for floating-point overflow, division by zero, and invalid operations exceptions. These are converted into SIGFPE signals; if the program has no SIGFPE handler, it terminates with a memory dump.

By default, IEEE 754 floating-point arithmetic is nonstop, and underflows are gradual. (See “Nonstandard Floating Point” on page 92 for a further explanation.)

(SPARC) Synonym for `-fns -ftrap=common`.

-fns [= {no, yes}]

(SPARC) Turns on the SPARC nonstandard floating-point mode.

The default is `-fns=no`, the SPARC standard floating-point mode. `-fns` is the same as `-fns=yes`.

Optional use of `=yes` or `=no` provides a way of toggling the `-fns` flag following some other macro flag that includes `-fns`, such as `-fast`. This flag enables the nonstandard floating point mode when a program begins execution. By default, the non-standard floating point mode will not be enabled automatically.

On some SPARC systems, the nonstandard floating point mode disables “gradual underflow,” causing tiny results to be flushed to zero rather than producing subnormal numbers. It also causes subnormal operands to be replaced silently by zero. On those SPARC systems that do not support gradual underflow and subnormal numbers in hardware, use of this option can significantly improve the performance of some programs.

When nonstandard mode is enabled, floating point arithmetic may produce results that do not conform to the requirements of the IEEE 754 standard. See the Numerical Computation Guide for more information.

This option is effective only on SPARC systems and only if used when compiling the main program. On x86 systems, the option is ignored.

`-fprecision=p`

(x86) `-fprecision={single, double, extended}`

Initializes the rounding-precision mode bits in the Floating-point Control Word to single (24 bits), double (53 bits), or extended (64 bits), respectively. The default floating-point rounding-precision mode is extended.

Note that on Intel, only the precision, not exponent, range is affected by the setting of floating-point rounding precision mode.

`-fround=r`

Sets the IEEE 754 rounding mode that is established at runtime during the program initialization.

r must be one of: `nearest`, `tozero`, `negative`, `positive`.

The default is `-fround=nearest`.

The meanings are the same as those for the `ieee_flags` subroutine.

When *r* is `tozero`, `negative`, or `positive`, this flag sets the rounding direction mode to round-to-zero, round-to-negative-infinity, or round-to-positive-infinity respectively when a program begins execution. When *r* is `nearest` or the `-fround` flag is not used, the rounding direction mode is not altered from its initial value (round-to-nearest by default).

This option is effective only if used when compiling the main program.

`-fsimple[=n]`

Allows the optimizer to make simplifying assumptions concerning floating-point arithmetic.

If *n* is present, it must be 0, 1, or 2. The defaults are:

- With no `-fsimple[=n]`, the compiler uses `-fsimple=0`
- With only `-fsimple`, no `=n`, the compiler uses `-fsimple=1`

`-fsimple=0`

Permits no simplifying assumptions. Preserve strict IEEE 754 conformance.

`-fsimple=1`

Allows conservative simplifications. The resulting code does not strictly conform to IEEE 754, but numeric results of most programs are unchanged.

With `-fsimple=1`, the optimizer can assume the following:

- IEEE 754 default rounding/trapping modes do not change after process initialization.
- Computations producing no visible result other than potential floating point exceptions may be deleted.
- Computations with Infinity or NaNs as operands need not propagate NaNs to their results; for example, $x*0$ may be replaced by 0.
- Computations do not depend on sign of zero.

With `-fsimple=1`, the optimizer is *not* allowed to optimize completely without regard to roundoff or exceptions. In particular, a floating-point computation cannot be replaced by one that produces different results with rounding modes held constant at runtime. The `-fast` macroflag includes `-fsimple=1`.

`-fsimple=2`

Permits aggressive floating point optimizations that may cause many programs to produce different numeric results due to changes in rounding. For example, `-fsimple=2` permits the optimizer to replace all computations of x/y in a given loop with $x*z$, where x/y is guaranteed to be evaluated at least once in the loop, $z=1/y$, and the values of y and z are known to have constant values during execution of the loop.

Even with `-fsimple=2`, the optimizer is not permitted to introduce a floating point exception in a program that otherwise produces none.

`-fsingle`

(`-Xt` and `-Xs` modes only) Causes the compiler to evaluate `float` expressions as single precision rather than double precision. This option has no effect if the compiler is used in either `-Xa` or `-Xc` modes, as `float` expressions are already evaluated as single precision.

`-fstore`

(*x86*) Causes the compiler to convert the value of a floating-point expression or function to the type on the left-hand side of an assignment, when that expression or function is assigned to a variable, or when the expression is cast to a shorter

floating-point type, rather than leaving the value in a register. Due to rounding and truncation, the results may be different from those that are generated from the register value. This is the default mode.

To turn off this option, use the `-nofstore` option.

`-fttrap=t`

Sets the IEEE 754 trapping mode in effect at startup.

t is a comma-separated list that consists of one or more of the following: `%all`, `%none`, `common`, `[no%]invalid`, `[no%]overflow`, `[no%]underflow`, `[no%]division`, `[no%]inexact`.

The default is `-fttrap=%none`.

This option sets the IEEE 754 trapping modes that are established at program initialization. Processing is left-to-right. The `common` exceptions, by definition, are `invalid`, `division by zero`, and `overflow`.

Example: `-fttrap=%all,no%inexact` means set all traps, except `inexact`.

The meanings are the same as for the `ieee_flags` subroutine, except that:

- `%all` turns on all the trapping modes.
- `%none`, the default, turns off all trapping modes.
- A `no%` prefix turns off that specific trapping mode.

If you compile one routine with `-fttrap=t`, compile all routines of the program with the same `-fttrap=t` option; otherwise, you can get unexpected results.

`-G`

Passes the option to the link editor to produce a shared object rather than a dynamically linked executable. This option is passed to `ld(1)`, and cannot be used with the `-dn` option.

`-g`

Produces additional symbol table information for the debugger.

This option invokes the incremental linker; see “`-xildoff`” on page 56 and “`-xildon`” on page 57. Invoke `ild` instead of `ld` unless you are using the `-G` or `-xildoff` options, or you are naming source files on the command line.

If you issue `-g`, and the optimization level is `-xO3` or lower, the compiler provides best-effort symbolic information with almost full optimization. Tail-call optimization and back-end inlining are disabled.

If you issue `-g` and the optimization level is `-xO4`, the compiler provides best-effort symbolic information with full optimization. Compile with the `-g` option to use the full capabilities of the Sun WorkShop 6 Performance Analyzer. While some performance analysis features do not require `-g`, you must compile with `-g` to view annotated source, some function level information, and compiler commentary messages. See the `analyzer(1)` man page and "Compiling Your Program for Data Collection and Analysis" in *Analyzing Program Performance With Sun WorkShop* for more information.

The commentary messages that are generated with `-g` describe the optimizations and transformations that the compiler made while compiling your program. Use the `er_src(1)` command to display the messages, which are interleaved with the source code.

For more information on debugging, see the *Debugging a Program With dbx* manual. For more information on optimization, see the *Analyzing Program Performance With Sun Workshop* manual.

-H

Prints to standard error, one per line, the path name of each file included during the current compilation. The display is indented so as to show which files are included by other files.

Here, the program `sample.c` includes the files, `stdio.h` and `math.h`; `math.h` includes the file, `floatingpoint.h`, which itself includes functions that use `sys/ieeefp.h`:

```
% cc -H sample.c
  /usr/include/stdio.h
  /usr/include/math.h
    /usr/include/floatingpoint.h
      /usr/include/sys/ieeefp.h
```

-h *name*

Assigns a name to a shared dynamic library as a way to have different versions of a library. In general, the *name* after `-h` should be the same as the file name given after the `-o` option. The space between `-h` and *name* is optional.

The linker assigns the specified *name* to the library and records the name in the library file as the *intrinsic* name of the library. If there is no *-hname* option, then no intrinsic name is recorded in the library file.

When the runtime linker loads the library into an executable file, it copies the intrinsic name from the library file into the executable, into a list of needed shared library files. Every executable has such a list. If there is no intrinsic name of a shared library, then the linker copies the path of the shared library file instead.

`-I [- | dir]`

`-I dir` adds *dir* to the list of directories that are searched for `#include` files with relative file names, that is, those not beginning with a / (slash). `-I` values accumulate. See “Include Files” on page 89 for a discussion of the search order used to find the include files.

`-I-` gives you more control over the algorithm that the compiler uses when searching for include files. `-I-` values do not accumulate. This section first describes the default search algorithms, then it describes the effect of `-I-` on these algorithms.

For more information on the search pattern of the compiler, see “Include Files” on page 89.

`-i`

Passes the option to the linker to ignore any `LD_LIBRARY_PATH` or `LD_LIBRARY_PATH_64` setting.

`-KPIC`

(SPARC) The `-KPIC` command is equivalent to `-xcode=pic32`. See also “`-xcode=v`” on page 52.

(x86) `-KPIC` is identical to `-Kpic`.

`-Kpic`

(SPARC) The `-Kpic` command is equivalent to `-xcode=pic13`. See “`-xcode=v`” on page 52.

(x86) Generate position-independent code for use in shared libraries (small model). Permits references to, at most, 2**11 unique external symbols.

-keeptmp

Retains temporary files created during compilation instead of deleting them automatically.

-Ldir

Adds *dir* to the list of directories searched for libraries by ld(1). This option and its arguments are passed to ld(1).

-lname

Links with object library *libname.so*, or *libname.a*. The order of libraries in the command-line is important, as symbols are resolved from left to right.

This option must follow the *sourcefile* arguments.

-mC

Removes duplicate strings from the *.comment* section of the object file. When you use the *-mC* flag, *mcs -c* is invoked.

-misalign

(SPARC) *-misalign* is equivalent to *-xmemalign=1i*. See “*-xmemalign=ab*” on page 62.

-misalign2

(SPARC) *-misalign2* is equivalent to *-xmemalign=2i*. See “*-xmemalign=ab*” on page 62.

`-mr [, string]`

`-mr` removes all strings from the `.comment` section. When you use this flag, `mcs -d -a` is invoked.

`-mr ,string` removes all strings from the `.comment` section and inserts *string* in that section of the object file. If *string* contains embedded blanks, it must be enclosed in quotation marks. A null *string* results in an empty `.comment` section. This option is passed as `-d -astring` to `mcs`.

`-mt`

Macro option that expands to `-D_REENTRANT -lthread`. If you are doing your own multithread coding, you must use this option in the compile and link steps. To obtain faster execution, this option requires a multiprocessor system. On a single-processor system, the resulting executable usually runs more slowly with this option.

`-native`

This option is a synonym for `-xtarget=native`.

`-nofstore`

(*x86*) Does not convert the value of a floating-point expression or function to the type on the left-hand side of an assignment, when that expression or function is assigned to a variable or is cast to a shorter floating-point type; rather, it leaves the value in a register. See also “`-fstore`” on page 29.

`-noqueue`

Instructs the compiler not to queue this compile request if a license is not available. Under normal circumstances, if no license is available, the compiler waits until one becomes available. With this option, the compiler returns immediately.

-O

Same as -xO2.

-o *filename*

Names the output file *filename* (as opposed to the default, `a.out`). *filename* cannot be the same as *sourcefile*, since `cc` does not overwrite the source file. This option and its arguments are passed to `ld(1)`.

-P

Runs the source file through the C preprocessor only. It then puts the output in a file with a `.i` suffix. Unlike `-E`, this option does not include preprocessor-type line number information in the output. See also the `-E` option.

-p

Prepares the object code to collect data for profiling with `prof(1)`. This option invokes a runtime recording mechanism that produces a `mon.out` file at normal termination.

-Q[y|n]

Emits or does not emit identification information to the output file. `-Qy` is the default.

If `-Qy` is used, identification information about each invoked compilation tool is added to the `.comment` section of output files, which is accessible with `mcs`. This option can be useful for software administration.

`-Qn` suppresses this information.

-qp

Same as `-p`.

-Rdir [: dir]

Passes a colon-separated list of directories used to specify library search directories to the runtime linker. If present and not null, it is recorded in the output object file and passed to the runtime linker.

If both `LD_RUN_PATH` and the `-R` option are specified, the `-R` option takes precedence.

-S

Directs `cc` to produce an assembly source file but not to assemble the program.

-S

Removes all symbolic debugging information from the output object file. This option cannot be specified with `-g`.

Passed to `ld(1)`.

-Uname

Undefines the preprocessor symbol *name*. This option removes any initial definition of the preprocessor symbol *name* created by `-D` on the same command line including those placed there by the command-line driver.

`-U` has no effect on any preprocessor directives in source files. You can give multiple `-U` options on the command line.

If the same *name* is specified for both `-D` and `-U` on the command line, *name* is undefined, regardless of the order the options appear. In the following example, `-U` undefines `__sun`:

```
cc -U__sun test.c
```

Preprocessor statements of the following form in `test.c` will not take effect because `__sun` is undefined.

```
#ifdef(__sun)
```

See “-Dname[=tokens]” on page 22 for a list of predefined symbols.

-V

Directs `cc` to print the name and version ID of each component as the compiler executes.

-v

Directs the compiler to perform stricter semantic checks and to enable other lint-like checks. For example, the code:

```
#include <stdio.h>
main(void)
{
    printf("Hello World.\n");
}
```

compiles and executes without problem. With `-v`, it still compiles; however, the compiler displays this warning:

```
"hello.c", line 5: warning: function has no return statement:
main
```

`-v` does not give all the warnings that `lint(1)` does. Try running the above example through `lint`.

-Wc, arg

Passes the argument `arg` to a specified component `c`. Each argument must be separated from the preceding only by a comma. All `-W` arguments are passed after the regular command-line arguments. A comma can be part of an argument by escaping it by an immediately preceding `\` (backslash) character. All `-W` arguments are passed after the regular command-line arguments.

For example, `-Wa, -o, objfile` passes `-o` and `objfile` to the assembler, in that order. Also, `-wl, -I, name` causes the linking phase to override the default name of the dynamic linker, `/usr/lib/ld.so.1`.

The order in which the argument(s) are passed to a tool with respect to the other specified command line options may change.

See “Components of the C Compilation System” on page 10 for a list of components. *c* can be one of the following:

- a Assembler: (*fbe*); (*gas*)
- c C code generator: (*cg*) (*SPARC*);
- d cc driver¹
- h Intermediate code translator (*ir2hf*)(*Intel*)
- i Inter-procedure analysis (*ube_ipa*)(*Intel*)
- l Link editor (*ld*)
- m *mcs*
- o Postoptimizer (*SPARC*)
- p Preprocessor (*cpp*)
- u C code generator (*ube*) (*Intel*)
- 0 Compiler (*acomp*) (*ssbd*, *SPARC*)
- 2 Optimizer: (*iropt*) (*SPARC*)

1. You cannot use `-wd` to pass the cc options listed in this chapter to the C compiler.

`-W`

Suppresses compiler warning messages.

This option overrides the `error_messages` pragma.

`-X[c | a | t | s]`

The `-x` (note uppercase *X*) options specify varying degrees of compliance to the ANSI/ISO C standard. `-xa` is the default mode.

See “K&R Sun C Incompatibilities with Sun ANSI/ISO C” on page 324 for a discussion of differences between ANSI/ISO C and K&R C.

`-xc`

(*c* = conformance) Issues errors and warnings for programs that use non-ANSI/ISO C constructs. This option is strictly conformant ANSI/ISO C, without K&R C compatibility extensions. The predefined macro `__STDC__` has a value of 1 with the `-xc` option.

`-xa`

(a = ANSI) This is the default compiler mode. ANSI C plus K&R C compatibility extensions, with semantic changes required by ANSI C. Where K&R C and ANSI C specify different semantics for the same construct, the compiler uses the ANSI C interpretation. If the `-xa` option is used in conjunction with the `-xtransition` option, the compiler issues warnings about the different semantics. The predefined macro `_STDC_` has a value of 0 with the `-xa` option.

`-xt`

(t = transition) This option uses ANSI/ISO C plus K&R C compatibility extensions *without* semantic changes required by ANSI/ISO C. Where K&R C and ANSI/ISO C specify different semantics for the same construct, the compiler uses the K&R C interpretation. If you use the `-xt` option in conjunction with the `-xtransition` option, the compiler issues warnings about the different semantics. The predefined macro `_STDC_` has a value of 0 with the `-xt` option.

`-xs`

(s = K&R C) Attempts to warn about all language constructs that have differing behavior between ANSI/ISO C and K&R C. The compiler language includes all features compatible with K&R C. This option invokes `cpp` for preprocessing. `_STDC_` is not defined in this mode. For more information regarding the effects of compiling with the `-xs` option, see “The Difference Between Sun C and ANSI/ISO C As Set By `-xs`” on page 332.

`-x386`

(x86) Optimizes for the 80386 processor.

`-x486`

(x86) Optimizes for the 80486 processor.

`-xa`

Inserts code to count how many times each basic block is executed. This option is the old style of basic block profiling for `tcov`. See “`-xprofile=p`” on page 68 for information on the new style of profiling and the `tcov(1)` man page for more details. See also *Analyzing Program Performance With Sun Workshop*.

`-xa` invokes a runtime recording mechanism that creates a `.d` file for every `.c` file at normal termination. The `.d` file accumulates execution data for the corresponding source file. `tcov(1)` can then be run on the source file to generate statistics about the program. Since this option entails some optimization, it is incompatible with `-g`.

If set at compile-time, the `TCOVDIR` environment variable specifies the directory where the `.d` files are located. If this variable is not set, the `.d` files remain in the same directory as the `.c` files.

The `-xprofile=tcov` and the `-xa` options are compatible in a single executable. That is, you can link a program that contains some files which have been compiled with `-xprofile=tcov`, and others with `-xa`. You cannot compile a single file with both options.

`-xalias_level[=l]`

The compiler uses the `-xalias_level` option to determine what assumptions it can make in order to perform optimizations using type-based alias-analysis. This option places the indicated alias level into effect for the translation units being compiled. If you do not issue the `-xalias_level` command, the compiler assumes `-xalias_level=basic`. If you specify `-xalias_level` without a value, the default is `-xalias_level=layout`.

Remember that if you issue the `-xalias_level` option but you fail to adhere to all of the assumptions and restrictions about aliasing described for any of the alias levels, the behavior of your program is undefined.

Replace *l* with one of the terms in the following table.

TABLE 2-4 The Levels of Alias-Disambiguation

Term	Meaning
any	The compiler assumes that all memory references can alias at this level. There is no type-based alias analysis at the level of <code>-xalias_level=any</code> .
basic	<p>If you use the <code>-xalias_level=basic</code> option, the compiler assumes that memory references that involve different C basic types do not alias each other. The compiler also assumes that references to all other types can alias each other as well as any C basic type. The compiler assumes that references using <code>char *</code> can alias any other type.</p> <p>For example, at the <code>-xalias_level=basic</code> level, the compiler assumes that a pointer variable of type <code>int *</code> is not going to access a float object. Therefore it is safe for the compiler to perform optimizations that assume a pointer of type <code>float *</code> will not alias the same memory that is referenced with a pointer of type <code>int *</code>.</p>
weak	<p>If you use the <code>-xalias_level=weak</code> option, the compiler assumes that any structure pointer can point to any structure type.</p> <p>Any structure or union type that contains a reference to any type that is either referenced in an expression in the source being compiled or is referenced from outside the source being compiled, must be declared prior to the expression in the source being compiled.</p> <p>You can satisfy this restriction by including all the header files of a program that contain types that reference any of the types of the objects referenced in any expression of the source being compiled.</p> <p>At the level of <code>-xalias_level=weak</code>, the compiler assumes that memory references that involve different C basic types do not alias each other. The compiler assumes that references using <code>char *</code> alias memory references that involve any other type.</p>
layout	<p>If you use the <code>-xalias_level=layout</code> option, the compiler assumes that memory references that involve types with the same sequence of types in memory can alias each other.</p> <p>The compiler assumes that two references with types that do not look the same in memory do not alias each other. The compiler assumes that any two memory accesses through different struct types alias if the initial members of the structures look the same in memory. However, at this level, you should not use a pointer to a struct to access some field of a dissimilar struct object that is beyond any of the common initial sequence of members that look the same in memory between the two structs. This is because the compiler assumes that such references do not alias each other.</p> <p>At the level of <code>-xalias_level=layout</code> the compiler assumes that memory references that involve different C basic types do not alias each other. The compiler assumes that references using <code>char *</code> can alias memory references involving any other type.</p>

TABLE 2-4 The Levels of Alias-Disambiguation (*Continued*)

Term	Meaning
<code>strict</code>	<p>If you use the <code>-xalias_level=strict</code> option, the compiler assumes that memory references, that involve types such as structs or unions, that are the same when tags are removed, can alias each other. Conversely, the compiler assumes that memory references involving types that are not the same even after tags are removed do not alias each other.</p> <p>However, any structure or union type that contains a reference to any type that is part of any object referenced in an expression in the source being compiled, or is referenced from outside the source being compiled, must be declared prior to the expression in the source being compiled.</p> <p>You can satisfy this restriction by including all the header files of a program that contain types that reference any of the types of the objects referenced in any expression of the source being compiled. At the level of <code>-xalias_level=strict</code> the compiler assumes that memory references that involve different C basic types do not alias each other. The compiler assumes that references using <code>char *</code> can alias any other type.</p>
<code>std</code>	<p>If you use the <code>-xalias_level=std</code> option, the compiler assumes that types and tags need to be the same to alias, however, references using <code>char *</code> can alias any other type. This rule is the same as the restrictions on the dereferencing of pointers that are found in the 1999 ISO C standard. Programs that properly use this rule will be very portable and should see good performance gains under optimization.</p>
<code>strong</code>	<p>If you use the <code>-xalias_level=strong</code> option, the same restrictions apply as at the <code>std</code> level, but additionally, the compiler assumes that pointers of type <code>char *</code> are used only to access an object of type <code>char</code>. Also, the compiler assumes that there are no interior pointers. An interior pointer is defined as a pointer that points to a member of a struct.</p>

`-xarch=isa`

Specify instruction set architecture (ISA).

Architectures that are accepted by `-xarch` keyword *isa* are shown in TABLE 2-5:

TABLE 2-5 `-xarch` ISA Keywords

Platform	Valid <code>-xarch</code> Keywords
SPARC	<code>generic</code> , <code>native</code> , <code>v7</code> , <code>v8a</code> , <code>v8</code> , <code>v8plus</code> , <code>v8plusa</code> , <code>v8plusb</code> , <code>v9</code> , <code>v9a</code> , <code>v9b</code>
x86	<code>generic</code> , <code>native</code> , <code>386</code> , <code>pentium_pro</code>

Note that although `-xarch` can be used alone, it is part of the expansion of the `-xtarget` option and may be used to override the `-xarch` value that is set by a specific `-xtarget` option. For example:

```
% cc -xtarget=ultra2 -xarch=v8plusb ...
```

overrides the `-xarch=v8` set by `-xtarget=ultra2`

If you use this option with optimization, the appropriate choice can provide good performance of the executable on the specified architecture. An inappropriate choice results in a binary program that is not executable on the intended target platform.

SPARC Only

The following table details the performance of an executable that is compiled with a given `-xarch` option and then executed by various SPARC processors. The purpose of this table is to help you identify the best `-xarch` option for your executable given a particular target machine. Start by identifying the range of machines that are of interest to you and then consider the cost of maintaining multiple binaries versus the benefit of extracting the last iota of performance from newer machines.

TABLE 2-6 `-xarch` Matrix

		Instruction Set of SPARC Machine:					
		V7	V8a	V8	V9 (Non-Sun Processor)	V9 (Sun processor)	V9b
<code>-xarch</code> compilation option	v7	N	S	S	S	S	S
	v8a	PD	N	S	S	S	S
	v8	PD	PD	N	S	S	S
	v8plus	NE	NE	NE	N	S	S
	v8plusa	NE	NE	NE	**	N	S
	v8plusb	NE	NE	NE	**	NE	N
	v9	NE	NE	NE	N	S	S
	v9a	NE	NE	NE	**	N	S
	v9b	NE	NE	NE	**	NE	N

** Note: An executable compiled with this instruction set may perform nominally on a V9 non-Sun processor chip or it may not execute at all. Check with your hardware vendor to make sure your executable can run on its target machine.

- N reflects Nominal performance. The program executes and takes full advantage of the processor's instruction set.
- S reflects Satisfactory performance. The program executes but may not exploit all available processor instructions.
- PD reflects Performance Degradation. The program executes, but depending on the instructions used, may experience slight to significant performance degradation. The degradation occurs when instructions that are not implemented by the processor are emulated by the kernel.
- NE means Not Executable. The program will not execute because the kernel does not emulate the instructions that are not implemented by the processor.

If you are going to compile your executable with the `v8plus` or `v8plusa` instruction set, consider compiling your executable with `v9` or `v9a` instead. The `v8plus` and `v8plusa` options were provided so that programs could take advantage of some SPARC V9 and UltraSPARC features prior to the availability of Solaris 7 with its support for 64-bit programs. Programs compiled with the `v8plus` or `v8plusa` option are not portable to SPARC V8 or older machines. Such programs can be re-compiled with `v9` or `v9a`, respectively, to take full advantage of all the features of SPARC V9 and UltraSPARC. *The V8+ Technical Specification* white paper, part number 802-7447-10, is available through your Sun representative and explains the limitations of `v8plus` and `v8plusa`.

- SPARC instruction set architectures V7, V8, and V8a are all binary compatible.
- Object binary files (`.o`) compiled with `v8plus` and `v8plusa` can be linked and can execute together, but only on a SPARC V8plusa compatible platform.
- Object binary files (`.o`) compiled with `v8plus`, `v8plusa`, and `v8plusb` can be linked and can execute together, but only on a SPARC V8plusb compatible platform.
- `-xarch` values `v9`, `v9a`, and `v9b` are only available on UltraSPARC 64-bit Solaris environments.
- Object binary files (`.o`) compiled with `v9` and `v9a` can be linked and can execute together, but will run only on a SPARC V9a compatible platform.
- Object binary files (`.o`) compiled with `v9`, `v9a`, and `v9b` can be linked and can execute together, but will run only on a SPARC V9b compatible platform.

For any particular choice, the generated executable may run much more slowly on earlier architectures. Also, although quad-precision (`REAL*16` and `long double`) floating-point instructions are available in many of these instruction set architectures, the compiler does not use these instructions in the code it generates.

The following table gives details for each of the `-xarch` keywords on SPARC platforms.

TABLE 2-7 `-xarch` Values for SPARC Platforms

<code>-xarch=</code>	Meaning
<code>generic</code>	<p>Compile for good performance on most systems. This is the default. This option uses the best instruction set for good performance on most processors without major performance degradation on any of them. With each new release, the definition of “best” instruction set may be adjusted, if appropriate.</p>
<code>native</code>	<p>Compile for good performance on this system. This is the default for the <code>-fast</code> option. The compiler chooses the appropriate setting for the current system processor it is running on.</p>
<code>v7</code>	<p>Compile for the SPARC-V7 ISA. Enables the compiler to generate code for good performance on the V7 ISA. This is equivalent to using the best instruction set for good performance on the V8 ISA, but without integer <code>mul</code> and <code>div</code> instructions, and the <code>fsmuld</code> instruction. Examples: SPARCstation 1, SPARCstation 2</p>
<code>v8a</code>	<p>Compile for the V8a version of the SPARC-V8 ISA. By definition, V8a means the V8 ISA, but without the <code>fsmuld</code> instruction. This option enables the compiler to generate code for good performance on the V8a ISA. Example: Any system based on the microSPARC I chip architecture</p>
<code>v8</code>	<p>Compile for the SPARC-V8 ISA. Enables the compiler to generate code for good performance on the V8 architecture. Example: SPARCstation 10</p>
<code>v8plus</code>	<p>Compile for the V8plus version of the SPARC-V9 ISA. By definition, <code>v8plus</code> means the V9 ISA, but limited to the 32-bit subset defined by the V8plus ISA specification, without the Visual Instruction Set (VIS), and without other implementation-specific ISA extensions.</p> <ul style="list-style-type: none"> • This option enables the compiler to generate code for good performance on the V8plus ISA. • The resulting object code is in SPARC-V8+ ELF32 format and only executes in a Solaris UltraSPARC environment—it does not run on a V7 or V8 processor. <p>Example: Any system based on the UltraSPARC chip architecture</p>

TABLE 2-7 `-xarch` Values for SPARC Platforms (*Continued*)

<code>-xarch=</code>	Meaning
<code>v8plusa</code>	<p>Compile for the V8plusa version of the SPARC-V9 ISA.</p> <p>By definition, <code>v8plusa</code> means the V8plus architecture, plus the Visual Instruction Set (VIS) version 1.0, and with UltraSPARC extensions.</p> <ul style="list-style-type: none">• This option enables the compiler to generate code for good performance on the UltraSPARC architecture, but limited to the 32-bit subset defined by the V8plus specification.• The resulting object code is in SPARC-V8+ ELF32 format and only executes in a Solaris UltraSPARC environment—it does not run on a V7 or V8 processor. <p>Example: Any system based on the UltraSPARC chip architecture</p>
<code>v8plusb</code>	<p>Compile for the V8plusb version of the SPARC-V8plus ISA with UltraSPARC III extensions.</p> <p>Enables the compiler to generate object code for the UltraSPARC architecture, plus the Visual Instruction Set (VIS) version 2.0, and with UltraSPARC III extensions.</p> <ul style="list-style-type: none">• The resulting object code is in SPARC-V8+ ELF32 format and executes only in a Solaris UltraSPARC III environment.• Compiling with this option uses the best instruction set for good performance on the UltraSPARC III architecture.
<code>v9</code>	<p>Compile for the SPARC-V9 ISA.</p> <p>Enables the compiler to generate code for good performance on the V9 SPARC architecture.</p> <ul style="list-style-type: none">• The resulting <code>.o</code> object files are in ELF64 format and can only be linked with other SPARC-V9 object files in the same format.• The resulting executable can only be run on an UltraSPARC processor running a 64-bit enabled Solaris operating environment with the 64-bit kernel.• <code>-xarch=v9</code> is only available when compiling in a 64-bit enabled Solaris environment.
<code>v9a</code>	<p>Compile for the SPARC-V9 ISA with UltraSPARC extensions.</p> <p>Adds to the SPARC-V9 ISA the Visual Instruction Set (VIS) and extensions specific to UltraSPARC processors, and enables the compiler to generate code for good performance on the V9 SPARC architecture.</p> <ul style="list-style-type: none">• The resulting <code>.o</code> object files are in ELF64 format and can only be linked with other SPARC-V9 object files in the same format.• The resulting executable can only be run on an UltraSPARC processor running a 64-bit enabled Solaris operating environment with the 64-bit kernel.• <code>-xarch=v9a</code> is only available when compiling in a 64-bit enabled Solaris operating environment.

TABLE 2-7 `-xarch` Values for SPARC Platforms (*Continued*)

<code>-xarch=</code>	Meaning
<code>v9b</code>	<p>Compile for the SPARC-V9 ISA with UltraSPARC III extensions.</p> <p>Adds UltraSPARC III extensions and VIS version 2.0 to the V9a version of the SPARC-V9 ISA. Compiling with this option uses the best instruction set for good performance in a Solaris UltraSPARC III environment.</p> <ul style="list-style-type: none"> • The resulting object code is in SPARC-V9 ELF64 format and can only be linked with other SPARC-V9 object files in the same format. • The resulting executable can only be run on an UltraSPARC III processor running a 64-bit enabled Solaris operating environment with the 64-bit kernel. • <code>-xarch=v9b</code> is only available when compiling in a 64-bit enabled Solaris operating environment.

x86 Only

TABLE 2-8 `-xarch` Values on x86

Value	Meaning
<code>generic</code>	Limits instruction set to the Intel x86 architecture and is the equivalent of the <code>386</code> option.
<code>native</code>	Compile for good performance on this system. This is the default for the <code>-fast</code> option. The compiler chooses the appropriate setting for the current system processor on which it is compiling.
<code>386</code>	Limits the instruction set to the Intel 386/486 architecture.
<code>pentium_pro</code>	Limits the instruction set to the <code>pentium_pro</code> architecture.

`-xautopar`

(*SPARC*) Turns on automatic parallelization for multiple processors. Does dependence analysis (analyze loops for inter-iteration data dependence) and loop restructuring. If optimization is not at `-xO3` or higher, optimization is raised to `-xO3` and a warning is emitted.

Avoid `-xautopar` if you do your own thread management.

To achieve faster execution, this option requires a multiple processor system. On a single-processor system, the resulting binary usually runs slower.

To determine how many processors you have, use the `psrinfo` command:

```
% psrinfo
0  on-line since 01/12/95 10:41:54
1  on-line since 01/12/95 10:41:54
3  on-line since 01/12/95 10:41:54
4  on-line since 01/12/95 10:41:54
```

To request a number of processors, set the `PARALLEL` environment variable. The default is 1.

- Do not request more processors than are available.
- If `N` is the number of processors on the machine, then for a one-user, multiprocessor system, try `PARALLEL=N-1`.

If you use `-xautopar` and compile and link in *one* step, then linking automatically includes the microtasking library and the threads-safe C runtime library. If you use `-xautopar` and compile and link in *separate* steps, then you must also link with `-xautopar`.

`-xbuiltin[=(%all|%none)]`

Use the `-xbuiltin[=(%all|%none)]` command when you want to improve the optimization of code that calls standard library functions. Many standard library functions, such as the ones defined in `math.h` and `stdio.h`, are commonly used by various programs. This command lets the compiler substitute intrinsic functions or inline system functions where profitable for performance.

The first default of this command is `-xbuiltin=%none`, which means no functions from the standard libraries are substituted or inlined. The first default applies when you do not specify `-xbuiltin`.

The second default of this command is `-xbuiltin%all`, which means the compiler substitutes intrinsics or inlines standard library functions as it determines the optimization benefit. The second default applies when you specify `-xbuiltin` but do not provide an argument.

If you compile with `-fast`, then `-xbuiltin` is set to `%all`.

`-xCC`

Accepts the C++-style comments. In particular, `//` can be used to indicate the start of a comment.

`-xc99=o`

The `-xc99` flag controls compiler recognition of the implemented features from the C99 standard (ISO/IEC 9899:1999, Programming Language - C).

`o` can be one of the following: `%all`, `%none`.

`-xc99=%none` turns off recognition of C99 features. `-xc99=%all` turns on recognition of supported C99 features.

Issuing `-xc99` without any arguments is the same as `-xc99=%all`.

For more information on the supported C99 features of this implementation, see “Supported Features of C99” on page 319.

`-xcache=c`

Defines the cache properties for use by the optimizer. `c` must be one of the following:

- `generic` (`SPARC`, `x86`)
- `s1/l1/a1`
- `s1/l1/a1:s2/l2/a2`
- `s1/l1/a1:s2/l2/a2:s3/l3/a3`

The `si/li/ai` are defined as follows:

- `si` The size of the data cache at level `i`, in kilobytes
- `li` The line size of the data cache at level `i`, in bytes
- `ai` The associativity of the data cache at level `i`

Although this option can be used alone, it is part of the expansion of the `-xtarget` option; its *primary use* is to override a value supplied by the `-xtarget` option.

This option specifies the cache properties that the optimizer can use. It does not guarantee that any particular cache property is used. The following table lists the `-xcache` values.

TABLE 2-9 The `-xcache` Values

Value	Meaning
<code>generic</code>	Define the cache properties for good performance on most x86 and SPARC architectures.
	This is the default value which directs the compiler to use cache properties for good performance on most x86 and SPARC processors, without major performance degradation on any of them.
	With each new release, these best timing properties will be adjusted, if appropriate.
<code>s1/l1/a1</code>	Define level 1 cache properties.
<code>s1/l1/a1:s2/l2/a2</code>	Define levels 1 and 2 cache properties.
<code>s1/l1/a1:s2/l2/a2:s3/l3/a3</code>	Define levels 1, 2, and 3 cache properties.

Example: `-xcache=16/32/4:1024/32/1` specifies the following:

Level 1 cache has:
 16K bytes
 32 bytes line size
 4-way associativity

Level 2 cache has:
 1024K bytes
 32 bytes line size
 Direct mapping associativity

`-xcg[89|92]`

(SPARC)

`-xcg89` is a macro for: `-xarch=v7 -xchip=old -xcache=64/32/1`.

`-xcg92` is a macro for: `-xarch=v8 -xchip=super`

`-xcache=16/32/4:1024/32/1`.

`-xchar_byte_order=o`

Produce an integer constant by placing the characters of a multi-character character-constant in the specified byte order. You can substitute one of the following values for *o*:

- **low**: place the characters of a multi-character character-constant in low-to-high byte order.
- **high**: place the characters of a multi-character character-constant in high-to-low byte order.
- **default**: place the characters of a multi-character character-constant in an order determined by the compilation mode `-X[c|a|t|s]`. For more information, see “Character Constants” on page 89.

`-xchip=c`

Specifies the target processor for use by the optimizer.

c must be one of the following: `generic`, `old`, `super`, `super2`, `micro`, `micro2`, `hyper`, `hyper2`, `powerup`, `ultra`, `ultra2`, `ultra2i`, `386`, `486`, `pentium`, `pentium_pro`, `603`, `604`.

Although this option can be used alone, it is part of the expansion of the `-xtarget` option; its *primary use* is to override a value supplied by the `-xtarget` option.

This option specifies timing properties by specifying the target processor.

Some effects are:

- The ordering of instructions, that is, scheduling
- The way the compiler uses branches
- The instructions to use in cases where semantically equivalent alternatives are available

TABLE 2-10 The `-xchip` Values

Value	Meaning
<code>generic</code>	Use timing properties for good performance on most x86 and SPARC architectures. This is the default value that directs the compiler to use the best timing properties for good performance on most processors, without major performance degradation on any of them.
<code>old</code>	Uses timing properties of pre-SuperSPARC processors.
<code>super</code>	Uses timing properties of the SuperSPARC processors.

TABLE 2-10 The `-xchip` Values (Continued)

Value	Meaning
<code>super2</code>	Uses timing properties of the SuperSPARC II processors.
<code>micro</code>	Uses timing properties of the microSPARC processors.
<code>micro2</code>	Uses timing properties of the microSPARC II processors.
<code>hyper</code>	Uses timing properties of the hyperSPARC processors.
<code>hyper2</code>	Uses timing properties of the hyperSPARC II processors.
<code>powerup</code>	Uses timing properties of the Weitek PowerUp processors.
<code>ultra</code>	Uses timing properties of the UltraSPARC processors.
<code>ultra2</code>	Uses timing properties of the UltraSPARC II processors.
<code>ultra2e</code>	Uses timing properties of the UltraSPARC Iie processors.
<code>ultra2i</code>	Uses timing properties of the UltraSPARC Iii processors.
<code>ultra3</code>	Uses timing properties of the UltraSPARC III processors.
<code>386</code>	Uses timing properties of the Intel 386 architecture.
<code>486</code>	Uses timing properties of the Intel 486 architecture.
<code>pentium</code>	Uses timing properties of the Intel pentium architecture.
<code>pentium_pro</code>	Uses timing properties of the Intel pentium_pro architecture.

`-xcode=v`

(SPARC) Specify code address space. *v* must be one of:

- `abs32` Generate 32-bit absolute addresses. Code + data + bss size is limited to 2^{32} bytes. This is the default on 32-bit architectures: `-xarch-generic`, `v7`, `v8`, `v8a`, `v8plus`, `v8plusa`
- `abs44` Generate 44-bit absolute addresses. Code + data + bss size is limited to 2^{44} bytes. Available only on 64-bit architectures: `-xarch=v9`, `v9a`

abs64	Generate 64-bit absolute addresses. Available only on 64-bit architectures: -xarch=v9, v9a
pic13	Generate position-independent code for use in shared libraries (small model). Equivalent to -Kpic. Permits references to at most 2^{11} unique external symbols on 32-bit architectures, 2^{10} on 64-bit architectures. The <code>-xcode=pic13</code> command is similar to <code>-xcode=pic32</code> , except that the size of the global offset table is limited to 8Kbytes.
pic32	Generate position-independent code for use in shared libraries (large model). Equivalent to -KPIC. Permits references to at most 2^{30} unique external symbols on 32-bit architectures, 2^{29} on 64-bit architectures. Each reference to a global datum is generated as a dereference of a pointer in the global offset table. Each function call is generated in pc-relative addressing mode through a procedure linkage table. With this option, the global offset table spans the range of 32-bit addresses in those rare cases where there are too many global data objects for <code>-xcode=pic32</code> .

The default is `-xcode=abs32` for SPARC V7 and V8, and `-xcode=abs64` for SPARC and UltraSPARC V9 (with `-xarch=v9 | v9a`).

When building shared dynamic libraries with `-xarch=v9` or `v9a` on 64-bit Solaris 7, you must specify `-xcode=pic13` or `-xcode=pic32`.

There are two nominal performance costs with `-xcode=pic13` and `-xcode=pic32`:

- A routine compiled with either `-xcode=pic13` or `-xcode=pic32` executes a few extra instructions upon entry to set a register to point at a table (`_GLOBAL_OFFSET_TABLE_`) used for accessing a shared library's global or static variables.
- Each access to a global or static variable involves an extra indirect memory reference through `_GLOBAL_OFFSET_TABLE_`. If the compile is done with `-xcode=pic32`, there are two additional instructions per global and static memory reference.

When considering the above costs, remember that the use of `-xcode=pic13` and `-xcode=pic32` can significantly reduce system memory requirements, due to the effect of library code sharing. Every page of code in a shared library compiled `-xcode=pic13` or `-xcode=pic32` can be shared by every process that uses the library. If a page of code in a shared library contains even a single non-pic (that is, absolute) memory reference, the page becomes nonsharable, and a copy of the page must be created each time a program using the library is executed.

The easiest way to tell whether or not a `.o` file has been compiled with `-xcode=pic13` or `-xcode=pic32` is with the `nm` command:

```
% nm file.o | grep _GLOBAL_OFFSET_TABLE_ U _GLOBAL_OFFSET_TABLE_
```

A .o file containing position-independent code contains an unresolved external reference to `_GLOBAL_OFFSET_TABLE_`, as indicated by the letter U.

To determine whether to use `-xcode=pic13` or `-xcode=pic32`, use `nm` to identify the number of distinct global and static variables used or defined in the library. If the size of `_GLOBAL_OFFSET_TABLE_` is under 8,192 bytes, you can use `-Kpic`. Otherwise, you must use `-xcode=pic32`.

`-xcrossfile=[n]`

(SPARC) Enables optimization and inlining across source files. If specified, *n* can be 0 or 1.

Normally the scope of the compiler's analysis is limited to each separate file on the command line. For example, `-xO4`'s automatic inlining is limited to subprograms defined and referenced within the same source file.

With `-xcrossfile`, the compiler analyzes all the files named on the command line as if they had been concatenated into a single source file. `-xcrossfile` is only effective when used with `-xO4` or `-xO5`.

The files produced from this compilation are interdependent due to possible inlining, and must be used as a unit when they are linked into a program. If any one routine is changed and the files recompiled, they must all be recompiled. As a result, using this option affects the construction of makefiles.

The default is `-xcrossfile=0`, and no cross-file optimizations are performed. `-xcrossfile` is equivalent to `-xcrossfile=1`.

`-xcsi`

Allows the C compiler to accept source code written in locales that do not conform to the ISO C source character code requirements. These locales include: `ja_JP.PCK`.

The compiler translation phases required to handle such locales may result in significantly longer compilation times. You should only use this option when you compile source files that contain source characters from one of these locales.

The compiler does not recognize source code written in locales that do not conform to the ISO C source character code requirements unless you issue `-xcsi`.

`-xdepend`

(SPARC) Analyzes loops for inter-iteration data dependencies and does loop restructuring. Loop restructuring includes loop interchange, loop fusion, scalar replacement, and elimination of “dead” array assignments. If optimization is not at `-xO3` or higher, optimization is raised to `-xO3` and a warning is issued.

Dependency analysis is also included with `-xautopar` or `-xparallel`. The dependency analysis is done at compile time.

Dependency analysis may help on single-processor systems. However, if you try `-xdepend` on single-processor systems, you should not use either `-xautopar` or `-xexplicitpar`. If either of them is on, then the `-xdepend` optimization is done for multiple-processor systems.

`-xe`

Performs only syntax and semantic checking on the source file, but does not produce any object or executable code.

`-xexplicitpar`

(SPARC) Generates parallelized code based on specification of `#pragma MP` directives. You do the dependency analysis: analyze and specify loops for inter-iteration data dependencies. The software parallelizes the specified loops. If optimization is not at `-xO3` or higher, optimization is raised to `-xO3` and a warning is issued. Avoid `-xexplicitpar` if you do your own thread management.

To get faster code, this option requires a multiprocessor system. On a single-processor system, the generated code usually runs slower.

If you identify a loop for parallelization, and the loop has dependencies, you can get incorrect results, possibly different ones with each run, and with no warnings. Do not apply an explicit parallel pragma to a reduction loop. The explicit parallelization is done, but the reduction aspect of the loop is not done, and the results can be incorrect.

In summary, to parallelize explicitly:

- Analyze the loops to find those that are safe to parallelize.
- Insert `#pragma MP` to parallelize a loop. See the “Explicit Parallelization and Pragmas” on page 129” for more information.
- Use the `-xexplicitpar` option.

The following is an example of inserting a parallel pragma immediately before the loop:

```
#pragma MP taskloop
  for (j=0; j<1000; j++){
    ...
  }
```

If you use `-xexplicitpar` and compile and link in *one* step, then linking automatically includes the microtasking library and the threads-safe C runtime library. If you use `-xexplicitpar` and compile and link in *separate* steps, then you must also *link* with `-xexplicitpar`.

Do not issue `-xexplicitpar` and `-xopenmp` together.

`-xF`

Enables performance analysis of the executable using the Analyzer. (See the `analyzer(1)` man pages.) Produces code that can be reordered at the function level. Each function in the file is placed in a separate section; for example, functions `foo()` and `bar()` are placed in the sections `.text%foo` and `.text%bar`, respectively. Function ordering in the executable can be controlled by using `-xF` in conjunction with the `-M` option to `ld` (see `ld(1)`). This option also causes the assembler to generate some debugging information in the object file, necessary for data collection.

`-xhelp=f`

Displays on-line help information.

f must be either `flags`, or `readme`.

`-xhelp=flags` displays a summary of the compiler options.

`-xhelp=readme` displays the README file.

`-xildoff`

Turns off the incremental linker and forces the use of `ld`. This option is the default if you do not use the `-g` option, or you do use the `-G` option, or any source files are present on the command line. Override this default by using the `-xildon` option.

`-xildon`

Turns on the incremental linker and forces the use of `ild` in incremental mode. This option is the default if you use the `-g` option, and you do not use the `-G` option, and there are no source files present on the command line. Override this default by using the `-xildoff` option.

`-xinline=list`

The format of the *list* for `-xinline` is as follows:

```
[ { %auto, func_name, no%func_name } [ , { %auto, func_name, no%func_name } ] . . . ]
```

`-xinline` tries to inline only those functions specified in the list. The list is comprised of either a comma-separated list of functions names, or a comma separated list of `no%func_name` values, or the value `%auto`. If you issue `no%func_name`, the compiler is not to inline the named function. If you issue `%auto`, the compiler is to attempt to automatically inline all functions in the source files.

If you are compiling with `-xO3`, you can use `-xinline` to increase optimization by inlining some or all functions. The `-xO3` level of optimization does not include inlining.

If you are compiling with `-xO4`, `-xinline` can decrease optimization by restricting inlining to only those routines in the list. With `-xO4`, the compiler normally tries to inline all references to functions defined in the source file. When you specify `xinline=` but do not name any functions or `%auto`, this indicates that none of the routines in the source file are to be inlined.

A function is not inlined if any of the following conditions apply. No warning is issued.

- Optimization is less than `-xO3`.
- The routine cannot be found.
- Inlining the routine does not look practicable to the optimizer.
- The source for the routine is not in the file being compiled (however, see `-xcrossfile`).

`-xipo[=a]`

Replace *a* with either 0 or 1. `-xipo` without any arguments is equivalent `-xipo=1`. `-xipo=0` is the default setting and turns off `-xipo`.

This compiler flag performs whole-program optimizations by invoking an interprocedural analysis pass. Unlike `-xcrossfile`, `-xipo` performs optimizations across all object files in the link step, and is not limited to just the source files of the compile command.

The `-xipo` option generates significantly larger object files due to the additional information needed to perform optimizations across files. However, this additional information does not become part of the final executable binary file. Any increase in the size of the executable program is due to the additional optimizations performed. The object files created in the compilation steps have additional analysis information compiled within them to permit crossfile optimizations to take place at the link step.

`-xipo` is particularly useful when compiling and linking large multi-file applications. Object files compiled with this flag have analysis information compiled within them that enables interprocedural analysis across source and pre-compiled program files.

However, analysis and optimization is limited to the object files compiled with `-xipo`, and does not extend to object files in the libraries.

`-xipo` is multiphased, so you need to specify `-xipo` for each step if you compile and link in separate steps.

In this example, compilation and linking occur in a single step:

```
cc -xipo -xO4 -o prog part1.c part2.c part3.c
```

The optimizer performs crossfile inlining across all three source files. This is done in the final link step, so the compilation of the source files need not all take place in a single compilation and could take place over a number of separate compilations, each specifying `-xipo`.

In this example, compilation and linking occur in separate steps:

```
cc -xipo -xO4 -c part1.c part2.c
cc -xipo -xO4 -c part3.c
cc -xipo -xO4 -o prog part1.o part2.o part3.o
```

A restriction is that libraries, even if compiled with `-xipo`, do not participate in crossfile interprocedural analysis, as this example shows:

```
cc -xipo -xO4 one.c two.c three.c
ar -r mylib.a one.o two.o three.o
...
cc -xipo -xO4 -o myprog main.c four.c mylib.a
```

Here interprocedural optimizations are performed between `one.c`, `two.c` and `three.c`, and between `main.c` and `four.c`, but not between `main.c` or `four.c` and the routines on `mylib.a`. (The first compilation may generate warnings about undefined symbols, but the interprocedural optimizations are performed because it is a compile and link step.)

Other important information about `-xipo`:

- It requires an optimization level of at least `-xO4`.
- It conflicts with `-xcrossfile`. If you use these together, the result is a compilation error.
- Objects that are compiled without `-xipo` can be linked freely with objects that are compiled with `-xipo`.

`-xlibmieee`

Forces IEEE 754 style return values for math routines in exceptional cases. In such cases, no exception message is printed, and you should not rely on `errno`.

`-xlibmil`

Inlines some library routines for faster execution. This option selects the appropriate assembly language inline templates for the floating-point option and platform for your system.

`-xlic_lib=sunperf`

(SPARC) Links in the Sun-supplied performance libraries.

`-xlicinfo`

Returns information about the licensing system. In particular, this option returns the name of the license server and the IDs of users who have checked out licenses. This option does not request compilation or check out a license.

`-xloopinfo`

(SPARC) Shows which loops are parallelized and which are not. Gives a short reason for not parallelizing a loop. The `-xloopinfo` option is valid only if `-xautopar`, or `-xparallel`, or `-xexplicitpar` is specified; otherwise, the compiler issues a warning.

To achieve faster execution, this option requires a multiprocessor system. On a single-processor system, the generated code usually runs slower.

`-xM`

Runs only the preprocessor on the named C programs, requesting that it generate makefile dependencies and send the result to the standard output (see `make(1)` for details about makefiles and dependencies).

For example:

```
#include <unistd.h>
void main(void)
{ }
```

generates this output:

```
e.o: e.c
e.o: /usr/include/unistd.h
e.o: /usr/include/sys/types.h
e.o: /usr/include/sys/machtypes.h
e.o: /usr/include/sys/select.h
e.o: /usr/include/sys/time.h
e.o: /usr/include/sys/types.h
e.o: /usr/include/sys/time.h
e.o: /usr/include/sys/unistd.h
```

-xM1

Collects dependencies like -xM, but excludes /usr/include files. For example:

```
more hello.c
#include<stdio.h>
main()
{
    (void)printf("hello\n");
}
cc -xM hello.c
hello.o: hello.c
hello.o: /usr/include/stdio.h
```

Compiling with -xM1 does not report header file dependencies:

```
cc -xM1 hello.c
hello.o: hello.c
```

-xM1 is not available under -Xs mode.

-xMerge

Merges data segments into text segments. Data initialized in the object file produced by this compilation is read-only and (unless linked with ld -N) is shared between processes.

-xmaxopt [=v]

where *v* is one of *off*, 1, 2, 3, 4, 5. This command limits the level of `pragma opt` to the level specified. The default value is `-xmaxopt=off` which causes `pragma opt` to be ignored. If you specify `-xmaxopt` without supplying an argument, that is the equivalent of specifying `-xmaxopt=5`.

`-xmemalign=ab`

Specify maximum assumed memory alignment and behavior of misaligned data accesses. There must be a value for both *a* (alignment) and *b* (behavior). *a* specifies the maximum assumed memory alignment and *b* specifies the behavior for misaligned memory accesses. The following table lists the alignment and behavior values for `-xmemalign`

TABLE 2-11 `-xmemalign` Alignment and Behavior Values

<i>a</i>		<i>b</i>	
1	Assume at most 1 byte alignment.	i	Interpret access and continue execution.
2	Assume at most 2 byte alignment.	s	Raise signal SIGBUS.
4	Assume at most 4 byte alignment.	f	Raise signal SIGBUS for alignments less or equal to 4, otherwise interpret access and continue execution.
8	Assume at most 8 byte alignment.		
16	Assume at most 16 byte alignment		

For memory accesses where the alignment is determinable at compile time, the compiler will generate the appropriate load/store instruction sequence for that alignment of data.

For memory accesses where the alignment cannot be determined at compile time, the compiler must assume an alignment to generate the needed load/store sequence.

The `-xmemalign` flag allows the user to specify the maximum memory alignment of data to be assumed by the compiler in these indeterminable situations. It also specifies the error behavior to be followed at run time when a misaligned memory access does take place.

Here are the default values for `-xmemalign`. The following default values only apply when no `-xmemalign` flag is present:

- `-xmemalign=4s` when `-xarch` has the value `generic`, `v7`, `v8`, `v8a`, `v8plus`, `v8plusa`.
- `-xmemalign=8s` when `-xarch` has the value `v9`, `v9a`.

Here is the default when `-xmemalign` flag is present but no value is given:

- `-xmemalign=1i` for all `-xarch` values.

The following table shows how you can use `-xmemalign` to handle different alignment situations.

TABLE 2-12 Examples of `-xmemalign`

Command	Situation
<code>-xmemalign=1s</code>	There are many misaligned accesses so trap handling is too slow.
<code>-xmemalign=8i</code>	There are occasional, intentional, misaligned accesses in code that is otherwise correct.
<code>-xmemalign=8s</code>	There should be no misaligned accesses in the program.
<code>-xmemalign=2s</code>	You want to check for possible odd-byte accesses.
<code>-xmemalign=2i</code>	You want to check for possible odd-byte access and you want the program to work.

`-xno lib`

Does not link any libraries by default; that is, no `-l` options are passed to `ld(1)`. Normally, the `cc` driver passes `-lc` to `ld`.

When you use `-xno lib`, you have to pass all the `-l` options yourself. For example:

```
% cc test.c -xno lib -Bstatic -lm -Bdynamic -lc
```

links `libm` statically and the other libraries dynamically.

`-xno libmil`

Does not inline math library routines. Use it after the `-fast` option. For example:

```
% cc -fast -xno libmil....
```

`-xO[1 | 2 | 3 | 4 | 5]`

Optimizes the object code; note the upper-case letter O. When `-xO` is used with the `-g` option, a limited amount of debugging is available. For more information, see “Debugging Optimized Code” in Chapter 1 of *Debugging a Program With dbx*.

The levels (1, 2, 3, 4, or 5) you can use with `-xO` differ according to the platform you are using.

(SPARC)

`-xO1`

Does basic local optimization (peephole).

`-xO2`

Does basic local and global optimization. This is induction variable elimination, local and global common subexpression elimination, algebraic simplification, copy propagation, constant propagation, loop-invariant optimization, register allocation, basic block merging, tail recursion elimination, dead code elimination, tail call elimination, and complex expression expansion.

The `-xO2` level does not assign global, external, or indirect references or definitions to registers. It treats these references and definitions as if they were declared `volatile`. In general, the `-xO2` level results in minimum code size.

`-xO3`

Performs like `-xO2`, but also optimizes references or definitions for external variables. Loop unrolling and software pipelining are also performed. This level does not trace the effects of pointer assignments. When compiling either device drivers, or programs that modify external variables from within signal handlers, you may need to use the `volatile` type qualifier to protect the object from optimization. In general, the `-xO3` level results in increased code size.

`-xO4`

Performs like `-xO3`, but also automatically inlines functions contained in the same file; this usually improves execution speed. If you want to control which functions are inlined, see “`-xinline=list`” on page 57.

This level traces the effects of pointer assignments, and usually results in increased code size.

-x05

Attempts to generate the highest level of optimization. Uses optimization algorithms that take more compilation time or that do not have as high a certainty of improving execution time. Optimization at this level is more likely to improve performance if it is done with profile feedback. See “-xprofile=p” on page 68.

(x86)

-x01

Preloads arguments from memory, cross-jumping (tail-merging), as well as the single pass of the default optimization.

-x02

Schedules both high- and low-level instructions and performs improved spill analysis, loop memory-reference elimination, register lifetime analysis, enhanced register allocation, and elimination of global common subexpressions.

-x03

Performs loop strength reduction, induction variable elimination, as well as the optimization done by level 2.

-x04

Performs loop unrolling, avoids creating stack frames when possible, and automatically inlines functions contained in the same file, as well as the optimization done by levels 2 and 3. Note that this optimization level can cause stack traces from `adb` and `dbx` to be incorrect.

-x05

Generates the highest level of optimization. Uses optimization algorithms that take more compilation time or that do not have as high a certainty of improving execution time. Some of these include generating local calling convention entry points for exported functions, further optimizing spill code and adding analysis to improve instruction scheduling.

If the optimizer runs out of memory, it tries to recover by retrying the current procedure at a lower level of optimization and resumes subsequent procedures at the original level specified in the command-line option.

If you optimize at -x03 or -x04 with very large procedures (thousands of lines of code in the same procedure), the optimizer may require a large amount of virtual memory. In such cases, machine performance may degrade.

For more information on debugging, see the *Debugging a Program With dbx* manual. For more information on optimization, see the *Analyzing Program Performance With Sun Workshop* manual.

`-xopenmp [=i]`

where *i* is one of `parallel`, `stubs`, or `none`. If you specify `-xopenmp` but do not include a value, the compiler assumes `-xopenmp=parallel`. If you do not specify `-xopenmp`, the compiler assumes `-xopenmp=none`.

`-xopenmp=parallel` enables recognition of OpenMP pragmas and applies to SPARC only. The optimization level under `-xopenmp=parallel` is `-xO3`. The compiler issues a warning if the optimization level of your program is changed from a lower level to `-xO3`. `-xopenmp=parallel` predefines the `_OPENMP` preprocessor token.

`-xopenmp=stubs` links with the stubs routines for the OpenMP API routines. Use this option if you need to compile your application to execute serially. `-xopenmp=stubs` also predefines the `_OPENMP` preprocessor token.

`-xopenmp=none` does not enable recognition of OpenMP pragmas, makes no change to the optimization level of your program, and does not predefine any preprocessor tokens.

Do not issue `-xopenmp`, `-xexplicitpar`, or `-xparallel` together.

For more information on how to compile a program that is OpenMP compliant, see “Parallelizing for OpenMP” on page 106.

For information that is specific to this implementation of OpenMP, see “Implementation-Specific Information of OpenMP” on page 335.

-xP

Prints prototypes for all K&R C functions defined in this module.

```
f()
{
}

main(argc,argv)
int argc;
char *argv[];
{
}
```

produces this output:

```
int f(void);
int main(int, char **);
```

-xparallel

(SPARC) Parallelizes loops both automatically by the compiler and explicitly specified by the programmer. The `-xparallel` option is a macro, and is equivalent to specifying all three of `-xautopar`, `-xdepend`, and `-xexplicitpar`. With explicit parallelization of loops, there is a risk of producing incorrect results. If optimization is not at `-xO3` or higher, optimization is raised to `-xO3` and a warning is issued.

Avoid `-xparallel` if you do your own thread management. Do not issue `-xparallel` if you are issuing `-xopenmp`. `-xparallel` sets `-xexplicitpar` which should not be issued if you specify `-xopenmp`.

To get faster code, this option requires a multiprocessor system. On a single-processor system, the generated code usually runs slower.

If you compile and link in *one* step, `-xparallel` links with the microtasking library and the threads-safe C runtime library. If you compile and link in *separate* steps, and you compile with `-xparallel`, then link with `-xparallel`

-xpentium

(x86) Optimizes for the Pentium processor.

`-xpg`

Prepares the object code to collect data for profiling with `gprof(1)`. It invokes a runtime recording mechanism that produces a `gmon.out` file at normal termination.

`-xprefetch=[val],val`

(SPARC) Enable prefetch instructions on those architectures that support prefetch, such as UltraSPARC II. (`-xarch=v8plus, v9plusa, v9, or v9a`)

Explicit prefetching should only be used under special circumstances that are supported by measurements.

`val` must be one of the following:

Value	Meaning
<code>auto</code>	Enable automatic generation of prefetch instructions
<code>no%auto</code>	Disable automatic generation of prefetch instructions
<code>explicit</code>	Enable explicit prefetch macros
<code>no%explicit</code>	Disable explicit prefetch macros
<code>yes</code>	Same as <code>-xprefetch=auto,explicit</code>
<code>no</code>	Same as <code>-xprefetch=no%auto,no%explicit</code>

If you do not specify `-xprefetch`, the default is `-xprefetch=no%auto,explicit`. If you specify `-xprefetch` without a value, that is equivalent to `-xprefetch=auto,explicit`.

The `sun_prefetch.h` header file provides the macros that you can use to specify explicit prefetch instructions. The prefetches are approximately at the place in the executable that corresponds to where the macros appear.

`-xprofile=p`

Collects data for a profile or uses a profile to optimize.

(SPARC) `p` must be `collect[:name]`, `use[:name]`, or `tcov`.

This option causes execution frequency data to be collected and saved during execution, then the data can be used in subsequent runs to improve performance. This option is only valid when you specify a level of optimization.

`collect[:name]`

Collects and saves execution frequency data for later use by the optimizer with `-xprofile=use`. The compiler generates code to measure statement execution-frequency.

The *name* is the name of the program that is being analyzed. This name is optional. If *name* is not specified, `a.out` is assumed to be the name of the executable.

You can set the environment variables `SUN_PROFDATA` and `SUN_PROFDATA_DIR` to control where a program compiled with `-xprofile=collect` stores the profile data. If set, the `-xprofile=collect` data is written to `SUN_PROFDATA_DIR/SUN_PROFDATA`.

These environment variables similarly control the path and names of the profile data files written by `tcov`, as described in the `tcov(1)` man page.

If these environment variables are not set, the profile data is written to `name.profile/feedback` in the current directory, where *name* is the name of the executable or the name specified in the `-xprofile=collect:name` flag. `-xprofile` does not append `.profile` to *name* if *name* already ends in `.profile`. If you run the program several times, the execution frequency data accumulates in the `feedback` file; that is, output from prior executions is not lost.

`use[:name]`

Uses execution frequency data to optimize strategically.

As with `collect:name`, the *name* is optional and may be used to specify the name of the program.

The program is optimized by using the execution frequency data previously generated and saved in the `feedback` files written by a previous execution of the program compiled with `-xprofile=collect`.

The source files and other compiler options must be exactly the same as those used for the compilation that created the compiled program that generated the `feedback` file. The same version of the compiler must be used for both the `collect` build and the `use` build as well. If compiled with `-xprofile=collect:name`, the same program name *name* must appear in the optimizing compilation: `-xprofile=use:name`.

`tcov`

Basic block coverage analysis using “new” style `tcov`.

The `-xprofile=tcov` option is the new style of basic block profiling for `tcov`. It has similar functionality to the `-xa` option, but correctly collects data for programs that have source code in header files. See “`-xa`” on page 39 for information on the old style of profiling, the `tcov(1)` man page, and *Analyzing Program Performance With Sun WorkShop* for more details.

Code instrumentation is performed similarly to that of the `-xa` option, but `.d` files are no longer generated. Instead, a single file is generated, the name of which is based on the final executable. For example, if the program is run out of `/foo/bar/myprog.profile`, the data file is stored in `/foo/bar/myprog.profile/myprog.tcovd`.

The `-xprofile=tcov` and the `-xa` options are compatible in a single executable. That is, you can link a program that contains some files that have been compiled with `-xprofile=tcov`, and others with `-xa`. You cannot compile a single file with both options.

When running `tcov`, you must pass it the `-x` option to make it use the new style of data. If not, `tcov` uses the old `.d` files, if any, by default for data, and produces unexpected output.

Unlike the `-xa` option, the `TCOVDIR` environment variable has no effect at compile-time. However, its value is used at program runtime. See `tcov(1)` and *Analyzing Program Performance With Sun WorkShop* for more details.

Note – `tcov`'s code coverage report can be unreliable if there is inlining of routines due to `-xO4` or `-xinline`.

`-xreduction`

(SPARC) Turns on reduction recognition during automatic parallelization. `-xreduction` must be specified with `-xautopar`, or `-xparallel`.

When reduction recognition is enabled, the compiler parallelizes reductions such as *dot* products, maximum and minimum finding. These reductions yield different roundoffs than obtained by unparallelized code.

`-xregs=r[,r...]`

(SPARC) Specifies the usage of registers for the generated code.

`r` is a comma-separated list that consists of one or more of the following: `[no%]appl`, `[no%]float`.

Example: `-xregs=appl,no%float`

TABLE 2-13 The `-xregs` Values

Value	Meaning
<code>appl</code>	Allows the use of the following registers: <code>g2, g3, g4</code> (<code>v8a, v8, v8plus, v8plusa, v8plusb</code>) <code>g2, g3</code> (<code>v9, v9a, v9b</code>) For more information on SPARC instruction sets, see “ <code>-xarch=isa</code> ” on page 42. In the SPARC ABI, these registers are described as <i>application</i> registers. Using these registers can increase performance because fewer load and store instructions are needed. However, such use can conflict with some old library programs written in assembly code.
<code>no%appl</code>	Does not use the <code>appl</code> registers.
<code>float</code>	Allows using the floating-point registers as specified in the SPARC ABI. You can use these registers even if the program contains no floating-point code.
<code>no%float</code>	Does not use the floating-point registers. With this option, a source program cannot contain any floating-point code.

The default is `-xregs=appl,float`.

`-xrestrict [=f]`

(SPARC) Treats pointer-valued function parameters as restricted pointers. *f* is `%all`, `%none`, or a comma-separated list of one or more function names: `{%all| %none| fn[fn...]}`.

If a function list is specified with this option, pointer parameters in the specified functions are treated as restricted; if `-xrestrict=%all` is specified, all pointer parameters in the entire C file are treated as restricted. Refer to “`_Restrict` Keyword” on page 86, for more information.

This command-line option can be used on its own, but it is best used with optimization. For example, the command:

```
%cc -xO3 -xrestrict=%all prog.c
```

treats all pointer parameters in the file `prog.c` as restricted pointers. The command:

```
%cc -xO3 -xrestrict=agc prog.c
```

treats all pointer parameters in the function `argc` in the file `prog.c` as restricted pointers.

The default is `%none`; specifying `-xrestrict` is equivalent to specifying `-xrestrict=%all`.

`-XS`

Disables Auto-Read for `dbx`. Use this option in case you cannot keep the `.o` files around. It passes the `-s` option to the assembler.

No Auto-Read is the older way of loading symbol tables. It places all symbol tables for `dbx` in the executable file. The linker links more slowly and `dbx` initializes more slowly.

Auto-Read is the newer and default way of loading symbol tables. With Auto-Read, the information is distributed in the `.o` files, so that `dbx` loads the symbol table information only if and when it is needed. Hence, the linker links faster, and `dbx` initializes faster.

With `-xs`, if you move the executables to another directory, then to use `dbx`, you can ignore the object (`.o`) files.

Without `-xs`, if you move the executables, you must move both the source files and the object (`.o`) files, or set the path with the `dbx pathmap` or `use` command.

`-xsafe=mem`

(*SPARC*) Allows the compiler to assume no memory-based traps occur.

This option grants permission to use the speculative load instruction on V9 machines. It is only effective when you specify `-xO5` optimization and `-xarch=v8plus|v8plusa|v9|v9a`.

`-xsb`

Generates extra symbol table information for the Source Browser. This option is not valid with the `-xs` mode of the compiler.

`-xsbfast`

Creates the database for the Source Browser. Does not compile source into an object file. This option is not valid with the `-xs` mode of the compiler.

`-xsfpcnst`

Represents unsuffixed floating-point constants as single precision, instead of the default mode of double precision. Not valid with `-xc`.

`-xspace`

Does no optimizations or parallelization of loops that increase code size.

Example: The compiler will not unroll loops or parallelize loops if it increases code size.

`-xstrcnst`

Inserts string literals into the read-only data section of the text segment instead of the default data segment.

`-xtarget=t`

Specifies the target system for instruction set and optimization.

The value of *t* must be one of the following: *native*, *generic*, *system-name* (*SPARC*, *x86*).

The `-fast` macro option includes `-xtarget=native` in its expansion.

The `-xtarget` option is a macro that permits a quick and easy specification of the `-xarch`, `-xchip`, and `-xcache` combinations that occur on real systems. The only meaning of `-xtarget` is in its expansion.

TABLE 2-14 The `-xtarget` Values

Value	Meaning
<code>native</code>	Gets the best performance on the host system.
<code>generic</code>	<p>The compiler generates code for the best performance on the host system. It determines the available architecture, chip, and cache properties of the machine on which the compiler is running.</p> <p>Gets the best performance for generic architecture, chip, and cache.</p> <p>The compiler expands <code>-xtarget=generic</code> to: <code>-xarch=generic -xchip=generic -xcache=generic</code></p> <p>This is the default value.</p>
<code>system-name</code>	<p>Gets the best performance for the specified system.</p> <p>You select a system name from TABLE 2-15 that lists the mnemonic encodings of the actual system name and numbers.</p>

The performance of some programs may benefit by providing the compiler with an accurate description of the target computer hardware. When program performance is critical, the proper specification of the target hardware could be very important. This is especially true when running on the newer SPARC processors. However, for most programs and older SPARC processors, the performance gain is negligible and a generic specification is sufficient.

Each specific value for `-xtarget` expands into a specific set of values for the `-xarch`, `-xchip`, and `-xcache` options. See TABLE 2-15 for the values. For example:

```
-xtarget=sun4/15 is equivalent to: -xarch=v8a -xchip=micro
-xcache=2/16/1
```

TABLE 2-15 `-xtarget` Expansions on SPARC

<code>-xtarget=</code>	<code>-xarch</code>	<code>-xchip</code>	<code>-xcache</code>
<code>generic</code>	<code>generic</code>	<code>generic</code>	<code>generic</code>
<code>cs6400</code>	<code>v8</code>	<code>super</code>	<code>16/32/4:2048/64/1</code>
<code>entr150</code>	<code>v8</code>	<code>ultra</code>	<code>16/32/1:512/64/1</code>

TABLE 2-15 -xtarget Expansions on SPARC (Continued)

-xtarget=	-xarch	-xchip	-xcache
entr2	v8plusa	ultra	16/32/1:512/64/1
entr2/1170	v8plusa	ultra	16/32/1:512/64/1
entr2/1200	v8plusa	ultra	16/32/1:512/64/1
entr2/2170	v8plusa	ultra	16/32/1:512/64/1
entr2/2200	v8plusa	ultra	16/32/1:512/64/1
entr3000	v8plusa	ultra	16/32/1:512/64/1
entr4000	v8plusa	ultra	16/32/1:512/64/1
entr5000	v8plusa	ultra	16/32/1:512/64/1
entr6000	v8plusa	ultra	16/32/1:512/64/1
sc2000	v8	super	16/32/4:2048/64/1
solb5	v7	old	128/32/1
solb6	v8	super	16/32/4:1024/32/1
ss1	v7	old	64/16/1
ss10	v8	super	16/32/4
ss10/20	v8	super	16/32/4
ss10/30	v8	super	16/32/4
ss10/40	v8	super	16/32/4
ss10/402	v8	super	16/32/4
ss10/41	v8	super	16/32/4:1024/32/1
ss10/412	v8	super	16/32/4:1024/32/1
ss10/50	v8	super	16/32/4
ss10/51	v8	super	16/32/4:1024/32/1
ss10/512	v8	super	16/32/4:1024/32/1
ss10/514	v8	super	16/32/4:1024/32/1
ss10/61	v8	super	16/32/4:1024/32/1
ss10/612	v8	super	16/32/4:1024/32/1
ss10/71	v8	super2	16/32/4:1024/32/1
ss10/712	v8	super2	16/32/4:1024/32/1
ss10/hs11	v8	hyper	256/64/1
ss10/hs12	v8	hyper	256/64/1

TABLE 2-15 -xtarget Expansions on SPARC (Continued)

-xtarget=	-xarch	-xchip	-xcache
ss10/hs14	v8	hyper	256/64/1
ss10/hs21	v8	hyper	256/64/1
ss10/hs22	v8	hyper	256/64/1
ss1000	v8	super	16/32/4:1024/32/1
sslplus	v7	old	64/16/1
ss2	v7	old	64/32/1
ss20	v8	super	16/32/4:1024/32/1
ss20/151	v8	hyper	512/64/1
ss20/152	v8	hyper	512/64/1
ss20/50	v8	super	16/32/4
ss20/502	v8	super	16/32/4
ss20/51	v8	super	16/32/4:1024/32/1
ss20/512	v8	super	16/32/4:1024/32/1
ss20/514	v8	super	16/32/4:1024/32/1
ss20/61	v8	super	16/32/4:1024/32/1
ss20/612	v8	super	16/32/4:1024/32/1
ss20/71	v8	super2	16/32/4:1024/32/1
ss20/712	v8	super2	16/32/4:1024/32/1
ss20/hs11	v8	hyper	256/64/1
ss20/hs12	v8	hyper	256/64/1
ss20/hs14	v8	hyper	256/64/1
ss20/hs21	v8	hyper	256/64/1
ss20/hs22	v8	hyper	256/64/1
ss2p	v7	powerup	64/32/1
ss4	v8a	micro2	8/16/1
ss4/110	v8a	micro2	8/16/1
ss4/85	v8a	micro2	8/16/1
ss5	v8a	micro2	8/16/1
ss5/110	v8a	micro2	8/16/1
ss5/85	v8a	micro2	8/16/1

TABLE 2-15 -xtarget Expansions on SPARC (*Continued*)

-xtarget=	-xarch	-xchip	-xcache
ss600/120	v7	old	64/32/1
ss600/140	v7	old	64/32/1
ss600/41	v8	super	16/32/4:1024/32/1
ss600/412	v8	super	16/32/4:1024/32/1
ss600/51	v8	super	16/32/4:1024/32/1
ss600/512	v8	super	16/32/4:1024/32/1
ss600/514	v8	super	16/32/4:1024/32/1
ss600/61	v8	super	16/32/4:1024/32/1
ss600/612	v8	super	16/32/4:1024/32/1
sselc	v7	old	64/32/1
ssipc	v7	old	64/16/1
ssipx	v7	old	64/32/1
sslc	v8a	micro	2/16/1
sslt	v7	old	64/32/1
sslx	v8a	micro	2/16/1
sslx2	v8a	micro2	8/16/1
ssslc	v7	old	64/16/1
ssvyger	v8a	micro2	8/16/1
sun4/110	v7	old	2/16/1
sun4/15	v8a	micro	2/16/1
sun4/150	v7	old	2/16/1
sun4/20	v7	old	64/16/1
sun4/25	v7	old	64/32/1
sun4/260	v7	old	128/16/1
sun4/280	v7	old	128/16/1
sun4/30	v8a	micro	2/16/1
sun4/330	v7	old	128/16/1
sun4/370	v7	old	128/16/1
sun4/390	v7	old	128/16/1
sun4/40	v7	old	64/16/1

TABLE 2-15 -xtarget Expansions on SPARC (Continued)

-xtarget=	-xarch	-xchip	-xcache
sun4/470	v7	old	128/32/1
sun4/490	v7	old	128/32/1
sun4/50	v7	old	64/32/1
sun4/60	v7	old	64/16/1
sun4/630	v7	old	64/32/1
sun4/65	v7	old	64/16/1
sun4/670	v7	old	64/32/1
sun4/690	v7	old	64/32/1
sun4/75	v7	old	64/32/1
ultra	v8plusa	ultra	16/32/1:512/64/1
ultra1/140	v8plusa	ultra	16/32/1:512/64/1
ultra1/170	v8plusa	ultra	16/32/1:512/64/1
ultra1/200	v8plusa	ultra	16/32/1:512/64/1
ultra2	v8plusa	ultra2	16/32/1:512/64/1
ultra2/1170	v8plusa	ultra	16/32/1:512/64/1
ultra2/1200	v8plusa	ultra	16/32/1:1024/64/1
ultra2/1300	v8plusa	ultra2	16/32/1:2048/64/1
ultra2/2170	v8plusa	ultra	16/32/1:512/64/1
ultra2/2200	v8plusa	ultra	16/32/1:1024/64/1
ultra2/2300	v8plusa	ultra2	16/32/1:2048/64/1
ultra2e	v8plusa	ultra2e	16/32/1:256/64/4
ultra2i	v8plusa	ultra2i	16/32/1:512/64/1
ultra3	v8plusa	ultra3	64/32/4:8192/512/1

The following table lists the -xtarget values for the Intel Architecture:

TABLE 2-16 -xtarget Expansions on Intel Architecture

-xtarget=	-xarch	-xchip	-xcache
generic	generic	generic	generic
386	386	386	generic

TABLE 2-16 -xtarget Expansions on Intel Architecture (Continued)

-xtarget=	-xarch	-xchip	-xcache
486	386	386	generic
pentium	386	pentium	generic
pentium_pro	pentium_pro	pentium_pro	generic

-xtemp=*dir*

Sets the directory for temporary files used by `cc` to *dir*. No space is allowed within this option string. Without this option, temporary files go into `/tmp`. `-xtemp` has precedence over the `TMPDIR` environment variable.

-xtime

Reports the time and resources used by each compilation component.

-xtransition

Issues warnings for the differences between K&R C and Sun ANSI/ISO C.

The `-xtransition` option issues warnings in conjunction with the `-Xa` and `-Xt` options. You can eliminate all warning messages about differing behavior through appropriate coding. The following warnings no longer appear unless you issue the `-xtransition` option:

- `\a` is ANSI C "alert" character
- `\x` is ANSI C hex escape
- bad octal digit
- base type is really *type tag: name*
- comment is replaced by "###"
- comment does not concatenate tokens
- declaration introduces new type in ANSI C: *type tag*
- macro replacement within a character constant
- macro replacement within a string literal
- no macro replacement within a character constant
- no macro replacement within a string literal
- operand treated as unsigned
- trigraph sequence replaced
- ANSI C treats constant as unsigned: *operator*
- semantics of *operator* change in ANSI C; use explicit cast

-xtrigraphs

The `-xtrigraphs` option determines whether the compiler recognizes trigraph sequences as defined by the ISO/ANSI C standard.

By default, the compiler assumes `-xtrigraphs=yes` and recognizes all trigraph sequences throughout the compilation unit.

If your source code has a literal string containing question marks (?) that the compiler is interpreting as a trigraph sequence, you can use the `-xtrigraph=no` suboption to turn off the recognition of trigraph sequences. The `-xtrigraphs=no` option turns off recognition of all trigraphs throughout the entire compilation unit.

Consider the following example source file named `trigraphs_demo.c`.

```
#include <stdio.h>

int main ()
{
    (void) printf("(\\?\\?) in a string appears as (??)\\n");
    return 0;
}
```

Here is the output if you compile this code with `-xtrigraphs=yes`.

```
example% cc -xtrigraphs=yes trigraphs_demo.c
example% a.out
(??) in a string appears as ( )
```

Here is the output if you compile this code with `-xtrigraphs=no`.

```
example% cc -xtrigraphs=no trigraphs_demo.c
example% a.out
(??) in a string appears as (??)
```

-xunroll=*n*

Suggests to the optimizer to unroll loops *n* times. *n* is a positive integer. When *n* is 1, it is a command, and the compiler unrolls no loops. When *n* is greater than 1, the `-xunroll=n` merely suggests to the compiler that it unroll loops *n* times.

`-xvector [= {yes | no}]`

Enable automatic generation of calls to the vector library functions.

`-xvector=yes` permits the compiler to transform math library calls within loops into single calls to the equivalent vector math routines when such transformations are possible. Such transformations could result in a performance improvement for loops with large loop counts.

If you do not issue `-xvector`, the default is `-xvector=no`. `-xvector=no` undoes a previously specified `-xvector=yes`. If you specify `-xvector` but do not supply a value, the default is `-xvector=yes`.

If you use `-xvector` on the command line without previously specifying `-xdepend`, `-xvector` triggers `-xdepend`. The `-xvector` option also raises the optimization level to `-x03` if optimization is not specified or optimization is set lower than `-x03`.

The compiler includes the `libmvec` libraries in the load step. If you compile and link with separate commands, be sure to use `-xvector` in the linking `cc` command.

`-xvpara`

(SPARC) Warns about loops that have `#pragma MP` directives specified when the loop may not be properly specified for parallelization. For example, when the optimizer detects data dependencies between loop iterations, it issues a warning.

Use `-xvpara` with the `-xexplicitpar` option or the `-xparallel` option and the `#pragma MP`. See “Explicit Parallelization and Pragmas” on page 129 for more information.

`-Yc, dir`

Specifies a new directory *dir* for the location of component *c*. *c* can consist of any of the characters representing components that are listed under the `-W` option.

If the location of a component is specified, then the new path name for the tool is *dir/tool*. If more than one `-Y` option is applied to any one item, then the last occurrence holds.

`-YA, dir`

Changes the default directory searched for components.

`-YI, dir`

Changes the default directory searched for `include` files.

`-YP, dir`

Changes the default directory for finding library files.

`-YS, dir`

Changes the default directory for startup object files.

`-Zll`

(*SPARC*) Creates the program database for `lock_lint`, but does not actually compile. Refer to the `lock_lint(1)` man page for more details.

Options Passed to the Linker

`cc` recognizes `-a`, `-e`, `-r`, `-t`, `-u`, and `-z` and passes these options and their arguments to `ld`. `cc` passes any unrecognized options to `ld` with a warning.

Sun ANSI/ISO C Compiler-Specific Information

The Sun ANSI/ISO C compiler is compatible with the C language described in the American National Standard for Programming Language-C, ANSI/ISO 9899-1990. This chapter documents those areas specific to the Sun ANSI/ISO C compiler and is organized into the following sections:

- “Environment Variables” on page 83
- “Global Behavior: Value Versus unsigned Preserving” on page 85
- “Keywords” on page 86
- “long long Data Type” on page 86
- “Constants” on page 88
- “Include Files” on page 89
- “Nonstandard Floating Point” on page 92
- “Preprocessing Directives and Names” on page 93

Environment Variables

TMPDIR

`cc` normally creates temporary files in the directory `/tmp`. You can specify another directory by setting the environment variable `TMPDIR` to the directory of your choice. However, if `TMPDIR` is not a valid directory, `cc` uses `/tmp`. The `-xtemp` option has precedence over the `TMPDIR` environment variable.

If you use a Bourne shell, type:

```
$ TMPDIR=dir; export TMPDIR
```

If you use a C shell, type:

```
% setenv TMPDIR dir
```

PARALLEL

(*SPARC*) Specifies the number of processors available to the program for multiprocessor execution. If the target machine has multiple processors, the threads can map to independent processors. Running the program leads to the creation of two threads that execute the parallelized portions of the program.

SUNPROF_DATA

Controls the name of the file in which the `-xprofile=collect` command stores execution-frequency data.

SUNPROF_DATA_DIR

Controls in which directory the `-xprofile=collect` command places the execution-frequency data-file.

SUNPRO_SB_INIT_FILE_NAME

The absolute path name of the directory containing the `.sbinit(5)` file. This variable is used only if the `-xsb` or `-xsbfast` flag is used.

SUNW_MP_THR_IDLE

Controls the status of each thread after it finishes its share of a parallel job. You can set `SUNW_MP_THR_IDLE` to either `spin` or `sleep [n s | n ms]`. The default is `spin`, which means the thread goes spin-waiting. The other choice, `sleep [n s | n ms]` puts the thread to sleep after spin-waiting for n units. The wait can be seconds (`s`) is the default unit) or milliseconds (`ms`) where `1s` means one second, `10 ms` means ten milliseconds. If a new job arrives before n units is reached, the thread stops spin-waiting and starts doing the new job. If `SUNW_MP_THR_IDLE` contains an illegal value or is not set, `spin` is used as the default.

Global Behavior: Value Versus unsigned Preserving

A program that depends on unsigned preserving arithmetic conversions behaves differently. This is considered to be the most serious change made by ANSI/ISO C.

In the first edition of K&R, *The C Programming Language* (Prentice-Hall, 1978), unsigned specified exactly one type; there were no unsigned chars, unsigned shorts, or unsigned longs, but most C compilers added these very soon thereafter.

In K&R C compilers, the unsigned preserving rule is used for promotions: when an unsigned type needs to be widened, it is widened to an unsigned type; when an unsigned type mixes with a signed type, the result is an unsigned type.

The other rule, specified by ANSI/ISO C, came to be called “value preserving,” in which the result type depends on the relative sizes of the operand types. When an unsigned char or unsigned short is widened, the result type is `int` if an `int` is large enough to represent all the values of the smaller type. Otherwise, the result type is unsigned `int`. The value preserving rule produces the least surprise arithmetic result for most expressions.

Only in the `-xt` and `-xs` modes does the compiler use the unsigned preserving promotions; in the other modes, `-xc` and `-xa`, the value preserving promotion rules are used. When the `-xtransition` option is used, the compiler warns about each expression whose behavior might depend on the promotion rules used.

Keywords

`__asm` Keyword

The `__asm` keyword (note the initial double-underscore) is a synonym for the `asm` keyword. If you use `asm`, rather than `__asm`, and compile in `-xc` mode, the compiler issues a warning. The compiler does not issue a warning if you use `__asm` in `-xc` mode. The `__asm` statement has the form:

```
__asm( "string" );
```

where *string* is a valid assembly language statement. The `__asm` statements must appear within function bodies.

`_Restrict` Keyword

The C compiler supports `_Restrict` keyword as an equivalent to the `restrict` keyword in the C99 standard. For more information on supported C99 features, see "Supported Features of C99" on page 319.

long long Data Type

The Sun ANSI/ISO C compiler includes the data types `long long`, and `unsigned long long`, which are similar to the data type `long`. `long long` can store 64 bits of information; `long` can store 32 bits of information. `long long` is not available in `-xc` mode.

Printing long long Data Types

To print or scan long long data types, prefix the conversion specifier with the letters "ll." For example, to print `llvar`, a variable of long long data type, in signed decimal format, use:

```
printf("%lld\n", llvar);
```

Usual Arithmetic Conversions

Some binary operators convert the types of their operands to yield a common type, which is also the type of the result. These are called the usual arithmetic conversions:

- If either operand is type `long double`, the other operand is converted to `long double`.
- Otherwise, if either operand has type `double`, the other operand is converted to `double`.
- Otherwise, if either operand has type `float`, the other operand is converted to `float`.
- Otherwise, the integral promotions are performed on both operands. Then, these rules are applied:
 - If either operand has type `unsigned long long int`, the other operand is converted to `unsigned long long int`.
 - If either operand has type `long long int`, the other operand is converted to `long long int`.
 - If either operand has type `unsigned long int`, the other operand is converted to `unsigned long int`.
 - Otherwise, if one operand has type `long int` and the other has type `unsigned int`, both operands are converted to `unsigned long int`.
 - Otherwise, if either operand has type `long int`, the other operand is converted to `long int`.
 - Otherwise, if either operand has type `unsigned int`, the other operand is converted to `unsigned int`.
 - Otherwise, both operands have type `int`.

Constants

This section contains information related to constants that is specific to the Sun ANSI/ISO C compiler.

Integral Constants

Decimal, octal, and hexadecimal integral constants can be suffixed to indicate type, as shown in the following table.

TABLE 3-1 Data Type Suffixes

Suffix	Type
u or U	unsigned
l or L	long
ll or LL	long long ¹
lu, LU, Lu, lU, ul, uL, Ul, or UL	unsigned long
llu, LLU, LLu, llU, ull, ULL, uLL, Ull	unsigned long long ¹

1. The long long and unsigned long long are not available in -Xc mode.

When assigning types to unsuffixed constants, the compiler uses the first of this list in which the value can be represented, depending on the size of the constant:

- int
- long int
- unsigned long int
- long long int
- unsigned long long int

Character Constants

A multiple-character constant that is not an escape sequence has a value derived from the numeric values of each character. For example, the constant `'123'` has a value of:

0	'3'	'2'	'1'
---	-----	-----	-----

or `0x333231`.

With the `-xs` option and in other, non-ANSI/ISO versions of C, the value is:

0	'1'	'2'	'3'
---	-----	-----	-----

or `0x313233`.

Include Files

To include any of the standard header files supplied with the C compilation system, use this format:

```
#include <stdio.h>
```

The angle brackets (`<>`) cause the preprocessor to search for the header file in the standard place for header files on your system, usually the `/usr/include` directory.

The format is different for header files that you have stored in your own directories:

```
#include "header.h"
```

For statements of the form `#include "foo.h"` (where quotation marks are used), the compiler searches for include files in the following order:

1. The current directory (that is, the directory containing the “including” file)
2. The directories named with `-I` options, if any

3. The `/usr/include` directory

If your header file is not in the same directory as the source files that include it, specify the path of the directory in which it is stored with the `-I` option to `cc`. Suppose, for instance, that you have included both `stdio.h` and `header.h` in the source file `mycode.c`:

```
#include <stdio.h>
#include "header.h"
```

Suppose further that `header.h` is stored in the directory `../defs`. The command:

```
% cc -I../defs mycode.c
```

directs the preprocessor to search for `header.h` first in the directory containing `mycode.c`, then in the directory `../defs`, and finally in the standard place. It also directs the preprocessor to search for `stdio.h` first in `../defs`, then in the standard place. The difference is that the current directory is searched only for header files whose names you have enclosed in quotation marks.

You can specify the `-I` option more than once on the `cc` command-line. The preprocessor searches the specified directories in the order they appear. You can specify multiple options to `cc` on the same command-line:

```
% cc -o prog -I../defs mycode.c
```

Using the `-I-` Option to Change the Search Algorithm

The new `-I-` option gives more control over the default search rules. When `-I-` appears in the command line:

- The compiler never searches the current directory, unless the directory is listed explicitly in a `-I` directive. This effect applies even for include statements of the form `#include "foo.h"`.
- For include statements of the form `#include "foo.h"`, the compiler searches for include files in the following order:
 - a. The directories named with `-I` options (both before and after `-I-`)
 - b. The `/usr/include` directory
- For include statements of the form `#include <foo.h>`, the compiler searches for include files in the following order:

- a. The directories named with `-I` that appear after `-I-` (that is, the compiler does not search the `-I` directories that appear before `-I-`)
- b. The `/usr/include` directory

The following example shows the results of using `-I-` when compiling `prog.c`.

<code>prog.c</code>	<pre>#include "a.h" #include <b.h> #include "c.h"</pre>
<code>c.h</code>	<pre>#ifndef _C_H_1 #define _C_H_1 int c1; #endif</pre>
<code>inc/a.h</code>	<pre>#ifndef _A_H #define _A_H #include "c.h" int a; #endif</pre>
<code>inc/b.h</code>	<pre>#ifndef _B_H #define _B_H #include <c.h> int b; #endif</pre>
<code>inc/c.h</code>	<pre>#ifndef _C_H_2 #define _C_H_2 int c2; #endif</pre>

The following command shows the default behavior of searching the current directory (the directory of the including file) for include statements of the form `#include "foo.h"`. When processing the `#include "c.h"` statement in `inc/a.h`, the preprocessor includes the `c.h` header file from the `inc` subdirectory. When

processing the `#include "c.h"` statement in `prog.c`, the preprocessor includes the `c.h` file from the directory containing `prog.c`. Note that the `-H` option instructs the compiler to print the paths of the included files.

```
example% cc -c -Iinc -H prog.c
inc/a.h
        inc/c.h
inc/b.h
        inc/c.h
c.h
```

The next command shows the effect of the `-I-` option. The preprocessor does not look in the including directory first when it processes statements of the form `#include "foo.h"`. Instead, it searches the directories named by the `-I` options in the order that they appear in the command line. When processing the `#include "c.h"` statement in `inc/a.h`, the preprocessor includes the `./c.h` header file instead of the `inc/c.h` header file.

```
example% cc -c -I. -I- -Iinc -H prog.c
inc/a.h
        ./c.h
inc/b.h
        inc/c.h
./c.h
```

For more information, see “`-I[-|dir]`” on page 32.

Nonstandard Floating Point

IEEE 754 floating-point default arithmetic is “nonstop.” Underflows are “gradual.” The following is a summary, see the *Numerical Computation Guide* for details.

Nonstop means that execution does not halt on occurrences like division by zero, floating-point overflow, or invalid operation exceptions. For example, consider the following, where x is zero and y is positive:

$$z = y / x;$$

By default, z is set to the value `+Inf`, and execution continues. With the `-fnonstd` option, however, this code causes an exit, such as a core dump.

Here is how gradual underflow works. Suppose you have the following code:

```
x = 10;
for (i = 0; i < LARGE_NUMBER; i++)
x = x / 10;
```

The first time through the loop, x is set to 1; the second time through, to 0.1; the third time through, to 0.01; and so on. Eventually, x reaches the lower limit of the machine's capacity to represent its value. What happens the next time the loop runs?

Let's say that the smallest number characterizable is $1.234567e-38$

The next time the loop runs, the number is modified by "stealing" from the mantissa and "giving" to the exponent so the new value is $1.23456e-39$ and, subsequently, $1.2345e-40$ and so on. This is known as "gradual underflow," which is the default behavior. In nonstandard mode, none of this "stealing" takes place; typically, x is simply set to zero.

Preprocessing Directives and Names

This section describes assertions, pragmas, and predefined names.

Assertions

A line of the form:

```
#assert predicate (token-sequence)
```

associates the *token-sequence* with the predicate in the assertion name space (separate from the space used for macro definitions). The predicate must be an identifier token.

```
#assert predicate
```

asserts that *predicate* exists, but does not associate any token sequence with it.

The compiler provides the following predefined predicates by default (not in `-Xc` mode):

```
#assert system (unix)
#assert machine (sparc)(SPARC)
#assert machine (i386)(Intel)
#assert cpu (sparc)(SPARC)
#assert cpu (i386)(Intel)
```

`lint` provides the following predefinition predicate by default (not in `-Xc` mode):

```
#assert lint (on)
```

Any assertion may be removed by using `#unassert`, which uses the same syntax as `assert`. Using `#unassert` with no argument deletes all assertions on the predicate; specifying an assertion deletes only that assertion.

An assertion may be tested in a `#if` statement with the following syntax:

```
#if #predicate(non-empty token-list)
```

For example, the predefined predicate `system` can be tested with the following line:

```
#if #system(unix)
```

which evaluates true.

Pragmas

Preprocessing lines of the form:

```
#pragma pp-tokens
```

specify implementation-defined actions.

The following `#pragmas` are recognized by the compilation system. The compiler ignores unrecognized `pragmas`. Using the `-v` option will give a warning for unrecognized `pragmas`.

`#pragma align integer (variable[, variable])`

The align pragma makes all the mentioned variables memory aligned to *integer* bytes, overriding the default. The following limitations apply:

- The *integer* value must be a power of 2 between 1 and 128; valid values are: 1, 2, 4, 8, 16, 32, 64, and 128.
- *variable* is a global or static variable; it cannot be an automatic variable.
- If the specified alignment is smaller than the default, the default is used.
- The pragma line must appear before the declaration of the variables which it mentions; otherwise, it is ignored.
- Any variable that is mentioned but not declared in the text following the pragma line is ignored. For example:

```
#pragma align 64 (aninteger, astring, astruct)
int aninteger;
static char astring[256];
struct astruct{int a; char *b};
```

`#pragma does_not_read_global_data (funcname [, funcname])`

This pragma asserts that the specified list of routines do not read global data directly or indirectly. This allows for better optimization of code around calls to such routines. In particular, assignment statements or stores could be moved around such calls.

This pragma is permitted only after the prototype for the specified functions are declared. If the assertion about global access is not true, then the behavior of the program is undefined.

`#pragma does_not_return (funcname [, funcname])`

This pragma is an assertion to the compiler backend that the calls to the specified routines will not return. This allows the optimizer to perform optimizations consistent with that assumption. For example, register life-times will terminate at the call sites which in turn allows more optimizations.

If the specified function does return, then the behavior of the program is undefined.

This pragma is permitted only after the prototype for the specified functions are declared as the following example shows:

```
extern void exit(int);
#pragma does_not_return(exit);

extern void __assert(int);
#pragma does_not_return(__assert);
```

`#pragma does_not_write_global_data` (*funcname* [, *funcname*])

This pragma asserts that the specified list of routines do not write global data directly or indirectly. This allows for better optimization of code around calls to such routines. In particular, assignment statements or stores could be moved around such calls.

This pragma is permitted only after the prototype for the specified functions are declared. If the assertion about global access is not true, then the behavior of the program is undefined.

`#pragma error_messages` (*on* | *off* | *default*, *tag... tag*)

The error message pragma provides control within the source program over the messages issued by the C compiler and lint. For the C compiler, the pragma has an effect on warning messages only. The `-w` option of the C compiler overrides this pragma by suppressing all warning messages.

- `#pragma error_messages` (*on*, *tag... tag*)

The `on` option ends the scope of any preceding `#pragma error_messages` option, such as the `off` option, and overrides the effect of the `-erroff` option.

- `#pragma error_messages` (*off*, *tag... tag*)

The `off` option prevents the C compiler or the lint program from issuing the given messages beginning with the token specified in the pragma. The scope of the pragma for any specified error message remains in effect until overridden by another `#pragma error_messages`, or the end of compilation.

- `#pragma error_messages` (*default*, *tag... tag*)

The `default` option ends the scope of any preceding `#pragma error_messages` directive for the specified tags.

`#pragma fini (f1[, f2...fn])`

Causes the implementation to call functions *f1* to *fn* (finalization functions) after it calls `main()` routine. Such functions are expected to be of type `void` and to accept no arguments, and are called either when a program terminates under program control or when the containing shared object is removed from memory. As with “initialization functions,” finalization functions are executed in the order processed by the link editors.

`#pragma ident string`

Places *string* in the `.comment` section of the executable.

`#pragma init (f1[, f2...fn])`

Causes the implementation to call functions *f1* to *fn* (initialization functions) before it calls `main()`. Such functions are expected to be of type `void` and to accept no arguments, and are called while constructing the memory image of the program at the start of execution. In the case of initializers in a shared object, they are executed during the operation that brings the shared object into memory, either program start-up or some dynamic loading operation, such as `dlopen()`. The only ordering of calls to initialization functions is the order in which they were processed by the link editors, both static and dynamic.

`#pragma [no_]inline (funcname[, funcname])`

This pragma controls the inlining of routine names listed in the argument of the pragma. The scope of this pragma is over the entire file. Only global inlining control is allowed, call-site specific control is not permitted by this pragma.

If you use `#pragma inline`, it provides a suggestion to the compiler to inline the calls in the current file that match the list of routines listed in the pragma. This suggestion may be ignored under certain cases. For example, the suggestion is ignored when the body of the function is in a different module and the `crossfile` option is not used.

If you use `#pragma no_inline`, it provides a suggestion to the compiler to not inline the calls in the current file that match the list of routines listed in the pragma.

Both `#pragma inline` and `#pragma no_inline` are permitted only after the prototype for the specified functions are declared as the following example shows:

```
static void foo(int);
static int bar(int, char *);
#pragma inline(foo, bar)
```

`#pragma int_to_unsigned` (*funcname*)

For a function that returns a type of unsigned, in `-xt` or `-xs` mode, changes the function return to be of type `int`.

(SPARC) `#pragma MP serial_loop`

Refer to “Serial Pragmas” on page 129 for details.

(SPARC) `#pragma MP serial_loop_nested`

Refer to “Serial Pragmas” on page 129 for details.

(SPARC) `#pragma MP taskloop`

Refer to “Parallel Pragma” on page 129 for details.

(SPARC) `#pragma nomemorydepend`

This pragma specifies that for any iteration of a loop, there are no memory dependences. That is, within any iteration of a loop there are no references to the same memory. This pragma will permit the compiler (pipeliner) to schedule instructions, more effectively, within a single iteration of a loop. If any memory dependences exist within any iteration of a loop, the results of executing the program are undefined. The pragma applies to the next `for` loop within the current block. The compiler takes advantage of this information at optimization level of 3 or above.

(SPARC) `#pragma no_side_effect(funcname[, funcname...])`

funcname specifies the name of a function within the current translation unit. The function must be declared prior to the pragma. The pragma must be specified prior to the function's definition. For the named function, *funcname*, the pragma declares that the function has no side effects of any kind. This means that *funcname* returns a result value that depends only on the passed arguments. In addition, *funcname* and any called descendants:

- Do not access for reading or writing any part of the program state visible in the caller at the point of the call.
- Do not perform I/O.
- Do not change any part of the program state not visible at the point of the call.

The compiler can use this information when doing optimizations using the function. If the function does have side effects, the results of executing a program which calls this function are undefined. The compiler takes advantage of this information at optimization level of 3 or above.

`#pragma opt level (funcname[, funcname])`

The value of *level* specifies the optimization level for the *funcname* subprograms. You can assign opt levels 0, 1, 2, 3, 4, 5. You can turn off optimization by setting *level* to 0. The *funcname* subprograms must be prototyped prior to the pragma.

The level of optimization for any function listed in the pragma is reduced to the value of `-xmaxopt`. The pragma is ignored when `-xmaxopt=off`.

`#pragma pack(n)`

Use `#pragma pack(n)`, to affect member packing of a structure. By default, members of a structure are aligned on their natural boundaries; one byte for a char, two bytes for a short, four bytes for an integer etc. If *n* is present, it must be a power of 2 specifying the strictest natural alignment for any structure member. Zero is not accepted.

You can use `#pragma pack(n)` to specify an alignment boundary for a structure member. For example, `#pragma pack(2)` aligns int, long, long long, float, double, long double, and pointers on two byte boundaries instead of their natural alignment boundaries.

If *n* is the same or greater than the strictest alignment on your platform, (four on Intel, eight on SPARC v8, and 16 on SPARC v9), the directive has the effect of natural alignment. Also, if *n* is omitted, member alignment reverts to the natural alignment boundaries.

The `#pragma pack(n)` directive applies to all structure definitions which follow it until the next `pack` directive. If the same structure is defined in different translation units with different packing, your program may fail in unpredictable ways. In particular, you should not use `#pragma pack(n)` prior to including a header that defines the interface of a precompiled library. The recommended usage of `#pragma pack(n)` is to place it in your program code immediately before any structure to be packed. Follow the packed structure immediately with `#pragma pack()`.

Note – If you use `#pragma pack` to align struct members on boundaries other than their natural boundaries, accessing these fields may lead to a bus error on SPARC. See “`-xmemalign=ab`” on page 62, for the optimal way to compile such programs.

(SPARC) `#pragma pipelooop(n)`

This pragma accepts a positive constant integer value, or 0, for the argument *n*. This pragma specifies that a loop is pipelinable and the minimum dependence distance of the loop-carried dependence is *n*. If the distance is 0, then the loop is effectively a Fortran-style `doall` loop and should be pipelined on the target processors. If the distance is greater than 0, then the compiler (pipeliner) will only try to pipeline *n* successive iterations. The pragma applies to the next `for` loop within the current block. The compiler takes advantage of this information at optimization level of 3 or above.

`#pragma rarely_called(funcname[, funcname])`

This pragma provides a hint to the compiler backend that the specified functions are called infrequently. This allows the compiler to perform profile-feedback style optimizations on the call-sites of such routines without the overhead of a profile-collections phase. Since this pragma is a suggestion, the compiler optimizer may not perform any optimizations based on this pragma.

The `#pragma rarely_called` preprocessor directive is only permitted after the prototype for the specified functions are declares. The following is an example of `#pragma rarely_called`:

```
extern void error (char *message);
#pragma rarely_called(error);
```

```
#pragma redefine_extname old_extname new_extname
```

This pragma causes every externally defined occurrence of the name *old_extname* in the object code to be replaced by *new_extname*. As a result, the linker only sees the name *new_extname* at link time. If #pragma redefine_extname is encountered after the first use of *old_extname*, as a function definition, an initializer, or an expression, the effect is undefined. (This pragma is not supported in -Xs mode.)

When #pragma redefine_extname is available, the compiler provides a definition of the predefined macro PRAGMA_REDEFINE_EXTNAME which lets you write portable code that works both with and without #pragma redefine_extname.

The purpose of #pragma redefine_extname is to allow an efficient means of redefining a function interface when the name of the function cannot be changed. For example, when the original function definition must be maintained in a library, for compatibility with existing programs, along with a new definition of the same function for use by new programs. This can be accomplished by adding the new

function definition to the library by a new name. Consequently, the header file that declares the function uses `#pragma redefine_extname` so that all of the uses of the function are linked with the new definition of that function.

```
#if    defined(__STDC__)

#ifdef __PRAGMA_REDEFINE_EXTNAME
extern int myroutine(const long *, int *);
#pragma redefine_extname myroutine __fixed_myroutine
#else /* __PRAGMA_REDEFINE_EXTNAME */

static int
myroutine(const long * arg1, int * arg2)
{
    extern int __myroutine(const long *, int*);
    return (__myroutine(arg1, arg2));
}
#endif /* __PRAGMA_REDEFINE_EXTNAME */

#else /* __STDC__ */

#ifdef __PRAGMA_REDEFINE_EXTNAME
extern int myroutine();
#pragma redefine_extname myroutine __fixed_myroutine
#else /* __PRAGMA_REDEFINE_EXTNAME */

static int
myroutine(arg1, arg2)
    long *arg1;
    int *arg2;
{
    extern int __fixed_myroutine();
    return (__fixed_myroutine(arg1, arg2));
}
#endif /* __PRAGMA_REDEFINE_EXTNAME */

#endif /* __STDC__ */
```

`#pragma returns_new_memory` (*funcname* [, *funcname*])

This pragma asserts that the return value of the specified functions does not alias with any memory at the call site. In effect, this call returns a new memory location. This information allows the optimizer to better track pointer values and clarify memory location. This results in improved scheduling, pipelining, and parallelization of loops. However, if the assertion is false, the behavior of the program is undefined.

This pragma is permitted only after the prototype for the specified functions are declared as the following example shows:

```
void *malloc(unsigned);  
#pragma returns_new_memory(malloc);
```

`#pragma unknown_control_flow (name[, name])`

In order to describe procedures that alter the flow graphs of their callers, the C compiler provides the `#pragma unknown_control_flow` directive. Typically, this directive accompanies declarations of functions like `setjmp()`. On Sun systems, the include file `<setjmp.h>` contains the following:

```
extern int setjmp();  
#pragma unknown_control_flow(setjmp);
```

Other functions with properties like those of `setjmp()` must be declared similarly.

In principle, an optimizer that recognizes this attribute could insert the appropriate edges in the control flow graph, thus handling function calls safely in functions that call `setjmp()`, while maintaining the ability to optimize code in unaffected parts of the flow graph.

(SPARC) `#pragma unroll (unroll_factor)`

This pragma accepts a positive constant integer value for the argument *unroll_factor*. The pragma applies to the next `for` loop within the current block. For unroll factor other than 1, this directive serves as a suggestion to the compiler that the specified loop should be unrolled by the given factor. The compiler will, when possible, use that unroll factor. When the unroll factor value is 1, this directive serves as a command which specifies to the compiler that the loop is not to be unrolled. The compiler takes advantage of this information at optimization level of 3 or above.

```
#pragma weak symbol1 [= symbol2]
```

Defines a weak global symbol. This pragma is used mainly in source files for building libraries. The linker does not produce an error message if it is unable to resolve a weak symbol.

```
#pragma weak symbol
```

defines *symbol* to be a weak symbol. The linker does not produce an error message if it does not find a definition for *symbol*.

```
#pragma weak symbol1 = symbol2
```

defines *symbol1* to be a weak symbol, which is an alias for the symbol *symbol2*. This form of the pragma can only be used in the same translation unit where *symbol2* is defined, either in the sourcefiles or one of its included headerfiles. Otherwise, a compilation error will result.

If your program calls but does not define *symbol1*, and *symbol1* is a weak symbol in a library being linked, the linker uses the definition from that library. However, if your program defines its own version of *symbol1*, then the program's definition is used and the weak global definition of *symbol1* in the library is not used. If the program directly calls *symbol2*, the definition from the library is used; a duplicate definition of *symbol2* causes an error.

Predefined Names

The following identifier is predefined as an object-like macro:

TABLE 3-2 Predefined Identifier

Identifier	Description
__STDC__	__STDC__ 1 -xc __STDC__ 0 -xa, -xt Not defined -xs

The compiler issues a warning if __STDC__ is undefined (#undef __STDC__).
__STDC__ is not defined in -xs mode.

Predefinitions (not valid in `-xc` mode):

- `sun`
- `unix`
- `sparc` (*SPARC*)
- `i386` (*Intel*)

The following predefinitions are valid in all modes:

- `__sun`
- `__unix`
- `__SUNPRO_C=0x530`
- `__'uname -s'_'uname -r'` (example: `__SunOS_5_7`)
- `__sparc` (*SPARC*)
- `__i386` (*Intel*)
- `__BUILTIN_VA_ARG_INCR`
- `__SVR4`
- `__sparcv9` (`-Xarch=v9, v9a`)

The compiler also predefines the object-like macro `__PRAGMA_REDEFINE_EXTNAME`

to indicate that the pragma will be recognized. The following is predefined in `-xa` and `-xt` modes only:

`__RESTRICT`

Parallelizing Sun ANSI/ISO C Code

The Sun ANSI/ISO C compiler can optimize code to run on SPARC shared-memory multiprocessor machines. The process is called *parallelizing*. The compiled code can execute in parallel using the multiple processors on the system. This chapter explains how you can take advantage of the compiler's parallelizing features and is organized into the following sections:

- "Overview" on page 105
- "Parallelizing For OpenMP" on page 106
- "Environment Variables" on page 107
- "Data Dependence and Interference" on page 109
- "Speedups" on page 116
- "Load Balance and Loop Scheduling" on page 121
- "Loop Transformations" on page 122
- "Aliasing and Parallelization" on page 126

Overview

The C compiler generates parallel code for those loops that it determines are safe to parallelize. Typically, these loops have iterations that are independent of each other. For such loops, it does not matter in what order the iterations are executed or if they are executed in parallel. Many, though not all, vector loops fall into this category.

Because of the way aliasing works in C, it is difficult to determine the safety of parallelization. To help the compiler, Sun ANSI/ISO C offers pragmas and additional pointer qualifications to provide aliasing information known to the programmer that the compiler cannot determine.

Example of Use

The following example illustrates how to enable and control parallelized C:

```
% cc -fast -xO4 -xautopar example.c -o example
```

This generates an executable called `example`, which can be executed normally. If you wish to take advantage of multiprocessor execution, see “`-xautopar`” on page 47.

Parallelizing for OpenMP

You can compile your code so that it complies with the OpenMP specification. For more information on the OpenMP specification for C, visit the web site at <http://www.openmp.org/specs/>.

To take advantage of the compiler’s OpenMP support, you need to do the following:

- Issue the compiler’s `-xopenmp` option. See “`-xopenmp [=i]`” on page 66.

Handling OpenMP Runtime Warnings

The OpenMP runtime system can issue warnings for non-fatal errors. Use the following function to register a call back function to handle these warnings:

```
int sunw_mp_register_warn(void (*func) (void *))
```

You can access the prototype for this function by issuing a `#include` preprocessor directive for `<sunw_mp_misc.h>`.

If you do not want to register a function, set the environment variable `SUNW_MP_WARN` to `TRUE` and the warning messages are sent to `stderr`. For more information on `SUNW_MP_WARN`, see “`SUNW_MP_WARN`” on page 108.

For information specific to this implementation of OpenMP, see “Implementation-Specific Information of OpenMP” on page 335.

Environment Variables

There are four environment variables that relate to parallelized C:

- PARALLEL
- SUNW_MP_THR_IDLE
- SUNW_MP_WARN
- STACKSIZE

PARALLEL

Set the PARALLEL environment variable if you can take advantage of multiprocessor execution. The PARALLEL environment variable specifies the number of processors available to the program. The following example shows that PARALLEL is set to two:

```
% setenv PARALLEL 2
```

If the target machine has multiple processors, the threads can map to independent processors. Running the program leads to the creation of two threads that execute the parallelized portions of the program.

SUNW_MP_THR_IDLE

Currently, the starting thread of a program creates bound threads. Once created, these bound threads participate in executing the parallel part of a program (parallel loop, parallel region, etc.) and keep spin-waiting while the sequential part of the program runs. These bound threads never sleep or stop until the program terminates. Having these threads spin-wait generally gives the best performance when a parallelized program runs on a dedicated system. However, threads that are spin-waiting use system resources.

Use the SUNW_MP_THR_IDLE environment variable to control the status of each thread after it finishes its share of a parallel job.

```
% setenv SUNW_MP_THR_IDLE value
```

You can substitute either `spin` or `sleep[n s|n ms]` for *value*. The default is `spin`, which means the thread should spin (or busy-wait) after completing a parallel task, until a new parallel task arrives.

The other choice, `sleep[n s | n ms]` puts the thread to sleep after spin-waiting n units. The wait unit can be seconds (`s`, the default unit) or milliseconds (`ms`), where `1s` means one second, `10ms` means ten milliseconds. `sleep` with no arguments puts the thread to sleep immediately after completing a parallel task. `sleep`, `sleep0`, `sleep0s`, and `sleep0ms` are all equivalent.

If a new job arrives before n units is reached, the thread stops spin-waiting and starts doing the new job. If `SUNW_MP_THR_IDLE` contains an illegal value or isn't set, `spin` is used as the default.

SUNW_MP_WARN

Set this environment variable to `TRUE` to print warning messages from OpenMP and other parallelization runtime-systems.

```
% setenv SUNW_MP_WARN TRUE
```

If you registered a function by using `sunw_mp_register_warn()` to handle warning messages, then `SUNW_MP_WARN` prints no warning messages, even if you set it to `TRUE`. If you did not register a function and set `SUNW_MP_WARN` to `TRUE`, `SUNW_MP_WARN` prints the warning messages to `stderr`. If you do not register a function and you do not set `SUNW_MP_WARN`, no warning messages are issued.. For more information on `sunw_mp_register_warn()` see "Handling OpenMP Runtime Warnings" on page 106.

STACKSIZE

The executing program maintains a main memory stack for the master thread and distinct stacks for each slave thread. Stacks are temporary memory address spaces used to hold arguments and automatic variables over subprogram invocations.

The default size of the main stack is about eight megabytes. Use the `limit` command to display the current main stack size as well as set it.

```
% limit
cputime unlimited
filesize unlimited
datasize 2097148 kbytes
stacksize 8192 kbytes <- current main stack size
coredumpsize 0 kbytes
descriptors 256
memorysize unlimited
% limit stacksize 65536 <- set main stack to 64Mb
```


Each slave thread of a multithreaded program has its own thread stack. This stack mimics the main stack of the master thread but is unique to the thread. The thread's private arrays and variables (local to the thread) are allocated on the thread stack. All slave threads have the same stack size, which is one megabyte for 32-bit applications and two megabytes for 64-bit applications by default. The size is set with the `STACKSIZE` environment variable:

```
% setenv STACKSIZE 8192 <- Set thread stack size to 8 Mb
```

Setting the thread stack size to a value larger than the default may be necessary for most parallelized code.

Sometimes the compiler may generate a warning message that indicates a bigger stack size is needed. However, it may not be possible to know just how large to set it, except by trial and error, especially if private/local arrays are involved. If the stack size is too small for a thread to run, the program will abort with a segmentation fault.

Keyword

The keyword `_Restrict` can be used with parallelized C. Refer to the section “`_Restrict` Keyword” on page 86 for details.

Data Dependence and Interference

The C compiler performs analysis on loops in programs to determine if it is safe to execute different iterations of the loops in parallel. The purpose of this analysis is to determine if any two iterations of the loop could interfere with each other. Typically this happens if one iteration of a variable could read a variable while another iteration is writing the very same variable. Consider the following program fragment:

CODE EXAMPLE 4-1 A Loop With Dependence

```
for (i=1; i < 1000; i++) {  
    sum = sum + a[i]; /* S1 */  
}
```

In CODE EXAMPLE 4-1 any two successive iterations, i and $i+1$, will write and read the same variable `sum`. Therefore, in order for these two iterations to execute in parallel some form of locking on the variable would be required. Otherwise it is not safe to allow the two iterations to execute in parallel.

However, the use of locks imposes overhead that might slowdown the program. The C compiler will not ordinarily parallelize the loop in CODE EXAMPLE 4-1. In CODE EXAMPLE 4-1 there is a data dependence between two iterations of the loop. Consider another example:

CODE EXAMPLE 4-2 A Loop Without Dependence

```
for (i=1; i < 1000; i++) {
    a[i] = 2 * a[i]; /* S1 */
}
```

In this case each iteration of the loop references a different array element. Therefore different iterations of the loop can be executed in any order. They may be executed in parallel without any locks because no two data elements of different iterations can possibly interfere.

The analysis performed by the compiler to determine if two different iterations of a loop could reference the same variable is called data dependence analysis. Data dependences prevent loop parallelization if one of the references writes to the variable. The dependence analysis performed by the compiler can have three outcomes:

- There is a dependence. In this case, it is not safe to execute the loop in parallel. CODE EXAMPLE 4-1 illustrates this case.
- There is no dependence. The loop may safely execute in parallel using an arbitrary number of processors. CODE EXAMPLE 4-2 illustrates this case.
- The dependence cannot be determined. The compiler assumes, for safety, that there might be a dependence that prevents parallel execution of the loop and will not parallelize the loop.

In CODE EXAMPLE 4-3, whether or not two iterations of the loop write to the same element of array `a` depends on whether or not array `b` contains duplicate elements. Unless the compiler can determine this fact, it assumes there is a dependence and does not parallelize the loop.

CODE EXAMPLE 4-3 A Loop That May or May Not Contain Dependencies

```
for (i=1; i < 1000; i++) {
    a[b[i]] = 2 * a[i];
}
```

Parallel Execution Model

The parallel execution of loops is performed by Solaris threads. The thread starting the initial execution of the program is called the master thread. At program start-up the master thread creates multiple slave threads as shown in the following figure. At the end of the program all the slave threads are terminated. Slave thread creation is performed exactly once to minimize the overhead.




FIGURE 4-1 Master and Slave Threads

After start-up, the master thread starts the execution of the program while slave threads wait idly. When the master thread encounters a parallel loop, different iterations of the loop are distributed among the slave and master threads which start the execution of the loop. After each thread finishes execution of its chunk it synchronizes with the remaining threads. This synchronization point is called a *barrier*. The master thread cannot continue executing the remainder of the program until all the threads have finished their work and reached the barrier. The slave threads go into a wait state after the barrier waiting for more parallel work, and the master thread continues to execute the program.

During this process, various overheads can occur:

- The overhead of synchronization and work distribution
- The overhead of barrier synchronization

In general, there may be some parallel loops for which the amount of useful work performed is not enough to justify the overhead. For such loops, there may be appreciable slowdown. In the following figure, a loop is parallelized. However the barriers, represented by horizontal bars, introduce significant overhead. The work between the barriers is performed serially or in parallel as indicated. The amount of time required to execute the loop in parallel is considerably less than the amount of time required to synchronize the master and slave threads at the barriers.




FIGURE 4-2 Parallel Execution of a Loop

Private Scalars and Private Arrays

There are some data dependences for which the compiler may still be able to parallelize a loop. Consider the following example.

CODE EXAMPLE 4-4 A Parallelizable Loop With Dependence

```
for (i=1; i < 1000; i++) {
    t = 2 * a[i];          /* S1 */
    b[i] = t;             /* S2 */
}
```

In this example, assuming that arrays *a* and *b* are non-overlapping arrays, there appears to be a data dependence in any two iterations due to the variable *t*. The following statements execute during iterations one and two.

CODE EXAMPLE 4-5 Iterations One and Two

```
t = 2*a[1]; /* 1 */
b[1] = t;   /* 2 */
t = 2*a[2]; /* 3 */
b[2] = t;   /* 4 */
```

Because statements one and three modify the variable `t`, the compiler cannot execute them in parallel. However, the value of `t` is always computed and used in the same iteration so the compiler can use a separate copy of `t` for each iteration. This eliminates the interference between different iterations due to such variables. In effect, we have made variable `t` as a private variable for each thread executing that iteration. This can be illustrated as follows:

CODE EXAMPLE 4-6 Variable `t` as a Private Variable for Each Thread

```
for (i=1; i < 1000; i++) {
    pt[i] = 2 * a[i];          /* S1 */
    b[i] = pt[i];            /* S2 */
}
```

CODE EXAMPLE 4-6 is essentially the same example as CODE EXAMPLE 4-3, but each scalar variable reference `t` is now replaced by an array reference `pt`. Each iteration now uses a different element of `pt`, and this results in eliminating any data dependencies between any two iterations. Of course one problem with this illustration is that it may lead to an extra large array. In practice, the compiler only allocates one copy of the variable for each thread that participates in the execution of the loop. Each such variable is, in effect, private to the thread.

The compiler can also privatize array variables to create opportunities for parallel execution of loops. Consider the following example:

CODE EXAMPLE 4-7 A Parallelizable Loop With an Array Variable

```
for (i=1; i < 1000; i++) {
    for (j=1; j < 1000; j++) {
        x[j] = 2 * a[i];      /* S1 */
        b[i][j] = x[j];      /* S2 */
    }
}
```

In CODE EXAMPLE 4-7, different iterations of the outer loop modify the same elements of array *x*, and thus the outer loop cannot be parallelized. However, if each thread executing the outer loop iterations has a private copy of the entire array *x*, then there would be no interference between any two iterations of the outer loop. This is illustrated as follows:

CODE EXAMPLE 4-8 A Parallelizable Loop Using a Privatized Array

```
for (i=1; i < 1000; i++) {
    for (j=1; j < 1000; j++) {
        px[i][j] = 2 * a[i];    /* S1 */
        b[i][j] = px[i][j];    /* S2 */
    }
}
```

As in the case of private scalars, it is not necessary to expand the array for all the iterations, but only up to the number of threads executing in the systems. This is done automatically by the compiler by allocating one copy of the original array in the private space of each thread.

Storeback

Privatization of variables can be very useful for improving the parallelism in the program. However, if the private variable is referenced outside the loop then the compiler needs to assure that it has the right value. Consider the following example:

CODE EXAMPLE 4-9 A Parallelized Loop Using Storeback

```
for (i=1; i < 1000; i++) {
    t = 2 * a[i];    /* S1 */
    b[i] = t;        /* S2 */
}
x = t;              /* S3 */
```

In CODE EXAMPLE 4-9 the value of *t* referenced in statement S3 is the final value of *t* computed by the loop. After the variable *t* has been privatized and the loop has finished executing, the right value of *t* needs to be stored back into the original variable. This is called storeback. This is done by copying the value of *t* on the final

iteration back to the original location of variable `t`. In many cases the compiler can do this automatically. But there are situations where the last value cannot be computed so easily:

CODE EXAMPLE 4-10 A Loop That Cannot Use Storeback

```
for (i=1; i < 1000; i++) {
    if (c[i] > x[i] ) {           /* C1 */
        t = 2 * a[i];           /* S1 */
        b[i] = t;               /* S2 */
    }
}
x = t*t;                         /* S3 */
```

For correct execution, the value of `t` in statement `S3` is not, in general, the value of `t` on the final iteration of the loop. It is in fact the last iteration for which the condition `C1` is true. Computing the final value of `t` is quite hard in the general cases. In cases like this the compiler will not parallelize the loop.

Reduction Variables

There are cases when there is a real dependence between iterations of a loop and the variables causing the dependence cannot simply be privatized. This can arise, for example, when values are being accumulated from one iteration to the next.

CODE EXAMPLE 4-11 A Loop That May or May Not Be Parallelized

```
for (i=1; i < 1000; i++) {
    sum += a[i]*b[i]; /* S1 */
}
```

In **CODE EXAMPLE 4-11**, the loop computes the vector product of two arrays into a common variable called `sum`. This loop cannot be parallelized in a simple manner. The compiler can take advantage of the associative nature of the computation in statement `S1` and allocate a private variable called `psum[i]` for each thread. Each copy of the variable `psum[i]` is initialized to 0. Each thread computes its own partial sum in its own copy of the variable `psum[i]`. Before crossing the barrier, all the partial sums are added onto the original variable `sum`. In this example, the variable `sum` is called a reduction variable because it computes a sum-reduction. However, one danger of promoting scalar variables to reduction variables is that the

manner in which rounded values are accumulated can change the final value of sum. The compiler performs this transformation only if you specifically give permission for it to do so.

Speedups

If the compiler does not parallelized a portion of a program where a significant amount of time is spent, then no speedup occurs. This is basically a consequence of Amdahls Law. For example, if a loop that accounts for five percent of the execution time of a program is parallelized, then the overall speedup is limited to five percent. However, there may not be any improvement depending on the size of the workload and parallel execution overheads.

As a general rule, the larger the fraction of program execution that is parallelized, the greater the likelihood of a speedup.

Each parallel loop incurs a small overhead during start-up and shutdown. The start overhead includes the cost of work distribution, and the shutdown overhead includes the cost of the barrier synchronization. If the total amount of work performed by the loop is not big enough then no speedup will occur. In fact the loop might even slow down. So if a large amount of program execution is accounted by a large number of short parallel loops, then the whole program may slow down instead of speeding up.

The compiler performs several loop transformations that try to increase the granularity of the loops. Some of these transformations are loop interchange and loop fusion. So in general, if the amount of parallelism in a program is small or is fragmented among small parallel regions, then the speedup is less.

Often scaling up a problem size improves the fraction of parallelism in a program. For example, consider a problem that consists of two parts: a quadratic part that is sequential, and a cubic part that is parallelizable. For this problem the parallel part of the workload grows faster than the sequential part. So at some point the problem will speedup nicely, unless it runs into resource limitations.

It is beneficial to try some tuning, experimentation with directives, problem sizes and program restructuring in order to achieve benefits from parallel C.

Amdahl's Law

Fixed problem-size speedup is generally governed by Amdahl's law. Amdahl's Law simply says that the amount of parallel speedup in a given problem is limited by the sequential portion of the problem. The following equation describes the speedup of a problem where F is the fraction of time spent in sequential region, and the remaining fraction of the time is spent uniformly among P processors. If the second term of the equation drops to zero, the total speedup is bounded by the first term, which remains fixed.

$$\frac{1}{S} = F + \frac{(1-F)}{P}$$

The following figure illustrates this concept diagrammatically. The darkly shaded portion represents the sequential part of the program, and remains constant for one, two, four, and eight processors, while the lightly shaded portion represents the parallel portion of the program that can be divided uniformly among arbitrary number of processors.




FIGURE 4-3 Fixed Problem Speedups

In reality, however, you may incur overheads due to communication and distribution of work to multiple processors. These overheads may or may not be fixed for arbitrary number of processors used.

FIGURE 4-4 illustrates the ideal speedups for a program continuing 0%, 2%, 5%, and 10% sequential portions. Here, no overhead is assumed.




FIGURE 4-4 Amdahl's Law Speedup Curve

Overheads

Once the overheads are incorporated in the model the speedup curves change dramatically. Just for the purposes of illustration we assume that overheads consist of two parts: a fixed part which is independent of the number of processors, and a non-fixed part that grows quadratically with the number of the processors used:

$$\frac{1}{S} = \frac{1}{F + \left(1 - \frac{F}{P}\right) + K_1 + K_2 P^2}$$

In this equation, K1 and K2 are some fixed factors. Under these assumptions the speedup curve is shown in the following figure. It is interesting to note that in this case the speedups peak out. After a certain point adding more processors is detrimental to performance as shown in the following figure.




FIGURE 4-5 Speedup Curve with Overheads

Gustafson's Law

Amdahls Law can be misleading for predicting parallel speedups in real problems. The fraction of time spent in sequential sections of the program sometimes depends on the problem size. That is, by scaling the problem size, you may improve the chances of speedup. The following example demonstrates this.

CODE EXAMPLE 4-12 Scaling the Problem Size May Improve Chances of Speedup

```
/*
 * initialize the arrays
 */
for (i=0; i < n; i++) {
    for (j=0; j < n; j++) {
        a[i][j] = 0.0;
        b[i][j] = ...
        c[i][j] = ...
    }
}
/*
 * matrix multiply
 */
for (i=0; i < n; i++) {
    for(j=0; j < n; j++) {
        for (k=0; k < n; k++) {
            a[i][j] = b[i][k]*c[k][j];
        }
    }
}
```

Assume an ideal overhead of zero and assume that only the second loop nest is executed in parallel. It is easy to see that for small problem sizes (i.e. small values of n), the sequential and parallel parts of the program are not so far from each other. However, as n grows larger, the time spent in the parallel part of the program grows faster than the time spent in the sequential part. For this problem, it is beneficial to increase the number of processors as the problem size increases.

Load Balance and Loop Scheduling

Loop scheduling is the process of distributing iterations of a parallel loop to multiple threads. In order to maximize the speedup, it is important that the work be distributed evenly among the threads while not imposing too much overhead. The compiler offers several types of scheduling for different situations.

Static or Chunk Scheduling

It is beneficial to divide the work evenly among the different threads on the system when the work performed by different iterations of a loop is the same. This approach is known as static scheduling.

CODE EXAMPLE 4-13 A Good Loop for Static Scheduling

```
for (i=1; i < 1000; i++) {  
    sum += a[i]*b[i];      /* S1 */  
}
```

Under static or chunk scheduling, each thread will get the same number of iterations. If there were 4 threads, then in the above example, each thread will get 250 iterations. Provided there are no interruptions and each thread progresses at the same rate, all the threads will complete at the same time.

Self Scheduling

Static scheduling will not achieve good load balance, in general, when the work performed by each iteration varies. In static scheduling, each thread grabs the same chunk of iterations. Each thread, except the master thread, upon completion of its chunk waits to participate in the next parallel loop execution. The master thread continues execution of the program. In self scheduling, each thread grabs a different small chunk of iteration and after completion of its assigned chunk, tries to acquire more chunks from the same loop.

Guided Self Scheduling

In guided self scheduling (GSS), each thread gets successively smaller number of chunks. In cases where the size of each iteration varies, GSS can help balance the load.

Loop Transformations

The compiler performs several loop restructuring transformations to help improve the parallelization of a loop in programs. Some of these transformations can also improve the single processor execution of loops as well. The transformations performed by the compiler are described below.

Loop Distribution

Often loops contain a few statements that cannot be executed in parallel and many statements that can be executed in parallel. Loop Distribution attempts to remove the sequential statements into a separate loop and gather the parallelizable statements into a different loop. This is illustrated in the following example:

CODE EXAMPLE 4-14 A Candidate for Loop Distribution

```
for (i=0; i < n; i++) {
    x[i] = y[i] + z[i]*w[i];           /* S1 */
    a[i+1] = (a[i-1] + a[i] + a[i+1])/3.0; /* S2 */
    y[i] = z[i] - x[i];               /* S3 */
}
```

Assuming that arrays *x*, *y*, *w*, *a*, and *z* do not overlap, statements S1 and S3 can be parallelized but statement S2 cannot be. Here is how the loop looks after it is split or distributed into two different loops:

CODE EXAMPLE 4-15 The Distributed Loop

```
/* L1: parallel loop */
for (i=0; i < n; i++) {
    x[i] = y[i] + z[i]*w[i];           /* S1 */
    y[i] = z[i] - x[i];               /* S3 */
}
/* L2: sequential loop */
for (i=0; i < n; i++) {
    a[i+1] = (a[i-1] + a[i] + a[i+1])/3.0; /* S2 */
}
```

After this transformation, loop L1 does not contain any statements that prevent the parallelization of the loop and may be executed in parallel. Loop L2, however, still has a non-parallelizable statement from the original loop.

Loop distribution is not always profitable or safe to perform. The compiler performs analysis to determine the safety and profitability of distribution.

Loop Fusion

If the granularity of a loop, or the work performed by a loop, is small, the performance gain from distribution may be insignificant. This is because the overhead of parallel loop start-up is too high compared to the loop workload. In such situations, the compiler uses loop fusion to combine several loops into a single parallel loop, and thus increase the granularity of the loop. Loop fusion is easy and safe when loops with identical trip counts are adjacent to each other. Consider the following example:

CODE EXAMPLE 4-16 Loops With Small Work Loads

```
/* L1: short parallel loop */
for (i=0; i < 100; i++) {
    a[i] = a[i] + b[i];           /* S1 */
}
/* L2: another short parallel loop */
for (i=0; i < 100; i++) {
    b[i] = a[i] * d[i];           /* S2 */
}
```

The two short parallel loops are next to each other, and can be safely combined as follows:

CODE EXAMPLE 4-17 The Two Loops Fused

```
/* L3: a larger parallel loop */
for (i=0; i < 100; i++) {
    a[i] = a[i] + b[i];      /* S1 */
    b[i] = a[i] * d[i];    /* S2 */
}
```

The new loop generates half the parallel loop execution overhead. Loop fusion can also help in other ways. For example if the same data is referenced in two loops, then combining them can improve the locality of reference.

However, loop fusion is not always safe to perform. If loop fusion creates a data dependence that did not exist before then the fusion may result in incorrect execution. Consider the following example:

CODE EXAMPLE 4-18 Unsafe Fusion Candidates

```
/* L1: short parallel loop */
for (i=0; i < 100; i++) {
    a[i] = a[i] + b[i];      /* S1 */
}
/* L2: a short loop with data dependence */
for (i=0; i < 100; i++) {
    a[i+1] = a[i] * d[i];   /* S2 */
}
```

If the loops in CODE EXAMPLE 4-18 are fused, a data dependence is created from statement S2 to S1. In effect, the value of `a[i]` in the right hand side of statement S1 is computed in statement S2. If the loops are not fused, this would not happen. The compiler performs safety and profitability analysis to determine if loop fusion should be done. Often, the compiler can fuse an arbitrary number of loops. Increasing the granularity in this manner can sometimes push a loop far enough up for it to be profitable for parallelization.

Loop Interchange

It is generally more profitable to parallelize the outermost loop in a nest of loops, since the overheads incurred are small. However, it is not always safe to parallelize the outermost loops due to dependences that might be carried by such loops. This is illustrated in the following:

CODE EXAMPLE 4-19 Nested Loop That Cannot be Parallelized

```
for (i=0; i <n; i++) {  
    for (j=0; j <n; j++) {  
        a[j][i+1] = 2.0*a[j][i-1];  
    }  
}
```

In this example, the loop with the index variable *i* cannot be parallelized, because of a dependency between two successive iterations of the loop. The two loops can be interchanged and the parallel loop (the *j*-loop) becomes the outer loop:

CODE EXAMPLE 4-20 The Loops Interchanged

```
for (j=0; j<n; j++) {  
    for (i=0; i<n; i++) {  
        a[j][i+1] = 2.0*a[j][i-1];  
    }  
}
```

The resulting loop incurs an overhead of parallel work distribution only once, while previously, the overhead was incurred *n* times. The compiler performs safety and profitability analysis to determine whether to perform loop interchange.

Aliasing and Parallelization

ANSI C aliasing can often prevent loops from getting parallelized. Aliasing occurs when there are two possible references to the same memory location. Consider the following example:

CODE EXAMPLE 4-21 A Loop With Two References to the Same Memory Location

```
void copy(float a[], float b[], int n) {
    int i;
    for (i=0; i < n; i++) {
        a[i] = b[i]; /* S1 */
    }
}
```

Since variables `a` and `b` are parameters, it is possible that `a` and `b` may be pointing to overlapping regions of memory. e.g if `copy` were called as follows:

```
copy (x[10], x[11], 20);
```

In the called routine, two successive iterations of the copy loop may be reading and writing the same element of the array `x`. However, if the routine `copy` were called as follows then there is no possibility of overlap in any of the 20 iterations of the loop:

```
copy (x[10], x[40], 20);
```

In general, it is not possible for the compiler to analyze this situation correctly without knowing how the routine is called. The compiler provides a keyword extension to ANSI/ISO C that lets you convey this kind of aliasing information. See “Restricted Pointers” on page 127 for more information.

Array and Pointer References

Part of the aliasing problem is that the C language can define array referencing and definition through pointer arithmetic. In order for the compiler to effectively parallelize loops, either automatically or explicitly with pragmas, all data that is laid out as an array must be referenced using C array reference syntax and not pointers.

If pointer syntax is used, the compiler cannot determine the relationship of the data between different iterations of a loop. Thus it will be conservative and not parallelize the loop.

Restricted Pointers

In order for a compiler to effectively perform parallel execution of a loop, it needs to determine if certain lvalues designate distinct regions of storage. Aliases are lvalues whose regions of storage are not distinct. Determining if two pointers to objects are aliases is a difficult and time consuming process because it could require analysis of the entire program. Consider function `vsq()` below:

CODE EXAMPLE 4-22 A Loop With Two Pointers

```
void vsq(int n, double * a, double * b) {
    int i;
    for (i=0; i<n; i++) {
        b[i] = a[i] * a[i];
    }
}
```

The compiler can parallelize the execution of the different iterations of the loops if it knows that pointers `a` and `b` access different objects. If there is an overlap in objects accessed through pointers `a` and `b` then it would be unsafe for the compiler to execute the loops in parallel. At compile time, the compiler does not know if the objects accessed by `a` and `b` overlap by simply analyzing the function `vsq()`; the compiler may need to analyze the whole program to get this information.

Restricted pointers are used to specify pointers which designate distinct objects so that the compiler can perform pointer alias analysis. To support restricted pointers, the keyword `_Restrict` is recognized by our ANSI C compiler as an extension. The following is an example of function `vsq()` in which function parameters are declared as restricted pointers:

```
void vsq(int n, double * _Restrict a, double * _Restrict b)
```

Pointers `a` and `b` are declared as restricted pointers, so the compiler knows that `a` and `b` point to distinct regions of storage. With this alias information, the compiler is able to parallelize the loop.

`_Restrict` is a type-qualifier, like `volatile`, and it shall only qualify pointer types. `_Restrict` is recognized as a keyword only for compilations modes `-Xa` (default) and `-xt`. For these two compilation modes, the compiler defines the macro

`__RESTRICT` so you can easily write portable code with restricted pointers. For example, the following code works on our compiler (all compilation modes) as well as on other compilers which do not support restricted pointers:

CODE EXAMPLE 4-23 Portable Code That Uses the `_Restrict` Keyword

```
#ifdef __RESTRICT
    #define restrict _Restrict
#else
    #define restrict
#endif
void vsq(int n, double * restrict a, double * restrict b)
{
    int i;
    for (i=0; i<n; i++)
        b[i] = a[i] * a[i];
}
```

Should restricted pointers become part of the ANSI/ISO C Standard, it is most likely that `restrict` will be the keyword. You may want to write code with restricted pointers that includes the following preprocessor directive as function `vsq()` does:

```
#define restrict _Restrict
```

This directive minimizes the changes you need to make now that `restrict` is a keyword in the ISO C Standard. We chose `_Restrict` as the keyword because it is in the implementors name space, so there is no conflict with identifiers in the users name space.

There are situations in which you may not want to change the source code. You can specify that pointer -valued function-parameters be treated as restricted pointers by using the following command line option:

```
-xrestrict=[func1, . . . ,funcn]
```

If a function list is specified, then pointer parameters in the specified functions are treated as restricted; otherwise, all pointer parameters in the entire C file are treated as restricted. For example, `-xrestrict=vsq`, qualifies the pointers `a` and `b` given in the first example of the function `vsq()` with the keyword `_Restrict`.

It is critical that you use `_Restrict` correctly. If pointers qualified as restricted pointers point to objects which are not distinct, the compiler can incorrectly parallelize loops resulting in undefined behavior. For example, assume that pointers `a` and `b` of function `vsq()` point to objects which overlap, such that `b[i]` and

`a[i+1]` are the same object. If `a` and `b` are not declared as restricted pointers the loops will be executed serially. If `a` and `b` are incorrectly qualified as restricted pointers the compiler may parallelize the execution of the loops, which is not safe, because `b[i+1]` should only be computed after `b[i]` had been computed.

Explicit Parallelization and Pragmas

Often, there is not enough information available for the compiler to make a decision on the legality or profitability of parallelization. Sun ANSI/ISO C supports pragmas that allow the programmer to effectively parallelize loops that otherwise would be too difficult or impossible for the compiler to handle.

Serial Pragmas

There are two serial pragmas, and both apply to “for” loops:

- `#pragma MP serial_loop`
- `#pragma MP serial_loop_nested`

The `#pragma MP serial_loop` pragma indicates to the compiler that the next for loop is not to be automatically parallelized.

The `#pragma MP serial_loop_nested` pragma indicates to the compiler that the next for loop and any for loops nested within the scope of this for loop are not to be automatically parallelized. The scope of the `serial_loop_nested` pragma does not extend beyond the scope of the loop to which it applies.

Parallel Pragma

There is one parallel pragma: `#pragma MP taskloop [options]`.

The `MP taskloop` pragma can, optionally, take one or more of the following arguments.

- `maxcpus` (*number_of_processors*)
- `private` (*list_of_private_variables*)
- `shared` (*list_of_shared_variables*)
- `readonly` (*list_of_readonly_variables*)
- `storeback` (*list_of_storeback_variables*)
- `savelast`
- `reduction` (*list_of_reduction_variables*)
- `schedtype` (*scheduling_type*)

Only one option can be specified per `MP taskloop` pragma; however, the pragmas are cumulative and apply to the next `for` loop encountered within the current block in the source code:

```
#pragma MP taskloop maxcpus(4)
#pragma MP taskloop shared(a,b)
#pragma MP taskloop storeback(x)
```

These options may appear multiple times prior to the `for` loop to which they apply. In case of conflicting options, the compiler will issue a warning message.

Nesting of for Loops

An `MP taskloop` pragma applies to the next `for` loop within the current block. There is no nesting of parallelized `for` loops by parallelized C.

Eligibility for Parallelizing

An `MP taskloop` pragma suggests to the compiler that, unless otherwise disallowed, the specified `for` loop should be parallelized.

Any `for` loop with irregular control flow and unknown loop iteration increment is ineligible for parallelization. For example, `for` loops containing `setjmp`, `longjmp`, `exit`, `abort`, `return`, `goto`, labels, and `break` should not be considered as candidates for parallelization.

Of particular importance is to note that `for` loops with inter-iteration dependencies can be eligible for explicit parallelization. This means that if a `MP taskloop` pragma is specified for such a loop the compiler will simply honor it, unless the `for` loop is disqualified. It is the user's responsibility to make sure that such explicit parallelization will not lead to incorrect results.

If both the `serial_loop` or `serial_loop_nested` and `taskloop` pragmas are specified for a `for` loop, the last one specified will prevail.

Consider the following example:

```
#pragma MP serial_loop_nested
  for (i=0; i<100; i++) {
    # pragma MP taskloop
      for (j=0; j<1000; j++) {
        ...
      }
    }
}
```

The `i` loop will not be parallelized but the `j` loop might be.

Number of Processors

`#pragma MP taskloop maxcpus (number_of_processors)` specifies the number of processors to be used for this loop, if possible.

The value of `maxcpus` must be a positive integer. If `maxcpus` equals 1, then the specified loop will be executed in serial. (Note that setting `maxcpus` to be 1 is equivalent to specifying the `serial_loop` pragma.) The smaller of the values of `maxcpus` or the interpreted value of the `PARALLEL` environment variable will be used. When the environment variable `PARALLEL` is not specified, it is interpreted as having the value 1.

If more than one `maxcpus` pragma is specified for a `for` loop, the last one specified will prevail.

Classifying Variables

A variable used in a loop is classified as being either a “private,” “shared,” “reduction,” or “readonly” variable. The variable will belong to only one of these classifications. A variable can only be classified as a reduction or readonly variable via an explicit pragma. See `#pragma MP taskloop reduction` and `#pragma MP taskloop readonly`. A variable can be classified as being either a “private or “shared” variable via an explicit pragma or through the following default scoping rules.

Default Scoping Rules for Private and Shared Variables

A private variable is one whose value is private to each processor processing some iterations of a `for` loop. In other words, the value assigned to a private variable in one iteration of a `for` loop is not propagated to other processors processing other iterations of that `for` loop. A shared variable, on the other hand, is a variable whose current value is accessible by all processors processing iterations of a `for` loop. The value assigned to a shared variable by one processor working on iterations of a loop may be seen by other processors working on other iterations of the loop. Loops being explicitly parallelized through use of `#pragma MP taskloop` directives, that contain references to shared variables, must ensure that such sharing of values does not cause any correctness problems (such as race conditions). No synchronization is provided by the compiler on updates and accesses to shared variables in an explicitly parallelized loop.

In analyzing explicitly parallelized loops, the compiler uses the following “default scoping rules” to determine whether a variable is private or shared:

- If a variable is not explicitly classified via a pragma, the variable will default to being classified as a shared variable if it is declared as a pointer or array, and is only referenced using array syntax within the loop. Otherwise, it will be classified as a private variable.
- The loop index variable is always treated as a private variable and is always a storeback variable.

It is *highly recommended* that all variables used in an explicitly parallelized `for` loop be explicitly classified as one of shared, private, reduction, or readonly, to avoid the “default scoping rules.”

Since the compiler does not perform any synchronization on accesses to shared variables, extreme care must be exercised before using an `MP taskloop` pragma for a loop that contains, for example, array references. If inter-iteration data dependencies exist in such an explicitly parallelized loop, then its parallel execution may give erroneous results. The compiler may or may not be able to detect such a potential problem situation and issue a warning message. In any case, the compiler will not disable the explicit parallelization of loops with potential shared variable problems.

Private Variables

```
#pragma MP taskloop private (list_of_private_variables)
```

Use this pragma to specify all the variables that should be treated as private variables for this loop. All other variables used in the loop that are not explicitly specified as shared, readonly, or reduction variables, are either shared or private as defined by the default scoping rules.

A private variable is one whose value is private to each processor processing some iterations of a loop. In other words, the value assigned to a private variable by one of the processors working on iterations of a loop is not propagated to other processors processing other iterations of that loop. A private variable has no initial value at the start of each iteration of a loop and must be set to a value within the iteration of a loop prior to its first use within that iteration. Execution of a program with a loop containing an explicitly declared private variable whose value is used prior to being set will result in undefined behavior.

Shared Variables

```
#pragma MP taskloop shared (list_of_shared_variables)
```

Use this pragma to specify all the variables that should be treated as shared variables for this loop. All other variables used in the loop that are not explicitly specified as private, readonly, storeback or reduction variables, are either shared or private as defined by the default scoping rules.

A shared variable is a variable whose current value is accessible by all processors processing iterations of a `for` loop. The value assigned to a shared variable by one processor working on iterations of a loop may be seen by other processors working on other iterations of the loop.

Read-only Variables

```
#pragma MP taskloop readonly (list_of_readonly_variables)
```

Read-only variables are a special class of shared variables that are not modified in any iteration of a loop. Use this pragma to indicate to the compiler that it may use a separate copy of that variable's value for each processor processing iterations of the loop.

Storeback Variables

```
#pragma MP taskloop storeback (list_of_storeback_variables)
```

Use this pragma to specify all the variables to be treated as storeback variables.

A storeback variable is one whose value is computed in a loop, and this computed value is then used after the termination of the loop. The last loop iteration values of storeback variables are available for use after the termination of the loop. Such a variable is a good candidate to be declared explicitly via this directive as a storeback variable when the variable is a private variable, whether by explicitly declaring the variable private or by the default scoping rules.

Note that the storeback operation for a storeback variable occurs at the last iteration of the explicitly parallelized loop, regardless of whether or not that iteration updates the value of the storeback variable. In other words the processor that processes the last iteration of a loop may not be the same processor that currently contains the last updated value for a storeback variable. Consider the following example:

```
#pragma MP taskloop private(x)
#pragma MP taskloop storeback(x)
    for (i=1; i <= n; i++) {
        if (...) {
            x=...
        }
    }
    printf ("%d", x);
```

In the previous example the value of the storeback variable `x` printed out via the `printf()` call may not be the same as that printed out by a serial version of the `i` loop, because in the explicitly parallelized case, the processor that processes the last

iteration of the loop (when $i==n$), which performs the storeback operation for x may not be the same processor that currently contains the last updated value for x . The compiler will attempt to issue a warning message to alert the user of such potential problems.

In an explicitly parallelized loop, variables referenced as arrays are not treated as storeback variables. Hence it is important to include them in the *list_of_storeback_variables* if such storeback operation is desired (for example, if the variables referenced as arrays have been declared as private variables).

Savelast

```
#pragma MP taskloop savelast
```

Use this pragma to specify all the private variables of a loop that you want to be treated as storeback variables. The syntax of this pragma is as follows:

```
#pragma MP taskloop savelast
```

It is often convenient to use this form, rather than list out each private variable of a loop when declaring each variable as storeback variables.

Reduction Variables

`#pragma MP taskloop reduction (list_of_reduction_variables)` specifies that all the variables appearing in the reduction list will be treated as reduction variables for the loop. A reduction variable is one whose partial values can be individually computed by each of the processors processing iterations of the loop, and whose final value can be computed from all its partial values. The presence of a list of reduction variables can facilitate the compiler in identifying that the loop is a reduction loop, allowing generation of parallel reduction code for it. Consider the following example:

```
#pragma MP taskloop reduction(x)
  for (i=0; i<n; i++) {
    x = x + a[i];
  }
```

the variable x is a (sum) reduction variable and the i loop is a (sum) reduction loop.

Scheduling Control

The Sun ANSI/ISO C compiler supports several pragmas that can be used in conjunction with the `taskloop` pragma to control the loop scheduling strategy for a given loop. The syntax for this pragma is:

```
#pragma MP taskloop schedtype (scheduling_type)
```

This pragma can be used to specify the specific *scheduling_type* to be used to schedule the parallelized loop. *Scheduling_type* can be one of the following:

- `static`

In static scheduling all the iterations of the loop are uniformly distributed among all the participating processors. Consider the following example:

```
#pragma MP taskloop maxcpus(4)
#pragma MP taskloop schedtype(static)
    for (i=0; i<1000; i++) {
        ...
    }
```

In the above example, each of the four processors will process 250 iterations of the loop.

- `self [(chunk_size)]`

In `self` scheduling, each participating processor processes a fixed number of iterations (called the “chunk size”) until all the iterations of the loop have been processed. The optional *chunk_size* parameter specifies the “chunk size” to be used. *Chunk_size* must be a positive integer constant, or variable of integral type. If specified as a variable *chunk_size* must evaluate to a positive integer value at the beginning of the loop. If this optional parameter is not specified or its value is not positive, the compiler will select the chunk size to be used. Consider the following example:

```
#pragma MP taskloop maxcpus(4)
#pragma MP taskloop schedtype(self(120))
    for (i=0; i<1000; i++) {
        ...
    }
```

In the above example, the number of iterations of the loop assigned to each participating processor, in order of work request, are:

120, 120, 120, 120, 120, 120, 120, 40.

- `gss [(min_chunk_size)]`

In guided self scheduling, each participating processor processes a variable number of iterations (called the “min chunk size”) until all the iterations of the loop have been processed. The optional *min_chunk_size* parameter specifies that each variable chunk size used must be at least *min_chunk_size* in size. *Min_chunk_size* must be a positive integer constant, or variable of integral type. If specified as a variable *min_chunk_size* must evaluate to a positive integer value at the beginning of the loop. If this optional parameter is not specified or its value is not positive, the compiler will select the chunk size to be used. Consider the following example:

```
#pragma MP taskloop maxcpus(4)
#pragma MP taskloop schedtype(gss(10))
for (i=0; i<1000; i++) {
    ...
}
```

In the above example, the number of iterations of the loop assigned to each participating processor, in order of work request, are:

250, 188, 141, 106, 79, 59, 45, 33, 25, 19, 14, 11, 10, 10, 10.

■ factoring [*min_chunk_size*]

In factoring scheduling, each participating processor processes a variable number of iterations (called the “min chunk size”) until all the iterations of the loop have been processed. The optional *min_chunk_size* parameter specifies that each variable chunk size used must be at least *min_chunk_size* in size. *Min_chunk_size* must be a positive integer constant, or variable of integral type. If specified as a variable *min_chunk_size* must evaluate to a positive integer value at the beginning of the loop. If this optional parameter is not specified or its value is not positive, the compiler will select the chunk size to be used. Consider the following example:

```
#pragma MP taskloop maxcpus(4)
#pragma MP taskloop schedtype(factoring(10))
for (i=0; i<1000; i++) {
    ...
}
```

In the above example, the number of iterations of the loop assigned to each participating processor, in order of work request, are:

125, 125, 125, 125, 62, 62, 62, 62, 32, 32, 32, 32, 16, 16, 16, 16, 10, 10, 10, 10, 10, 10

Incremental Link Editor (`ild`)

This chapter describes `ild`, `ild`-specific features, example messages, and `ild` options. This chapter is organized into the following sections:

- “Introduction” on page 137
- “Overview of Incremental Linking” on page 138
- “How to Use `ild`” on page 138
- “How `ild` Works” on page 140
- “What `ild` Cannot Do” on page 141
- “Reasons for Full Relinks” on page 142
- “`ild` Options” on page 145
- “Options Passed to `ild` from the Compilation System” on page 150
- “Environment” on page 152
- “`ild` Options not Supported by `ild`” on page 154
- “Additional Unsupported Commands” on page 155
- “Files That `ild` Uses” on page 156

Introduction

`ild` is an incremental version of the Link Editor `ld`, and replaces `ld` for linking programs. Use `ild` to complete the edit, compile, link, and debug loop efficiently and more quickly. You can avoid relinking entirely by using the *fix and continue* feature of `dbx` which allows you to work without relinking. However, if you need to relink, the process can be faster if you use `ild`. For more information on *fix and continue*, see Chapter 11 in *Debugging a Program With dbx*.

`ild` links incrementally so you can insert modified object code into an executable file that you created earlier, without relinking unmodified object files. The time required to relink depends upon the amount of code modified. Linking your application on every build does not require the same amount of time; small changes in code can be relinked very quickly.

On the initial link, `ild` requires about the same amount of time that `ld` requires, but subsequent `ild` links can be much faster than an `ld` link. The cost of the reduced link time is an increase in the size of the executable.

Overview of Incremental Linking

When you use `ild` in place of `ld`, the initial link causes the various text, data, bss, exception table sections, etc., to be padded with additional space for future expansion (see FIGURE 5-1). Additionally, all relocation records and the global symbol table are saved into a new persistent state region in the executable file. On subsequent incremental links, `ild` uses timestamps to determine which object files have changed and patches the changed object code into a previously built executable. That is, previous versions of the object files are invalidated and the new object files are loaded into the space vacated, or into the pad sections of the executable when needed. All references to symbols in invalidated object files are patched to point to the correct new object files.

`ild` does not support all `ld` command options. If `ild` is passed a command option that it does not support, `ild` directly invokes `/usr/ccs/bin/ld` to perform the link. See “`ld` Options Not Supported by `ild`” on page 154 for more information on commands that are not supported by the Incremental Linker.

How to Use `ild`

`ild` is invoked automatically by the compilation system in place of `ld` under certain conditions. When you invoke a compilation system, you are invoking a compiler driver. When you pass certain options to the driver, the driver uses `ild`. The compiler driver reads the options from the command line and executes various programs in the correct order and adds files from the list of arguments that are passed.

For example, `cc` first runs `acompile` (the front-end of the compiler), then `acompile` runs the optimizing code generator, then `cc` does the same thing for the other source files listed on the command line. The driver can then generate a call to either `ild` or `ld`, depending on the options, passing it all of the files just compiled, plus other files and libraries needed to make the program complete

The following figures shows an example of incremental linking.

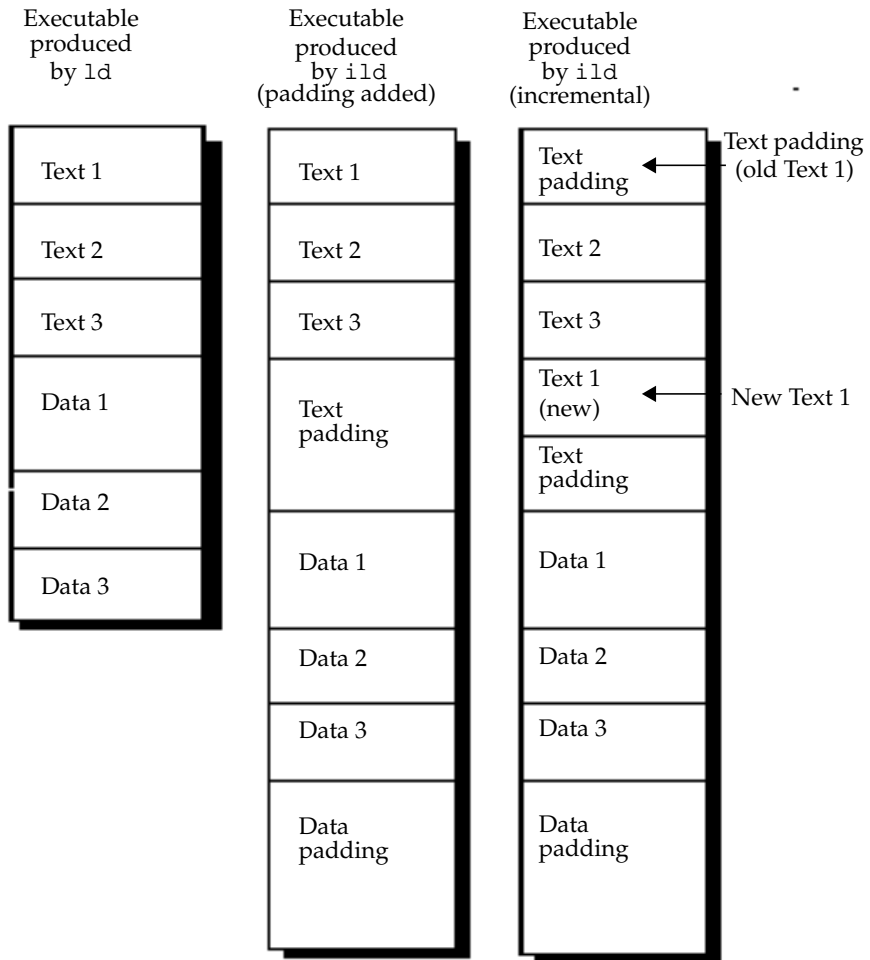


FIGURE 5-1 An Example of Incremental Linking

The following compilation system options control whether a link step is performed by `ild` or `ld`:

- `-xildon` Always use `ild`
- `-xildoff` Always use `ld`

Note – If `-xildon` and `-xildoff` are both present, the last command listed is used by the linker.

- `-g` When neither `-xildoff` or `-G` are given, use `ild` for link-only invocations (no source files on the command line). For a complete explanation of `-g`, see “`-g`” on page 30.
- `-G` Prevents the `-g` option from having any effect on linker selection. For a complete explanation of `-G`, see “`-G`” on page 30.

When you use the `-g` option to invoke debugging, and you have the default Makefile structure (which includes compile-time options such as `-g` on the link command line), you use `ild` automatically when doing development.

How `ild` Works

On an initial link, `ild` saves information about:

- All of the object files looked at.
- The symbol table for the executable produced.
- All symbolic references not resolved at compile time.

Initial `ild` links take about as much time as an `ld` link.

On incremental links, `ild`:

- Determines which files have changed.
- Relinks the modified object files.
- Uses stored information to modify changed symbolic references in the rest of the program.

Incremental `ild` links are much faster than `ld` links.

In general, you do one initial link and all subsequent links are incremental.

For example, `ild` saves a list of all places where symbol `foo` is referenced in your code. If you do an incremental link that changes the value of `foo`, `ild` must change the value of all references to `foo`.

`ild` spreads out the components of the program and each section of the executable has padding added to it. Padding makes the executable modules larger than when they were linked by `ld`. As object files increase in size during successive incremental links, the padding can become exhausted. If this occurs, `ild` displays a message and does a complete full relink of the executable.

For example, as FIGURE 5-1 shows, each of the three columns shows the sequence of text and data in a linked executable program. The left column shows text and data in an executable linked by `ld`. The center column shows the addition of text and data padding in an executable linked by `ild`. Assume that a change is made to the source file for Text 1 that causes the Text section to grow without affecting the size of the

other sections. The right column shows that the original location of Text 1 has been replaced by Text padding (Text 1 has been invalidated). Text 1 has been moved to occupy a portion of the Text padding space.

To produce a smaller nonincremental executable, run the compiler driver (for example, `cc` or `CC`) with the `-xildoff` option, and `ld` is invoked to produce a more compact executable.

The resulting executable from `ild` can be debugged by `dbx` because `dbx/Debugger` understands the padding that `ild` inserts between programs.

For any command-line option that `ild` does not understand, `ild` invokes `ld`. `ild` is compatible with `ld` (in `/usr/ccs/bin/ld`). See “`ild` Options” on page 145, for details.

There are no special or extra files used by `ild`.

What `ild` Cannot Do

When `ild` is invoked to create shared objects, `ild` invokes `ld` to create the link.

Performance of `ild` may suffer greatly if you change a high percentage of object files. `ild` automatically does an full relink when it detects that a high percentage of files have been changed.

Do not use `ild` to produce the final production code for shipment. `ild` makes the file larger because parts of the program have been spread out due to padding. Because of the padding and additional time required to link, it is recommended that you do not use the `-xildon` option for production code. (Use `-xildoff` on the link line if `-g` is present.)

`ild` may not link small programs much faster, and the increase in size of the executable is greater than that for larger programs.

Third-party tools that work on executables may have unexpected results on `ild`-produced binaries.

Any program that modifies an executable, for example `strip` or `mcs`, might affect the ability of `ild` to perform an incremental link. When this happens, `ild` issues a message and performs a full relink. For more information on a full relink, see “Reasons for Full Relinks” on page 142.

Reasons for Full Relinks

The following section explains under which circumstances `ild` calls `ld` to complete a link.

`ild` Deferred-Link Messages

The message `'ild: calling ld to finish link' . . .` means that `ild` cannot complete the link, and is deferring the link request to `ld` for completion. By default, these messages are displayed as needed. You can suppress these messages by using the `-z i_quiet` option.

The following message is suppressed if `ild` is implicitly requested (`-g`), but is displayed if `-xildon` is on the command line. This message is displayed in all cases if you use the `-z i_verbose` option, and never displayed if you use the `-z i_quiet` option.

```
ild: calling ld to finish link -- cannot handle shared libraries
in archive library name
```

Here are further examples of `-z i_verbose` messages:

```
ild: calling ld to finish link -- cannot handle keyword Keyword
```

```
ild: calling ld to finish link -- cannot handle -d Keyword
```

```
ild: calling ld to finish link -- cannot handle -z keyword
```

```
ild: calling ld to finish link -- cannot handle argument keyword
```

`ild` Relink Messages

The message `'ild: (Performing full relink)'` . . . means that for some reason `ild` cannot do an incremental link and must do a full relink. This is not an error. It is to inform you that this link will take longer than an incremental link (see “How `ild` Works” on page 140, for more details). `ild` messages can be controlled by `ild` options `-z i_quiet` and `-z i_verbose`. Some messages have a verbose mode with more descriptive text.

You can suppress all of these messages by using the `ild` option `-z i_quiet`. If the default message has a verbose mode, the message ends with an ellipsis (`[...]`) indicating more information is available. You can view the additional information by using the `-z i_verbose` option. Example messages are shown with the `-z i_verbose` option selected.

Example 1: Internal Free Space Exhausted

The most common of the full relink messages is the internal free space exhausted message:

```
# This creates test1.o
# This creates a.out with
minimal debugging information.
# A one-line compile and link
puts all debugging information
into a.out.
```

```
$ cat test1.c
int main() { return 0; }
$ rm a.out
$ cc -xildon -c -g test1.c
$ cc -xildon -z i_verbose -g test1.o
$ cc -xildon -z i_verbose -g test1.c

ild: (Performing full relink) internal free
space in output file exhausted (sections)
$
```

These commands show that going from a one-line compile to a two-line compile causes debugging information to grow in the executable. This growth causes `ild` to run out of space and do an full relink.

Example 2: Running `strip`

Another problem arises when you run `strip`. Continuing from Example 1:

```
# Strip a.out
# Try to do an incremental
link
```

```
$ strip a.out
$ cc -xildon -z i_verbose -g test1.c

ild: (Performing full relink) a.out has been
altered since the last incremental link --
maybe you ran strip or mcs on it?
$
```

Example 3: ild Version

When a new version of ild is run on an executable created by an older version of ild, you see the following error message:

```
# Assume old_executable was  
created by an earlier  
version of ild
```

```
$ cc -xildon -z i_verbose foo.o -o old_executable  
  
ild: (Performing full relink) an updated ild  
has been installed since a.out was last linked  
(2/16)
```

Note – The numbers (2/16) are used only for internal reporting.

Example 4: Too Many Files Changed

Sometimes ild determines that it will be faster to do a full relink than an incremental link. For example:

```
$ rm a.out  
$ cc -xildon -z i_verbose \  
    x0.o x1.o x2.o x3.o x4.o x5.o x6.o x7.o x8.o test2.o  
$ touch x0.o x1.o x2.o x3.o x4.o x5.o x6.o x7.o x8.o  
$ cc -xildon -z i_verbose \  
    x0.o x1.o x2.o x3.o x4.o x5.o x6.o x7.o x8.o test2.o  
ild: (Performing full relink) too many files changed
```

Here, use of the touch command causes ild to determine that files x0.o through x8.o have changed and that a full relink will be faster than incrementally relinking all nine object files.

Example 5: Full Relink

There are certain conditions that can cause a full relink on the next link, as compared to the previous examples that cause a full relink on this link.

The next time you try to link that program, you see the message:

```
# ild detects previous  
error and does a full  
relink
```

```
$ cc -xildon -z i_verbose broken.o  
ild: (Performing full relink) cannot do incremental  
relink due to problems in the previous link
```

A full relink occurs.

Example 6: New Working Directory

```
# initial link with cwd  
equal to /tmp
```

```
# incremental link, cwd  
is now /tmp/junk
```

```
% cd /tmp  
% cat y.c  
    int main(){ return 0;}  
% cc -c y.c  
% rm -f a.out  
% cc -xildon -z i_verbose y.o -o a.out  
  
% mkdir junk  
% mv y.o y.c a.out junk  
% cd junk  
% cc -xildon -z i_verbose y.o -o a.out  
  
ild: (Performing full relink) current directory has  
changed from '/tmp' to '/tmp/junk'  
%
```

ild Options

This section describes the linker control options directly accepted by the compilation system and linker options that may be passed through the compilation system to `ild`.

`-a`

In static mode only, produce an executable object file; give errors for undefined references. This is the default behavior for static mode.

`-B dynamic | static`

Options governing library inclusion. Option `-Bdynamic` is valid in dynamic mode only. These options can be specified any number of times on the command line as toggles: if the `-Bstatic` option is given, no shared objects are accepted until `-Bdynamic` is seen. See option “`-lx`” on page 147.

`-d y|n`

When `-dy` (the default) is specified, `ild` uses dynamic linking; when `-dn` is specified, `ild` uses static linking. See option “`-B dynamic | static`” on page 146.

`-e epsym`

Set the entry point address for the output file to be that of the symbol `epsym`.

`-g`

The compilation systems invoke `ild` in place of `ld` when the `-g` option (output debugging information) is given, unless any of the following are true:

- The `-G` option (produce a shared library) is given
- The `-xildoff` option is present
- Any source files are named on the command line

`-I name`

When building an executable, use `name` as the path name of the interpreter to be written into the program header. The default in static mode is no interpreter; in dynamic mode, the default is the name of the runtime linker, `/usr/lib/ld.so.1`. Either case may be overridden by `-Iname`. `exec` only loads this interpreter when it loads `a.out` and will pass control to the interpreter rather than to `a.out` directly.

-i

Ignores `LD_LIBRARY_PATH` setting. This option is useful when an `LD_LIBRARY_PATH` setting is in effect to influence the runtime library search, which would interfere with the link editing being performed. (This also applies to the setting of `LD_LIBRARY_PATH_64`).

-L*path*

Adds *path* to the library search directories. `ild` searches for libraries first in any directories specified by the `-L` options, and then in the standard directories. This option is useful only if it precedes the `-l` options to which it applies on the command line. You can use the environment variable `LD_LIBRARY_PATH` and `LD_LIBRARY_PATH_64` to supplement the library search path (see “`LD_LIBRARY_PATH`” on page 152).

-l*x*

Searches a library `libx.so` or `libx.a`, the conventional names for shared object and archive libraries, respectively. In dynamic mode, unless the `-Bstatic` option is in effect, `ild` searches each directory specified in the library search path for a file `libx.so` or `libx.a`. The directory search stops at the first directory containing either. `ild` chooses the file ending in `.so` if `-l` expands to two files whose names are of the form `libx.so` and `libx.a`. If no `libx.so` is found, then `ild` accepts `libx.a`. In static mode, or when the `-Bstatic` option is in effect, `ild` selects only the file ending in `.a`. A library is searched when its name is encountered, so the placement of `-l` is significant.

-m

Produce a memory map or listing of the input/output sections on the standard output.

-o *outfile*

Produces an output object file named *outfile*. The name of the default object file is `a.out`.

-Q y | n

Under `-Qy`, an `ident` string is added to the `.comment` section of the output file to identify the version of the link editor used to create the file. This results in multiple `ld ident`s when there have been multiple linking steps, such as when using `ld -r`. This is identical with the default action of the `cc` command. Option `-Qn` suppresses version identification.

-Rpath

This option gives a colon-separated list of directories that specifies library search directories to the runtime linker. If present and not null, *path* is recorded in the output object file and passed to the runtime linker. Multiple instances of this option are concatenated and separated by a colon.

-S

Strips symbolic information from the output file. Any debugging information and associated relocation entries are removed. Except for relocatable files or shared objects, the symbol table and string table sections are also removed from the output object file.

-t

Turn off the warning about multiply defined symbols that are not the same size.

-u *symname*

Enter *symname* as an undefined symbol in the symbol table. This is useful for loading entirely from an archive library, since initially the symbol table is empty and an unresolved reference is needed to force the loading of the first routine. The placement of this option on the command line is significant; it must be placed before the library that defines the symbol.

-V

Output a message about the version of `ld` being used.

`-xildoff`

Incremental linker off. Force the use of bundled `ld`. This is the default if `-g` is not being used, or `-G` is being used. You can override this default with `-xildon`.

`-xildon`

Incremental linker. Force the use of `ild` in incremental mode. This is the default if `-g` is being used. You can override this default with `-xildoff`.

`-YP, dirlist`

(cc only) Changes the default directories used for finding libraries. Option *dirlist* is a colon-separated path list.

Note – `ild` uses the “`-z name`” form for special options. The *i_* prefix to the `-z` options identifies those options peculiar to `ild`.

`-z allextact | defaultextract | weakextract`

Alter the extraction criteria of objects from any archives that follow. By default archive members are extracted to satisfy undefined references and to promote tentative definitions with data definitions. Weak symbol references do not trigger extraction. Under `-z allextact`, all archive members are extracted from the archive. Under `-z weakextract`, weak references trigger archive extraction. `-z defaultextract` provides a means of returning to the default following use of the former extract options.

`-z defs`

Forces a fatal error if any undefined symbols remain at the end of the link. This is the default when building an executable. It is also useful when building a shared object to assure that the object is self-contained, that is, that all its symbolic references are resolved internally.

`-z i_dryrun`

(`ild` only.) Prints the list of files that would be linked by `ild` and exits.

`-z i_full`

(`ild` only.) Does a complete relink in incremental mode.

`-z i_noincr`

(`ild` only.) Runs `ild` in nonincremental mode (not recommended for customer use — used for testing only).

`-z i_quiet`

(`ild` only) Turns off all `ild` relink messages.

`-z i_verbose`

(`ild` only) Expands on default information on some `ild` relink messages.

`-z nodefs`

Allows undefined symbols. This is the default when building a shared object. When used with executables, the behavior of references to such “undefined symbols” is unspecified.

Options Passed to `ild` From the Compilation System

The following options are accepted by `ild`, but you must use the form:

`-Wl, arg, arg` (for `cc`), or `-Qoption ld arg, arg` (for others),

to pass them to `ild` via the compilation system

`-a`

In static mode only, produces an executable object file; gives errors for undefined references. This is the default behavior for static mode. Option `-a` cannot be used with the `-r` option.

`-e epsym`

Sets the entry point address for the output file to be that of the symbol *epsym*.

`-I name`

When building an executable, uses *name* as the path name of the interpreter to be written into the program header. The default in static mode is no interpreter; in dynamic mode, the default is the name of the runtime linker, `/usr/lib/ld.so.1`. Either case can be overridden by `-I name`. The `exec` system call loads this interpreter when it loads the `a.out` and passes control to the interpreter rather than to the `a.out` directly.

`-m`

Produces a memory map or listing of the input/output sections on the standard output.

`-t`

Turn off the warning about symbols that are defined more than once and that are not the same size.

`-u symname`

Enters *symname* as an undefined symbol in the symbol table. This is useful for loading entirely from an archive library because, initially, the symbol table is empty and an unresolved reference is needed to force the loading of the first routine. The placement of this option on the command line is significant; it must be placed before the library that defines the symbol.

Environment

LD_LIBRARY_PATH

A list of directories which is searched for the libraries that are specified with the `-l` option. Multiple directories are separated by a colon. In the most general case, it contains two directory lists separated by a semicolon:

```
dirlist1; dirlist2
```

If `ild` is called with any number of occurrences of `-L`, as in:

```
ild ...-Lpath1 ... -Lpathn ...
```

then the search path ordering is:

```
dirlist1 path1 ... pathn dirlist2 LIBPATH
```

When the list of directories does not contain a semicolon, it is interpreted as follows:

```
dirlist2
```

LD_LIBRARY_PATH is also used to specify library search directories to the runtime linker. That is, if LD_LIBRARY_PATH exists in the environment, the runtime linker searches the directories named in it, before its default directory, for shared objects to be linked with the program at execution.

Note – When running a `set-user-ID` or `set-group-ID` program, the runtime linker searches only for libraries in `/usr/lib`. It also searches for any full pathname specified within the executable. A full pathname is the result of a runpath being specified when the executable was constructed. Any library dependencies specified as relative pathnames are silently ignored.

LD_LIBRARY_PATH_64

On Solaris 7 and Solaris 8, this environment variable is similar to `LD_LIBRARY_PATH` but overrides it when searching for 64-bit dependencies.

When you run Solaris 7 or Solaris 8 on a SPARC processor and link in 32-bit mode, `LD_LIBRARY_PATH_64` is ignored. If only `LD_LIBRARY_PATH` is defined, it is used for both 32-bit and 64-bit linking. If both `LD_LIBRARY_PATH` and `LD_LIBRARY_PATH_64` are defined, the 32-bit linking will be done using `LD_LIBRARY_PATH` and the 64-bit linking will be done using `LD_LIBRARY_PATH_64`.

LD_OPTIONS

A default set of options to `ild`. `LD_OPTIONS` is interpreted by `ild` as though its value had been placed on the command line immediately following the name used to invoke `ild`, as in:

```
ild $LD_OPTIONS ... other-arguments ...
```

LD_PRELOAD

A list of shared objects that are to be interpreted by the runtime linker. The specified shared objects are linked in after the program being executed and before any other shared objects that the program references.

Note – When running a `set-user-ID` or `set-group-ID` program, this option is silently ignored.

LD_RUN_PATH

An alternative mechanism for specifying a runpath to the link editor (see the `-R` option). If both `LD_RUN_PATH` and the `-R` option are specified, the `-R` is used.

LD_DEBUG

(not supported by `ild`) Provide a list of tokens that cause the runtime linker to print debugging information to the standard error. The special token `help` indicates the full list of tokens available.

Note – Environment variable names beginning with the characters ‘LD_’ are reserved for possible future enhancements to ld. Environment variable-names beginning with the characters ‘LLD_’ are reserved for possible future enhancements to ild.

ld Options not Supported by ild

If ild determines that a command line option is not implemented, ild directly invokes /usr/css/bin/ld to perform the link.

The following options, which may be given to the compilation system, are not supported by ild.

-B *symbolic*

In dynamic mode only, when building a shared object, bind references to global symbols to their definitions within the object, if definitions are available. Normally, references to global symbols within shared objects are not bound until runtime, even if definitions are available, so that definitions of the same symbol in an executable or other shared objects can override the object’s own definition. ld issues warnings for undefined symbols unless -z *defs* overrides.

-b

In dynamic mode only, when creating an executable, does not do special processing for relocations that reference symbols in shared objects. Without the -b option, the link editor creates special position-independent relocations for references to functions defined in shared objects and arranges for data objects defined in shared objects to be copied into the memory image of the executable by the runtime linker. With the -b option, the output code can be more efficient, but it is less sharable.

-G

In dynamic mode only, produces a shared object. Undefined symbols are allowed.

`-h name`

In dynamic mode only, when building a shared object, records *name* in the object's dynamic section. Option *name* is recorded in executables that are linked with this object rather than the object's UNIX System file name. Accordingly, *name* is used by the runtime linker as the name of the shared object to search for at runtime.

`-z muldefs`

Allows multiple symbol definitions. By default, multiple symbol definitions occurring between relocatable objects result in a fatal error condition. This option suppresses the error condition, and allows the first symbol definition to be taken.

`-z text`

In dynamic mode only, forces a fatal error if any relocations against non-writable, allocatable sections remain.

Additional Unsupported Commands

In addition, the following options that may be passed directly to `ld`, are not supported by `ild`:

`-D token,token, ...`

Prints debugging information as specified by each token, to the standard error. The special token *help* indicates the full list of tokens available.

`-F name`

Useful only when building a shared object. Specifies that the symbol table of the shared object is used as a "filter" on the symbol table of the shared object specified by *name*.

-M *mapfile*

Reads *mapfile* as a text file of directives to `ld`. See *SunOS 5.3 Linker and Libraries Manual* for a description of mapfiles.

-r

Combines relocatable object files to produce one relocatable object file. `ld` does not complain about unresolved references. This option cannot be used in dynamic mode or with `-a`.

Files That `ld` Uses

`libx.a` libraries

`a.out` output file

`LIBPATH` usually `/usr/lib`

lint Source Code Checker

This chapter explains how you can use the `lint` program to check your C code for errors that may cause a compilation failure or unexpected results at runtime. In many cases, `lint` warns you about incorrect, error-prone, or nonstandard code that the compiler does not necessarily flag.

The `lint` program issues every error and warning message produced by the C compiler. It also issues warnings about potential bugs and portability problems. Many messages issued by `lint` can assist you in improving your program's effectiveness, including reducing its size and required memory.

The `lint` program uses the same locale as the compiler and the output from `lint` is directed to `stderr`. This chapter is organized into the following sections:

- "Basic and Enhanced `lint` Modes" on page 157
- "Using `lint`" on page 158
- "The `lint` Options" on page 160
- "`lint` Messages" on page 173
- "`lint` Directives" on page 177
- "`lint` Reference and Examples" on page 181

See Chapter 7 for more information on and examples of how to use `lint` to check code before you perform type-based alias-disambiguation.

Basic and Enhanced `lint` Modes

The `lint` program operates in two modes:

- *Basic*, which is the default
- *Enhanced*, which includes everything done by basic `lint`, as well as additional, detailed analysis of code

In both basic and enhanced modes, `lint` compensates for separate and independent compilation in C by flagging inconsistencies in definition and use across files, including any libraries you have used. In a large project environment especially, where the same function may be used by different programmers in hundreds of separate modules of code, `lint` can help discover bugs that otherwise might be difficult to find. A function called with one less argument than expected, for example, looks at the stack for a value the call has never pushed, with results correct in one condition, incorrect in another, depending on whatever happens to be in memory at that stack location. By identifying dependencies like this one and dependencies on machine architecture as well, `lint` can improve the reliability of code run on your machine or someone else's.

In enhanced mode, `lint` provides more detailed reporting than in basic mode. In basic mode, `lint`'s capabilities include:

- Structure and flow analysis of the source program
- Constant propagations and constant expression evaluations
- Analysis of control flow and data flow
- Analysis of data types usage

In enhanced mode, `lint` can detect these problems:

- Unused `#include` directives, variables, and procedures
- Memory usage after its deallocation
- Unused assignments
- Usage of a variable value before its initialization
- Deallocation of nonallocated memory
- Usage of pointers when writing in constant data segments
- Nonequivalent macro redefinitions
- Unreached code
- Conformity of the usage of value types in unions
- Implicit casts of actual arguments.

Using `lint`

Invoke the `lint` program and its options from the command line. To invoke `lint` in the basic mode, use the following command:

```
% lint file1.c file2.c
```

Enhanced `lint` is invoked with the `-Nlevel` or `-Ncheck` option. For example, you can invoke enhanced `lint` as follows:

```
% lint -Nlevel=3 file1.c file2.c
```

`lint` examines code in two *passes*. In the first pass, `lint` checks for error conditions within C source files; in the second pass, it checks for inconsistencies across C source files. This process is invisible to the user unless `lint` is invoked with `-c`:

```
% lint -c file1.c file2.c
```

That command directs `lint` to execute the first pass only and collect information relevant to the second—about inconsistencies in definition and use across `file1.c` and `file2.c`—in intermediate files named `file1.ln` and `file2.ln`:

```
% ls
file1.c
file1.ln
file2.c
file2.ln
```

This way, the `-c` option to `lint` is analogous to the `-c` option to `cc`, which suppresses the link editing phase of compilation. Generally speaking, `lint`'s command-line syntax closely follows `cc`'s.

When the `.ln` files are linted:

```
% lint file1.ln file2.ln
```

the second pass is executed. `lint` processes any number of `.c` or `.ln` files in their command-line order. Thus,

```
% lint file1.ln file2.ln file3.c
```

directs `lint` to check `file3.c` for errors internal to it and all three files for consistency.

`lint` searches directories for included header files in the same order as `cc`. You can use the `-I` option to `lint` as you would the `-I` option to `cc`. See “Include Files” on page 89

You can specify multiple options to `lint` on the same command line. Options can be concatenated unless one of the options takes an argument or if the option has more than one letter:

```
% lint -cp -Idir1 -Idir2 file1.c file2.c
```

That command directs `lint` to:

- Execute the first pass only
- Perform additional portability checks
- Search the specified directories for included header files

`lint` has many options you can use to direct `lint` to perform certain tasks and report on certain conditions.

The `lint` Options

The `lint` program is a static analyzer. It cannot evaluate the runtime consequences of the dependencies it detects. Certain programs, for instance, may contain hundreds of unreachable `break` statements that are of little importance, but which `lint` flags nevertheless. This is one example where the `lint` command-line options and directives—special comments embedded in the source text—come in:

- You can invoke `lint` with the `-b` option to suppress all the error messages about unreachable `break` statements.
- You can precede any unreachable statement with the comment `/*NOTREACHED*/` to suppress the diagnostic for that statement.

The `lint` options are listed below alphabetically. Several `lint` options relate to suppressing `lint` diagnostic messages. These options are also listed in TABLE 6-6, following the alphabetized options, along with the specific messages they suppress. The options for invoking enhanced `lint` begin with `-N`.

`lint` recognizes many `cc` command-line options, including `-A`, `-D`, `-E`, `-g`, `-H`, `-O`, `-P`, `-U`, `-Xa`, `-Xc`, `-Xs`, `-Xt`, and `-Y`, although `-g` and `-O` are ignored. Unrecognized options are warned about and ignored.

`-#`

Turns on verbose mode, showing each component as it is invoked.

`-###`

Shows each component as it is invoked, but does not actually execute it.

`-a`

Suppresses certain messages. Refer to TABLE 6-6.

`-b`

Suppresses certain messages. Refer to TABLE 6-6.

`-C filename`

Creates a `.ln` file with the file name specified. These `.ln` files are the product of `lint`'s first pass only. *filename* can be a complete path name.

`-c`

Creates a `.ln` file consisting of information relevant to `lint`'s second pass for every `.c` file named on the command line. The second pass is not executed.

`-dirout=dir`

Specifies the directory *dir* where the `lint` output files (`.ln` files) will be placed. This option affects the `-c` option.

`-err=warn`

`-err=warn` is a macro for `-errwarn=%all`. See “`-errwarn=t`” on page 166.

`-errchk=l(, l)`

Check structural arguments passed by value; Check portability to environment for which the size of long integers and pointers is 64 bits.

l is a comma-separated list of checks that consists of one or more of the following:

`%all`

Perform all of `errchk`'s checks.

`%none`

Perform none of `errchk`'s checks. This is the default.

`locfmtchk`

Use this option when you want `lint` to check `printf`-like format strings during its first pass. Regardless of whether or not you use `-errchk=locfmtchk`, `lint` always checks for `printf`-like format strings in its second pass.

`longptr64`

Check portability to environment for which the size of long integers and pointers is 64 bits and the size of plain integers is 32 bits. Check assignments of pointer expressions and long integer expressions to plain integers, even when explicit cast is used.

`no%longptr64`

Perform none of `errchk`'s `longptr64` checks.

The values may be a comma separated list, for example
`-errchk=longptr64,structarg`.

The default is `-errchk=%none`. Specifying `-errchk` is equivalent to specifying `-errchk=%all`.

no%structarg

Perform none of `errchk`'s *structarg* checks.

parentheses

Use this option to enhance the maintainability of code. If `-errchk=parentheses` returns a warning, consider using additional parentheses to clearly signify the precedence of operations within the code.

signext

This option produces error messages when the normal ANSI/ISO C value-preserving rules allow the extension of the sign of a signed-integral value in an expression of unsigned-integral type. This option only produces error messages when you specify `-errchk=longptr64` as well.

sizematch

Issues a warning when a larger integer is assigned to a smaller integer. These warnings are also issued for assignment between same size integers that have different signs (unsigned int = signed int).

structarg

Check structural arguments passed by value and report the cases when formal parameter type is not known.

`-errfmt=f`

Specifies the format of `lint` output. *f* can be one of the following: `macro`, `simple`, `src`, or `tab`.

TABLE 6-1 The `-errfmt` Values

Value	Meaning
<code>macro</code>	Displays the source code, the line number, and the place of the error, with macro unfolding
<code>simple</code>	Displays the line number and the place number, in brackets, of the error, for one-line (simple) diagnostic messages. Similar to the <code>-s</code> option, but includes error-position information
<code>src</code>	Displays the source code, the line number, and the place of the error (no macro unfolding)
<code>tab</code>	Displays in tabular format. This is the default.

The default is `-errfmt=tab`. Specifying `-errfmt` is equivalent to specifying `-errfmt=tab`.

If more than one format is specified, the last format specified is used, and `lint` warns about the unused formats.

`-errhdr=h`

Enables the reporting of certain messages for header files when used with `-Ncheck`. *h* is a comma-separated list that consists of one or more of the following: `dir`, `no%dir`, `%all`, `%none`, `%user`.

TABLE 6-2 The `-errhdr` Values

Value	Meaning
<code>dir</code>	Checks header files used in the directory <i>dir</i>
<code>no%dir</code>	Does not check header files used in the directory <i>dir</i>
<code>%all</code>	Checks all used header files
<code>%none</code>	Does not check header files. This is the default.
<code>%user</code>	Checks all used user header files, that is, all header files except those in <code>/usr/include</code> and its subdirectories, as well as those supplied by the compiler

The default is `-errhdr=%none`. Specifying `-errhdr` is equivalent to specifying `-errhdr=%user`.

Examples:

```
% lint -errhdr=inc1 -errhdr=./inc2
```

checks used header files in directories `inc1` and `./inc2`.

```
% lint -errhdr=%all,no%./inc
```

checks all used header files except those in the directory `./inc`.

`-erroff=tag(, tag)`

Suppresses or enables `lint` error messages.

t is a comma-separated list that consists of one or more of the following: *tag*, *no%tag*, *%all*, *%none*.

TABLE 6-3 The `-erroff` Values

Value	Meaning
<i>tag</i>	Suppresses the message specified by this <i>tag</i> . You can display the tag for a message by using the <code>-errtags=yes</code> option.
<i>no%tag</i>	Enables the message specified by this <i>tag</i>
<i>%all</i>	Suppresses all messages
<i>%none</i>	Enables all messages. This is the default.

The default is `-erroff=%none`. Specifying `-erroff` is equivalent to specifying `-erroff=%all`.

Examples:

```
% lint -erroff=%all,no%E_ENUM_NEVER_DEF,no%E_STATIC_UNUSED
```

prints only the messages “enum never defined” and “static unused”, and suppresses other messages.

```
% lint -erroff=E_ENUM_NEVER_DEF,E_STATIC_UNUSED
```

suppresses only the messages “enum never defined” and “static unused”.

`-errtags=a`

Displays the message tag for each error message. *a* can be either yes or no. The default is `-errtags=no`. Specifying `-errtags` is equivalent to specifying `-errtags=yes`.

Works with all `-errfmt` options.

`-errwarn=t`

If the indicated warning message is issued, `lint` exits with a failure status. *t* is a comma-separated list that consists of one or more of the following: *tag*, `no%tag`, `%all`, `%none`. Order is important; for example `%all,no%tag` causes `lint` to exit with a fatal status if any warning except *tag* is issued. The following table list the `-errwarn` values:

TABLE 6-4 `-errwarn` Values

<i>tag</i>	Cause <code>lint</code> to exit with a fatal status if the message specified by this <i>tag</i> is issued as a warning message. Has no effect if <i>tag</i> is not issued.
<code>no%tag</code>	Prevent <code>lint</code> from exiting with a fatal status if the message specified by <i>tag</i> is issued only as a warning message. Has no effect if <i>tag</i> is not issued. Use this option to revert a warning message that was previously specified by this option with <i>tag</i> or <code>%all</code> from causing <code>lint</code> to exit with a fatal status when issued as a warning message.
<code>%all</code>	Cause <code>lint</code> to exit with a fatal status if any warning messages are issued. <code>%all</code> can be followed by <code>no%tag</code> to exempt specific warning messages from this behavior.
<code>%none</code>	Prevents any warning message from causing <code>lint</code> to exit with a fatal status should any warning message be issued.

The default is `-errwarn=%none`. If you specify `-errwarn` alone, it is equivalent to `-errwarn=%all`.

-F

Prints the path names as supplied on the command line rather than only their base names when referring to the `.c` files named on the command line.

-fd

Reports about old-style function definitions or declarations.

-flagsrc=*file*

Executes `lint` with options contained in the file *file*. Multiple options can be specified in *file*, one per line.

-h

Suppresses certain messages. Refer to TABLE 6-6.

-I*dir*

Searches the directory *dir* for included header files.

-k

Alter the behavior of `/* LINTED [message] */` directives or `NOTE(LINTED(message))` annotations. Normally, `lint` suppresses warning messages for the code following these directives. Instead of suppressing the messages, `lint` prints an additional message containing the comment inside the directive or annotation.

-L*dir*

Searches for a `lint` library in the directory *dir* when used with `-l`.

-lx

Accesses the lint library `llib-lx.ln`.

-m

Suppresses certain messages. Refer to TABLE 6-6.

-Ncheck=c

Checks header files for corresponding declarations; checks macros. *c* is a comma-separated list of checks that consists of one or more of the following: `macro`, `extern`, `%all`, `%none`, `no%macro`, `no%extern`.

TABLE 6-5 The -Ncheck Values

Value	Meaning
<code>macro</code>	Checks for consistency of macro definitions across files
<code>extern</code>	Checks for one-to-one correspondence of declarations between source files and their associated header files (for example, for <code>file1.c</code> and <code>file1.h</code>). Ensure that there are neither extraneous nor missing <code>extern</code> declarations in a header file.
<code>%all</code>	Performs all of -Ncheck's checks
<code>%none</code>	Performs none of -Ncheck's checks. This is the default.
<code>no%macro</code>	Performs none of -Ncheck's macro checks
<code>no%extern</code>	Performs none of -Ncheck's extern checks

The default is `-Ncheck=%none`. Specifying `-Ncheck` is equivalent to specifying `-Ncheck=%all`.

Values may be combined with a comma, for example, `-Ncheck=extern,macro`.

Example:

```
% lint -Ncheck=%all,no%macro
```

performs all checks except macro checks.

`-Nlevel=n`

Specifies the level of analysis for reporting problems. This option allows you to control the amount of detected errors. The higher the level, the longer the verification time. *n* is a number: 1, 2, 3, or 4. The default is `-Nlevel=2`. Specifying `-Nlevel` is equivalent to specifying `-Nlevel=4`.

`-Nlevel=1`

Analyzes single procedures. Reports unconditional errors that occur on some program execution paths. Does not do global data and control flow analysis.

`-Nlevel=2`

The default. Analyzes the whole program, including global data and control flow. Reports unconditional errors that occur on some program execution paths.

`-Nlevel=3`

Analyzes the whole program, including constant propagation, cases when constants are used as actual arguments, as well as the analysis performed under `-Nlevel=2`.

Verification of a C program at this analysis level takes two to four times longer than at the preceding level. The extra time is required because lint assumes partial interpretation of the program by creating sets of possible values for program variables. These sets of variables are created on the basis of constants and conditional statements that contain constant operands available in the program. The sets form the basis for creating other sets (a form of constant propagation). Sets received as the result of the analysis are evaluated for correctness according to the following algorithm:

If a correct value exists among all possible values of an object, then that correct value is used as the basis for further propagation; otherwise an error is diagnosed.

`-Nlevel=4`

Analyzes the whole program, and reports conditional errors that could occur when certain program execution paths are used, as well as the analysis performed under `-Nlevel=3`.

At this analysis level, there are additional diagnostic messages. The analysis algorithm generally corresponds to the analysis algorithm of `-Nlevel=3` with the exception that any invalid values now generate an error message. The amount of

time required for analysis at this level can increase as much as two orders (about 20 to 100 times more slowly). In this case the extra time required is directly proportional to the program complexity as characterized by recursion, conditional statements etc. As a result of this, it may be difficult to use this level of analysis for a program that exceeds 100,000 lines.

-n

Suppresses checks for compatibility with the default `lint` standard C library.

-oX

Causes `lint` to create a `lint` library with the name `llib-lx.ln`. This library is created from all the `.ln` files that `lint` used in its second pass. The `-c` option nullifies any use of the `-o` option. To produce a `llib-lx.ln` without extraneous messages, you can use the `-x` option. The `-v` option is useful if the source file(s) for the `lint` library are just external interfaces. The `lint` library produced can be used later if `lint` is invoked with `-lx`.

By default, you create libraries in `lint`'s basic format. If you use `lint`'s enhanced mode, the library created will be in enhanced format, and can only be used in enhanced mode.

-p

Enables certain messages relating to portability issues.

-Rfile

Write a `.ln` file to *file*, for use by `cxref(1)`. This option disables the enhanced mode, if it is switched on.

-S

Converts compound messages into simple ones.

`-u`

Suppresses certain messages. Refer to TABLE 6-6. This option is suitable for running `lint` on a subset of files of a larger program.

`-V`

Writes the product name and releases to standard error.

`-v`

Suppresses certain messages. Refer to TABLE 6-6.

`-wfile`

Write a `.ln` file to *file*, for use by `cf1ow(1)`. This option disables the enhanced mode, if it is switched on.

`-X`

Suppresses certain messages. Refer to TABLE 6-6.

`-XCC=a`

Accepts C++-style comments. In particular, `//` can be used to indicate the start of a comment. *a* can be either `yes` or `no`. The default is `-XCC=no`. Specifying `-XCC` is equivalent to specifying `-XCC=yes`.

`-Xalias_level[=l]`

where *l* is one of `any`, `basic`, `weak`, `layout`, `strict`, `std`, or `strong`. See “The Levels of Alias-Disambiguation” on page 41 for a detailed explanation of the different levels of disambiguation.

If you do not specify `-Xalias_level`, the default of the flag is `-Xalias_level=any`. This means that there is no type-based alias-analysis. If you specify `-Xalias_level` but do not supply a level, the default is `-Xalias_level=layout`.

Be sure to run lint at a level of disambiguation that is no more strict than the level at which you ran the compiler. If you run lint at a level of disambiguation that is more strict than the level at which you compiled, the results will be difficult to interpret and possibly misleading.

See Chapter 7 for a detailed explanation of disambiguation as well as a list of pragmas designed to help with disambiguation.

`-Xarch=v9`

Predefines the `__sparcv9` macro and searches for v9 versions of lint libraries.

`-Xexplicitpar=a`

(SPARC) Directs lint to recognize `#pragma MP` directives. *a* can be either `yes` or `no`. The default is `-Xexplicitpar=no`. Specifying `-Xexplicitpar` is equivalent to specifying `-Xexplicitpar=yes`.

`-Xkeeptmp=a`

Keeps temporary files created during linting instead of deleting them automatically. *a* can be either `yes` or `no`. The default is `-Xkeeptmp=no`. Specifying `-Xkeeptmp` is equivalent to specifying `-Xkeeptmp=yes`.

`-Xtemp=dir`

Sets the directory for temporary files to *dir*. Without this option, temporary files go into `/tmp`.

`-Xtime=a`

Reports the execution time for each lint pass. *a* can be either `yes` or `no`. The default is `-Xtime=no`. Specifying `-Xtime` is equivalent to specifying `-Xtime=yes`.

`-Xtransition=a`

Issues warnings for the differences between K&R C and Sun ANSI/ISO C. *a* can be either `yes` or `no`. The default is `-Xtransition=no`. Specifying `-Xtransition` is equivalent to specifying `-Xtransition=yes`.

`-y`

Treats every `.c` file named on the command line as if it begins with the directive `/* LINTLIBRARY */` or the annotation `NOTE(LINTLIBRARY)`. A `lint` library is normally created using the `/* LINTLIBRARY */` directive or the `NOTE(LINTLIBRARY)` annotation.

lint Messages

Most of `lint`'s messages are simple, one-line statements printed for each occurrence of the problem they diagnose. Errors detected in included files are reported multiple times by the compiler, but only once by `lint`, no matter how many times the file is included in other source files. Compound messages are issued for inconsistencies across files and, in a few cases, for problems within them as well. A single message describes every occurrence of the problem in the file or files being checked. When use of a `lint` filter (see “`lint` Libraries” on page 186) requires that a message be printed for each occurrence, compound diagnostics can be converted to the simple type by invoking `lint` with the `-s` option.

`Lint`'s messages are written to `stderr`.

The Error and Warning Messages File, located in `/opt/SUNWspro/READMEs/c_lint_messages`, contains all the C compiler error and warning messages and all the `lint` program's messages. Many of the messages are self-explanatory. You can obtain a description of the messages and, in many cases, code examples, by searching the text file for a string from the message that was generated.

Options to Suppress Messages

You can use several `lint` options to suppress `lint` diagnostic messages. Messages can be suppressed with the `-erroff` option, followed by one or more *tags*. These mnemonic tags can be displayed with the `-errtags=yes` option.

The following table lists the options that suppress lint messages.

TABLE 6-6 lint Options to Suppress Messages

Option	Messages Suppressed
-a	assignment causes implicit narrowing conversion conversion to larger integral type may sign-extend incorrectly
-b	statement not reached (unreachable break and empty statements)
-h	assignment operator "=" found where equality operator "==" was expected constant operand to op: "!" fallthrough on case statements pointer cast may result in improper alignment precedence confusion possible; parenthesize statement has no consequent: if statement has no consequent: else
-m	declared global, could be static
-erroff= <i>tag</i>	One or more lint messages specified by <i>tag</i>
-u	name defined but never used name used but not defined
-v	arguments unused in function
-x	name declared but never used or defined

lint Message Formats

The lint program can, with certain options, show precise source file lines with pointers to the line position where the error occurred. The option enabling this feature is `-errfmt=f`. Under this option, lint provides the following information:

- Source line(s) and position(s)
- Macro unfolding
- Error-prone stack(s)

For example, the following program, `Test1.c`, contains an error.

```
1 #include <string.h>
2 static void cpv(char *s, char* v, unsigned n)
3 { int i;
4   for (i=0; i<=n; i++){
5     *v++ = *s++;}
6 }
7 void main(int argc, char* argv[])
8 {
9   if (argc != 0){
10    cpv(argv[0], argc, strlen(argv[0]));}
11}
```

Using `lint` on `Test1.c` with the option:

```
% lint -errfmt=src -Nlevel=2 Test1.c
```

produces the following output:

```
static void cpv(char *s, char* v, unsigned n)
|          ^ line 2, Test1.c
|
|          cpv(argv[0], argc, strlen(argv[0]));
|          ^ line 10, Test1.c
warning: improper pointer/integer combination: arg #2
|
| static void cpv(char *s, char* v, unsigned n)
|          ^ line 2, Test1.c
|
| cpv(argv[0], argc, strlen(argv[0]));
|          ^ line 10, Test1.c
|
|          *v++ = *s++;
|          ^ line 5, Test1.c
warning: use of a pointer produced in a questionable way
v defined at Test1.c(2)::Test1.c(5)
call stack:
main()          ,Test1.c(10)
cpv()           ,Test1.c(5)
```

The first warning indicates two source lines that are contradictory. The second warning shows the call stack, with the control flow leading to the error.

Another program, `Test2.c`, contains a different error:

```
1 #define AA(b) AR[b+1]
2 #define B(c,d) c+AA(d)
3
4 int x=0;
5
6 int AR[10]={1,2,3,4,5,6,77,88,99,0};
7
8 main()
9 {
10  int y=-5, z=5;
11  return B(y,z);
12 }
```

Using `lint` on `Test2.c` with the option:

```
% lint -errfmt=macro Test2.c
```

produces the following output, showing the steps of macro substitution:

```
| return B(y,z);
|         ^ line 11, Test2.c
|
| #define B(c,d) c+AA(d)
|                 ^ line 2, Test2.c
|
| #define AA(b) AR[b+1]
|                         ^ line 1, Test2.c
error: undefined symbol: l
|
|         return B(y,z);
|                 ^ line 11, Test2.c
|
| #define B(c,d) c+AA(d)
|                 ^ line 2, Test2.c
|
| #define AA(b) AR[b+1]
|                         ^ line 1, Test2.c
variable may be used before set: l
lint: errors in Test2.c; no output created
lint: pass2 not run - errors in Test2.c
```

lint Directives

Predefined Values

The following predefinitions are valid in all modes:

- `__sun`
- `__unix`
- `__lint`
- `__SUNPRO_C=0x530`
- `__'uname -s'_'uname -r'` (example: `__SunOS_5_7`)
- `__RESTRICT` (`-Xa` and `-Xt` modes only)
- `__sparc` (*SPARC*)
- `__i386` (*Intel*)
- `__BUILTIN_VA_ARG_INCR`
- `__SVR4`
- `__sparcv9` (`-Xarch=v9`)

These predefinitions are not valid in `-Xc` mode:

- `sun`
- `unix`
- `sparc` (*SPARC*)
- `i386` (*Intel*)
- `lint`

Directives

`lint` directives in the form of `/*...*/` are supported for existing annotations, but will not be supported for future annotations. Directives in the form of source code annotations, `NOTE(...)`, are recommended for all annotations.

Specify `lint` directives in the form of source code annotations by including the file `note.h`, for example:

```
#include <note.h>
```

Lint shares the Source Code Annotations scheme with several other tools. When you install the Sun C compiler, you also automatically install the file `/usr/lib/note/SUNW_SPRO-lint`, which contains the names of all the annotations that LockLint

understands. However, the Sun C source code checker, `lint`, also checks all the files in `/usr/lib/note` and `/opt/SUNWspr/<current-release>/note` for all valid annotations.

You may specify a location other than `/usr/lib/note` by setting the environment variable `NOTEPATH`, as in:

```
setenv NOTEPATH $NOTEPATH:other_location
```

The following table lists the `lint` directives along with their actions.

TABLE 6-7 `lint` Directives

Directive	Action
<code>NOTE(ALIGNMENT(<i>fname</i>,<i>n</i>))</code> where <i>n</i> =1, 2, 4, 8, 16, 32, 64, 128	Makes <code>lint</code> set the following function result alignment in <i>n</i> bytes. For example, <code>malloc()</code> is defined as returning a <code>char*</code> or <code>void*</code> when in fact it really returns pointers that are word, or even doubleword, aligned. Suppresses the following message: <ul style="list-style-type: none"> • improper alignment
<code>NOTE(ARGSUSED(<i>n</i>))</code> <code>/*ARGSUSEDn*/</code>	This directive acts like the <code>-v</code> option for the next function. Suppresses the following message for every argument but the first <i>n</i> in the function definition it precedes. Default is 0. For the <code>NOTE</code> format, <i>n</i> must be specified. <ul style="list-style-type: none"> • argument unused in function
<code>NOTE(ARGUNUSED(<i>par_name</i>[,<i>par_name</i>...]))</code>	Makes <code>lint</code> not check the mentioned arguments for usage (this option acts only for the next function). Suppresses the following message for every argument listed in <code>NOTE</code> or directive. <ul style="list-style-type: none"> • argument unused in function
<code>NOTE(CONSTCOND)</code> <code>/*CONSTCOND*/</code>	Suppresses complaints about constant operands for the conditional expression. Suppresses the following messages for the constructs it precedes. Also <code>NOTE(CONSTANTCONDITION)</code> or <code>/* CONSTANTCONDITION */</code> . constant in conditional context constant operands to <code>op</code> : <code>!"</code> logical expression always false: <code>op "&&"</code> logical expression always true: <code>op " "</code>

TABLE 6-7 lint Directives (Continued)

Directive	Action
NOTE(EMPTY) /*EMPTY*/	<p>Suppresses complaints about a null statement consequent on an if statement. This directive should be placed after the test expression, and before the semicolon. This directive is supplied to support empty if statements when a valid else statement follows. It suppresses messages on an empty else consequent.</p> <p>Suppresses the following messages when inserted between the controlling expression of the if and semicolon.</p> <ul style="list-style-type: none"> • statement has no consequent: else when inserted between the else and semicolon; • statement has no consequent: if
NOTE(FALLTHRU) /*FALLTHRU*/	<p>Suppresses complaints about a fall through to a case or default labelled statement. This directive should be placed immediately preceding the label.</p> <p>Suppresses the following message for the case statement it precedes. Also NOTE(FALLTHROUGH) or /* FALLTHROUGH */.</p> <ul style="list-style-type: none"> • fallthrough on case statement
NOTE(LINTED (msg)) /*LINTED [msg]*/	<p>Suppresses any intra-file warning except those dealing with unused variables or functions. This directive should be placed on the line immediately preceding where the lint warning occurred. The -k option alters the way in which lint handles this directive. Instead of suppressing messages, lint prints an additional message, if any, contained in the comments. This directive is useful in conjunction with the -s option for post-lint filtering.</p> <p>When -k is not invoked, suppresses every warning pertaining to an intra-file problem, except:</p> <ul style="list-style-type: none"> • argument unused in function • declarations unused in block • set but not used in function • static unused • variable not used in function <p>for the line of code it precedes. msg is ignored.</p>
NOTE(LINTLIBRARY) /*LINTLIBRARY*/	<p>When -o is invoked, writes to a library .ln file only definitions in the .c file it heads. This directive suppresses complaints about unused functions and function arguments in this file.</p>

TABLE 6-7 lint Directives (*Continued*)

Directive	Action
NOTE(NOTREACHED) /*NOTREACHED*/	<p>At appropriate points, stops comments about unreachable code. This comment is typically placed just after calls to functions such as <code>exit(2)</code>.</p> <p>Suppresses the following messages for the closing curly brace it precedes at the end of the function.</p> <ul style="list-style-type: none">• <code>statement not reached</code> <p>for the unreachable statements it precedes;</p> <ul style="list-style-type: none">• <code>fallthrough</code> on case statement <p>for the case it precedes that cannot be reached from the preceding case;</p> <ul style="list-style-type: none">• <code>function falls off bottom without returning value</code>
NOTE(PRINTFLIKE(<i>n</i>)) NOTE(PRINTFLIKE(<i>fun_name,n</i>)) /*PRINTFLIKE <i>n</i> */	<p>Treats the <i>n</i>th argument of the function definition it precedes as a <code>[fs]printf()</code> format string and issues the following messages for mismatches between the remaining arguments and the conversion specifications. <code>lint</code> issues these warnings by default for errors in the calls to <code>[fs]printf()</code> functions provided by the standard C library.</p> <p>For the NOTE format, <i>n</i> must be specified.</p> <ul style="list-style-type: none">• <code>malformed format strings</code> <p>for invalid conversion specifications in that argument, and function argument type inconsistent with format;</p> <ul style="list-style-type: none">• <code>too few arguments for format</code>• <code>too many arguments for format</code>

TABLE 6-7 lint Directives (Continued)

Directive	Action
NOTE(PROTOLIB(<i>n</i>)) /*PROTOLIB <i>n</i> */	When <i>n</i> is 1 and NOTE(LINTLIBRARY) or /*LINTLIBRARY */ is used, writes to a library .ln file only function prototype declarations in the .c file it heads. The default is 0, which cancels the process. For the NOTE format, <i>n</i> must be specified.
NOTE(SCANFLIKE(<i>n</i>)) NOTE(SCANLIKE(<i>fun_name</i> , <i>n</i>)) /*SCANFLIKE <i>n</i> */	Same as NOTE(PRINTFLIKE(<i>n</i>)) or /*PRINTFLIKE <i>n</i> */, except that the <i>n</i> th argument of the function definition is treated as a [fs]scanf() format string. By default, lint issues warnings for errors in the calls to [fs]scanf() functions provided by the standard C library. For the NOTE format, <i>n</i> must be specified.
NOTE(VARARGS(<i>n</i>)) NOTE(VARARGS(<i>fun_name</i> , <i>n</i>)) /*VARARGS <i>n</i> */	Suppresses the usual checking for variable numbers of arguments in the following function declaration. The data types of the first <i>n</i> arguments are checked; a missing <i>n</i> is taken to be 0. The use of the ellipsis (...) terminator in the definition is suggested in new or updated code. For the function whose definition it precedes, suppresses the following message for calls to the function with <i>n</i> or more arguments. For the NOTE format, <i>n</i> must be specified. <ul style="list-style-type: none"> • functions called with variable number of arguments

lint Reference and Examples

This section provides reference information on lint, including checks performed by lint, lint libraries, and lint filters.

Diagnostics Performed by lint

lint-specific diagnostics are issued for three broad categories of conditions: inconsistent use, nonportable code, and questionable constructs. In this section, we review examples of lint's behavior in each of these areas, and suggest possible responses to the issues they raise.

Consistency Checks

Inconsistent use of variables, arguments, and functions is checked within files as well as across them. Generally speaking, the same checks are performed for prototype uses, declarations, and parameters as `lint` checks for old-style functions. If your program does not use function prototypes, `lint` checks the number and types of parameters in each call to a function more strictly than the compiler. `lint` also identifies mismatches of conversion specifications and arguments in `[fs]printf()` and `[fs]scanf()` control strings.

Examples:

- Within files, `lint` flags non-void functions that “fall off the bottom” without returning a value to the invoking function. In the past, programmers often indicated that a function was not meant to return a value by omitting the return type: `fun() {}`. That convention means nothing to the compiler, which regards `fun()` as having the return type `int`. Declare the function with the return type `void` to eliminate the problem.
- Across files, `lint` detects cases where a non-void function does not return a value, yet is used for its value in an expression—and the opposite problem, a function returning a value that is sometimes or always ignored in subsequent calls. When the value is *always* ignored, it may indicate an inefficiency in the function definition. When it is *sometimes* ignored, it's probably bad style (typically, not testing for error conditions). If you need not check the return values of string functions like `strcat()`, `strcpy()`, and `sprintf()`, or output functions like `printf()` and `putchar()`, cast the offending calls to `void`.
- `lint` identifies variables or functions that are declared but not used or defined; used, but not defined; or defined, but not used. When `lint` is applied to some, but not all files of a collection to be loaded together, it produces error messages about functions and variables that are:
 - Declared in those files, but defined or used elsewhere
 - Used in those files, but defined elsewhere
 - Defined in those files, but used elsewhere

Invoke the `-x` option to suppress the first complaint, `-u` to suppress the latter two.

Portability Checks

Some nonportable code is flagged by `lint` in its default behavior, and a few more cases are diagnosed when `lint` is invoked with `-p` or `-xc`. The latter causes `lint` to check for constructs that do not conform to the ANSI/ISO C standard. For the messages issued under `-p` and `-xc`, see “`lint` Libraries” on page 186.

Examples:

- In some C language implementations, character variables that are not explicitly declared `signed` or `unsigned` are treated as signed quantities with a range typically from `-128` to `127`. In other implementations, they are treated as nonnegative quantities with a range typically from `0` to `255`. So the test:

```
char c;  
c = getchar();  
if (c == EOF) ...
```

where `EOF` has the value `-1`, always fails on machines where character variables take on nonnegative values. `lint` invoked with `-p` checks all comparisons that imply a *plain* `char` may have a negative value. However, declaring `c` as a `signed char` in the above example eliminates the diagnostic, not the problem. That's because `getchar()` must return all possible characters and a distinct `EOF` value, so a `char` cannot store its value. We cite this example, perhaps the most common one arising from implementation-defined sign-extension, to show how a thoughtful application of `lint`'s portability option can help you discover bugs not related to portability. In any case, declare `c` as an `int`.

- A similar issue arises with bit-fields. When constant values are assigned to bit-fields, the field may be too small to hold the value. On a machine that treats bit-fields of type `int` as unsigned quantities, the values allowed for `int x:3` range from `0` to `7`, whereas on machines that treat them as signed quantities, they range from `-4` to `3`. However, a three-bit field declared type `int` cannot hold the value `4` on the latter machines. `lint` invoked with `-p` flags all bit-field types other than `unsigned int` or `signed int`. These are the only *portable* bit-field types. The compiler supports `int`, `char`, `short`, and `long` bit-field types that may be `unsigned`, `signed`, or *plain*. It also supports the `enum` bit-field type.
- Bugs can arise when a larger-sized type is assigned to a smaller-sized type. If significant bits are truncated, accuracy is lost:

```
short s;  
long l;  
s = l;
```

`lint` flags all such assignments by default; the diagnostic can be suppressed by invoking the `-a` option. Bear in mind that you may be suppressing other diagnostics when you invoke `lint` with this or any other option. Check the list in “`lint` Libraries” on page 186 for the options that suppress more than one diagnostic.

- A cast of a pointer to one object type to a pointer to an object type with stricter alignment requirements may not be portable. `lint` flags:

```
int *fun(y)
char *y;
{
    return(int *)y;
}
```

because, on most machines, an `int` cannot start on an arbitrary byte boundary, whereas a `char` can. You can suppress the diagnostic by invoking `lint` with `-h`, although, again, you may be disabling other messages. Better still, eliminate the problem by using the generic pointer `void *`.

- ANSI/ISO C leaves the order of evaluation of complicated expressions undefined. That is, when function calls, nested assignment statements, or the increment and decrement operators cause side effects when a variable is changed as a by-product of the evaluation of an expression, the order in which the side effects take place is highly machine-dependent. By default, `lint` flags any variable changed by a side effect and used elsewhere in the same expression:

```
int a[10];
main()
{
    int i = 1;
    a[i++] = i;
}
```

In this example, the value of `a[1]` may be 1 if one compiler is used, 2 if another. The bitwise logical operator `&` can give rise to this diagnostic when it is mistakenly used in place of the logical operator `&&`:

```
if ((c = getchar()) != EOF & c != '0')
```

Questionable Constructs

`lint` flags a miscellany of legal constructs that may not represent what the programmer intended. Examples:

- An unsigned variable always has a nonnegative value. So the test:

```
unsigned x;
if (x < 0) ...
```

always fails. The test:

```
unsigned x;  
if (x > 0) ...
```

is equivalent to:

```
if (x != 0) ...
```

This may not be the intended action. `lint` flags questionable comparisons of unsigned variables with negative constants or 0. To compare an unsigned variable to the bit pattern of a negative number, cast it to unsigned:

```
if (u == (unsigned) -1) ...
```

Or use the `U` suffix:

```
if (u == -1U) ...
```

- `lint` flags expressions without side effects that are used in a context where side effects are expected—that is, where the expression may not represent what the programmer intends. It issues an additional warning whenever the equality operator is found where the assignment operator is expected—that is, where a side effect is expected:

```
int fun()  
{  
    int a, b, x, y;  
    (a = x) && (b == y);  
}
```

- `lint` cautions you to parenthesize expressions that mix both the logical and bitwise operators (specifically, `&`, `|`, `^`, `<<`, `>>`), where misunderstanding of operator precedence may lead to incorrect results. Because the precedence of bitwise `&`, for example, falls below logical `==`, the expression:

```
if (x & a == 0) ...
```

is evaluated as:

```
if (x & (a == 0)) ...
```

which is most likely not what you intended. Invoking `lint` with `-h` disables the diagnostic.

lint Libraries

You can use `lint` libraries to check your program for compatibility with the library functions you have called in it—the declaration of the function return type, the number and types of arguments the function expects, and so on. The standard `lint` libraries correspond to libraries supplied by the C compilation system, and generally are stored in a standard place on your system. By convention, `lint` libraries have names of the form `llib-lx.ln`.

The `lint` standard C library, `llib-1c.ln`, is appended to the `lint` command line by default; checks for compatibility with it can be suppressed by invoking the `-n` option. Other `lint` libraries are accessed as arguments to `-l`. That is:

```
% lint -lx file1.c file2.c
```

directs `lint` to check the usage of functions and variables in `file1.c` and `file2.c` for compatibility with the `lint` library `llib-lx.ln`. The library file, which consists only of definitions, is processed exactly as are ordinary source files and ordinary `.ln` files, except that functions and variables used inconsistently in the library file, or defined in the library file but not used in the source files, elicit no complaints.

To create your own `lint` library, insert the directive `NOTE(LINTLIBRARY)` at the head of a C source file, then invoke `lint` for that file with the `-o` option and the library name given to `-l`:

```
% lint -ox file1.c file2.c
```

causes only definitions in the source files headed by `NOTE(LINTLIBRARY)` to be written to the file `llib-lx.ln`. (Note the analogy of `lint -o` to `cc -o`.) A library can be created from a file of function prototype declarations in the same way, except that both `NOTE(LINTLIBRARY)` and `NOTE(PROTOLIB(n))` must be inserted at the head of the declarations file. If `n` is 1, prototype declarations are written to a library

.ln file just as are old-style definitions. If *n* is 0, the default, the process is cancelled. Invoking `lint` with `-y` is another way of creating a `lint` library. The command line:

```
% lint -y -ox file1.c file2.c
```

causes each source file named on that line to be treated as if it begins with `NOTE(LINTLIBRARY)`, and only its definitions to be written to `llib-1x.ln`.

By default, `lint` searches for `lint` libraries in the standard place. To direct `lint` to search for a `lint` library in a directory other than the standard place, specify the path of the directory with the `-L` option:

```
% lint -Ldir -1x file1.c file2.c
```

In enhanced mode, `lint` produces .ln files which store additional information than .ln files produced in basic mode. In enhanced mode, `lint` can read and understand all .ln files generated by either basic or enhanced `lint` modes. In basic mode, `lint` can read and understand .ln files generated only using basic `lint` mode.

By default, `lint` uses libraries from the `/usr/lib` directory. These libraries are in the basic `lint` format. You can run a `makefile` once, and create enhanced `lint` libraries in a new format, which will enable enhanced `lint` to work more effectively. To run the `makefile` and create the new libraries, enter the command:

```
% cd /opt/SUNWspro/WS6/src/lintlib; make
```

where `/opt/SUNWspro/WS6` is the installation directory. After the `makefile` is run, `lint` will use the new libraries in enhanced mode, instead of the libraries in the `/usr/lib` directory.

The specified directory is searched before the standard place.

lint Filters

A `lint` filter is a project-specific post-processor that typically uses an `awk` script or similar program to read the output of `lint` and discard messages that your project has deemed as *not* identifying real problems—string functions, for instance, returning values that are sometimes or always ignored. `lint` filters generate customized diagnostic reports when `lint` options and directives do not provide sufficient control over output.

Two options to `lint` are particularly useful in developing a filter:

- Invoking `lint` with `-s` causes compound diagnostics to be converted into simple, one-line messages issued for each occurrence of the problem diagnosed. The easily parsed message format is suitable for analysis by an `awk` script.
- Invoking `lint` with `-k` causes certain comments you have written in the source file to be printed in output, and can be useful both in documenting project decisions and specifying the post-processor's behavior. In the latter instance, if the comment identifies an expected `lint` message, and the reported message is the same, the message can be filtered out. To use `-k`, insert on the line preceding the code you wish to comment the `NOTE(LINTED(msg))` directive, where `msg` refers to the comment to be printed when `lint` is invoked with `-k`.

Refer to the list of directives in TABLE 6-7 for an explanation of what `lint` does when `-k` is *not* invoked for a file containing `NOTE(LINTED(msg))`.

Type-Based Alias Analysis

This document explains how to use the `-xalias_level` option and several new pragmas to enable the compiler to perform type-based alias analysis and optimizations. You use these extensions to express type-based information about the way pointers are used in your C program. The C compiler uses this information, in turn, to do a significantly better job of alias disambiguation for pointer-based memory references in your program.

This document contains the following sections:

- “Introduction to Type-Based Analysis” on page 189
- “Using Pragmas For Finer Control” on page 190
- “Checking With `lint`” on page 193
- “Examples of Memory Reference Constraints” on page 196

See “`-xalias_level[=l]`” on page 40 for a detailed explanation of this command’s syntax. Also, see “`-xalias_level[=l]`” on page 171 for an explanation of the `lint` program’s type-based alias-analysis capabilities.

Introduction to Type-Based Analysis

You can use the `-xalias_level` option to specify one of seven alias levels. Each level specifies a certain set of properties about the way you use pointers in your C program.

As you compile with higher levels of the `-xalias_level` option, the compiler makes increasingly extensive assumptions about the pointers in your code. You have greater programming freedom when the compiler makes fewer assumptions. However, the optimizations that result from these narrow assumptions may not result in significant runtime performance improvement. If you code in accordance

with the compiler assumptions of the more advanced levels of the `-xalias_level` option, there is a greater chance that the resulting optimizations will enhance runtime performance.

The `-xalias_level` option specifies which alias level applies to each translation unit. For cases where more detail is beneficial, you can use new pragmas to override whatever alias levels are in effect so that you can explicitly specify the aliasing relationships between individual types or pointer variables in the translation unit. These pragmas are most useful when the pointer usage in a translation unit is covered by one of the available alias levels, but a few specific pointer variables are used in an irregular way that is not allowed by one of the available levels.

Using Pragmas for Finer Control

For cases in which type-based analysis can benefit from more detail, you can use the following pragmas to override the alias level in effect and specify the aliasing relationships between individual types or pointer variables in the translation unit. These pragmas provide the most benefit when the use of pointers in a translation unit is consistent with one of the available alias levels, but a few specific pointer variables are used in an irregular way not allowed by one of the available levels.

Note – You must declare the named type or variable prior to the pragma or a warning message is issued and the pragma is ignored. The results of the program are undefined if the pragma appears after the first memory reference to which its meaning applies.

The following terms are used in the pragma definitions.

Term	Meaning
<i>level</i>	Any of the alias levels listed under “ <code>-xalias_level[=l]</code> ” on page 40.
<i>type</i>	Any of the following: <ul style="list-style-type: none">• <code>char</code>, <code>short</code>, <code>int</code>, <code>long</code>, <code>long long</code>, <code>float</code>, <code>double</code>, <code>long double</code>• <code>void</code>, which denotes all pointer types• <code>typedef name</code>, which is the name of a defined type from a <code>typedef</code> declaration• <code>struct name</code>, which is the keyword <code>struct</code> followed by a <i>struct tag</i> name• <code>union</code>, which is the keyword <code>union</code> followed by a <i>union tag</i> name
<i>pointer_name</i>	The name of any variable of pointer type in the translation unit.

`#pragma alias_level level list`

Replace *level* with one of the seven alias levels: *any*, *basic*, *weak*, *layout*, *strict*, *std*, or *strong*. You can replace *list* with either a single type or a comma-delimited list of types, or you can replace *list* with either a single pointer or a comma-delimited list of pointers. For example, you can issue `#pragma alias_level` as follows:

- `#pragma alias_level level (type [, type])`
- `#pragma alias_level level (pointer [, pointer])`

This pragma specifies that the indicated alias level applies either to all of the memory references of the translation unit for the listed types, or to all of the dereferences of the translation unit where any of the named pointer variables are being dereferenced.

If you specify more than one alias level to be applied to a particular dereference, the level that is applied by the pointer name, if any, has precedence over all other levels. The level applied by the type name, if any, has precedence over the level applied by the option. In the following example, the *std* level applies to *p* if the program is compiled with `#pragma alias_level` set higher than *any*.

```
typedef int * int_ptr;
int_ptr p;
#pragma alias_level strong (int_ptr)
#pragma alias_level std (p)
```

`#pragma alias (type, type [, type]...)`

This pragma specifies that all the listed types alias each other. In the following example, the compiler assumes that the indirect access `*pt` aliases the indirect access `*pf`.

```
#pragma alias (int, float)
int *pt;
float *pf;
```

`#pragma alias (pointer, pointer [, pointer]...)`

This pragma specifies that at the point of any dereference of any of the named pointer variables, the pointer value being dereferenced can point to the same object as any of the other named pointer variables. This pragma overrides the aliasing

assumptions of any applied alias levels. In the following example, the compiler assumes that the indirect access `*p` aliases the indirect access `*q` regardless of the types of the two pointers.

```
#pragma alias(p, q)
```

```
#pragma may_point_to (pointer, variable  
[, variable]...)
```

This pragma specifies that at the point of any dereference of the named pointer variable, the pointer value being dereferenced can point to the objects that are contained in any of the named variables. This pragma overrides the aliasing assumptions of any applied alias levels. In the following example, the compiler assumes that the indirect access `*p` aliases the direct accesses `a`, `b`, and `c`.

```
#pragma alias may_point_to(p, a, b, c)
```

```
#pragma noalias (type, type [, type]...)
```

This pragma specifies that the listed types do not alias each other. In the following example, the compiler assumes that the indirect access `*p` does not alias the indirect access `*ps`.

```
struct S {  
    float f;  
    ...} *ps;  
  
#pragma noalias(int, struct S)  
int *p;
```

```
#pragma noalias (pointer, pointer  
[, pointer]...)
```

This pragma specifies that at the point of any dereference of any of the named pointer variables, the pointer value being dereferenced does not point to the same object as any of the other named pointer variables. This pragma overrides all other

applied alias levels. In the following example, the compiler assumes that the indirect access `*p` does not alias the indirect access `*q` regardless of the types of the two pointers.

```
#pragma noalias(p, q)
```

```
#pragma may_not_point_to (pointer, variable  
[, variable]...)
```

This pragma specifies that at the point of any dereference of the named pointer variable, the pointer value being dereferenced does not point to the objects that are contained in any of the named variables. This pragma overrides all other applied alias levels. In the following example, the compiler assumes that the indirect access `*p` does not alias the direct accesses `a`, `b`, or `c`.

```
#pragma may_not_point_to(p, a, b, c)
```

Checking With `lint`

The `lint` program recognizes the same levels of type-based alias-disambiguation as the compiler's `-xalias_level` command. The `lint` program also recognizes the pragmas related to type-based alias-disambiguation documented in this chapter. For a detailed explanation of the `lint -xalias_level` command, see "`-xalias_level[=l]`" on page 171.

There are four situations that `lint` detects and for which it generates warnings:

- Casting a scalar pointer to a struct pointer
- Casting a void pointer to a struct pointer
- Casting a structure field to a scalar pointer
- Casting a struct pointer to a struct pointer at the level of `-xalias_level=strict` without explicit aliasing.

Struct Pointer Cast of Scalar Pointer

In the following example, the pointer `p` of type `integer`, is cast as a pointer of type `struct foo`. With lint `-Xalias_level=weak` (or higher), this generates a warning.

```
struct foo {
    int a;
    int b;
};

struct foo *f;
int *p;

void main()
{
    f = (struct foo *)p; /* struct pointer cast of scalar pointer
warning */
}
```

Struct Pointer Cast of Void Pointer

In the following example, the pointer `vp` of type `void`, is cast as pointer of type `struct foo`. With lint `-Xalias_level=weak` (or higher), this generates a warning.

```
struct foo {
    int a;
    int b;
};

struct foo *f;
void *vp;

void main()
{
    f = (struct foo *)vp; /* struct pointer cast of void pointer
warning */
}
```

Cast of Struct Field to Structure Pointer

In the following example, the address of structure member `foo.b` is being cast as an integer pointer and then assigned to `p`. With lint `-Xalias_level=weak` (or higher), this generates a warning.

```
struct foo {
    int a;
    int b;
};

struct foo *f;
int *p;

void main()
{
    p = (int *)&f->b; /* cast of struct field to structure pointer
warning */
}
```

Explicit Aliasing Required

In the following example, the pointer `f1`, of type `struct fooa`, is being cast as a pointer of type `struct foob`. With lint `-Xalias_level=strict` (or higher) such a cast requires explicit aliasing. Use `#pragma alias (struct fooa, struct foob)` before the assignment to `f1` and lint stops generating the warning.

```
struct fooa {
    int a;
};

struct foob {
    int b;
};

struct fooa *f1;
struct foob *f2;

void main()
{
    f1 = (struct fooa *)f2; /* explicit aliasing required warning
*/
}
```

Examples of Memory Reference Constraints

This section provides examples of code that are likely to appear in your source files. Each example is followed by a discussion of the compiler's assumptions about the code as dictated by the applied level of type-based analysis.

Consider the following code which can compile with different levels of aliasing to demonstrate the aliasing relationship of the shown types.

CODE EXAMPLE 7-1

```
struct foo {
    int f1;
    short f2;
    short f3;
    int f4;
} *fp;

struct bar {
    int b1;
    int b2;
    int b3;
} *bp;

int *ip;
short *sp;
```

If CODE EXAMPLE 7-1 is compiled with the `-xalias_level=any` option, the compiler considers the following indirect accesses as aliases to each other:

`*ip, *sp, *fp, *bp, fp->f1, fp->f2, fp->f3, fp->f4, bp->b1, bp->b2, bp->b3`

If CODE EXAMPLE 7-1 is compiled with the `-xalias_level=basic` option, the compiler considers the following indirect accesses as aliases to each other:

`*ip, *bp, fp->f1, fp->f4, bp->b1, bp->b2, bp->b3`

Additionally, `*sp, fp->f2, and fp->f3` can alias each other, and `*sp` and `*fp` can alias each other.

However, under `-xalias_level=basic`, the compiler assumes the following:

- `*ip` does not alias `*sp`.
- `*ip` does not alias `fp->f2` and `fp->f3`.

- `*sp` does not alias `fp->f1`, `fp->f4`, `bp->b1`, `bp->b2`, and `bp->b3`.

The compiler makes these assumptions because the access types of the two indirect accesses are different basic types.

If CODE EXAMPLE 7-1 is compiled with the `-xalias_level=weak` option, the compiler assumes the following alias information:

- `*ip` can alias `*fp`, `fp->f1`, `fp->f4`, `*bp`, `bp->b1`, `bp->b2`, and `bp->b3`.
- `*sp` can alias `*fp`, `fp->f2` and `fp->f3`.
- `fp->f1` can alias `bp->b1`.
- `fp->f4` can alias `bp->b3`.

The compiler assumes that `fp->fp1` does not alias `bp->b2` because `f1` is a field with offset 0 in a structure, whereas `b2` is a field with a 4-byte offset in a structure. Similarly, the compiler assumes that `fp->f1` does not alias `bp->b3`, and `fp->f4` does not alias either `bp->b1` or `bp->b2`.

If CODE EXAMPLE 7-1 is compiled with the `-xalias_level=layout` option, the compiler assumes the following information:

- `*ip` can alias `*fp`, `*bp`, `fp->f1`, `fp->f4`, `bp->b1`, `bp->b2`, and `bp->b3`.
- `*sp` can alias `*fp`, `fp->f2`, and `fp->f3`.
- `fp->f1` can alias `bp->b1` and `*bp`.
- `*fp` and `*bp` can alias each other.

`fp->f4` does not alias `bp->b3` because `f4` and `b3` are not corresponding fields in the common initial sequence of `foo` and `bar`.

If CODE EXAMPLE 7-1 is compiled with the `-xalias_level=strict` option, the compiler assumes the following alias information:

- `*ip` can alias `*fp`, `fp->f1`, `fp->f4`, `*bp`, `bp->b1`, `bp->b2`, and `bp->b3`.
- `*sp` can alias `*fp`, `fp->f2`, and `fp->f3`.

With `-xalias_level=strict`, the compiler assumes that `*fp`, `*bp`, `fp->f1`, `fp->f2`, `fp->f3`, `fp->f4`, `bp->b1`, `bp->b2`, and `bp->b3` do not alias each other because `foo` and `bar` are not the same when field names are ignored. However, `fp` aliases `fp->f1` and `bp` aliases `bp->b1`.

If CODE EXAMPLE 7-1 is compiled with the `-xalias_level=std` option, the compiler assumes the following alias information:

- `*ip` can alias `*fp`, `fp->f1`, `fp->f4`, `*bp`, `bp->b1`, `bp->b2`, and `bp->b3`.
- `*sp` can alias `*fp`, `fp->f2`, and `fp->f3`.

However, `fp->f1` does not alias `bp->b1`, `bp->b2`, or `bp->b3` because `foo` and `bar` are not the same when field names are considered.

If CODE EXAMPLE 7-1 is compiled with the `-xalias_level=strong` option, the compiler assumes the following alias information:

- `*ip` does not alias `fp->f1`, `fp->f4`, `bp->b1`, `bp->b2`, and `bp->b3` because a pointer, such as `*ip`, should not point to the interior of a structure.
- Similarly, `*sp` does not alias `fp->f1` or `fp->f3`.
- `*ip` does not alias `*fp`, `*bp`, and `*sp` due to differing types.
- `*sp` does not alias `*fp`, `*bp`, and `*ip` due to differing types.

Consider the following example source code that can compile with different levels of aliasing to demonstrate the aliasing relationship of the shown types.

CODE EXAMPLE 7-2

```
struct foo {
    int f1;
    int f2;
    int f3;
} *fp;

struct bar {
    int b1;
    int b2;
    int b3;
} *bp;
```

If CODE EXAMPLE 7-2 is compiled with the `-xalias_level=any` option, the compiler assumes the following alias information:

`*fp`, `*bp`, `fp->f1`, `fp->f2`, `fp->f3`, `bp->b1`, `bp->b2` and `bp->b3` all can alias each other because any two memory accesses alias each other at the level of `-xalias_level=any`.

If CODE EXAMPLE 7-2 is compiled with the `-xalias_level=basic` option, the compiler assumes the following alias information:

`*fp`, `*bp`, `fp->f1`, `fp->f2`, `fp->f3`, `bp->b1`, `bp->b2` and `bp->b3` all can alias each other. Any two field accesses using pointers `*fp` and `*bp` can alias each other in this example because all the structure fields are the same basic type.

If CODE EXAMPLE 7-2 is compiled with the `-xalias_level=weak` option, the compiler assumes the following alias information:

- `*fp` and `*bp` can alias each other.
- `fp->f1` can alias `bp->b1`, `*bp` and `*fp`.
- `fp->f2` can alias `bp->b2`, `*bp` and `*fp`.
- `fp->f3` can alias `bp->b3`, `*bp` and `*fp`.

However, `-xalias_level=weak` imposes the following restrictions:

- `fp->f1` does not alias `bp->b2` or `bp->b3` because `f1` has an offset of zero, which is different from that of `b2` (four bytes) and `b3` (eight bytes).
- `fp->f2` does not alias `bp->b1` or `bp->b3` because `f2` has an offset of four bytes, which is different from `b1` (zero bytes) and `b3` (eight bytes).
- `fp->f3` does not alias `bp->b1` or `bp->b2` because `f3` has an offset of eight bytes, which is different from `b1` (zero bytes) and `b2` (four bytes).

If aCODE EXAMPLE 7-2 is compiled with the `-xalias_level=layout` options, the compiler assumes the following alias information:

- `*fp` and `*bp` can alias each other.
- `fp->f1` can alias `bp->b1`, `*bp`, and `*fp`.
- `fp->f2` can alias `bp->b2`, `*bp`, and `*fp`.
- `fp->f3` can alias `bp->b3`, `*bp`, and `*fp`.

However, `-xalias_level=layout` imposes the following restrictions:

- `fp->f1` does not alias `bp->b2` or `bp->b3` because field `f1` corresponds to field `b1` in the common initial sequence of `f00` and `bar`.
- `fp->f2` does not alias `bp->b1` or `bp->b3` because `f2` corresponds to field `b2` in the common initial sequence of `f00` and `bar`.
- `fp->f3` does not alias `bp->b1` or `bp->b2` because `f3` corresponds to field `b3` in the common initial sequence of `f00` and `bar`.

If CODE EXAMPLE 7-2 is compiled with the `-xalias_level=strict` option, the compiler assumes the following alias information:

- `*fp` and `*bp` can alias each other.
- `fp->f1` can alias `bp->b1`, `*bp`, and `*fp`.
- `fp->f2` can alias `bp->b2`, `*bp`, and `*fp`.
- `fp->f3` can alias `bp->b3`, `*bp`, and `*fp`.

However, `-xalias_level=strict` imposes the following restrictions:

- `fp->f1` does not alias `bp->b2` or `bp->b3` because field `f1` corresponds to field `b1` in the common initial sequence of `f00` and `bar`.
- `fp->f2` does not alias `bp->b1` or `bp->b3` because `f2` corresponds to field `b2` in the common initial sequence of `f00` and `bar`.
- `fp->f3` does not alias `bp->b1` or `bp->b2` because `f3` corresponds to field `b3` in the common initial sequence of `f00` and `bar`.

If CODE EXAMPLE 7-2 is compiled with the `-xalias_level=std` option, the compiler assumes the following alias information:

`fp->f1`, `fp->f2`, `fp->f3`, `bp->b1`, `bp->b2`, and `bp->b3` do not alias each other.

If CODE EXAMPLE 7-2 is compiled with the `-xalias_level=strong` option, the compiler assumes the following alias information:

`fp->f1`, `fp->f2`, `fp->f3`, `bp->b1`, `bp->b2`, and `bp->b3` do not alias each other.

Consider the following example source code that demonstrates that certain levels of aliasing cannot handle interior pointers. For a definition of interior pointers see TABLE 2-4.

CODE EXAMPLE 7-3

```
struct foo {
    int f1;
    struct bar *f2;
    struct bar *f3;
    int f4;
    int f5;
    struct bar fb[10];
} *fp;

struct bar
    struct bar *b2;
    struct bar *b3;
    int b4;
} *bp;

bp=(struct bar*)&fp->f2;
```

The dereference in CODE EXAMPLE 7-3 is not supported by `weak`, `layout`, `strict`, or `std`. After the pointer assignment `bp=(struct bar*)&fp->f2`, the following pair of memory accesses touches the same memory locations:

- `fp->f2` and `bp->b2` access the same memory location
- `fp->f3` and `bp->b3` access the same memory location
- `fp->f4` and `bp->b4` access the same memory location

However, under options `weak`, `layout`, `strict`, and `std`, the compiler assumes that `fp->f2` and `bp->b2` do not alias. The compiler makes this assumption because `b2` has an offset of zero, which is different from the offset of `f2` (four bytes), and `foo` and `bar` do not have a common initial sequence. Similarly, the compiler also assumes that `bp->b3` does not alias `fp->f3`, and `bp->b4` does not alias `fp->f4`.

Thus, the pointer assignment `bp=(struct bar*)(&fp->f2)` creates a situation in which the compiler's assumptions about alias information are incorrect. This may lead to incorrect optimization.

Try compiling after you make the modifications shown in the following example.

```
struct foo {
    int f1;
    struct bar fb; /* Modified line */
#define f2 fb.b2 /* Modified line */
#define f3 fb.b3 /* Modified line */
#define f4 fb.b4 /* Modified line */
    int f5;
    struct bar fb[10];
} *fp;

struct bar
    struct bar *b2;
    struct bar *b3;
    int b4;
} *bp;

bp=(struct bar*)(&fp->f2);
```

After the pointer assignment `bp=(struct bar*)(&fp->f2)`, the following pair of memory accesses touches the same memory locations:

- `fp->f2` and `bp->b2`
- `fp->f3` and `bp->b3`
- `fp->f4` and `bp->b4`

By examining the changes shown in the preceding code example, you can see that the expression `fp->f2` is another form of the expression `fp->fb.b2`. Because `fp->fb` is of type `bar`, `fp->f2` accesses the `b2` field of `bar`. Furthermore, `bp->b2` also accesses the `b2` field of `bar`. Therefore, the compiler assumes that `fp->f2` aliases `bp->b2`. Similarly, the compiler assumes that `fp->f3` aliases `bp->b3`, and `fp->f4` aliases `bp->b4`. As a result, the aliasing assumed by the compiler matches the actual aliases caused by the pointer assignment.

Consider the following example source code.

CODE EXAMPLE 7-4

```
struct foo {
    int f1;
    int f2;
} *fp;

struct bar {
    int b1;
    int b2;
} *bp;

struct cat {
    int c1;
    struct foo cf;
    int c2;
    int c3;
} *cp;

struct dog {
    int d1;
    int d2;
    struct bar db;
    int d3;
} *dp;
```

If CODE EXAMPLE 7-4 is compiled with the `-xalias_level=weak` option, the compiler assumes the following alias information:

- `fp->f1` can alias `bp->b1`, `cp->c1`, `dp->d1`, `cp->cf.f1`, and `df->db.b1`.
 - `fp->f2` can alias `bp->b2`, `cp->cf.f1`, `dp->d2`, `cp->cf.f2`, `df->db.b2`, `cp->c2`.
 - `bp->b1` can alias `fp->f1`, `cp->c1`, `dp->d1`, `cp->cf.f1`, and `df->db.b1`.
 - `bp->b2` can alias `fp->f2`, `cp->cf.f1`, `dp->d2`, `cp->cf.f1`, and `df->db.b2`.
- `fp->f2` can alias `cp->c2` because `*dp` can alias `*cp` and `*fp` can alias `dp->db`.
- `cp->c1` can alias `fp->f1`, `bp->b1`, `dp->d1`, and `dp->db.b1`.
 - `cp->cf.f1` can alias `fp->f1`, `fp->f2`, `bp->b1`, `bp->b2`, `dp->d2`, and `dp->d1`.
- `cp->cf.f1` does not alias `dp->db.b1`.
- `cp->cf.f2` can alias `fp->f2`, `bp->b2`, `dp->db.b1`, and `dp->d2`.
 - `cp->c2` can alias `dp->db.b2`.
- `cp->c2` does not alias `dp->db.b1` and `cp->c2` does not alias `dp->d3`.

With respect to offsets, `cp->c2` can alias `db->db.b1` only if `*dp` aliases `cp->cf`. However, if `*dp` aliases `cp->cf`, then `dp->db.b1` must alias beyond the end of `foo.cf`, which is prohibited by object restrictions. Therefore, the compiler assumes that `cp->c2` cannot alias `db->db.b1`.

`cp->c3` can alias `dp->d3`.

Notice that `cp->c3` does not alias `dp->db.b2`. These memory references do not alias because the offsets of the fields of the types involved in the dereferences differ and do not overlap. Based on this, the compiler assumes they cannot alias.

- `dp->d1` can alias `fp->f1`, `bp->b1`, and `cp->c1`.
- `dp->d2` can alias `fp->f2`, `bp->b2`, and `cp->cf.f1`.
- `dp->db.b1` can alias `fp->f1`, `bp->b1`, and `cp->c1`.
- `dp->db.b2` can alias `fp->f2`, `bp->b2`, `cp->c2`, and `cp->cf.f1`.
- `dp->d3` can alias `cp->c3`.

Notice that `dp->d3` does not alias `cp->cf.f2`. These memory references do not alias because the offsets of the fields of the types involved in the dereferences differ and do not overlap. Based on this, the compiler assumes they cannot alias.

If CODE EXAMPLE 7-4 is compiled with the `-xalias_level=layout` option, the compiler assumes only the following alias information:

- `fp->f1`, `bp->b1`, `cp->c1` and `dp->d1` all can alias each other.
- `fp->f2`, `bp->b2` and `dp->d2` all can alias each other.
- `fp->f1` can alias `cp->cf.f1` and `dp->db.b1`.
- `bp->b1` can alias `cp->cf.f1` and `dp->db.b1`.
- `fp->f2` can alias `cp->cf.f2` and `dp->db.b2`.
- `bp->b2` can alias `cp->cf.f2` and `dp->db.b2`.

If CODE EXAMPLE 7-4 is compiled with the `-xalias_level=strict` option, the compiler assumes only the following alias information:

- `fp->f1` and `bp->b1` can alias each other.
- `fp->f2` and `bp->b2` can alias each other.
- `fp->f1` can alias `cp->cf.f1` and `dp->db.b1`.
- `bp->b1` can alias `cp->cf.f1` and `dp->db.b1`.
- `fp->f2` can alias `cp->cf.f2` and `dp->db.b2`.
- `bp->b2` can alias `cp->cf.f2` and `dp->db.b2`.

If CODE EXAMPLE 7-4 is compiled with the `-xalias_level=std` option, the compiler assumes only the following alias information:

- `fp->f1` can alias `cp->cf.f1`.
- `bp->b1` can alias `dp->db.b1`.
- `fp->f2` can alias `cp->cf.f2`.
- `bp->b2` can alias `dp->db.b2`.

Consider the following example source code.

CODE EXAMPLE 7-5

```
struct foo {
    short f1;
    short f2;
    int   f3;
} *fp;

struct bar {
    int b1;
    int b2;
} *bp;

union moo {
    struct foo u_f;
    struct bar u_b;
} u;
```

Here are the compiler's assumptions based on the following alias levels:

- If CODE EXAMPLE 7-5 is compiled with the `-xalias_level=weak` option, `fp->f3` and `bp->b2` can alias each other.
- If CODE EXAMPLE 7-5 is compiled with the `-xalias_level=layout` option, no fields can alias each other.
- If CODE EXAMPLE 7-5 is compiled with the `-xalias_level=strict` option, `fp->f3` and `bp->b2` can alias each other.
- If CODE EXAMPLE 7-5 is compiled with the `-xalias_level=std` option, no fields can alias each other.

Consider the following example source code.

CODE EXAMPLE 7-6

```
struct bar;

struct foo {
    struct foo *ffp;
    struct bar *fbp;
} *fp;

struct bar {
    struct bar *bbp;
    long       b2;
} *bp;
```

Here are the compiler's assumptions based on the following alias levels:

- If CODE EXAMPLE 7-6 is compiled with the `-xalias_level=weak` option, only `fp->ffp` and `bp->bbp` can alias each other.
- If CODE EXAMPLE 7-6 is compiled with the `-xalias_level=layout` option, only `fp->ffp` and `bp->bbp` can alias each other.
- If CODE EXAMPLE 7-6 is compiled with the `-xalias_level=strict` option, no fields can alias because the two struct types are still different even after their tags are removed.
- If CODE EXAMPLE 7-6 is compiled with the `-xalias_level=std` option, no fields can alias because the two types and the tags are not the same.

Consider the following example source code:

CODE EXAMPLE 7-7

```

struct foo;
struct bar;
#pragma alias (struct foo, struct bar)

struct foo {
    int f1;
    int f2;
} *fp;

struct bar {
    short b1;
    short b2;
    int b3;
} *bp;

```

The pragma in this example tells the compiler that `foo` and `bar` are allowed to alias each other. The compiler makes the following assumptions about alias information:

- `fp->f1` can alias with `bp->b1`, `bp->b2`, and `bp->b3`
- `fp->f2` can alias with `bp->b1`, `bp->b2`, and `bp->b3`

Transitioning to ANSI/ISO C

This chapter contains the following sections:

- “Basic Modes” on page 207
- “A Mixture of Old- and New-Style Functions” on page 208
- “Functions with Varying Arguments” on page 212
- “Promotions: Unsigned Versus Value Preserving” on page 215
- “Tokenization and Preprocessing” on page 219
- “`const` and `volatile`” on page 224
- “Multibyte Characters and Wide Characters” on page 228
- “Standard Headers and Reserved Names” on page 231
- “Internationalization” on page 234
- “Grouping and Evaluation in Expressions” on page 238
- “Incomplete Types” on page 241
- “Compatible and Composite Types” on page 244

Basic Modes

The ANSI/ISO C compiler allows both old-style and new-style C code. The following `-X` (note case) options provide varying degrees of compliance to the ANSI/ISO C standard. `-Xa` is the default mode.

`-Xa`

(a = ANSI) ANSI/ISO C plus K&R C compatibility extensions, with semantic changes required by ANSI/ISO C. Where K&R C and ANSI/ISO C specify different semantics for the same construct, the compiler issues warnings about the conflict and uses the ANSI/ISO C interpretation. This is the default mode.

-Xc

(c = conformance) Maximally conformant ANSI/ISO C, without K&R C compatibility extensions. The compiler issues errors and warnings for programs that use non-ANSI/ISO C constructs.

-Xs

(s = K&R C) The compiled language includes all features compatible with pre-ANSI/ISO K&R C. The compiler warns about all language constructs that have differing behavior between ANSI/ISO C and K&R C.

-Xt

(t = transition) ANSI/ISO C plus K&R C compatibility extensions, *without* semantic changes required by ANSI/ISO C. Where K&R C and ANSI/ISO C specify different semantics for the same construct, the compiler issues warnings about the conflict and uses the K&R C interpretation.

A Mixture of Old- and New-Style Functions

ANSI/ISO C's most sweeping change to the language is the function prototype borrowed from the C++ language. By specifying for each function the number and types of its parameters, not only does every regular compile get the benefits of argument and parameter checks (similar to those of `lint`) for each function call, but arguments are automatically converted (just as with an assignment) to the type expected by the function. ANSI/ISO C includes rules that govern the mixing of old- and new-style function declarations since there are many, many lines of existing C code that could and should be converted to use prototypes.

Writing New Code

When you write an entirely new program, use new-style function declarations (function prototypes) in headers and new-style function declarations and definitions in other C source files. However, if there is a possibility that someone will port the

code to a machine with a pre-ANSI/ISO C compiler, we suggest you use the macro `__STDC__` (which is defined only for ANSI/ISO C compilation systems) in both header and source files. Refer to “Mixing Considerations” on page 209 for an example.

An ANSI/ISO C-conforming compiler must issue a diagnostic whenever two incompatible declarations for the same object or function are in the same scope. If all functions are declared and defined with prototypes, and the appropriate headers are included by the correct source files, all calls should agree with the definition of the functions. This protocol eliminates one of the most common C programming mistakes.

Updating Existing Code

If you have an existing application and want the benefits of function prototypes, there are a number of possibilities for updating, depending on how much of the code you would like to change:

1. Recompile without making any changes.

Even with no coding changes, the compiler warns you about mismatches in parameter type and number when invoked with the `-v` option.

2. Add function prototypes just to the headers.

All calls to global functions are covered.

3. Add function prototypes to the headers and start each source file with function prototypes for its local (static) functions.

All calls to functions are covered, but doing this requires typing the interface for each local function twice in the source file.

4. Change all function declarations and definitions to use function prototypes.

For most programmers, choices 2 and 3 are probably the best cost/benefit compromise. Unfortunately, these options are precisely the ones that require detailed knowledge of the rules for mixing old and new styles.

Mixing Considerations

For function prototype declarations to work with old-style function definitions, both must specify functionally identical interfaces or have *compatible types* using ANSI/ISO C’s terminology.

For functions with varying arguments, there can be no mixing of ANSI/ISO C's ellipsis notation and the old-style `varargs()` function definition. For functions with a fixed number of parameters, the situation is fairly straightforward: just specify the types of the parameters as they were passed in previous implementations.

In K&R C, each argument was converted just before it was passed to the called function according to the default argument promotions. These promotions specified that all integral types narrower than `int` were promoted to `int` size, and any `float` argument was promoted to `double`, hence simplifying both the compiler and libraries. Function prototypes are more expressive—the specified parameter type is what is passed to the function.

Thus, if a function prototype is written for an existing (old-style) function definition, there should be no parameters in the function prototype with any of the following types:

<code>char</code>	<code>signed char</code>	<code>unsigned char</code>	<code>float</code>
<code>short</code>	<code>signed short</code>	<code>unsigned short</code>	

There still remain two complications with writing prototypes: `typedef` names and the promotion rules for narrow unsigned types.

If parameters in old-style functions were declared using `typedef` names, such as `off_t` and `ino_t`, it is important to know whether or not the `typedef` name designates a type that is affected by the default argument promotions. For these two, `off_t` is a `long`, so it is appropriate to use in a function prototype; `ino_t` used to be an `unsigned short`, so if it were used in a prototype, the compiler issues a diagnostic because the old-style definition and the prototype specify different and incompatible interfaces.

Just what should be used instead of an `unsigned short` leads us into the final complication. The one biggest incompatibility between K&R C and the ANSI/ISO C compiler is the promotion rule for the widening of `unsigned char` and `unsigned short` to an `int` value. (See “Promotions: Unsigned Versus Value Preserving” on page 215.) The parameter type that matches such an old-style parameter depends on the compilation mode used when you compile:

- `-Xs` and `-Xt` should use `unsigned int`
- `-Xa` and `-Xc` should use `int`

The best approach is to change the old-style definition to specify either `int` or `unsigned int` and use the matching type in the function prototype. You can always assign its value to a local variable with the narrower type, if necessary, after you enter the function.

Watch out for the use of id's in prototypes that may be affected by preprocessing. Consider the following example:

```
#define status 23
void my_exit(int status); /* Normally, scope begins */
                          /* and ends with prototype */
```

Do not mix function prototypes with old-style function declarations that contain narrow types.

```
void foo(unsigned char, unsigned short);
void foo(i, j) unsigned char i; unsigned short j; {...}
```

Appropriate use of `__STDC__` produces a header file that can be used for both the old and new compilers:

```
header.h:
struct s { /* . . . */ };
#ifdef __STDC__
    void errmsg(int, ...);
    struct s *f(const char *);
    int g(void);
#else
    void errmsg();
    struct s *f();
    int g();
#endif
```

The following function uses prototypes and can still be compiled on an older system:

```
struct s *
#ifdef __STDC__
    f(const char *p)
#else
    f(p) char *p;
#endif
{
    /* . . . */
}
```

Here is an updated source file (as with choice 3 above). The local function still uses an old-style definition, but a prototype is included for newer compilers:

```
source.c:
#include "header.h"
    typedef /* . . . */ MyType;
#ifdef __STDC__
    static void del(MyType *);
    /* . . . */
    static void
    del(p)
    MyType *p;
    {
    /* . . . */
    }
    /* . . . */
```

Functions With Varying Arguments

In previous implementations, you could not specify the parameter types that a function expected, but ANSI/ISO C encourages you to use prototypes to do just that. To support functions such as `printf()`, the syntax for prototypes includes a special ellipsis (`...`) terminator. Because an implementation might need to do unusual things to handle a varying number of arguments, ANSI/ISO C requires that all declarations and the definition of such a function include the ellipsis terminator.

Since there are no names for the “...” part of the parameters, a special set of macros contained in `stdarg.h` gives the function access to these arguments. Earlier versions of such functions had to use similar macros contained in `varargs.h`.

Let us assume that the function we wish to write is an error handler called `errmsg()` that returns `void`, and whose only fixed parameter is an `int` that specifies details about the error message. This parameter can be followed by a file name, a line number, or both, and these are followed by format and arguments, similar to those of `printf()`, that specify the text of the error message.

To allow our example to compile with earlier compilers, we make extensive use of the macro `__STDC__` which is defined only for ANSI/ISO C compilation systems. Thus, the function's declaration in the appropriate header file is:

```
#ifdef __STDC__
    void errmsg(int code, ...);
#else
    void errmsg();
#endif
```

The file that contains the definition of `errmsg()` is where the old and new styles can get complex. First, the header to include depends on the compilation system:

```
#ifdef __STDC__
#include <stdarg.h>
#else
#include <varargs.h>
#endif
#include <stdio.h>
```

`stdio.h` is included because we call `fprintf()` and `vfprintf()` later.

Next comes the definition for the function. The identifiers `va_alist` and `va_dcl` are part of the old-style `varargs.h` interface.

```
void
#ifdef __STDC__
errmsg(int code, ...)
#else
errmsg(va_alist) va_dcl /* Note: no semicolon! */
#endif
{
    /* more detail below */
}
```

Since the old-style variable argument mechanism did not allow us to specify any fixed parameters, we must arrange for them to be accessed before the varying portion. Also, due to the lack of a name for the “...” part of the parameters, the new `va_start()` macro has a second argument—the name of the parameter that comes just before the “...” terminator.

As an extension, Sun ANSI/ISO C allows functions to be declared and defined with no fixed parameters, as in:

```
int f(...);
```

For such functions, `va_start()` should be invoked with an empty second argument, as in:

```
va_start(ap,)
```

The following is the body of the function:

```
{
    va_list ap;
    char *fmt;
#ifdef __STDC__
    va_start(ap, code);
#else
    int code;
    va_start(ap);
    /* extract the fixed argument */
    code = va_arg(ap, int);
#endif
    if (code & FILENAME)
        (void)fprintf(stderr, "\\\"%s\\\": ", va_arg(ap, char *));
    if (code & LINENUMBER)
        (void)fprintf(stderr, "%d: ", va_arg(ap, int));
    if (code & WARNING)
        (void)fputs("warning: ", stderr);
    fmt = va_arg(ap, char *);
    (void)vfprintf(stderr, fmt, ap);
    va_end(ap);
}
```

Both the `va_arg()` and `va_end()` macros work the same for the old-style and ANSI/ISO C versions. Because `va_arg()` changes the value of `ap`, the call to `vfprintf()` cannot be:

```
(void)vfprintf(stderr, va_arg(ap, char *), ap);
```

The definitions for the macros `FILENAME`, `LINENUMBER`, and `WARNING` are presumably contained in the same header as the declaration of `errmsg()`.

A sample call to `errmsg()` could be:

```
errmsg(FILENAME, "<command line>", "cannot open: %s\n",
argv[optind]);
```

Promotions: Unsigned Versus Value Preserving

The following information appears in the Rationale section that accompanies the draft C Standard: “QUIET CHANGE”. A program that depends on unsigned preserving arithmetic conversions will behave differently, probably without complaint. This is considered to be the most serious change made by the Committee to a widespread current practice.

This section explores how this change affects our code.

Background

According to K&R, *The C Programming Language* (First Edition), unsigned specified exactly one type; there were no unsigned chars, unsigned shorts, or unsigned longs, but most C compilers added these very soon thereafter. Some compilers did not implement unsigned long but included the other two. Naturally, implementations chose different rules for type promotions when these new types mixed with others in expressions.

In most C compilers, the simpler rule, “unsigned preserving,” is used: when an unsigned type needs to be widened, it is widened to an unsigned type; when an unsigned type mixes with a signed type, the result is an unsigned type.

The other rule, specified by ANSI/ISO C, is known as “value preserving,” in which the result type depends on the relative sizes of the operand types. When an unsigned char or unsigned short is widened, the result type is int if an int is large enough to represent all the values of the smaller type. Otherwise, the result type is unsigned int. The value preserving rule produces the least surprise arithmetic result for most expressions.

Compilation Behavior

Only in the transition or pre-ANSI/ISO modes (`-xt` or `-xs`) does the ANSI/ISO C compiler use the unsigned preserving promotions; in the other two modes, conforming (`-xc`) and ANSI/ISO (`-xa`), the value preserving promotion rules are used.

First Example: The Use of a Cast

In the following code, assume that an unsigned char is smaller than an int.

```
int f(void)
{
    int i = -2;
    unsigned char uc = 1;

    return (i + uc) < 17;
}
```

The code above causes the compiler to issue the following warning when you use the `-xtransition` option:

```
line 6: warning: semantics of "<" change in ANSI/ISO C; use
explicit cast
```

The result of the addition has type `int` (value preserving) or `unsigned int` (unsigned preserving), but the bit pattern does not change between these two. On a two's-complement machine, we have:

```
    i:   111...110 (-2)
+   uc:  000...001 ( 1)
=====
        111...111 (-1 or UINT_MAX)
```

This bit representation corresponds to `-1` for `int` and `UINT_MAX` for `unsigned int`. Thus, if the result has type `int`, a signed comparison is used and the less-than test is true; if the result has type `unsigned int`, an unsigned comparison is used and the less-than test is false.

The addition of a cast serves to specify which of the two behaviors is desired:

```
value preserving:
    (i + (int)uc) < 17
unsigned preserving:
    (i + (unsigned int)uc) < 17
```

Since differing compilers chose different meanings for the same code, this expression can be ambiguous. The addition of a cast is as much to help the reader as it is to eliminate the warning message.

Bit-fields

The same situation applies to the promotion of bit-field values. In ANSI/ISO C, if the number of bits in an `int` or unsigned `int` bit-field is less than the number of bits in an `int`, the promoted type is `int`; otherwise, the promoted type is unsigned `int`. In most older C compilers, the promoted type is unsigned `int` for explicitly unsigned bit-fields, and `int` otherwise.

Similar use of casts can eliminate situations that are ambiguous.

Second Example: Same Result

In the following code, assume that both unsigned `short` and unsigned `char` are narrower than `int`.

```
int f(void)
{
    unsigned short us;
    unsigned char uc;
    return uc < us;
}
```

In this example, both automatics are either promoted to `int` or to unsigned `int`, so the comparison is sometimes unsigned and sometimes signed. However, the C compiler does not warn you because the result is the same for the two choices.

Integral Constants

As with expressions, the rules for the types of certain integral constants have changed. In K&R C, an unsuffixed decimal constant had type `int` only if its value fit in an `int`; an unsuffixed octal or hexadecimal constant had type `int` only if its value fit in an unsigned `int`. Otherwise, an integral constant had type `long`. At times, the value did not fit in the resulting type. In ANSI/ISO C, the constant type is the first type encountered in the following list that corresponds to the value:

unsuffixed decimal:

`int, long, unsigned long`

unsuffixed octal or hexadecimal:

`int, unsigned int, long, unsigned long`

U suffixed:

`unsigned int, unsigned long`

L suffixed:

`long, unsigned long`

UL suffixed:

`unsigned long`

The ANSI/ISO C compiler warns you, when you use the `-xtransition` option, about any expression whose behavior might change according to the typing rules of the constants involved. The old integral constant typing rules are used only in the transition mode; the ANSI/ISO and conforming modes use the new rules.

Third Example: Integral Constants

In the following code, assume `ints` are 16 bits.

```
int f(void)
{
    int i = 0;

    return i > 0xffff;
}
```

Because the hexadecimal constant's type is either `int` (with a value of `-1` on a two's-complement machine) or an unsigned `int` (with a value of `65535`), the comparison is true in `-xs` and `-xt` modes, and false in `-xa` and `-xc` modes.

Again, an appropriate cast clarifies the code and suppresses a warning:

```
-xt, -xs modes:
    i > (int)0xffff

-Xa, -Xc modes:
    i > (unsigned int)0xffff
    or
    i > 0xffffU
```

The `U` suffix character is a new feature of ANSI/ISO C and probably produces an error message with older compilers.

Tokenization and Preprocessing

Probably the least specified part of previous versions of C concerned the operations that transformed each source file from a bunch of characters into a sequence of tokens, ready to parse. These operations included recognition of white space (including comments), bundling consecutive characters into tokens, handling preprocessing directive lines, and macro replacement. However, their respective ordering was never guaranteed.

ANSI/ISO C Translation Phases

The order of these translation phases is specified by ANSI/ISO C:

1. Every trigraph sequence in the source file is replaced. ANSI/ISO C has exactly nine trigraph sequences that were invented solely as a concession to deficient character sets, and are three-character sequences that name a character not in the ISO 646-1983 character set:

TABLE 8-1 Trigraph Sequences

Trigraph Sequence	Converts to	Trigraph Sequence	Converts to
??=	#	??<	{
??-	~	??>	}
??([??/	\
??)]	??'	^
??!			

These sequences must be understood by ANSI/ISO C compilers, but we do not recommend their use. The ANSI/ISO C compiler warns you, when you use the `-xtransition` option, whenever it replaces a trigraph while in transition (`-xt`) mode, even in comments. For example, consider the following:

```
/* comment *??/  
/* still comment? */
```

The `??/` becomes a backslash. This character and the following newline are removed. The resulting characters are:

```
/* comment */* still comment? */
```

The first `/` from the second line is the end of the comment. The next token is the `*`.

1. Every backslash/new-line character pair is deleted.
2. The source file is converted into preprocessing tokens and sequences of white space. Each comment is effectively replaced by a space character.
3. Every preprocessing directive is handled and all macro invocations are replaced. Each `#included` source file is run through the earlier phases before its contents replace the directive line.
4. Every escape sequence (in character constants and string literals) is interpreted.
5. Adjacent string literals are concatenated.

6. Every preprocessing token is converted into a regular token; the compiler properly parses these and generates code.
7. All external object and function references are resolved, resulting in the final program.

Old C Translation Phases

Previous C compilers did not follow such a simple sequence of phases, nor were there any guarantees for when these steps were applied. A separate preprocessor recognized tokens and white space at essentially the same time as it replaced macros and handled directive lines. The output was then completely retokenized by the compiler proper, which then parsed the language and generated code.

Because the tokenization process within the preprocessor was a moment-by-moment operation and macro replacement was done as a character-based, not token-based, operation, the tokens and white space could have a great deal of variation during preprocessing.

There are a number of differences that arise from these two approaches. The rest of this section discusses how code behavior may change due to line splicing, macro replacement, stringizing, and token pasting, which occur during macro replacement.

Logical Source Lines

In K&R C, backslash/new-line pairs were allowed only as a means to continue a directive, a string literal, or a character constant to the next line. ANSI/ISO C extended the notion so that a backslash/new-line pair can continue anything to the next line. The result is a logical source line. Therefore, any code that relied on the separate recognition of tokens on either side of a backslash/new-line pair does not behave as expected.

Macro Replacement

The macro replacement process has never been described in detail prior to ANSI/ISO C. This vagueness spawned a great many divergent implementations. Any code that relied on anything fancier than manifest constant replacement and simple function-like macros was probably not truly portable. This manual cannot uncover all the differences between the old C macro replacement implementation and the ANSI/ISO C version. Nearly all uses of macro replacement with the exception of

token pasting and stringizing produce exactly the same series of tokens as before. Furthermore, the ANSI/ISO C macro replacement algorithm can do things not possible in the old C version. For example,

```
#define name (*name)
```

causes any use of name to be replaced with an indirect reference through name. The old C preprocessor would produce a huge number of parentheses and stars and eventually produce an error about macro recursion.

The major change in the macro replacement approach taken by ANSI/ISO C is to require macro arguments, other than those that are operands of the macro substitution operators # and ##, to be expanded recursively prior to their substitution in the replacement token list. However, this change seldom produces an actual difference in the resulting tokens.

Using Strings

Note – In ANSI/ISO C, the examples below marked with a † produce a warning about use of old features, when you use the `-xtransition` option. Only in the transition mode (`-Xt` and `-Xs`) is the result the same as in previous versions of C.

In K&R C, the following code produced the string literal "x y!":

```
#define str(a) "a!" †
str(x y)
```

Thus, the preprocessor searched inside string literals and character constants for characters that looked like macro parameters. ANSI/ISO C recognized the importance of this feature, but could not condone operations on parts of tokens. In ANSI/ISO C, all invocations of the above macro produce the string literal "a!". To achieve the old effect in ANSI/ISO C, we make use of the # macro substitution operator and the concatenation of string literals.

```
#define str(a) #a "!"
str(x y)
```

The above code produces the two string literals "x y" and "!" which, after concatenation, produces the identical "x y!".

There is no direct replacement for the analogous operation for character constants. The major use of this feature was similar to the following:

```
#define CNTL(ch) (037 & 'ch') †  
CNTL(L)
```

which produced

```
(037 & 'L')
```

which evaluates to the ASCII control-L character. The best solution we know of is to change all uses of this macro to:

```
#define CNTL(ch) (037 & (ch))  
CNTL('L')
```

This code is more readable and more useful, as it can also be applied to expressions.

Token Pasting

In K&R C, there were at least two ways to combine two tokens. Both invocations in the following produced a single identifier `x1` out of the two tokens `x` and `1`.

```
#define self(a) a  
#define glue(a,b) a/**/b †  
self(x)1  
glue(x,1)
```

Again, ANSI/ISO C could not sanction either approach. In ANSI/ISO C, both the above invocations would produce the two separate tokens `x` and `1`. The second of the above two methods can be rewritten for ANSI/ISO C by using the `##` macro substitution operator:

```
#define glue(a,b) a ## b  
glue(x, 1)
```

`#` and `##` should be used as macro substitution operators only when `__STDC__` is defined. Since `##` is an actual operator, the invocation can be much freer with respect to white space in both the definition and invocation.

There is no direct approach to effect the first of the two old-style pasting schemes, but since it put the burden of the pasting at the invocation, it was used less frequently than the other form.

const and volatile

The keyword `const` was one of the C++ features that found its way into ANSI/ISO C. When an analogous keyword, `volatile`, was invented by the ANSI/ISO C Committee, the “type qualifier” category was created. This category still remains one of the more nebulous parts of ANSI/ISO C.

Types, Only for lvalue

`const` and `volatile` are part of an identifier’s type, not its storage class. However, they are often removed from the topmost part of the type when an object’s value is fetched in the evaluation of an expression—exactly at the point when an lvalue becomes an rvalue. These terms arise from the prototypical assignment “L=R”; in which the left side must still refer directly to an object (an lvalue) and the right side need only be a value (an rvalue). Thus, only expressions that are lvalues can be qualified by `const` or `volatile` or both.

Type Qualifiers in Derived Types

The type qualifiers may modify type names and derived types. Derived types are those parts of C’s declarations that can be applied over and over to build more and more complex types: pointers, arrays, functions, structures, and unions. Except for functions, one or both type qualifiers can be used to change the behavior of a derived type.

For example,

```
const int five = 5;
```

declares and initializes an object with type `const int` whose value is not changed by a correct program. The order of the keywords is not significant to C. For example, the declarations:

```
int const five = 5;
```

and

```
const five = 5;
```

are identical to the above declaration in its effect.

The declaration

```
const int *pci = &five;
```

declares an object with type pointer to `const int`, which initially points to the previously declared object. The pointer itself does not have a qualified type—it points to a qualified type, and can be changed to point to essentially any `int` during program execution. `pci` cannot be used to modify the object to which it points unless a cast is used, as in the following:

```
*(int *)pci = 17;
```

If `pci` actually points to a `const` object, the behavior of this code is undefined.

The declaration

```
extern int *const cpi;
```

says that somewhere in the program there exists a definition of a global object with type `const` pointer to `int`. In this case, `cpi`'s value will not be changed by a correct program, but it can be used to modify the object to which it points. Notice that `const` comes after the `*` in the above declaration. The following pair of declarations produces the same effect:

```
typedef int *INT_PTR;  
extern const INT_PTR cpi;
```

These declarations can be combined as in the following declaration in which an object is declared to have type `const` pointer to `const int`:

```
const int *const cpci;
```

`const` Means `readonly`

In hindsight, `readonly` would have been a better choice for a keyword than `const`. If one reads `const` in this manner, declarations such as:

```
char *strcpy(char *, const char *);
```

are easily understood to mean that the second parameter is only used to read character values, while the first parameter overwrites the characters to which it points. Furthermore, despite the fact that in the above example, the type of `cpci` is a pointer to a `const int`, you can still change the value of the object to which it points through some other means, unless it actually points to an object declared with `const int` type.

Examples of `const` Usage

The two main uses for `const` are to declare large compile-time initialized tables of information as unchanging, and to specify that pointer parameters do not modify the objects to which they point.

The first use potentially allows portions of the data for a program to be shared by other concurrent invocations of the same program. It may cause attempts to modify this invariant data to be detected immediately by means of some sort of memory protection fault, since the data resides in a read-only portion of memory.

The second use helps locate potential errors before generating a memory fault during that demo. For example, functions that temporarily place a null character into the middle of a string are detected at compile time, if passed a pointer to a string that cannot be so modified.

volatile Means Exact Semantics

So far, the examples have all used `const` because it's conceptually simpler. But what does `volatile` really mean? To a compiler writer, it has one meaning: take no code generation shortcuts when accessing such an object. In ANSI/ISO C, it is a programmer's responsibility to declare every object that has the appropriate special properties with a `volatile` qualified type.

Examples of volatile Usage

The usual four examples of `volatile` objects are:

- An object that is a memory-mapped I/O port
- An object that is shared between multiple concurrent processes
- An object that is modified by an asynchronous signal handler
- An automatic storage duration object declared in a function that calls `setjmp`, and whose value is changed between the call to `setjmp` and a corresponding call to `longjmp`

The first three examples are all instances of an object with a particular behavior: its value can be modified at any point during the execution of the program. Thus, the seemingly infinite loop:

```
flag = 1;
while (flag);
```

is valid as long as `flag` has a `volatile` qualified type. Presumably, some asynchronous event sets `flag` to zero in the future. Otherwise, because the value of `flag` is unchanged within the body of the loop, the compilation system is free to change the above loop into a truly infinite loop that completely ignores the value of `flag`.

The fourth example, involving variables local to functions that call `setjmp`, is more involved. The fine print about the behavior of `setjmp` and `longjmp` notes that there are no guarantees about the values for objects matching the fourth case. For the most desirable behavior, it is necessary for `longjmp` to examine every stack frame between the function calling `setjmp` and the function calling `longjmp` for saved register values. The possibility of asynchronously created stack frames makes this job even harder.

When an automatic object is declared with a `volatile` qualified type, the compilation system knows that it has to produce code that exactly matches what the programmer wrote. Therefore, the most recent value for such an automatic object is always in memory and not just in a register, and is guaranteed to be up-to-date when `longjmp` is called.

Multibyte Characters and Wide Characters

At first, the internationalization of ANSI/ISO C affected only library functions. However, the final stage of internationalization—multibyte characters and wide characters—also affected the language proper.

Asian Languages Require Multibyte Characters

The basic difficulty in an Asian-language computer environment is the huge number of ideograms needed for I/O. To work within the constraints of usual computer architectures, these ideograms are encoded as sequences of bytes. The associated operating systems, application programs, and terminals understand these byte sequences as individual ideograms. Moreover, all of these encodings allow intermixing of regular single-byte characters with the ideogram byte sequences. Just how difficult it is to recognize distinct ideograms depends on the encoding scheme used.

The term “multibyte character” is defined by ANSI/ISO C to denote a byte sequence that encodes an ideogram, no matter what encoding scheme is employed. All multibyte characters are members of the “extended character set.” A regular single-byte character is just a special case of a multibyte character. The only requirement placed on the encoding is that no multibyte character can use a null character as part of its encoding.

ANSI/ISO C specifies that program comments, string literals, character constants, and header names are all sequences of multibyte characters.

Encoding Variations

The encoding schemes fall into two camps. The first is one in which each multibyte character is self-identifying, that is, any multibyte character can simply be inserted between any pair of multibyte characters.

The second scheme is one in which the presence of special shift bytes changes the interpretation of subsequent bytes. An example is the method used by some character terminals to get in and out of line-drawing mode. For programs written in multibyte characters with a shift-state-dependent encoding, ANSI/ISO C requires that each comment, string literal, character constant, and header name must both begin and end in the unshifted state.

Wide Characters

Some of the inconvenience of handling multibyte characters would be eliminated if all characters were of a uniform number of bytes or bits. Since there can be thousands or tens of thousands of ideograms in such a character set, a 16-bit or 32-bit sized integral value should be used to hold all members. (The full Chinese alphabet includes more than 65,000 ideograms!) ANSI/ISO C includes the typedef name `wchar_t` as the implementation-defined integral type large enough to hold all members of the extended character set.

For each wide character, there is a corresponding multibyte character, and vice versa; the wide character that corresponds to a regular single-byte character is required to have the same value as its single-byte value, including the null character. However, there is no guarantee that the value of the macro `EOF` can be stored in a `wchar_t`, just as `EOF` might not be representable as a `char`.

Conversion Functions

ANSI/ISO C provides five library functions that manage multibyte characters and wide characters:

TABLE 8-2 Multibyte Character Conversion Functions

<code>mblen()</code>	length of next multibyte character
<code>mbtowc()</code>	convert multibyte character to wide character
<code>wctomb()</code>	convert wide character to multibyte character
<code>mbstowcs()</code>	convert multibyte character string to wide character string
<code>wcstombs()</code>	convert wide character string to multibyte character string

The behavior of all of these functions depends on the current locale. (See “The `setlocale()` Function” on page 235.)

It is expected that vendors providing compilation systems targeted to this market supply many more string-like functions to simplify the handling of wide character strings. However, for most application programs, there is no need to convert any multibyte characters to or from wide characters. Programs such as `diff`, for example, read in and write out multibyte characters, needing only to check for an exact byte-for-byte match. More complicated programs, such as `grep`, that use regular expression pattern matching, may need to understand multibyte characters, but only the common set of functions that manages the regular expression needs this knowledge. The program `grep` itself requires no other special multibyte character handling.

C Language Features

To give even more flexibility to the programmer in an Asian-language environment, ANSI/ISO C provides wide character constants and wide string literals. These have the same form as their non-wide versions, except that they are immediately prefixed by the letter `L`:

`'x'` regular character constant

`'¥'` regular character constant

`L'x'` wide character constant

`L'¥'` wide character constant

`"abc¥xyz"` regular string literal

`L"abcxyz"` wide string literal

Multibyte characters are valid in both the regular and wide versions. The sequence of bytes necessary to produce the ideogram `¥` is encoding-specific, but if it consists of more than one byte, the value of the character constant `'¥'` is implementation-defined, just as the value of `'ab'` is implementation-defined. Except for escape sequences, a regular string literal contains exactly the bytes specified between the quotes, including the bytes of each specified multibyte character.

When the compilation system encounters a wide character constant or wide string literal, each multibyte character is converted into a wide character, as if by calling the `mbtowc()` function. Thus, the type of `L'¥'` is `wchar_t`; the type of `abc¥xyz` is array of `wchar_t` with length eight. Just as with regular string literals, each wide string literal has an extra zero-valued element appended, but in these cases, it is a `wchar_t` with value zero.

Just as regular string literals can be used as a shorthand method for character array initialization, wide string literals can be used to initialize `wchar_t` arrays:

```
wchar_t *wp = L"a¥z";
wchar_t x[] = L"a¥z";
wchar_t y[] = {L'a', L'¥', L'z', 0};
wchar_t z[] = {'a', L'¥', 'z', '\0'};
```

In the above example, the three arrays `x`, `y`, and `z`, and the array pointed to by `wp`, have the same length. All are initialized with identical values.

Finally, adjacent wide string literals are concatenated, just as with regular string literals. However, adjacent regular and wide string literals produce undefined behavior. A compiler is not required to produce an error if it does not accept such concatenations.

Standard Headers and Reserved Names

Early in the standardization process, the ANSI/ISO Standards Committee chose to include library functions, macros, and header files as part of ANSI/ISO C. While this decision was necessary for the writing of truly portable C programs, a side effect is the basis of some of the most negative comments about ANSI/ISO C from the public—a large set of reserved names.

This section presents the various categories of reserved names and some rationale for their reservations. At the end is a set of rules to follow that can steer your programs clear of any reserved names.

Balancing Process

To match existing implementations, the ANSI/ISO C committee chose names like `printf` and `NULL`. However, each such name reduced the set of names available for free use in C programs.

On the other hand, before standardization, implementors felt free to add both new keywords to their compilers and names to headers. No program could be guaranteed to compile from one release to another, let alone port from one vendor's implementation to another.

As a result, the Committee made a hard decision: to restrict all conforming implementations from including any extra names, except those with certain forms. It is this decision that causes most C compilation systems to be almost conforming. Nevertheless, the Standard contains 32 keywords and almost 250 names in its headers, none of which necessarily follow any particular naming pattern.

Standard Headers

The standard headers are:

TABLE 8-3 Standard Headers

<code>assert.h</code>	<code>locale.h</code>	<code>stddef.h</code>
<code>ctype.h</code>	<code>math.h</code>	<code>stdio.h</code>
<code>errno.h</code>	<code>setjmp.h</code>	<code>stdlib.h</code>
<code>float.h</code>	<code>signal.h</code>	<code>string.h</code>
<code>limits.h</code>	<code>stdarg.h</code>	<code>time.h</code>

Most implementations provide more headers, but a strictly conforming ANSI/ISO C program can only use these.

Other standards disagree slightly regarding the contents of some of these headers. For example, POSIX (IEEE 1003.1) specifies that `fdopen` is declared in `stdio.h`. To allow these two standards to coexist, POSIX requires the macro `_POSIX_SOURCE` to be `#defined` prior to the inclusion of any header to guarantee that these additional names exist. In its *Portability Guide*, X/Open has also used this macro scheme for its extensions. X/Open's macro is `_XOPEN_SOURCE`.

ANSI/ISO C requires the standard headers to be both self-sufficient and idempotent. No standard header needs any other header to be `#included` before or after it, and each standard header can be `#included` more than once without causing problems. The Standard also requires that its headers be `#included` only in safe contexts, so that the names used in the headers are guaranteed to remain unchanged.

Names Reserved for Implementation Use

The Standard places further restrictions on implementations regarding their libraries. In the past, most programmers learned not to use names like `read` and `write` for their own functions on UNIX Systems. ANSI/ISO C requires that only names reserved by the Standard be introduced by references within the implementation.

Thus, the Standard reserves a subset of all possible names for implementations to use. This class of names consists of identifiers that begin with an underscore and continue with either another underscore or a capital letter. The class of names contains all names matching the following regular expression:

```
_[_A-Z][0-9_a-zA-Z]*
```

Strictly speaking, if your program uses such an identifier, its behavior is undefined. Thus, programs using `_POSIX_SOURCE` (or `_XOPEN_SOURCE`) have undefined behavior.

However, undefined behavior comes in different degrees. If, in a POSIX-conforming implementation you use `_POSIX_SOURCE`, you know that your program's undefined behavior consists of certain additional names in certain headers, and your program still conforms to an accepted standard. This deliberate loophole in the ANSI/ISO C standard allows implementations to conform to seemingly incompatible specifications. On the other hand, an implementation that does not conform to the POSIX standard is free to behave in any manner when encountering a name such as `_POSIX_SOURCE`.

The Standard also reserves all other names that begin with an underscore for use in header files as regular file scope identifiers and as tags for structures and unions, but not in local scopes. The common practice of having functions named `_filbuf` and `_doprnt` to implement hidden parts of the library is allowed.

Names Reserved for Expansion

In addition to all the names explicitly reserved, the Standard also reserves (for implementations and future standards) names matching certain patterns:

TABLE 8-4 Names Reserved for Expansion

File	Reserved Name Pattern
<code>errno.h</code>	<code>E[0-9A-Z].*</code>
<code>ctype.h</code>	<code>(to is)[a-z].*</code>
<code>locale.h</code>	<code>LC_[A-Z].*</code>
<code>math.h</code>	<i>current function names</i> [<code>fl</code>]
<code>signal.h</code>	<code>(SIG SIG_[A-Z]).*</code>
<code>stdlib.h</code>	<code>str[a-z].*</code>
<code>string.h</code>	<code>(str mem wcs)[a-z].*</code>

In the above lists, names that begin with a capital letter are macros and are reserved only when the associated header is included. The rest of the names designate functions and cannot be used to name any global objects or functions.

Names Safe to Use

There are four simple rules you can follow to keep from colliding with any ANSI/ISO C reserved names:

- #include all system headers at the top of your source files (except possibly after a #define of `_POSIX_SOURCE` or `_XOPEN_SOURCE`, or both).
- Do not define or declare any names that begin with an underscore.
- Use an underscore or a capital letter somewhere within the first few characters of all file scope tags and regular names. Beware of the `va_` prefix found in `stdarg.h` or `varargs.h`.
- Use a digit or a non-capital letter somewhere within the first few characters of all macro names. Almost all names beginning with an `E` are reserved if `errno.h` is #included.

These rules are just a general guideline to follow, as most implementations will continue to add names to the standard headers by default.

Internationalization

The section “Multibyte Characters and Wide Characters” on page 228 introduced the internationalization of the standard libraries. This section discusses the affected library functions and gives some hints on how programs should be written to take advantage of these features.

Locales

At any time, a C program has a current locale—a collection of information that describes the conventions appropriate to some nationality, culture, and language. Locales have names that are strings. The only two standardized locale names are “C” and “”. Each program begins in the “C” locale, which causes all library functions to behave just like they have historically. The “” locale is the implementation’s best guess at the correct set of conventions appropriate to the program’s invocation. “C” and “” can cause identical behavior. Other locales may be provided by implementations.

For the purposes of practicality and expediency, locales are partitioned into a set of categories. A program can change the complete locale, or just one or more categories. Generally, each category affects a set of functions disjoint from the functions affected by other categories, so temporarily changing one category for a little while can make sense.

The `setlocale()` Function

The `setlocale()` function is the interface to the program's locale. In general, any program that uses the invocation country's conventions should place a call such as:

```
#include <locale.h>
/*...*/
setlocale(LC_ALL, "");
```

early in the program's execution path. This call causes the program's current locale to change to the appropriate local version, since `LC_ALL` is the macro that specifies the entire locale instead of one category. The following are the standard categories:

<code>LC_COLLATE</code>	sorting information
<code>LC_CTYPE</code>	character classification information
<code>LC_MONETARY</code>	currency printing information
<code>LC_NUMERIC</code>	numeric printing information
<code>LC_TIME</code>	date and time printing information

Any of these macros can be passed as the first argument to `setlocale()` to specify that category.

The `setlocale()` function returns the name of the current locale for a given category (or `LC_ALL`) and serves in an inquiry-only capacity when its second argument is a null pointer. Thus, code similar to the following can be used to change the locale or a portion thereof for a limited duration:

```
#include <locale.h>
/*...*/
char *oloc;
/*...*/
oloc = setlocale(LC_category, NULL);
if (setlocale(LC_category, "new") != 0)
{
    /* use temporarily changed locale */
    (void)setlocale(LC_category, oloc);
}
```

Most programs do not need this capability.

Changed Functions

Wherever possible and appropriate, existing library functions were extended to include locale-dependent behavior. These functions came in two groups:

- Those declared by the `ctype.h` header (character classification and conversion), and
- Those that convert to and from printable and internal forms of numeric values, such as `printf()` and `strtod()`.

All `ctype.h` predicate functions, except `isdigit()` and `isxdigit()`, can return nonzero (true) for additional characters when the `LC_CTYPE` category of the current locale is other than "C". In a Spanish locale, `isalpha('ñ')` should be true. Similarly, the character conversion functions, `tolower()` and `toupper()`, should appropriately handle any extra alphabetic characters identified by the `isalpha()` function. The `ctype.h` functions are almost always macros that are implemented using table lookups indexed by the character argument. Their behavior is changed by resetting the table(s) to the new locale's values, and therefore there is no performance impact.

Those functions that write or interpret printable floating values can change to use a decimal-point character other than period (.) when the `LC_NUMERIC` category of the current locale is other than "C". There is no provision for converting any numeric values to printable form with thousands separator-type characters. When converting from a printable form to an internal form, implementations are allowed to accept such additional forms, again in other than the "C" locale. Those functions that make use of the decimal-point character are the `printf()` and `scanf()` families,

`atof()`, and `strtod()`. Those functions that are allowed implementation-defined extensions are `atof()`, `atoi()`, `atol()`, `strtod()`, `strtol()`, `strtoul()`, and the `scanf()` family.

New Functions

Certain locale-dependent capabilities were added as new standard functions. Besides `setlocale()`, which allows control over the locale itself, the Standard includes the following new functions:

<code>localeconv()</code>	numeric/monetary conventions
<code>strcoll()</code>	collation order of two strings
<code>strxfrm()</code>	translate string for collation
<code>strxfrm()</code>	translate string for collation

In addition, there are the multibyte functions `mblen()`, `mbtowc()`, `mbstowcs()`, `wctomb()`, and `wcstombs()`.

The `localeconv()` function returns a pointer to a structure containing information useful for formatting numeric and monetary information appropriate to the current locale's `LC_NUMERIC` and `LC_MONETARY` categories. This is the only function whose behavior depends on more than one category. For numeric values, the structure describes the decimal-point character, the thousands separator, and where the separator(s) should be located. There are fifteen other structure members that describe how to format a monetary value.

The `strcoll()` function is analogous to the `strcmp()` function, except that the two strings are compared according to the `LC_COLLATE` category of the current locale. The `strxfrm()` function can also be used to transform a string into another, such that any two such after-translation strings can be passed to `strcmp()`, and get an ordering analogous to what `strcoll()` would have returned if passed the two pre-translation strings.

The `strftime()` function provides formatting similar to that used with `sprintf()` of the values in a `struct tm`, along with some date and time representations that depend on the `LC_TIME` category of the current locale. This function is based on the `asctime()` function released as part of UNIX System V Release 3.2.

Grouping and Evaluation in Expressions

One of the choices made by Dennis Ritchie in the design of C was to give compilers a license to rearrange expressions involving adjacent operators that are mathematically commutative and associative, even in the presence of parentheses. This is explicitly noted in the appendix in the *The C Programming Language* by Kernighan and Ritchie. However, ANSI/ISO C does not grant compilers this same freedom.

This section discusses the differences between these two definitions of C and clarifies the distinctions between an expression's side effects, grouping, and evaluation by considering the expression statement from the following code fragment.

```
int i, *p, f(void), g(void);
/*...*/
i = *++p + f() + g();
```

Definitions

The side effects of an expression are its modifications to memory and its accesses to `volatile` qualified objects. The side effects in the above expression are the updating of `i` and `p` and any side effects contained within the functions `f()` and `g()`.

An expression's grouping is the way values are combined with other values and operators. The above expression's grouping is primarily the order in which the additions are performed.

An expression's evaluation includes everything necessary to produce its resulting value. To evaluate an expression, all specified side effects must occur anywhere between the previous and next sequence point, and the specified operations are performed with a particular grouping. For the above expression, the updating of `i` and `p` must occur after the previous statement and by the `;` of this expression statement; the calls to the functions can occur in either order, any time after the previous statement, but before their return values are used. In particular, the operators that cause memory to be updated have no requirement to assign the new value before the value of the operation is used.

The K&R C Rearrangement License

The K&R C rearrangement license applies to the above expression because addition is mathematically commutative and associative. To distinguish between regular parentheses and the actual grouping of an expression, the left and right curly braces designate grouping. The three possible groupings for the expression are:

```
i = { { *++p + f() } + g() };  
i = { *++p + { f() + g() } };  
i = { { *++p + g() } + f() };
```

All of these are valid given K&R C rules. Moreover, all of these groupings are valid even if the expression were written instead, for example, in either of these ways:

```
i = *++p + ( f() + g() );  
i = ( g() + *++p ) + f();
```

If this expression is evaluated on an architecture for which either overflows cause an exception, or addition and subtraction are not inverses across an overflow, these three groupings behave differently if one of the additions overflows.

For such expressions on these architectures, the only recourse available in K&R C was to split the expression to force a particular grouping. The following are possible rewrites that respectively enforce the above three groupings:

```
i = *++p; i += f(); i += g()  
i = f(); i += g(); i += *++p;  
i = *++p; i += g(); i += f();
```

The ANSI/ISO C Rules

ANSI/ISO C does not allow operations to be rearranged that are mathematically commutative and associative, but that are not actually so on the target architecture. Thus, the precedence and associativity of the ANSI/ISO C grammar completely describes the grouping for all expressions; all expressions must be grouped as they are parsed. The expression under consideration is grouped in this manner:

```
i = { { *++p + f() } + g() };
```

This code still does not mean that `f()` must be called before `g()`, or that `p` must be incremented before `g()` is called.

In ANSI/ISO C, expressions need not be split to guard against unintended overflows.

The Parentheses

ANSI/ISO C is often erroneously described as honoring parentheses or evaluating according to parentheses due to an incomplete understanding or an inaccurate presentation.

Since ANSI/ISO C expressions simply have the grouping specified by their parsing, parentheses still only serve as a way of controlling how an expression is parsed; the natural precedence and associativity of expressions carry exactly the same weight as parentheses.

The above expression could have been written as:

```
i = ((*(++p)) + f()) + g();
```

with no different effect on its grouping or evaluation.

The As If Rule

There were several reasons for the K&R C rearrangement rules:

- The rearrangements provide many more opportunities for optimizations, such as compile-time constant folding.
- The rearrangements do not change the result of integral-typed expressions on most machines.
- Some of the operations are both mathematically and computationally commutative and associative on all machines.

The ANSI/ISO C Committee eventually became convinced that the rearrangement rules were intended to be an instance of the *as if* rule when applied to the described target architectures. ANSI/ISO C's *as if* rule is a general license that permits an implementation to deviate arbitrarily from the abstract machine description as long as the deviations do not change the behavior of a valid C program.

Thus, all the binary bitwise operators (other than shifting) are allowed to be rearranged on any machine because there is no way to notice such regroupings. On typical two's-complement machines in which overflow wraps around, integer expressions involving multiplication or addition can be rearranged for the same reason.

Therefore, this change in C does not have a significant impact on most C programmers.

Incomplete Types

The ANSI/ISO C standard introduced the term “incomplete type” to formalize a fundamental, yet misunderstood, portion of C, implicit from its beginnings. This section describes incomplete types, where they are permitted, and why they are useful.

Types

ANSI/ISO separates C's types into three distinct sets: function, object, and incomplete. Function types are obvious; object types cover everything else, except when the size of the object is not known. The Standard uses the term “object type” to specify that the designated object must have a known size, but it is important to know that incomplete types other than `void` also refer to an object.

There are only three variations of incomplete types: `void`, arrays of unspecified length, and structures and unions with unspecified content. The type `void` differs from the other two in that it is an incomplete type that cannot be completed, and it serves as a special function return and parameter type.

Completing Incomplete Types

An array type is completed by specifying the array size in a following declaration in the same scope that denotes the same object. When an array without a size is declared and initialized in the same declaration, the array has an incomplete type only between the end of its declarator and the end of its initializer.

An incomplete structure or union type is completed by specifying the content in a following declaration in the same scope for the same tag.

Declarations

Certain declarations can use incomplete types, but others require complete object types. Those declarations that require object types are array elements, members of structures or unions, and objects local to a function. All other declarations permit incomplete types. In particular, the following constructs are permitted:

- Pointers to incomplete types
- Functions returning incomplete types
- Incomplete function parameter types
- `typedef` names for incomplete types

The function return and parameter types are special. Except for `void`, an incomplete type used in such a manner must be completed by the time the function is defined or called. A return type of `void` specifies a function that returns no value, and a single parameter type of `void` specifies a function that accepts no arguments.

Since array and function parameter types are rewritten to be pointer types, a seemingly incomplete array parameter type is not actually incomplete. The typical declaration of `main`'s `argv`, namely, `char *argv[]`, as an unspecified length array of character pointers, is rewritten to be a pointer to character pointers.

Expressions

Most expression operators require complete object types. The only three exceptions are the unary `&` operator, the first operand of the comma operator, and the second and third operands of the `?:` operator. Most operators that accept pointer operands also permit pointers to incomplete types, unless pointer arithmetic is required. The list includes the unary `*` operator. For example, given:

```
void *p
```

`&*p` is a valid subexpression that makes use of this.

Justification

Why are incomplete types necessary? Ignoring `void`, there is only one feature provided by incomplete types that C has no other way to handle, and that has to do with forward references to structures and unions. If one has two structures that need pointers to each other, the only way to do so is with incomplete types:

```
struct a { struct b *bp; };
struct b { struct a *ap; };
```

All strongly typed programming languages that have some form of pointer and heterogeneous data types provide some method of handling this case.

Examples

Defining `typedef` names for incomplete structure and union types is frequently useful. If you have a complicated bunch of data structures that contain many pointers to each other, having a list of `typedefs` to the structures up front, possibly in a central header, can simplify the declarations.

```
typedef struct item_tag Item;
typedef union note_tag Note;
typedef struct list_tag List;
. . .
struct item_tag { . . . };
. . .
struct list_tag {
    struct list_tag {
};
```

Moreover, for those structures and unions whose contents should not be available to the rest of the program, a header can declare the tag without the content. Other parts of the program can use pointers to the incomplete structure or union without any problems, unless they attempt to use any of its members.

A frequently used incomplete type is an external array of unspecified length. Generally, it is not necessary to know the extent of an array to make use of its contents.

Compatible and Composite Types

With K&R C, and even more so with ANSI/ISO C, it is possible for two declarations that refer to the same entity to be other than identical. The term “compatible type” is used in ANSI/ISO C to denote those types that are “close enough”. This section describes compatible types as well as “composite types”—the result of combining two compatible types.

Multiple Declarations

If a C program were only allowed to declare each object or function once, there would be no need for compatible types. Linkage, which allows two or more declarations to refer to the same entity, function prototypes, and separate compilation all need such a capability. Separate translation units (source files) have different rules for type compatibility from within a single translation unit.

Separate Compilation Compatibility

Since each compilation probably looks at different source files, most of the rules for compatible types across separate compiles are structural in nature:

- Matching scalar (integral, floating, and pointer) types must be compatible, as if they were in the same source file.
- Matching structures, unions, and enums must have the same number of members. Each matching member must have a compatible type (in the separate compilation sense), including bit-field widths.
- Matching structures must have the members in the same order. The order of union and enum members does not matter.
- Matching enum members must have the same value.

An additional requirement is that the names of members, including the lack of names for unnamed members, match for structures, unions, and enums, but not necessarily their respective tags.

Single Compilation Compatibility

When two declarations in the same scope describe the same object or function, the two declarations must specify compatible types. These two types are then combined into a single composite type that is compatible with the first two. More about composite types later.

The compatible types are defined recursively. At the bottom are type specifier keywords. These are the rules that say that `unsigned short` is the same as `unsigned short int`, and that a type without type specifiers is the same as one with `int`. All other types are compatible only if the types from which they are derived are compatible. For example, two qualified types are compatible if the qualifiers, `const` and `volatile`, are identical, and the unqualified base types are compatible.

Compatible Pointer Types

For two pointer types to be compatible, the types they point to must be compatible and the two pointers must be identically qualified. Recall that the qualifiers for a pointer are specified after the `*`, so that these two declarations

```
int *const cpi;
int *volatile vpi;
```

declare two differently qualified pointers to the same type, `int`.

Compatible Array Types

For two array types to be compatible, their element types must be compatible. If both array types have a specified size, they must match, that is, an incomplete array type (see “Incomplete Types” on page 241) is compatible both with another incomplete array type and an array type with a specified size.

Compatible Function Types

To make functions compatible, follow these rules:

- For two function types to be compatible, their return types must be compatible. If either or both function types have prototypes, the rules are more complicated.

- For two function types with prototypes to be compatible, they also must have the same number of parameters, including use of the ellipsis (. . .) notation, and the corresponding parameters must be parameter-compatible.
- For an old-style function definition to be compatible with a function type with a prototype, the prototype parameters must *not* end with an ellipsis (. . .). Each of the prototype parameters must be parameter-compatible with the corresponding old-style parameter, after application of the default argument promotions.
- For an old-style function declaration (not a definition) to be compatible with a function type with a prototype, the prototype parameters must not end with an ellipsis (. . .). All of the prototype parameters must have types that would be unaffected by the default argument promotions.
- For two types to be parameter-compatible, the types must be compatible after the top-level qualifiers, if any, have been removed, and after a function or array type has been converted to the appropriate pointer type.

Special Cases

signed int behaves the same as int, except possibly for bit-fields, in which a plain int may denote an unsigned-behaving quantity.

Another interesting note is that each enumeration type must be compatible with some integral type. For portable programs, this means that enumeration types are separate types. In general, the ANSI/ISO C standard views them in that manner.

Composite Types

The construction of a composite type from two compatible types is also recursively defined. The ways compatible types can differ from each other are due either to incomplete arrays or to old-style function types. As such, the simplest description of the composite type is that it is the type compatible with both of the original types, including every available array size and every available parameter list from the original types.

Converting Applications

This chapter provides the information you need for writing code for the 32-bit or the 64-bit compilation environment. This chapter is organized into the following sections:

- “Overview of the Data Model Differences” on page 248
- “Implementing Single Source Code” on page 249
- “Converting to the LP64 Data Type Model” on page 253

Once you try to write or modify code for both the 32-bit and 64-bit compilation environments, you face two basic issues:

- Data type consistency between the different data-type models
- Interaction between the applications using different data-type models

Maintaining a single code-source with as few `#ifdefs` as possible is usually better than maintaining multiple source trees. Therefore, this chapter provides guidelines for writing code that works correctly in both 32-bit and 64-bit compilation environments. In some cases, the conversion of current code requires only a recompilation and relinking with the 64-bit libraries. However, for those cases where code changes are required, this chapter discusses the tools and strategies that make conversion easier.

The rest of this chapter provides the following information:

- “Overview of the Data Model Differences” on page 248 introduces the terminology that describes the 32-bit and 64-bit environments and provides an overview of some basic differences.
- “Implementing Single Source Code” on page 249 describe some of the available resources that you can use to write single-source code that supports 32-bit and 64-bit compilation.
- “Converting to the LP64 Data Type Model” on page 253 illustrates some of the more common problems you are likely to encounter when you convert code and where appropriate, shows the corresponding `lint` warnings.
- “Other Considerations” on page 261 provides general tips for troubleshooting code after you have made modifications.
- Finally, the “Checklist for Getting Started” on page 263 helps you get started.

Overview of the Data Model Differences

The biggest difference between the 32-bit and the 64-bit compilation environments is the change in data-type models.

The C data-type model for 32-bit applications is the ILP32 model, so named because integers, longs, and pointers are 32-bit data types. The LP64 data model, so named because longs and pointers grow to 64-bits, is the creation of a consortium of companies across the industry. The remaining C types int, long long, short, and char are the same in both data-type models.

Regardless of the data-type model, the standard relationship between C integral types holds true:

```
sizeof (char) <= sizeof (short) <= sizeof (int) <= sizeof (long)
```

The following table lists the basic C data types and their corresponding sizes in bits for both the ILP32 and LP64 data models.

TABLE 9-1 Data Type Size for ILP32 and LP64

C Data Type	LP32	LP64
char	8	8
short	16	16
int	32	32
long	32	64
long long	64	64
pointer	32	64
enum	32	32
float	32	32
double	64	64
long double	128	128

It is not unusual for current 32-bit applications to assume that integers, pointers, and longs are the same size. Because the size of longs and pointers change in the LP64 data model, you need to be aware that this change alone can cause many ILP32 to LP64 conversion problems.

In addition, it becomes very important to examine declarations and casts; how expressions are evaluated can be affected when the types change. The effects of standard C conversion rules are influenced by the change in data-type sizes. To adequately show what you intend, you need to explicitly declare the types of constants. You can also use casts in expressions to make certain that the expression is evaluated the way you intend. This is particularly true in the case of sign extension, where explicit casting is essential for demonstrating intent.

Implementing Single Source Code

The following sections describe some of the available resources that you can use to write single-source code that supports 32-bit and 64-bit compilation.

Derived Types

Use the system derived types to make code safe for both the 32-bit and the 64-bit compilation environment. In general, it is good programming practice to use derived types to allow for change. When you use derived data-types, only the system derived types need to change due to data model changes, or due to a port.

The system include files `<sys/types.h>` and `<inttypes.h>` contain constants, macros, and derived types that are helpful in making applications 32-bit and 64-bit safe.

`<sys/types.h>`

Include `<sys/types.h>` in an application source file to gain access to the definition of `_LP64` and `_ILP32`. This header also contains a number of basic derived types that should be used whenever appropriate. In particular, the following are of special interest:

- `clock_t` represents the system times in clock ticks.
- `dev_t` is used for device numbers.
- `off_t` is used for file sizes and offsets.
- `ptrdiff_t` is the signed integral type for the result of subtracting two pointers.
- `size_t` reflects the size, in bytes, of objects in memory.
- `ssize_t` is used by functions that return a count of bytes or an error indication.
- `time_t` counts time in seconds.

All of these types remain 32-bit quantities in the ILP32 compilation environment and grow to 64-bit quantities in the LP64 compilation environment.

<inttypes.h>

The include file <inttypes.h> provides constants, macros, and derived types that help you make your code compatible with explicitly sized data items, independent of the compilation environment. It contains mechanisms for manipulating 8-bit, 16-bit, 32-bit, and 64-bit objects. The file is part of an ANSI/ISO C proposal and tracks the ISO/JTC1/SC22/WG14 C committee's working draft for the revision of the current ISO C standard, ISO/IEC 9899:1990 Programming language - C. The following is a discussion of the basic features provided by <inttypes.h>:

- Fixed-width integer types.
- Helpful types such as `uintptr_t`
- Constant macros
- Limits
- Format string macros

The following sections provide more information about the basic features of <inttypes.h>.

Fixed-Width Integer Types

The fixed-width integer types that <inttypes.h> provides, include signed integer types, such as `int8_t`, `int16_t`, `int32_t`, `int64_t`, and unsigned integer types such as, `uint8_t`, `uint16_t`, `uint32_t`, and `uint64_t`.

Derived types defined as the smallest integer types that can hold the specified number of bits include `int_least8_t`,..., `int_least64_t`, `uint_least8_t`,..., `uint_least64_t`.

It is safe to use an integer for such operations as loop counters and file descriptors; it is also safe to use a long for an array index. However, do not use these fixed-width types indiscriminately. Use fixed-width types for explicit binary representations of the following:

- On-disk data
- Over the data wire
- Hardware registers
- Binary interface specifications
- Binary data structures

Helpful Types Such as `uintptr_t`

The <inttypes.h> file includes signed and unsigned integer types large enough to hold a pointer. These are given as `intptr_t` and `uintptr_t`. In addition, <inttypes.h> provides `intmax_t` and `uintmax_t` which are the longest (in bits) signed and unsigned integer types available.

Use the `uintptr_t` type as the integral type for pointers instead of a fundamental type such as unsigned long. Even though an unsigned long is the same size as a pointer in both the ILP32 and LP64 data models, using `uintptr_t` means that only the definition of `uintptr_t` is effected if the data model changes. This makes your code portable to many other systems. It is also a more clear way to express your intentions in C.

The `intptr_t` and `uintptr_t` types are extremely useful for casting pointers when you want to perform address arithmetic. Use `intptr_t` and `uintptr_t` types instead of long or unsigned long for this purpose.

Constant Macros

Use the macros `INT8_C(c)`, ..., `INT64_C(c)`, `UINT8_C(c)`, ..., `UINT64_C(c)` to specify the size and sign of a given constant. Basically, these macros place an `l`, `ul`, `ll`, or `ull` at the end of the constant, if necessary. For example, `INT64_C(1)` appends `ll` to the constant 1 for ILP32 and an `l` for LP64.

Use the `INTMAX_C(c)` and `UINTMAX_C(c)` macros to make a constant the biggest type. These macros can be very useful for specifying the type of constants described in “Converting to the LP64 Data Type Model” on page 253.

Limits

The limits defined by `<inttypes.h>` are constants that specify the minimum and maximum values of various integer types. This includes minimum and maximum values for each of the fixed-width types such as `INT8_MIN`, ..., `INT64_MIN`, `INT8_MAX`, ..., `INT64_MAX`, and their unsigned counterparts.

The `<inttypes.h>` file also provides the minimum and maximum for each of the least-sized types. These include `INT_LEAST8_MIN`, ..., `INT_LEAST64_MIN`, `INT_LEAST8_MAX`, ..., `INT_LEAST64_MAX`, as well as their unsigned counterparts.

Finally, `<inttypes.h>` defines the minimum and maximum value of the largest supported integer types. These include `INTMAX_MIN` and `INTMAX_MAX` and their corresponding unsigned versions.

Format String Macros

The `<inttypes.h>` file also includes the macros that specify the `printf(3S)` and `scanf(3S)` format specifiers. Essentially, these macros prepend the format specifier with an `l` or `ll` to identify the argument as a long or long long, given that the number of bits in the argument is built into the name of the macro.

There are macros for `printf(3S)` that print both the smallest and largest integer types in decimal, octal, unsigned, and hexadecimal formats as the following example shows:

```
int64_t i;
printf("i =%" PRIx64 "\n", i);
```

Similarly, there are macros for `scanf(3S)` that read both the smallest and largest integer types in decimal, octal, unsigned, and hexadecimal formats.

```
uint64_t u;
scanf("%" SCNu64 "\n", &u);
```

Do not use these macros indiscriminately. They are best used in conjunction with the fixed-width types discussed in “Fixed-Width Integer Types” on page 250.

Tools

Sun WorkShop includes an enhanced version of the `lint` program that detects potential 64-bit problems. In addition, the `-v` option to the C compiler performs additional and more strict semantic checks. The `-v` option also enables certain lint-like checks on the named files.

When you enhance code to be 64-bit safe, use the header files present in the Solaris 7 operating environment because these files have the correct definition of the derived types and data structures for the 64-bit compilation environment.

lint

Use `lint` to check code that is written for both the 32-bit and the 64-bit compilation environment. Issue the `-errchk=longptr64` option to generate LP64 warnings. Also use the `-errchk=longptr64` flag which checks portability to an environment for which the size of long integers and pointers is 64 bits and the size of plain integers is 32 bits. The `-errchklongptr64` flag checks assignments of pointer expressions and long integer expressions to plain integers, even when explicit casts are used.

Use the `-xarch=v9` option of `lint` when you want to check code that you intend to run in the 64-bit compilation environment only.

When `lint` generates warnings, it prints the line number of the offending code, a message that describes the problem, and whether or not a pointer is involved. The warning message also indicates the sizes of the involved data types. When you know a pointer is involved and you know the size of the data types, you can find specific 64-bit problems and avoid the pre-existing problems between 32-bit and smaller types.

Be aware, however, that even though `lint` gives warnings about potential 64-bit problems, it cannot detect all problems. Also, in many cases, code that is intentional and correct for the application generates a warning.

You can suppress the warning for a given line of code by placing a comment of the form `/*LINTED*/` on the previous line. This is useful when you want `lint` to ignore certain lines of code such as casts and assignments. Exercise extreme care when you use the `/*LINTED*/` comment because it can mask real problems. Refer to the `lint` man page for more information.

Converting to the LP64 Data Type Model

The examples that follow illustrate some of the more common problems you are likely to encounter when you convert code. Where appropriate, the corresponding `lint` warnings are shown.

Integer and Pointer Size Change

Since integers and pointers are the same size in the ILP32 compilation environment, some code relies on this assumption. Pointers are often cast to `int` or `unsigned int` for address arithmetic. Instead, cast your pointers to `long` because `long` and pointers are the same size in both ILP32 and LP64 data-type models. Rather than explicitly using `unsigned long`, use `uintptr_t` instead because it expresses your intent more closely and makes the code more portable, insulating it against future changes. Consider the following example:

```
char *p;
p = (char *) ((int)p & PAGEOFFSET);
%
warning: conversion of pointer loses bits
```

Here is the modified version:

```
char *p;  
p = (char *) ((uintptr_t)p & PAGEOFFSET);
```

Integer and Long Size Change

Because integers and longs are never really distinguished in the ILP32 data-type model, your existing code probably uses them indiscriminately. Modify any code that uses integers and longs interchangeably so it conforms to the requirements of both the ILP32 and LP64 data-type models. While an integer and a long are both 32-bits in the ILP32 data-type model, a long is 64 bits in the LP64 data-type model. Consider the following example:

```
int waiting;  
long w_io;  
long w_swap;  
...  
waiting = w_io + w_swap;  
  
%  
warning: assignment of 64-bit integer to 32-bit integer
```

Sign Extension

Sign extension is a common problem when you convert to the 64-bit compilation environment because the type conversion and promotion rules are somewhat obscure. To prevent sign extension problems, use explicit casting to achieve the intended results.

To understand why sign extension occurs, it helps to understand the conversion rules for ANSI/ISO C. The conversion rules that seem to cause the most sign extension problems between the 32-bit and the 64-bit compilation environment come into effect during the following operations:

- Integral promotion

You can use a `char`, `short`, `enumerated type`, or `bit-field`, whether signed or unsigned, in any expression that calls for an integer.

If an integer can hold all possible values of the original type, the value is converted to an integer; otherwise, the value is converted to an unsigned integer.

- Conversion between signed and unsigned integers

When an integer with a negative sign is promoted to an unsigned integer of the same or larger type, it is first promoted to the signed equivalent of the larger type, then converted to the unsigned value.

When the following example is compiled as a 64-bit program, the `addr` variable becomes sign-extended, even though both `addr` and `a.base` are unsigned types.

```
%cat test.c
struct foo {
  unsigned int base:19, rehash:13;
};

main(int argc, char *argv[])
{
  struct foo a;
  unsigned long addr;

  a.base = 0x40000;
  addr = a.base << 13; /* Sign extension here! */
  printf("addr 0x%lx\n", addr);

  addr = (unsigned int)(a.base << 13); /* No sign extension here! */
  printf("addr 0x%lx\n", addr);
}
```

This sign extension occurs because the conversion rules are applied as follows:

- `a.base` is converted from an unsigned `int` to an `int` because of the integral promotion rule. Thus, the expression `a.base << 13` is of type `int`, but no sign extension has yet occurred.
- The expression `a.base << 13` is of type `int`, but it is converted to a `long` and then to an unsigned `long` before being assigned to `addr`, because of signed and unsigned integer promotion rules. The sign extension occurs when it is converted from an `int` to a `long`.

```
% cc -o test64 -xarch=v9 test.c
% ./test64
addr 0xffffffff80000000
addr 0x80000000
%
```

When this same example is compiled as a 32-bit program it does not display any sign extension:

```
cc -o test test.c
%test

addr 0x80000000
addr 0x80000000
```

For a more detailed discussion of the conversion rules, refer to the ANSI/ISO C standard. Also included in this standard are useful rules for ordinary arithmetic conversions and integer constants.

Pointer Arithmetic Instead of Address Arithmetic

In general, using pointer arithmetic works better than address arithmetic because pointer arithmetic is independent of the data model, whereas address arithmetic might not be. Also, you can usually simplify your code by using pointer arithmetic. Consider the following example:

```
int *end;
int *p;
p = malloc(4 * NUM_ELEMENTS);
end = (int *)((unsigned int)p + 4 * NUM_ELEMENTS);

%
warning: conversion of pointer loses bits
```

Here is the modified version:

```
int *end;
int *p;
p = malloc(sizeof (*p) * NUM_ELEMENTS);
end = p + NUM_ELEMENTS;
```

Structures

Check the internal data structures in an applications for holes. Use extra padding between fields in the structure to meet alignment requirements. This extra padding is allocated when long or pointer fields grow to 64 bits for the LP64 data-type model.

In the 64-bit compilation environment on SPARC platforms, all types of structures are aligned to the size of the largest quantity within them. When you repack a structure, follow the simple rule of moving the long and pointer fields to the beginning of the structure. Consider the following structure definition:

```
struct bar {
    int i;
    long j;
    int k;
    char *p;
}; /* sizeof (struct bar) = 32 */
```

Here is the same structure with the long and pointer data types defined at the beginning of the structure:

```
struct bar {
    char *p;
    long j;
    int i;
    int k;
}; /* sizeof (struct bar) = 24 */
```

Unions

Be sure to check unions because their fields can change size between the ILP32 and the LP64 data-type models.

```
typedef union {
    double _d;
    long _l[2];
} llx_t;
```

Here is the modified version

```
typedef union {
    double _d;
    int _l[2];
} llx_t;
```

Type Constants

A lack of precision can cause the loss of data in some constant expressions. Be explicit when you specify the data types in your constant expression. Specify the type of each integer constant by adding some combination of {u,U,l,L}. You can also use casts to specify the type of a constant expression. Consider the following example:

```
int i = 32;
long j = 1 << i; /* j will get 0 because RHS is integer */
                /* expression */
```

Here is the modified version:

```
int i = 32;
long j = 1L << i;
```

Beware of Implicit Declarations

The C compiler assumes that any function or variable that is used in a module and not defined or declared externally is an integer. Any longs and pointers used in this way are truncated by the compiler's implicit integer declaration. Place the appropriate extern declaration for the function or variable in a header and not in the C module. Include this header in any C module that uses the function or variable. If this is a function or variable defined by the system headers, you still need to include the proper header in the code. Consider the following example:

```
int
main(int argc, char *argv[])
{
    char *name = getlogin()
    printf("login = %s\n", name);
    return (0);
}

%
warning: improper pointer/integer combination: op "="
warning: cast to pointer from 32-bit integer
implicitly declared to return int
getlogin      printf
```

The proper headers are now in the modified version

```
#include <unistd.h>
#include <stdio.h>

int
main(int argc, char *argv[])
{
    char *name = getlogin();
    (void) printf("login = %s\n", name);
    return (0);
}
```

sizeof() is an Unsigned Long

In the LP64 data-type model, `sizeof()` has the effective type of an unsigned long. Occasionally, `sizeof()` is passed to a function expecting an argument of type `int`, or assigned or cast to an integer. In some cases, this truncation causes loss of data.

```
long a[50];
unsigned char size = sizeof (a);

%
warning: 64-bit constant truncated to 8 bits by assignment
warning: initializer does not fit or is out of range: 0x190
```

Use Casts to Show Your Intentions

Relational expressions can be tricky because of conversion rules. You should be very explicit about how you want the expression to be evaluated by adding casts wherever necessary.

Check Format String Conversion Operation

Make sure the format strings for `printf(3S)`, `sprintf(3S)`, `scanf(3S)`, and `sscanf(3S)` can accommodate long or pointer arguments. For pointer arguments, the conversion operation given in the format string should be `%p` to work in both the 32-bit and 64-bit compilation environments.

```
char *buf;
struct dev_info *devi;
...
(void) sprintf(buf, "di%x", (void *)devi);

%
warning: function argument (number) type inconsistent with format
sprintf (arg 3)      void *: (format) int
```

Here is the modified version

```
char *buf;
struct dev_info *devi;
...
(void) sprintf(buf, 'di%p', (void *)devi);
```

For long arguments, the long size specification, `l`, should be prepended to the conversion operation character in the format string. Furthermore, check to be sure that the storage pointed to by `buf` is large enough to contain 16 digits.

```
size_t nbytes;
u_long align, addr, raddr, alloc;
printf("kalloca:%d%%d from heap got%x.%x returns%x\n",
nbytes, align, (int)raddr, (int)(raddr + alloc), (int)addr);

%
warning: cast of 64-bit integer to 32-bit integer
warning: cast of 64-bit integer to 32-bit integer
warning: cast of 64-bit integer to 32-bit integer
```


Here is the modified version

```
size_t nbytes;
u_long align, addr, raddr, alloc;
printf("kalloca:%lu%%%lu from heap got%lx.%lx returns%lx\n",
nbytes, align, raddr, raddr + alloc, addr);
```

Other Considerations

The remaining guidelines highlight common problems encountered when converting an application to a full 64-bit program.

Derived Types That Have Grown in Size

A number of derived types have changed to now represent 64-bit quantities in the 64-bit application compilation environment. This change does not affect 32-bit applications; however, any 64-bit applications that consume or export data described by these types need to be re-evaluated. An example of this is in applications that directly manipulate the `utmp(4)` or `utmpx(4)` files. For correct operation in the 64-bit application environment, do not attempt to directly access these files. Instead, use the `getutxent(3C)` and related family of functions.

Check for Side Effects of Changes

Be aware that a type change in one area can result in an unexpected 64-bit conversion in another area. For example, check all the callers of a function that previously returned an `int` and now returns an `ssize_t`.

Check Whether Literal Uses of `long` Still Make Sense

A variable that is defined as a `long` is 32 bits in the ILP32 data-type model and 64 bits in the LP64 data-type model. Where it is possible, avoid problems by redefining the variable and use a more portable derived type.

Related to this, a number of derived types have changed under the LP64 data-type model. For example, `pid_t` remains a long in the 32-bit environment, but under the 64-bit environment, a `pid_t` is an `int`.

Use `#ifdef` for Explicit 32-bit Versus 64-bit Prototypes

In some cases, specific 32-bit and 64-bit versions of an interface are unavoidable. You can distinguish these by specifying the `_LP64` or `_ILP32` feature test macros in the headers. Similarly, code that runs in 32-bit and 64-bit environments needs to utilize the appropriate `#ifdefs`, depending on the compilation mode.

Calling Convention Changes

When you pass structures by value and compile the code for SPARC V9, the structure is passed in registers rather than as a pointer to a copy if it is small enough. This can cause problems if you try to pass structures between C code and hand-written assembly code.

Floating point parameters work in a similar fashion; some floating point values passed by value are passed in floating point registers.

Algorithm Changes

After your code is safe for the 64-bit environment, review your code again to verify that the algorithms and data structures still make sense. The data types are larger, so data structures might use more space. The performance of your code might change as well. Given these concerns, you might need to modify your code appropriately.

Checklist for Getting Started

Use the following checklist to help you convert your code to 64-bit.

- Review all data structures and interfaces to verify that these are still valid in the 64-bit environment.
- Include `<sys/types.h>` (or at a minimum, `<sys/isa_defs.h>`) in your code to pull in the `_ILP32` or `_LP64` definitions as well as many basic derived types.
- Move function prototypes and external declarations with non-local scope to headers and include these headers in your code.
- Run lint using the `-errchk=longptr64` and `-D__sparcv9` flags and review each warning individually. Keep in mind that not all warnings require a change to the code. Depending on the changes, run lint again in both 32-bit and 64-bit modes.
- Compile code as both 32-bit and 64-bit, unless the application is being provided only as 64-bit.
- Test the application by executing the 32-bit version on the 32-bit operating system, and the 64-bit version on the 64-bit operating system. You can also test the 32-bit version on the 64-bit operating system.

cscope: Interactively Examining a C Program

`cscope` is an interactive program that locates specified elements of code in C, `lex`, or `yacc` source files. With `cscope`, you can search and edit your source files more efficiently than you could with a typical editor. That's because `cscope` supports function calls—when a function is being called, when it is doing the calling—as well as C language identifiers and keywords.

This chapter is a tutorial on the `cscope` browser provided with this release and is organized into the following sections:

- “The `cscope` Process” on page 265
 - “Basic Use” on page 266
 - “Unknown Terminal Type Error” on page 284
-

The `cscope` Process

When `cscope` is called for a set of C, `lex`, or `yacc` source files, it builds a symbol cross-reference table for the functions, function calls, macros, variables, and preprocessor symbols in those files. You can then query that table about the locations of symbols you specify. First, it presents a menu and asks you to choose the type of search you would like to have performed. You may, for instance, want `cscope` to find all the functions that call a specified function.

When `cscope` has completed this search, it prints a list. Each list entry contains the name of the file, the number of the line, and the text of the line in which `cscope` has found the specified code. In our case, the list also includes the names of the functions that call the specified function. You now have the option of requesting another search or examining one of the listed lines with the editor. If you choose the latter, `cscope` invokes the editor for the file in which the line appears, with the cursor on that line. You can now view the code in context and, if you wish, edit the file as any other file. You can then return to the menu from the editor to request a new search.

Because the procedure you follow depends on the task at hand, there is no single set of instructions for using `cscope`. For an extended example of its use, review the `cscope` session described in the next section. It shows how you can locate a bug in a program without learning all the code.

Basic Use

Suppose you are given responsibility for maintaining the program `prog`. You are told that an error message, out of storage, sometimes appears just as the program starts up. Now you want to use `cscope` to locate the parts of the code that are generating the message. Here is how you do it.

Step 1: Set Up the Environment

`cscope` is a screen-oriented tool that can only be used on terminals listed in the Terminal Information Utilities (`terminfo`) database. Be sure you have set the `TERM` environment variable to your terminal type so that `cscope` can verify that it is listed in the `terminfo` database. If you have not done so, assign a value to `TERM` and export it to the shell as follows:

In a Bourne shell, type:

```
$ TERM=term_name; export TERM
```

In a C shell, type:

```
% setenv TERM term_name
```

You may now want to assign a value to the `EDITOR` environment variable. By default, `cscope` invokes the `vi` editor. (The examples in this chapter illustrate `vi` usage.) If you prefer not to use `vi`, set the `EDITOR` environment variable to the editor of your choice and export `EDITOR`, as follows:

In a Bourne shell, type:

```
$ EDITOR=emacs; export EDITOR
```

In a C shell, type:

```
% setenv EDITOR emacs
```

You may have to write an interface between `cscope` and your editor. For details, see “Command-Line Syntax for Editors” on page 283.

If you want to use `cscope` only for browsing (without editing), you can set the `VIEWER` environment variable to `pg` and export `VIEWER`. `cscope` will then invoke `pg` instead of `vi`.

An environment variable called `VPATH` can be set to specify directories to be searched for source files. See “View Paths” on page 278.

Step 2: Invoke the `cscope` Program

By default, `cscope` builds a symbol cross-reference table for all the C, `lex`, and `yacc` source files in the current directory, and for any included header files in the current directory or the standard place. So, if all the source files for the program to be browsed are in the current directory, and if its header files are there or in the standard place, invoke `cscope` without arguments:

```
% cscope
```

To browse through selected source files, invoke `cscope` with the names of those files as arguments:

```
% cscope file1.c file2.c file3.h
```

For other ways to invoke `cscope`, see “Command-Line Options” on page 275.

`cscope` builds the symbol cross-reference table the first time it is used on the source files for the program to be browsed. By default, the table is stored in the file `cscope.out` in the current directory. On a subsequent invocation, `cscope` rebuilds the cross-reference only if a source file has been modified or the list of source files is different. When the cross-reference is rebuilt, the data for the unchanged files is copied from the old cross-reference, which makes rebuilding faster than the initial build, and reduces startup time for subsequent invocations.

Step 3: Locate the Code

Now let's return to the task we undertook at the beginning of this section: to identify the problem that is causing the error message out of storage to be printed. You have invoked `cscope`, the cross-reference table has been built. The `cscope` menu of tasks appears on the screen.

The `cscope` Menu of Tasks:

```
% cscope

cscope      Press the ? key for help

Find this C symbol:
Find this global definition:
Find functions called by this function:
Find functions calling this function:
Find this text string:
Change this text string:
Find this egrep pattern:
Find this file:
Find files #including this file:
```

Press the Return key to move the cursor down the screen (with wraparound at the bottom of the display), and `^p` (Control-p) to move the cursor up; or use the up (`ua`) and down (`da`) arrow keys. You can manipulate the menu and perform other tasks with the following single-key commands:

TABLE 10-1 `cscope` Menu Manipulation Commands

Tab	Move to the next input field.
Return	Move to the next input field.
<code>^n</code>	Move to the next input field.
<code>^p</code>	Move to the previous input field.
<code>^y</code>	Search with the last text typed.
<code>^b</code>	Move to the previous input field and search pattern.
<code>^f</code>	Move to the next input field and search pattern.
<code>^c</code>	Toggle ignore/use letter case when searching. For example, a search for <code>FILE</code> matches <code>file</code> and <code>File</code> when ignoring the letter case.
<code>^r</code>	Rebuild cross-reference.

TABLE 10-1 cscope Menu Manipulation Commands (*Continued*)

!	Start an interactive shell. Type ^d to return to cscope.
^l	Redraw the screen.
?	Display the list of commands.
^d	Exit cscope.

If the first character of the text for which you are searching matches one of these commands, you can escape the command by entering a \ (backslash) before the character.

Now move the cursor to the fifth menu item, `Find this text string`, enter the text out of storage, and press the Return key.

cscope Function: Requesting a Search for a Text String:

```
$ cscope
cscope      Press the ? key for help

Find this C symbol
Find this global definition
Find functions called by this function
Find functions calling this function
Find this text string:  out of storage
Change this text string
Find this egrep pattern
Find this file
Find files #including this file
```

Note – Follow the same procedure to perform any other task listed in the menu except the sixth, `Change this text string`. Because this task is slightly more complex than the others, there is a different procedure for performing it. For a description of how to change a text string, see “Examples” on page 279.

cscope searches for the specified text, finds one line that contains it, and reports its finding.

cscope Function: Listing Lines Containing the Text String:

```
Text string: out of storage

File Line
1 alloc.c 63 (void) fprintf(stderr, "\n%s: out of storage\n",
argv0);

Find this C symbol:
Find this global definition:
Find functions called by this function:
Find functions calling this function:
Find this text string:
Change this text string:
Find this egrep pattern:
Find this file:
Find files #including this file:
```

After `cscope` shows you the results of a successful search, you have several options. You may want to change one of the lines or examine the code surrounding it in the editor. Or, if `cscope` has found so many lines that a list of them does not fit on the screen at once, you may want to look at the next part of the list. The following table shows the commands available after `cscope` has found the specified text:

TABLE 10-2 Commands for Use After an Initial Search

1 - 9	Edit the file referenced by this line. The number you type corresponds to an item in the list of lines printed by <code>cscope</code> .
Space	Display the next set of matching lines.
+	Display the next set of matching lines.
^v	Display the next set of matching lines.
-	Display the previous set of matching lines.
^e	Edit the displayed files in order.
>	Append the list of lines being displayed to a file.
	Pipe all lines to a shell command.

Again, if the first character of the text for which you are searching matches one of these commands, you can escape the command by entering a backslash before the character.

Now examine the code around the newly found line. Enter 1 (the number of the line in the list). The editor is invoked with the file `alloc.c` with the cursor at the beginning of line 63 of `alloc.c`.

`cscope` Function: Examining a Line of Code:

```
{
    return(alloctest(realloc(p, (unsigned) size)));
}

/* check for memory allocation failure */

static char *
alloctest(p)
char *p;
{
    if (p == NULL) {
        (void) fprintf(stderr, "\n%s: out of storage\n", argv0);
        exit(1);
    }
    return(p);
}
~
~
~
~
~
~
~
~
"alloc.c" 67 lines, 1283 characters
```

You can see that the error message is generated when the variable `p` is `NULL`. To determine how an argument passed to `alloctest()` could have been `NULL`, you must first identify the functions that call `alloctest()`.

Exit the editor by using normal quit conventions. You are returned to the menu of tasks. Now type `alloctest` after the fourth item, Find functions calling this function.

cscope Function: Requesting a List of Functions That Call `allocstest()`:

```
Text string: out of storage

File Line
1 alloc.c 63(void)fprintf(stderr,"\n%s:  out of storage\n",argv0);

Find this C symbol:
Find this global definition:
Find functions called by this function:
Find functions calling this function:  allocstest
Find this text string:
Change this text string:
Find this egrep pattern:
Find this file:
Find files #including this file:
```

cscope finds and lists three such functions.

cscope Function: Listing Functions That Call `allocstest()`:

```
Functions calling this function: allocstest
File Function Line
1 alloc.c mymalloc 33 return(allocstest(malloc((unsigned) size)));
2 alloc.c mycalloc 43 return(allocstest(calloc((unsigned) nelem,
(unsigned) size)));
3 alloc.c myrealloc 53 return(allocstest(realloc(p, (unsigned)
size)));

Find this C symbol:
Find this global definition:
Find functions called by this function:
Find functions calling this function:
Find this text string:
Change this text string:
Find this egrep pattern:
Find this file:
Find files #including this file:
```

Now you want to know which functions call `mymalloc()`. cscope finds ten such functions. It lists nine of them on the screen and instructs you to press the space bar to see the rest of the list.

cscope Function: Listing Functions That Call mymalloc():

```
Functions calling this function: mymalloc
```

File	Function	Line
1 alloc.c	stralloc	24 return(strcpy(mymalloc (strlen(s) + 1), s));
2 crossref.c	crossref	47 symbol = (struct symbol *)mymalloc (msymbols * sizeof(struct symbol));
3 dir.c	makevpsrkdirs	63 srcdirs = (char **) mymalloc (nsrkdirs * sizeof(char*));
4 dir.c	addinmdir	167 incdirs = (char **)mymalloc (sizeof(char *));
5 dir.c	addinmdir	168 incnames = (char **) mymalloc(sizeof(char *));
6 dir.c	addsrcfile	439 p = (struct listitem *) mymalloc (sizeof(struct listitem));
7 display.c	dispinit	87 displine = (int *) mymalloc (mdisprefs * sizeof(int));
8 history.c	addcmd	19 h = (struct cmd *) mymalloc (sizeof(struct cmd));
9 main.c	main	212 s = mymalloc((unsigned) (strlen(reffile) +strlen(home) + 2));

```
* 9 more lines - press the space bar to display more *
```

```
Find this C symbol:
```

```
Find this global definition:
```

```
Find functions called by this function:
```

```
Find functions calling this function:
```

```
Find this text string:
```

```
Change this text string:
```

```
Find this egrep pattern:
```

```
Find this file:
```

```
Find files #including this file:
```

Because you know that the error message out of storage is generated at the beginning of the program, you can guess that the problem may have occurred in the function `dispinit()` (display initialization).

To view `dispinit()`, the seventh function on the list, type 7.

cscope Function: Viewing `dispinit()` in the Editor:

```
void
dispinit()
{
    /* calculate the maximum displayed reference lines */
    lastdispline = FLDLINE - 4;
    mdisprefs = lastdispline - REFLINE + 1;
    if (mdisprefs > 9) {
        mdisprefs = 9;
    }
    /* allocate the displayed line array */
    dispiline = (int *) mymalloc(mdisprefs * sizeof(int));
}
^L/* display a page of the references */

void
display()
{
    char file[PATHLEN + 1]; /* file name */
    char function[PATLEN + 1]; /* function name */
    char linenum[NUMLLEN + 1]; /* line number */
    int screenline; /* screen line number */
    int width; /* source line display width */
    register int i, j;
}
"display.c" 622 lines, 14326 characters
```

`mymalloc()` failed because it was called either with a very large number or a negative number. By examining the possible values of `FLDLINE` and `REFLINE`, you can see that there are situations in which the value of `mdisprefs` is negative, that is, in which you are trying to call `mymalloc()` with a negative number.

Step 4: Edit the Code

On a windowing terminal, you may have multiple windows of arbitrary size. The error message out of storage might have appeared as a result of running `prog` in a window with too few lines. In other words, that may have been one of the situations in which `mymalloc()` was called with a negative number. Now you want to be sure that when the program aborts in this situation in the future, it does so after printing the more meaningful error message screen too small. Edit the function `dispinit()` as follows.

cscope Function: Correcting the Problem:

```
void
dispinit()
{
/* calculate the maximum displayed reference lines */
    lastdispline = FLDLINE - 4;
    mdisprefs = lastdispline - REFLINE + 1;
    if (mdisprefs > 9) {
        mdisprefs = 9;
    }
/* allocate the displayed line array */
    displine = (int *) mymalloc(mdisprefs * sizeof(int));
}
^L/* display a page of the references */

void
display()
{
    char file[PATHLEN + 1]; /* file name */
    char function[PATLEN + 1]; /* function name */
    char linenum[NUMLEN + 1]; /* line number */
    int screenline; /* screen line number */
    int width; /* source line display width */
    register int i, j;
"display.c" 622 lines, 14326 characters
}
```

You have fixed the problem we began investigating at the beginning of this section. Now if `prog` is run in a window with too few lines, it does not simply fail with the unedifying error message out of storage. Instead, it checks the window size and generates a more meaningful error message before exiting.

Command-Line Options

As noted, `cscope` builds a symbol cross-reference table for the `C`, `lex`, and source files in the current directory by default. That is,

```
% cscope
```

is equivalent to:

```
% cscope *. [chly]
```

We have also seen that you can browse through selected source files by invoking `cscope` with the names of those files as arguments:

```
% cscope file1.c file2.c file3.h
```

`cscope` provides command-line options with greater flexibility in specifying source files to be included in the cross-reference. When you invoke `cscope` with the `-s` option and any number of directory names (separated by commas):

```
% cscope -s dir1,dir2,dir3
```

`cscope` builds a cross-reference for all the source files in the specified directories as well as the current directory. To browse through all of the source files whose names are listed in *file* (file names separated by spaces, tabs, or new-lines), invoke `cscope` with the `-i` option and the name of the file containing the list:

```
% cscope -i file
```

If your source files are in a directory tree, use the following commands to browse through all of them:

```
% find . -name '*.chly' -print | sort > file  
% cscope -i file
```

If this option is selected, however, `cscope` ignores any other files appearing on the command-line.

The `-I` option can be used for `cscope` in the same way as the `-I` option to `cc`. See “Include Files” on page 89.

You can specify a cross-reference file other than the default `cscope.out` by invoking the `-f` option. This is useful for keeping separate symbol cross-reference files in the same directory. You may want to do this if two programs are in the same directory, but do not share all the same files:

```
% cscope -f admin.ref admin.c common.c aux.c libs.c  
% cscope -f delta.ref delta.c common.c aux.c libs.c
```


In this example, the source files for two programs, `admin` and `delta`, are in the same directory, but the programs consist of different groups of files. By specifying different symbol cross-reference files when you invoke `cscope` for each set of source files, the cross-reference information for the two programs is kept separate.

You can use the `-pn` option to specify that `cscope` display the path name, or part of the path name, of a file when it lists the results of a search. The number you give to `-p` stands for the last n elements of the path name you want to be displayed. The default is 1, the name of the file itself. So if your current directory is `home/common`, the command:

```
% cscope -p2
```

causes `cscope` to display `common/file1.c`, `common/file2.c`, and so forth when it lists the results of a search.

If the program you want to browse contains a large number of source files, you can use the `-b` option, so that `cscope` stops after it has built a cross-reference; `cscope` does not display a menu of tasks. When you use `cscope -b` in a pipeline with the `batch(1)` command, `cscope` builds the cross-reference in the background:

```
% echo 'cscope -b' | batch
```

Once the cross-reference is built, and as long as you have not changed a source file or the list of source files in the meantime, you need only specify:

```
% cscope
```

for the cross-reference to be copied and the menu of tasks to be displayed in the normal way. You can use this sequence of commands when you want to continue working without having to wait for `cscope` to finish its initial processing.

The `-d` option instructs `cscope` not to update the symbol cross-reference. You can use it to save time if you are sure that no such changes have been made; `cscope` does not check the source files for changes.

Note – Use the `-d` option with care. If you specify `-d` under the erroneous impression that your source files have not been changed, `cscope` refers to an outdated symbol cross-reference in responding to your queries.

Check the `cscope(1)` man page for other command-line options.

View Paths

As we have seen, `cscope` searches for source files in the current directory by default. When the environment variable `VPATH` is set, `cscope` searches for source files in directories that comprise your view path. A view path is an ordered list of directories, each of which has the same directory structure below it.

For example, suppose you are part of a software project. There is an *official* set of source files in directories below `/fs1/ofc`. Each user has a home directory (`/usr/you`). If you make changes to the software system, you may have copies of just those files you are changing in `/usr/you/src/cmd/prog1`. The official versions of the entire program can be found in the directory `/fs1/ofc/src/cmd/prog1`.

Suppose you use `cscope` to browse through the three files that comprise `prog1`, namely, `f1.c`, `f2.c`, and `f3.c`. You would set `VPATH` to `/usr/you` and `/fs1/ofc` and export it, as in:

In a Bourne shell, type:

```
$ VPATH=/usr/you:/fs1/ofc; export VPATH
```

In a C shell, type:

```
% setenv VPATH /usr/you:/fs1/ofc
```

You then make your current directory `/usr/you/src/cmd/prog1`, and invoke `cscope`:

```
% cscope
```

The program locates all the files in the view path. In case duplicates are found, `cscope` uses the file whose parent directory appears earlier in `VPATH`. Thus, if `f2.c` is in your directory, and all three files are in the official directory, `cscope` examines `f2.c` from your directory, and `f1.c` and `f3.c` from the official directory.

The first directory in `VPATH` must be a prefix of the directory you will be working in, usually `$HOME`. Each colon-separated directory in `VPATH` must be absolute: it should begin at `/`.

cscope and Editor Call Stacks

`cscope` and editor calls can be stacked. That is, when `cscope` puts you in the editor to view a reference to a symbol and there is another reference of interest, you can invoke `cscope` again from within the editor to view the second reference without exiting the current invocation of either `cscope` or the editor. You can then back up by exiting the most recent invocation with the appropriate `cscope` and editor commands.

Examples

This section presents examples of how `cscope` can be used to perform three tasks: changing a constant to a preprocessor symbol, adding an argument to a function, and changing the value of a variable. The first example demonstrates the procedure for changing a text string, which differs slightly from the other tasks on the `cscope` menu. That is, once you have entered the text string to be changed, `cscope` prompts you for the new text, displays the lines containing the old text, and waits for you to specify which of these lines you want it to change.

Changing a Constant to a Preprocessor Symbol

Suppose you want to change a constant, `100`, to a preprocessor symbol, `MAXSIZE`. Select the sixth menu item, `Change this text string`, and enter `\100`. The `1` must be escaped with a backslash because it has a special meaning (item `1` on the menu) to `cscope`. Now press Return. `cscope` prompts you for the new text string. Type `MAXSIZE`.

`cscope` Function: Changing a Text String:

```
cscope Press the ? key for help

Find this C symbol:
Find this global definition:
Find functions called by this function:
Find functions calling this function:
Find this text string:
Change this text string: \100
Find this egrep pattern:
Find this file:
Find files #including this file:
To: MAXSIZE
```

`cscope` displays the lines containing the specified text string, and waits for you to select those in which you want the text to be changed.

`cscope` Function: Prompting for Lines to be Changed:

```
cscope Press the ? key for help

Find this C symbol:
Find this global definition:
Find functions called by this function:
Find functions calling this function:
Find this text string:
Change this text string: \100
Find this egrep pattern:
Find this file:
Find files #including this file:
To: MAXSIZE
```

You know that the constant 100 in lines 1, 2, and 3 of the list (lines 4, 26, and 8 of the listed source files) should be changed to `MAXSIZE`. You also know that 0100 in `read.c` and 100.0 in `err.c` (lines 4 and 5 of the list) should not be changed. You select the lines you want changed with the following single-key commands:

TABLE 10-3 Commands for Selecting Lines to Be Changed

1-9	Mark or unmark the line to be changed.
*	Mark or unmark all displayed lines to be changed.
Space	Display the next set of lines.
+	Display the next set of lines.
-	Display the previous set of lines.
a	Mark all lines to be changed.
^d	Change the marked lines and exit.
Esc	Exit without changing the marked lines.

In this case, enter 1, 2, and 3. The numbers you type are not printed on the screen. Instead, `cscope` marks each list item you want to be changed by printing a `>` (greater than) symbol after its line number in the list.

cscope Function: Marking Lines to be Changed:

```
Change "100" to "MAXSIZE"

File Line
1>init.c 4 char s[100];
2>init.c 26 for (i = 0; i < 100; i++)
3>find.c 8 if (c < 100) {
4 read.c 12 f = (bb & 0100);
5 err.c 19 p = total/100.0; /* get percentage */

Find this C symbol:
Find this global definition:
Find functions called by this function:
Find functions calling this function:
Find this text string:
Change this text string:
Find this egrep pattern:
Find this file:
Find files #including this file:
Select lines to change (press the ? key for help):
```

Now type `^d` to change the selected lines. `cscope` displays the lines that have been changed and prompts you to continue.

cscope Function: Displaying Changed Lines of Text:

```
Changed lines:

char s[MAXSIZE];
for (i = 0; i < MAXSIZE; i++)
if (c < MAXSIZE) {

Press the RETURN key to continue:
```

When you press Return in response to this prompt, `cscope` redraws the screen, restoring it to its state before you selected the lines to be changed.

The next step is to add the `#define` for the new symbol `MAXSIZE`. Because the header file in which the `#define` is to appear is not among the files whose lines are displayed, you must escape to the shell by typing `!`. The shell prompt appears at the bottom of the screen. Then enter the editor and add the `#define`.

cscope Function: Exiting to the Shell:

```
Text string: 100

File Line
1 init.c 4 char s[100];
2 init.c 26 for (i = 0; i < 100; i++)
3 find.c 8 if (c < 100) {
4 read.c 12 f = (bb & 0100);
5 err.c 19 p = total/100.0;/* get percentage */

Find this C symbol:
Find this global definition:
Find functions called by this function:
Find functions calling this function:
Find this text string:
Change this text string:
Find this egrep pattern:
Find this file:
Find files #including this file:
$ vi defs.h
```

To resume the cscope session, quit the editor and type ^d to exit the shell.

Adding an Argument to a Function

Adding an argument to a function involves two steps: editing the function itself and adding the new argument to every place in the code where the function is called.

First, edit the function by using the second menu item, Find this global definition. Next, find out where the function is called. Use the fourth menu item, Find functions calling this function, to obtain a list of all the functions that call it. With this list, you can either invoke the editor for each line found by entering the list number of the line individually, or invoke the editor for all the lines automatically by typing ^e. Using cscope to make this kind of change ensures that none of the functions you need to edit are overlooked.

Changing the Value of a Variable

At times, you may want to see how a proposed change affects your code.

Suppose you want to change the value of a variable or preprocessor symbol. Before doing so, use the first menu item, `Find this C symbol`, to obtain a list of references that are affected. Then use the editor to examine each one. This step helps you predict the overall effects of your proposed change. Later, you can use `cscope` in the same way to verify that your changes have been made.

Command-Line Syntax for Editors

`cscope` invokes the `vi` editor by default. You can override the default setting by assigning your preferred editor to the `EDITOR` environment variable and exporting `EDITOR`, as described in “Step 1: Set Up the Environment” on page 266. However, `cscope` expects the editor it uses to have a command-line syntax of the form:

```
% editor +linenum filename
```

as does `vi`. If the editor you want to use does not have this command-line syntax, you must write an interface between `cscope` and the editor.

Suppose you want to use `ed`. Because `ed` does not allow specification of a line number on the command-line, you cannot use it to view or edit files with `cscope` unless you write a shell script that contains the following line:

```
/usr/bin/ed $2
```

Let’s name the shell script `myedit`. Now set the value of `EDITOR` to your shell script and export `EDITOR`:

In a Bourne shell, type:

```
$ EDITOR=myedit; export EDITOR
```

In a C shell, type:

```
% setenv EDITOR myedit
```

When `cscope` invokes the editor for the list item you have specified, say, line 17 in `main.c`, it invokes your shell script with the command-line:

```
% myedit +17 main.c
```

`myedit` then discards the line number (`$1`) and calls `ed` correctly with the file name (`$2`). Of course, you are not moved automatically to line 17 of the file and must execute the appropriate `ed` commands to display and edit the line.

Unknown Terminal Type Error

If you see the error message:

```
Sorry, I don't know how to deal with your "term" terminal
your terminal may not be listed in the Terminal Information Utilities (terminfo)
database that is currently loaded. Make sure you have assigned the correct value to
TERM. If the message reappears, try reloading the Terminal Information Utilities.
```

If this message is displayed:

```
Sorry, I need to know a more specific terminal type than
"unknown"
```

set and export the `TERM` variable as described in “Step 1: Set Up the Environment” on page 266.

ANSI/ISO C Data Representations

This appendix describes how ANSI C represents data in storage and the mechanisms for passing arguments to functions. It is intended as a guide to programmers who want to write or use modules in languages other than C and have those modules interface with C code. This appendix is organized into the following sections:

- “Storage Allocation” on page 285
 - “Data Representations” on page 286
 - “Argument-Passing Mechanism” on page 294
-

Storage Allocation

The following table shows the data types and how they are represented.

Note – Storage allocated on the stack (identifiers with internal, or automatic, linkage) should be limited to two gigabytes or less.

TABLE A-1 Storage Allocation for Data Types

Data Type	Internal Representation
char elements	A single 8-bit byte aligned on a byte boundary.
short integers	Halfword (two bytes or 16 bits), aligned on a two-byte boundary
int	32 bits on v8 (four bytes or one word), aligned on a four-byte boundary 32 bits on v9 (four bytes or one word), aligned on a four-byte boundary

TABLE A-1 Storage Allocation for Data Types (*Continued*)

Data Type	Internal Representation
long	32 bits on v8 (four bytes or one word), aligned on a four-byte boundary 64 bits on v9 (eight bytes or two words) aligned on an eight-byte boundary)
long long ¹	(SPARC) 64 bits (eight bytes or two words), aligned on an eight-byte boundary (Intel) 64 bits (eight bytes or two words), aligned on a four-byte boundary
float	32 bits (four bytes or one word), aligned on a four-byte boundary. A float has a sign bit, 8-bit exponent, and 23-bit fraction.
double	64 bits (eight bytes or two words), aligned on an eight-byte boundary (SPARC) or aligned on a four-byte boundary (Intel). A double element has a sign bit, an 11-bit exponent and a 52-bit fraction.
long double	v8 (SPARC) 128 bits (16 bytes or four words), aligned on an eight-byte boundary. A long double element has a sign bit, a 15-bit exponent and a 112-bit fraction. v9 (SPARC) 128 bits (16 bytes or four words), aligned on a 16 byte boundary. A long double element has a sign bit, a 15-bit exponent and a 112-bit fraction. (Intel) 96 bits (12 bytes or three words) aligned on a four-byte boundary. A long double element has a sign bit, a 16-bit exponent, and a 64-bit fraction. 16 bits are unused.

1. long long is not available in -Xc mode.

Data Representations

Bit numberings of any given data element depend on the architecture in use: SPARCstation™ machines use bit 0 as the least significant bit, with byte 0 being the most significant byte. The tables in this section describe the various representations.

Integer Representations

Integer types used in ANSI C are short, int, long, and long long:

TABLE A-2 Representation of short

Bits	Content
8 - 15	Byte 0 (<i>SPARC</i>) Byte 1 (<i>Intel</i>)
0 - 7	Byte 1 (<i>SPARC</i>) Byte 0 (<i>Intel</i>)

TABLE A-3 Representation of int

Bits	Content
24 - 31	Byte 0 (<i>SPARC</i>) Byte 3 (<i>Intel</i>)
16 - 23	Byte 1 (<i>SPARC</i>) Byte 2 (<i>Intel</i>)
8 - 15	Byte 2 (<i>SPARC</i>) Byte 1 (<i>Intel</i>)
0 - 7	Byte 3 (<i>SPARC</i>) Byte 0 (<i>Intel</i>)

TABLE A-4 Representation of long on Intel and SPARC v8 versus SPARC v9

Bits	Content
24 - 31	Byte 0 (<i>SPARC</i>) v8 Byte 4 (<i>SPARC</i>) v9 Byte 3 (<i>Intel</i>)
16 - 23	Byte 1 (<i>SPARC</i>) v8 Byte 5 (<i>SPARC</i>) v9 Byte 2 (<i>Intel</i>)
8 - 15	Byte 2 (<i>SPARC</i>) v8 Byte 6 (<i>SPARC</i>) v9 Byte 1 (<i>Intel</i>)
0 - 7	Byte 3 (<i>SPARC</i>) v8 Byte 7 (<i>SPARC</i>) v9 Byte 0 (<i>Intel</i>)

TABLE A-5 Representation of long long¹

Bits	Content
56 - 63	Byte 0 (<i>SPARC</i>) Byte 7 (<i>Intel</i>)
48 - 55	Byte 1 (<i>SPARC</i>) Byte 6 (<i>Intel</i>)
40 - 47	Byte 2 (<i>SPARC</i>) Byte 5 (<i>Intel</i>)
32 - 39	Byte 3 (<i>SPARC</i>) Byte 4 (<i>Intel</i>)
24 - 31	Byte 4 (<i>SPARC</i>) Byte 3 (<i>Intel</i>)
16 - 23	Byte 5 (<i>SPARC</i>) Byte 2 (<i>Intel</i>)
8 - 15	Byte 6 (<i>SPARC</i>) Byte 1 (<i>Intel</i>)
0 - 7	Byte 7 (<i>SPARC</i>) Byte 0 (<i>Intel</i>)

1. long long is not available in -Xc mode.

Floating-Point Representations

float, double, and long double data elements are represented according to the ANSI/ISO IEEE 754-1985 standard. The representation is:

$$(-1)^s e^{-bias} \times 2^{j:f}$$

where:

- s = sign
- e = biased exponent
- j is the leading bit, determined by the value of e . In the case of long double (*Intel*), the leading bit is explicit; in all other cases, it is implicit.
- f = fraction
- u means that the bit can be either 0 or 1.

The following tables show the position of the bits.

TABLE A-6 float Representation

Bits	Name
31	Sign
23 - 30	Exponent
0 - 22	Fraction

TABLE A-7 double Representation

Bits	Name
63	Sign
52 - 62	Exponent
0 - 51	Fraction

TABLE A-8 long double Representation (*SPARC*)

Bits	Name
127	Sign
112 - 126	Exponent
0 - 111	Fraction

TABLE A-9 long double Representation (*Intel*)

Bits	Name
80 - 95	Unused
79	Sign
64 - 78	Exponent
63	Leading bit
0 - 62	Fraction

For further information, refer to the *Numerical Computation Guide*.

Exceptional Values

float and double numbers are said to contain a “hidden,” or implied, bit, providing for one more bit of precision than would otherwise be the case. In the case of long double, the leading bit is implicit (*SPARC*) or explicit (*Intel*); this bit is 1 for normal numbers, and 0 for subnormal numbers.

TABLE A-10 float Representations

normal number ($0 < e < 255$):	$(-1)^{\text{Sign}2} (\text{exponent} - 127)1.f$
subnormal number ($e=0, f!=0$):	$(-1)^{\text{Sign}2} (-126)0.f$
zero ($e=0, f=0$):	$(-1)^{\text{Sign}n}0.0$
signaling NaN	$s=u, e=255(\text{max}); f=.0uuu-uu$; at least one bit must be nonzero
quiet NaN	$s=u, e=255(\text{max}); f=.1uuu-uu$
Infinity	$s=u, e=255(\text{max}); f=.0000-00$ (all zeroes)

TABLE A-11 double Representations

normal number ($0 < e < 2047$):	$(-1)^{\text{Sign}2} (\text{exponent} - 1023)1.f$
subnormal number ($e=0, f!=0$):	$(-1)^{\text{Sign}2} (-1022)0.f$
zero ($e=0, f=0$):	$(-1)^{\text{Sign}n}0.0$
signaling NaN	$s=u, e=2047(\text{max}); f=.0uuu-uu$; at least one bit must be nonzero
quiet NaN	$s=u, e=2047(\text{max}); f=.1uuu-uu$
Infinity	$s=u, e=2047(\text{max}); f=.0000-00$ (all zeroes)

TABLE A-12 long double Representations

normal number ($0 < e < 32767$):	$(-1)^{\text{Sign}2} (\text{exponent} - 16383)1.f$
subnormal number ($e=0, f!=0$):	$(-1)^{\text{Sign}2} (-16382)0.f$
zero ($e=0, f=0$):	$(-1)^{\text{Sign}n}0.0$
signaling NaN	$s=u, e=32767(\text{max}); f=.0uuu-uu$; at least one bit must be nonzero
quiet NaN	$s=u, e=32767(\text{max}); f=.1uuu-uu$
Infinity	$s=u, e=32767(\text{max}); f=.0000-00$ (all zeroes)

Hexadecimal Representation of Selected Numbers

The following tables show the hexadecimal representations.

TABLE A-13 Hexadecimal Representation of Selected Numbers (*SPARC*)

Value	float	double	long double
+0	00000000	0000000000000000	00000000000000000000000000000000
-0	80000000	8000000000000000	80000000000000000000000000000000
+1.0	3F800000	3FF0000000000000	3FFF0000000000000000000000000000
-1.0	BF800000	BFF0000000000000	BFFF0000000000000000000000000000
+2.0	40000000	4000000000000000	40000000000000000000000000000000
+3.0	40400000	4008000000000000	40080000000000000000000000000000
+Infinity	7F800000	7FF0000000000000	7FFF0000000000000000000000000000
-Infinity	FF800000	FFF0000000000000	FFFF0000000000000000000000000000
NaN	7FBFFFFF	7FF7FFFFFFFFFFFF	7FFF7FFFFFFFFFFFFFFFFFFFFFFFFFFFFF

TABLE A-14 Hexadecimal Representation of Selected Numbers (*Intel*)

Value	float	double	long double
+0	00000000	0000000000000000	000000000000000000000000
-0	80000000	0000000080000000	800000000000000000000000
+1.0	3F800000	000000003FF00000	3FFF80000000000000000000
-1.0	BF800000	00000000BFF00000	BFFF80000000000000000000
+2.0	40000000	0000000040000000	400080000000000000000000
+3.0	40400000	0000000040080000	4000C0000000000000000000
+Infinity	7F800000	000000007FF00000	7FFF80000000000000000000
-Infinity	FF800000	00000000FFF00000	FFFF80000000000000000000
NaN	7FBFFFFF	FFFFFFFF7FF7FFFF	7FFF7FFFFFFFFFFFFFFFFFFFFFFFFF

For further information, refer to the *Numerical Computation Guide*.

Pointer Representation

A pointer in C occupies four bytes. The NULL value pointer is equal to zero.

Array Storage

Arrays are stored with their elements in a specific storage order. The elements are actually stored in a linear sequence of storage elements.

C arrays are stored in row-major order; the last subscript in a multidimensional array varies the fastest.

String data types are simply arrays of `char` elements. The maximum number of characters allowed in a string literal or wide string literal (after concatenation) is 4,294,967,295. See “Storage Allocation” on page 285 for information on the size limit of storage allocated on the stack.

TABLE A-15 Array Types and Storage

Type	Maximum Number of Elements for SPARC and Intel	Maximum Number of Elements for SPARC V9
<code>char</code>	4,294,967,295	2,305,843,009,213,693,951
<code>short</code>	2,147,483,647	1,152,921,504,606,846,975
<code>int</code>	1,073,741,823	576,460,752,303,423,487
<code>long</code>	1,073,741,823	288,230,376,151,711,743
<code>float</code>	1,073,741,823	576,460,752,303,423,487
<code>double</code>	536,870,911	288,230,376,151,711,743
<code>long double</code>	268,435,451	144,115,188,075,855,871
<code>long long</code> ¹	536,870,911	288,230,376,151,711,743

1. Not valid in `-Xc` mode

Static and global arrays can accommodate many more elements.

Arithmetic Operations on Exceptional Values

This section describes the results derived from applying the basic arithmetic operations to combinations of exceptional and ordinary floating-point values. The information that follows assumes that no traps or any other exception actions are taken.

The following tables explain the abbreviations:

TABLE A-16 Abbreviation Usage

Abbreviation	Meaning
Num	Subnormal or normal number
Inf	Infinity (positive or negative)
NaN	Not a number
Uno	Unordered

The tables that follow describe the types of values that result from arithmetic operations performed with combinations of different types of operands.

TABLE A-17 Addition and Subtraction Results

Left Operand	Right Operand			
	0	Num	Inf	NaN
0	0	Num	Inf	NaN
Num	Num	See Note	Inf	NaN
Inf	Inf	Inf	See Note	NaN
NaN	NaN	NaN	NaN	NaN

Note – Num + Num could be Inf, rather than Num, when the result is too large (overflow). Inf + Inf = NaN when the infinities are of opposite sign.

TABLE A-18 Multiplication Results

Left Operand	Right Operand			
	0	Num	Inf	NaN
0	0	0	NaN	NaN
Num	0	Num	Inf	NaN
Inf	NaN	Inf	Inf	NaN
NaN	NaN	NaN	NaN	NaN

TABLE A-19 Division Results

Left Operand	Right Operand			
	0	Num	Inf	NaN
0	NaN	0	0	NaN
Num	Inf	Num	0	NaN
Inf	Inf	Inf	NaN	NaN
NaN	NaN	NaN	NaN	NaN

TABLE A-20 Comparison Results

Left Operand	Right Operand			
	0	+Num	+Inf	NaN
0	=	<	<	Uno
+Num	>	The result of the comparison	<	Uno
+Inf	>	>	=	Uno
NaN	Uno	Uno	Uno	Uno

Note – NaN compared with NaN is unordered, and results in inequality. +0 compares equal to -0.

Argument-Passing Mechanism

This section describes how arguments are passed in ANSI/ISO C.

All arguments to C functions are passed by value.

Actual arguments are passed in the reverse order from which they are declared in a function declaration.

Actual arguments which are expressions are evaluated before the function reference. The result of the expression is then placed in a register or pushed onto the stack.

(SPARC)

Functions return integer results in register %o0, float results in register %f0, and double results in registers %f0 and %f1.

long long¹ integers are *passed* in registers with the higher word order in %oN, and the lower order word in %o(N+1). In-register results are *returned* in %i0 and %i1, with similar ordering.

All arguments, except doubles and long doubles, are passed as four-byte values. A double is passed as an eight-byte value. The first six four-byte values (double counts as 8) are passed in registers %o0 through %o5. The rest are passed onto the stack. Structures are passed by making a copy of the structure and passing a pointer to the copy. A long double is passed in the same manner as a structure.

Upon return from a function, it is the responsibility of the caller to pop arguments from the stack. Registers described are as seen by the caller.

(Intel)

Functions return integer results in register %eax.

long long results are *returned* in registers %edx and %eax. Functions return float, double, and long double results in register %st(0).

All arguments, except structs, unions, long longs, doubles and long doubles, are passed as four-byte values; a long long is passed as an eight-byte value, a double is passed as an eight-byte value, and a long double is passed as a 12-byte value.

structs and unions are copied onto the stack. The size is rounded up to a multiple of four bytes. Functions returning structs and unions are passed a hidden first argument, pointing to the location into which the returned struct or union is stored.

Upon return from a function, it is the responsibility of the caller to pop arguments from the stack, except for the extra argument for struct and union returns that is popped by the called function.

1. Not available in -Xc mode.

Implementation-Defined ISO/IEC C Behavior

The ISO/IEC 9899:1990, Programming Languages - C standard specifies the form and establishes the interpretation of programs written in C. However, this standard leaves a number of issues as implementation-defined, that is, as varying from compiler to compiler. This chapter details these areas. They can be readily compared to the ISO/IEC 9899:1990 standard itself:

- Each issue uses the same section text as found in the ISO standard.
- Each issue is preceded by its corresponding section number in the ISO standard.

This appendix is organized into the following sections:

- “Translation (G.3.1)” on page 298
- “Environment (G.3.2)” on page 298
- “Identifiers (G.3.3)” on page 299
- “Characters(G.3.4)” on page 299
- “Integers(G.3.5)” on page 301
- “Floating-Point(G.3.6)” on page 303
- “Arrays and Pointers(G.3.7)” on page 304
- “Registers(G.3.8)” on page 305
- “Structures, Unions, Enumerations, and Bit-Fields(G.3.9)” on page 305
- “Qualifiers(G.3.10)” on page 307
- “Declarators(G.3.11)” on page 307
- “Statements(G.3.12)” on page 307
- “Preprocessing Directives(G.3.13)” on page 308
- “Library Functions(G.3.14)” on page 309
- “Locale-Specific Behavior(G.4)” on page 316

Implementation Compared to the ANSI/ ISO Standard

Translation (G.3.1)

The numbers in parentheses correspond to section numbers in the ISO/IEC 9899:1990 standard.

(5.1.1.3) Identification of diagnostics:

Error messages have the following format:

filename, line *line number*: *message*

Warning messages have the following format:

filename, line *line number*: *warning message*

Where:

- *filename* is the name of the file containing the error or warning
- *line number* is the number of the line on which the error or warning is found
- *message* is the diagnostic message

Environment (G.3.2)

(5.1.2.2.1) Semantics of arguments to main:

```
int main (int argc, char *argv[])
{
    ....
}
```

`argc` is the number of command-line arguments with which the program is invoked with. After any shell expansion, `argc` is always equal to at least 1, the name of the program.

`argv` is an array of pointers to the command-line arguments.

(5.1.2.3) What constitutes an interactive device:

An interactive device is one for which the system library call `isatty()` returns a nonzero value.

Identifiers (G.3.3)

(6.1.2) The number of significant initial characters (beyond 31) in an identifier without external linkage:

The first 1,023 characters are significant. Identifiers are case-sensitive.

(6.1.2) The number of significant initial characters (beyond 6) in an identifier with external linkage:

The first 1,023 characters are significant. Identifiers are case-sensitive.

Characters(G.3.4)

(5.2.1)The members of the source and execution character sets, except as explicitly specified in the Standard:

Both sets are identical to the ASCII character sets, plus locale-specific extensions.

(5.2.1.2)The shift states used for the encoding of multibyte characters:

There are no shift states.

(5.2.4.2.1)The number of bits in a character in the execution character set:

There are 8 bits in a character for the ASCII portion; locale-specific multiple of 8 bits for locale-specific extended portion.

(6.1.3.4)The mapping of members of the source character set (in character and string literals) to members of the execution character set:

Mapping is identical between source and execution characters.

(6.1.3.4)The value of an integer character constant that contains a character or escape sequence not represented in the basic execution character set or the extended character set for a wide character constant:

It is the numerical value of the rightmost character. For example, '\q' equals 'q'. A warning is emitted if such an escape sequence occurs.

(3.1.3.4)The value of an integer character constant that contains more than one character or a wide character constant that contains more than one multibyte character:

A multiple-character constant that is not an escape sequence has a value derived from the numeric values of each character.

(6.1.3.4)The current locale used to convert multibyte characters into corresponding wide characters (codes) for a wide character constant:

The valid locale specified by LC_ALL, LC_CTYPE, or LANG environment variable.

(6.2.1.1)Whether a plain char has the same range of values as signed char or unsigned char:

A char is treated as a signed char (SPARC) (Intel) .

Integers(G.3.5)

(6.1.2.5)The representations and sets of values of the various types of integers:

TABLE B-1 Representations and Sets of Values of Integers

Integer	Bits	Minimum	Maximum
char (SPARC) (Intel)	8	-128	127
signed char	8	-128	127
unsigned char	8	0	255
short	16	-32768	32767
signed short	16	-32768	32767
unsigned short	16	0	65535
int	32	-2147483648	2147483647
signed int	32	-2147483648	2147483647
unsigned int	32	0	4294967295
long (SPARC) v8	32	-2147483648	2147483647
long (SPARC) v9	64	-9223372036854775808	9223372036854775807
signed long (SPARC)v8	32	-2147483648	2147483647
signed long (SPARC) v9	64	-9223372036854775808	9223372036854775807
unsigned long (SPARC) v8	32	0	4294967295
unsigned long (SPARC) v9	64	0	18446744073709551615
long long ¹	64	-9223372036854775808	9223372036854775807
signed long long1	64	-9223372036854775808	9223372036854775807
unsigned long long1	64	0	18446744073709551615

1. Not valid in -Xc mode

(6.2.1.2)The result of converting an integer to a shorter signed integer, or the result of converting an unsigned integer to a signed integer of equal length, if the value cannot be represented:

When an integer is converted to a shorter `signed` integer, the low order bits are copied from the longer integer to the shorter `signed` integer. The result may be negative.

When an unsigned integer is converted to a `signed` integer of equal size, the low order bits are copied from the unsigned integer to the `signed` integer. The result may be negative.

(6.3)The results of bitwise operations on signed integers:

The result of a bitwise operation applied to a `signed` type is the bitwise operation of the operands, including the `sign` bit. Thus, each bit in the result is set if—and only if—each of the corresponding bits in both of the operands is set.

(6.3.5)The sign of the remainder on integer division:

The result is the same sign as the dividend; thus, the remainder of $-23/4$ is -3 .

(6.3.7)The result of a right shift of a negative-valued signed integral type:

The result of a right shift is a `signed` right shift.

Floating-Point(G.3.6)

(6.1.2.5)The representations and sets of values of the various types of floating-point numbers:

TABLE B-2 Values for a float

<i>float</i>	
Bits	32
Min	1.17549435E-38
Max	3.40282347E+38
Epsilon	1.19209290E-07

TABLE B-3 Values for a double

<i>double</i>	
Bits	64
Min	2.2250738585072014E-308
Max	1.7976931348623157E+308
Epsilon	2.2204460492503131E-16

TABLE B-4 Values for long double

<i>long double</i>	
Bits	128 (SPARC) 80 (Intel)
Min	3.362103143112093506262677817321752603E-4932 (SPARC) 3.3621031431120935062627E-4932 (Intel)
Max	1.189731495357231765085759326628007016E+4932 (SPARC) 1.1897314953572317650213E4932 (Intel)
Epsilon	1.925929944387235853055977942584927319E-34 (SPARC) 1.0842021724855044340075E-19 (Intel)

(6.2.1.3)The direction of truncation when an integral number is converted to a floating-point number that cannot exactly represent the original value:

Numbers are rounded to the nearest value that can be represented.

(6.2.1.4)The direction of truncation or rounding when a floating- point number is converted to a narrower floating-point number:

Numbers are rounded to the nearest value that can be represented.

Arrays and Pointers(G.3.7)

(6.3.3.4, 7.1.1)The type of integer required to hold the maximum size of an array; that is, the type of the `sizeof` operator, `size_t`:

`unsigned int` as defined in `stddef.h`.

`unsigned long` for `-Xarch=v9`

(6.3.4)The result of casting a pointer to an integer, or vice versa:

The bit pattern does not change for pointers and values of type `int`, `long`, `unsigned int` and `unsigned long`.

(6.3.6, 7.1.1)The type of integer required to hold the difference between two pointers to members of the same array, `ptrdiff_t`:

`int` as defined in `stddef.h`.

`long` for `-Xarch=v9`

Registers(G.3.8)

(6.5.1)The extent to which objects can actually be placed in registers by use of the `register` storage-class specifier:

The number of effective register declarations depends on patterns of use and definition within each function and is bounded by the number of registers available for allocation. Neither the compiler nor the optimizer is required to honor register declarations.

Structures, Unions, Enumerations, and Bit-Fields(G.3.9)

(6.3.2.3)A member of a union object is accessed using a member of a different type:

The bit pattern stored in the union member is accessed, and the value interpreted, according to the type of the member by which it is accessed.

(6.5.2.1)The padding and alignment of members of structures.

TABLE B-5 Padding and Alignment of Structure Members

Type	Alignment Boundary	Byte Alignment
char	Byte	1
short	Halfword	2
int	Word	4
long (SPARC) v8	Word	4
long (SPARC) v9	Doubleword	8
float (SPARC)	Word	4
double (SPARC)	Doubleword (SPARC) Word (Intel)	8 (SPARC) 4 (Intel)
long double (SPARC) v8	Doubleword (SPARC) Word (Intel)	8 (SPARC) 4 (Intel)
long double (SPARC) v9	Quadword	16
pointer (SPARC) v8	Word	4
pointer (SPARC) v9	Quadword	8
long long ¹	Doubleword (SPARC) Word (Intel)	8 (SPARC) 4 (Intel)

1. Not available in -Xc mode.

Structure members are padded internally, so that every element is aligned on the appropriate boundary.

Alignment of structures is the same as its more strictly aligned member. For example, a `struct` with only `chars` has no alignment restrictions, whereas a `struct` containing a `double` would be aligned on an 8-byte boundary.

(6.5.2.1)Whether a plain `int` bit-field is treated as a signed `int` bit-field or as an unsigned `int` bit-field:

It is treated as an unsigned `int`.

(6.5.2.1)The order of allocation of bit-fields within an `int`:

Bit-fields are allocated within a storage unit from high-order to low-order.

(6.5.2.1) Whether a bit-field can straddle a storage-unit boundary:

Bit-fields do not straddle storage-unit boundaries.

(6.5.2.2) The integer type chosen to represent the values of an enumeration type:

This is an `int`.

Qualifiers(G.3.10)

(6.5.5.3) What constitutes an access to an object that has volatile-qualified type:

Each reference to the name of an object constitutes one access to the object.

Declarators(G.3.11)

(6.5.4) The maximum number of declarators that may modify an arithmetic, structure, or union type:

No limit is imposed by the compiler.

Statements(G.3.12)

(6.6.4.2) The maximum number of case values in a switch statement:

No limit is imposed by the compiler.

Preprocessing Directives(G.3.13)

(6.8.1)Whether the value of a single-character character constant in a constant expression that controls conditional inclusion matches the value of the same character constant in the execution character set:

A character constant within a preprocessing directive has the same numeric value as it has within any other expression.

(6.8.1)Whether such a character constant may have a negative value:

Character constants in this context may have negative values (*SPARC*) (*Intel*) .

(6.8.2)The method for locating includable source files:

A file whose name is delimited by `< >` is searched for first in the directories named by the `-I` option, and then in the standard directory. The standard directory is `/usr/include`, unless the `-YI` option is used to specify a different default location.

A file whose name is delimited by quotes is searched for first in the directory of the source file that contains the `#include`, then in directories named by the `-I` option, and last in the standard directory.

If a file name enclosed in `< >` or double quotes begins with a `/` character, the file name is interpreted as a path name beginning in the root directory. The search for this file begins in the root directory only.

(6.8.2)The support of quoted names for includable source files:

Quoted file names in `include` directives are supported.

(6.8.2)The mapping of source file character sequences:

Source file characters are mapped to their corresponding ASCII values.

(6.8.6)The behavior on each recognized #pragma directive:

The following pragmas are supported. See “Pragmas” on page 94 for more information.

- `align integer (variable[, variable])`
- `does_not_read_global_data (funcname [, funcname])`
- `does_not_return (funcname[, funcname])`
- `does_not_write_global_data (funcname[, funcname])`
- `error_messages (on|off|default, tag1[tag2... tagn])`
- `fini (f1[, f2..., fn])`
- `ident string`
- `init (f1[, f2..., fn])`
- `inline (funcname[, funcname])`
- `int_to_unsigned (funcname)`
- `MP serial_loop`
- `MP serial_loop_nested`
- `MP taskloop`
- `no_inline (funcname[, funcname])`
- `nomemorydepend`
- `no_side_effect (funcname[, funcname])`
- `opt_level (funcname[, funcname])`
- `pack(n)`
- `pipelooop(n)`
- `rarely_called (funcname[, funcname])`
- `redefine_extname old_extname new_extname`
- `returns_new_memory (funcname[, funcname])`
- `unknown_control_flow (name[, name])`
- `unroll (unroll_factor)`
- `weak (symbol1 [= symbol2])`

(6.8.8)The definitions for `__DATE__` and `__TIME__` when, respectively, the date and time of translation are not available:

These macros are always available from the environment.

Library Functions(G.3.14)

(7.1.6)The null pointer constant to which the macro `NULL` expands:

`NULL` equals 0.

(7.2) The diagnostic printed by and the termination behavior of the `assert` function:

The diagnostic is:

Assertion failed: *statement*. file *filename*, line *number*

Where:

- *statement* is the statement which failed the assertion
- *filename* is the name of the file containing the failure
- *line number* is the number of the line on which the failure occurs

(7.3.1) The sets of characters tested for by the `isalnum`, `isalpha`, `iscntrl`, `islower`, `isprint`, and `isupper` functions:

TABLE B-6 Character Sets Tested by `isalpha`, `islower`, Etc.

<code>isalnum</code>	ASCII characters A-Z, a-z and 0-9
<code>isalpha</code>	ASCII characters A-Z and a-z, plus locale-specific single-byte letters
<code>iscntrl</code>	ASCII characters with value 0-31 and 127
<code>islower</code>	ASCII characters a-z
<code>isprint</code>	Locale-specific single-byte printable characters
<code>isupper</code>	ASCII characters A-Z

(7.5.1) The values returned by the mathematics functions on domain errors:

TABLE B-7 Values Returned on Domain Errors

Error	Math Functions	Compiler Modes	
		-Xs, -Xt	-Xa, -Xc
DOMAIN	$\text{acos}(x >1)$	0.0	0.0
DOMAIN	$\text{asin}(x >1)$	0.0	0.0
DOMAIN	$\text{atan2}(+0,+0)$	0.0	0.0
DOMAIN	$y0(0)$	-HUGE	-HUGE_VAL
DOMAIN	$y0(x<0)$	-HUGE	-HUGE_VAL
DOMAIN	$y1(0)$	-HUGE	-HUGE_VAL
DOMAIN	$y1(x<0)$	-HUGE	-HUGE_VAL
DOMAIN	$yn(n,0)$	-HUGE	-HUGE_VAL
DOMAIN	$yn(n,x<0)$	-HUGE	-HUGE_VAL
DOMAIN	$\log(x<0)$	-HUGE	-HUGE_VAL
DOMAIN	$\log_{10}(x<0)$	-HUGE	-HUGE_VAL
DOMAIN	$\text{pow}(0,0)$	0.0	1.0
DOMAIN	$\text{pow}(0,\text{neg})$	0.0	-HUGE_VAL
DOMAIN	$\text{pow}(\text{neg},\text{non-integral})$	0.0	NaN
DOMAIN	$\text{sqrt}(x<0)$	0.0	NaN
DOMAIN	$\text{fmod}(x,0)$	x	NaN
DOMAIN	$\text{remainder}(x,0)$	NaN	NaN
DOMAIN	$\text{acosh}(x<1)$	NaN	NaN
DOMAIN	$\text{atanh}(x >1)$	NaN	NaN

(7.5.1) Whether the mathematics functions set the integer expression `errno` to the value of the macro `ERANGE` on underflow range errors:

Mathematics functions, except `scalbn`, set `errno` to `ERANGE` when underflow is detected.

(7.5.6.4) Whether a domain error occurs or zero is returned when the `fmod` function has a second argument of zero:

In this case, it returns the first argument with domain error.

(7.7.1.1) The set of signals for the `signal` function:

The following table shows the semantics for each signal as recognized by the `signal` function:

TABLE B-8 Semantics for `signal` Signals

Signal	No.	Default	Event
SIGHUP	1	Exit	hangup
SIGINT	2	Exit	interrupt
SIGQUIT	3	Core	quit
SIGILL	4	Core	illegal instruction (not reset when caught)
SIGTRAP	5	Core	trace trap (not reset when caught)
SIGIOT	6	Core	IOT instruction
SIGABRT	6	Core	Used by abort
SIGEMT	7	Core	EMT instruction
SIGFPE	8	Core	floating point exception
SIGKILL	9	Exit	kill (cannot be caught or ignored)
SIGBUS	10	Core	bus error
SIGSEGV	11	Core	segmentation violation
SIGSYS	12	Core	bad argument to system call
SIGPIPE	13	Exit	write on a pipe with no one to read it
SIGALRM	14	Exit	alarm clock
SIGTERM	15	Exit	software termination signal from kill
SIGUSR1	16	Exit	user defined signal 1
SIGUSR2	17	Exit	user defined signal 2
SIGCLD	18	Ignore	child status change
SIGCHLD	18	Ignore	child status change alias
SIGPWR	19	Ignore	power-fail restart
SIGWINCH	20	Ignore	window size change

TABLE B-8 Semantics for `signal` Signals (Continued)

Signal	No.	Default	Event
SIGURG	21	Ignore	urgent socket condition
SIGPOLL	22	Exit	pollable event occurred
SIGIO	22	Exit	socket I/O possible
SIGSTOP	23	Stop	stop (cannot be caught or ignored)
SIGTSTP	24	Stop	user stop requested from tty
SIGCONT	25	Ignore	stopped process has been continued
SIGTTIN	26	Stop	background tty read attempted
SIGTTOU	27	Stop	background tty write attempted
SIGVTALRM	28	Exit	virtual timer expired
SIGPROF	29	Exit	profiling timer expired
SIGXCPU	30	Core	exceeded cpu limit
SIGXFSZ	31	Core	exceeded file size limit
SIGWAITINGT	32	Ignore	process's lwps are blocked

(7.7.1.1) The default handling and the handling at program startup for each signal recognized by the signal function:

See above.

(7.7.1.1) If the equivalent of `signal(sig, SIG_DFL);` is not executed prior to the call of a signal handler, the blocking of the signal that is performed:

The equivalent of `signal(sig, SIG_DFL)` is always executed.

(7.7.1.1) Whether the default handling is reset if the SIGILL signal is received by a handler specified to the signal function:

Default handling is not reset in SIGILL.

(7.9.2) Whether the last line of a text stream requires a terminating new-line character:

The last line does not need to end in a newline.

(7.9.2) Whether space characters that are written out to a text stream immediately before a new-line character appear when read in:

All characters appear when the stream is read.

(7.9.2) The number of null characters that may be appended to data written to a binary stream:

No null characters are appended to a binary stream.

(7.9.3) Whether the file position indicator of an append mode stream is initially positioned at the beginning or end of the file:

The file position indicator is initially positioned at the end of the file.

(7.9.3) Whether a write on a text stream causes the associated file to be truncated beyond that point:

A write on a text stream does not cause a file to be truncated beyond that point unless a hardware device forces it to happen.

(7.9.3) The characteristics of file buffering:

Output streams, with the exception of the standard error stream (`stderr`), are by default-buffered if the output refers to a file, and line-buffered if the output refers to a terminal. The standard error output stream (`stderr`) is by default unbuffered.

A buffered output stream saves many characters, and then writes the characters as a block. An unbuffered output stream queues information for immediate writing on the destination file or terminal immediately. Line-buffered output queues each line of output until the line is complete (a newline character is requested).

(7.9.3) Whether a zero-length file actually exists:

A zero-length file does exist since it has a directory entry.

(7.9.3) The rules for composing valid file names:

A valid file name can be from 1 to 1,023 characters in length and can use all character except the characters `null` and `/` (slash).

(7.9.3) Whether the same file can be open multiple times:

The same file can be opened multiple times.

(7.9.4.1) The effect of the `remove` function on an open file:

The file is deleted on the last call which closes the file. A program cannot open a file which has already been removed.

(7.9.4.2) The effect if a file with the new name exists prior to a call to the `rename` function:

If the file exists, it is removed and the new file is written over the previously existing file.

(7.9.6.1) The output for `%p` conversion in the `fprintf` function:

The output for `%p` is equivalent to `%x`.

(7.9.6.2) The input for `%p` conversion in the `fscanf` function:

The input for `%p` is equivalent to `%x`.

(7.9.6.2) The interpretation of a - character that is neither the first nor the last character in the scan list for %[conversion in the fscanf function:

The - character indicates an inclusive range; thus, [0-9] is equivalent to [0123456789].

Locale-Specific Behavior(G.4)

(7.12.1) The local time zone and Daylight Savings Time:

The local time zone is set by the environment variable TZ.

(7.12.2.1) The era for the clock function

The era for the clock is represented as clock ticks with the origin at the beginning of the execution of the program.

The following characteristics of a hosted environment are locale-specific:

(5.2.1) The content of the execution character set, in addition to the required members:

Locale-specific (no extension in C locale).

(5.2.2) The direction of printing:

Printing is always left to right.

(7.1.1) The decimal-point character:

Locale-specific ("," in C locale).

(7.3) The implementation-defined aspects of character testing and case mapping functions:

Same as 4.3.1.

(7.11.4.4) The collation sequence of the execution character set:

Locale-specific (ASCII collation in C locale).

(7.12.3.5) The formats for time and date:

Locale-specific. Formats for the C locale are shown in the tables below.

The names of the months are:

TABLE B-9 Names of Months

January	May	September
February	June	October
March	July	November
April	August	December

The names of the days of the week are:

TABLE B-10 Days and Abbreviated Days of the Week

Days		Abbreviated Days	
Sunday	Thursday	Sun	Thu
Monday	Friday	Mon	Fri
Tuesday	Saturday	Tue	Sat
Wednesday		Wed	

The format for time is:

`%H:%M:%S`

The format for date is:

`%m/%d/%y`

The formats for AM and PM designation are: AM PM

Supported Features of C99

This appendix details the supported features of the ISO/IEC 9899:1999, Programming Language - C standard. The following features are supported:

- C99 keywords: ISO/IEC 9899:1999, subclause 6.4.1 Keywords

The C99 standard introduces the following new keywords:

- `inline`
- `restrict`
- `_Imaginary`
- `_Complex`
- `_Bool`

The compiler issues a warning if you use these keywords as identifiers while compiling with `-xc99=%none`. Without `-xc99=%none` the compiler issues a warning or error messages for use of these keywords as identifiers depending on the context.

- `__func__` support: ISO/IEC 9899:1999, 6.4.2.2 Predefined identifiers
- Macros with a variable number of arguments: ISO/IEC 9899:1999, subclause 6.10.3 Macro replacement
- Variable length arrays (VLA): ISO/IEC 9899:1999, 6.7.5.2 Array declarators

VLAs are allocated on the stack as if by calling the `alloca` function. Their lifetime, regardless of their scope, is the same as any data allocated on the stack by calling `alloca`; until the function returns. The space allocated is freed when the stack is released upon returning from the function in which the VLA is allocated.

Not all constraints are yet enforced for variable length arrays. Constraint violations lead to undefined results.

- `inline` specifier for static functions: ISO/IEC 9899:1999, 6.7.4 Function specifiers
The C99 function-specifier `inline` has been added. `inline` is fully functional for functions with internal linkage. For functions defined with external linkage use of the `inline` function-specifier creates an inline definition only, no external definition of the function is created. Thus pointers to inline functions with external linkage are unique to each translation unit and will not compare equal.
- `//` introduced comments: ISO/IEC 9899:1999, 6.4.9 Comments

Performance Tuning (*SPARC*)

This appendix describes performance tuning on SPARC platforms and is organized into the following sections:

- “Limits” on page 321
- “libfast.a Library” on page 322

Limits

Some parts of the C library cannot be optimized for speed, even though doing so would benefit most applications. Some examples:

- Integer arithmetic routines—Current SPARC V8 processors support integer multiplication and division instructions. However, if standard C library routines were to use these instructions, programs running on V7 SPARC processors would either run slowly due to kernel emulation overhead, or might break altogether. Hence, integer multiplication and division instructions cannot be used in the standard C library routines.
- Doubleword memory access—Block copy and move routines, such as `memmove()` and `bcopy()`, could run considerably faster if they used SPARC doubleword load and store instructions (`ldd` and `std`). Some memory-mapped devices, such as frame buffers, do not support 64-bit access; nevertheless, these devices are expected to work correctly with `memmove()` and `bcopy()`. Hence, `ldd` and `std` cannot be used in the standard C library routines.
- Memory allocation algorithms—The C library routines `malloc()` and `free()` are typically implemented as a compromise between speed, space, and insensitivity to coding errors in old UNIX programs. Memory allocators based on “buddy system” algorithms typically run faster than the standard library version, but tend to use more space.

libfast.a Library

The library `libfast.a` provides speed-tuned versions of standard C library functions. Because it is an optional library, it can use algorithms and data representations that may not be appropriate for the standard C library, even though they improve the performance of most applications.

Use profiling to determine whether the routines in the following checklist are important to the performance of your application, then use this checklist to decide whether `libfast.a` benefits the performance:

- *Do* use `libfast.a` if performance of integer multiplication or division is important, even if a single binary version of the application must run on both V7 and V8 SPARC platforms. The important routines are: `.mul`, `.div`, `.rem`, `.umul`, `.udiv`, and `.urem`.
- *Do* use `libfast.a` if performance of memory allocation is important, and the size of the most commonly allocated blocks is close to a power of two. The important routines are: `malloc()`, `free()`, `realloc()`.
- *Do* use `libfast.a` if performance of block move or fill routines is important. The important routines are: `bcopy()`, `bzero()`, `memcpy()`, `memmove()`, and `memset()`.
- *Do not* use `libfast.a` if the application requires user mode, memory-mapped access to an I/O device that does not support 64-bit memory operations.
- *Do not* use `libfast.a` if the application is multithreaded.

When linking the application, add the option `-lfast` to the `cc` command used at link time. The `cc` command links the routines in `libfast.a` ahead of their counterparts in the standard C library.

The Differences Between K&R Sun C and Sun ANSI/ISO C

This appendix describes the differences between the previous K&R Sun C and Sun ANSI/ISO C. This appendix is organized into the following sections:

- “K&R Sun C Incompatibilities with Sun ANSI/ISO C” on page 324
- “The Difference Between Sun C and ANSI/ISO C as Set by -xs” on page 332
- “Keywords” on page 334

For more information see “Standards Conformance” on page 9.

K&R Sun C Incompatibilities with Sun ANSI/ISO C

TABLE E-1 K&R Sun C Incompatibilities With Sun ANSI/ISO C

Topic	Sun C (K&R)	Sun ANSI/ISO C
envp argument to main()	Allows envp as third argument to main().	Allows this third argument; however, this usage is not strictly conforming to the ANSI/ISO C standard.
Keywords	Treats the identifiers <code>const</code> , <code>volatile</code> , and <code>signed</code> as ordinary identifiers.	<code>const</code> , <code>volatile</code> , and <code>signed</code> are keywords.
extern and static functions declarations inside a block	Promotes these function declarations to file scope.	The ANSI/ISO standard does not guarantee that block scope function declarations are promoted to file scope.
Identifiers	Allows dollar signs (\$) in identifiers.	\$ not allowed.
long float types	Accepts long float declarations and treats these as double(s).	Does not accept these declarations.
Multi-character character-constants	<code>int mc = 'abcd';</code> yields: <code>abcd</code>	<code>int mc = 'abcd';</code> yields: <code>dcba</code>
Integer constants	Accepts 8 or 9 in octal escape sequences.	Does not accept 8 or 9 in octal escape sequences.
Assignment operators	Treats the following operator pairs as two tokens, and as a consequence, permits white space between them: <code>*=, /=, %=, +=, -=, <<=,</code> <code>>>=, &=, ^=, =</code>	Treats them as single tokens, and therefore disallows white space in between.
Unsigned preserving semantics for expressions	Supports unsigned preserving, that is, unsigned <code>char/shorts</code> are converted into unsigned <code>int(s)</code> .	Supports value-preserving, that is, unsigned <code>char/short(s)</code> are converted into <code>int(s)</code> .

TABLE E-1 K&R Sun C Incompatibilities With Sun ANSI/ISO C (*Continued*)

Topic	Sun C (K&R)	Sun ANSI/ISO C
Single/double precision calculations	Promotes the operands of floating point expressions to double. Functions which are declared to return floats always promote their return values to doubles.	Allows operations on floats to be performed in single precision calculations. Allows float return types for these functions.
Name spaces of struct/union members	Allows struct, union, and arithmetic types using member selection operators ('.', '->') to work on members of other struct(s) or unions.	Requires that every unique struct/union have its own unique name space.
A cast as an lvalue	Supports casts as lvalue(s). For example: <pre>(char *)ip = &char;</pre>	Does not support this feature.
Implied int declarations	Supports declarations without an explicit type specifier. A declaration such as <code>num;</code> is treated as implied <code>int</code> . For example: <pre>num; /*num implied as an int*/ int num2; /* num2 explicitly*/ /* declared an int */</pre>	The <code>num;</code> declaration (without the explicit type specifier <code>int</code>) is not supported, and generates a syntax error.
Empty declarations	Allows empty declarations, such as: <pre>int;</pre>	Except for tags, disallows empty declarations.
Type specifiers on type definitions	Allows type specifiers such as <code>unsigned</code> , <code>short</code> , <code>long</code> on typedefs declarations. For example: <pre>typedef short small; unsigned small x;</pre>	Does not allow type specifiers to modify typedef declarations.
Types allowed on bit fields	Allows bit fields of all integral types, including unnamed bit fields. The ABI requires support of unnamed bit fields and the other integral types.	Supports bitfields only of the type <code>int</code> , <code>unsigned int</code> and <code>signed int</code> . Other types are undefined.

TABLE E-1 K&R Sun C Incompatibilities With Sun ANSI/ISO C (*Continued*)

Topic	Sun C (K&R)	Sun ANSI/ISO C
Treatment of tags in incomplete declarations	<p> Ignores the incomplete type declaration. In the following example, <code>f1</code> refers to the outer struct:</p> <pre> struct x { . . . } s1; {struct x; struct y {struct x f1; } s2; struct x { . . . };}</pre>	<p>In an ANSI/ISO-conforming implementation, an incomplete struct or union type specifier hides an enclosing declaration with the same tag.</p>
Mismatch on struct/union/enum declarations	<p> Allows a mismatch on the struct/enum/union type of a tag in nested struct/union declarations. In the following example, the second declaration is treated as a struct:</p> <pre> struct x { . . . }s1; {union x s2; . . .}</pre>	<p> Treats the inner declaration as a new declaration, hiding the outer tag.</p>
Labels in expressions	<p> Treats labels as <code>(void *)</code> lvalues.</p>	<p> Does not allow labels in expressions.</p>
switch condition type	<p> Allows <code>float(s)</code> and <code>double(s)</code> by converting them to <code>int(s)</code>.</p>	<p> Evaluates only integral types (<code>int</code>, <code>char</code>, and enumerated) for the switch condition type.</p>
Syntax of conditional inclusion directives	<p> The preprocessor ignores trailing tokens after an <code>#else</code> or <code>#endif</code> directive.</p>	<p> Disallows such constructs.</p>
Token-pasting and the ## preprocessor operator	<p> Does not recognize the <code>##</code> operator. Token-pasting is accomplished by placing a comment between the two tokens being pasted:</p> <pre> #define PASTE(A,B) A/*any comment*/B</pre>	<p> Defines <code>##</code> as the preprocessor operator that performs token-pasting, as shown in this example:</p> <pre> #define PASTE(A,B) A##B</pre> <p> Furthermore, the Sun ANSI/ISO C preprocessor doesn't recognize the Sun C method. Instead, it treats the comment between the two tokens as white space.</p>

TABLE E-1 K&R Sun C Incompatibilities With Sun ANSI/ISO C (Continued)

Topic	Sun C (K&R)	Sun ANSI/ISO C
Preprocessor rescanning	The preprocessor recursively substitutes: <pre>#define F(X) X(arg) F(F)</pre> <p>yields</p> <pre>arg(arg)</pre>	A macro is not replaced if it is found in the replacement list during the rescan: <pre>#define F(X)X(arg) F(F)</pre> <p>yields:</p> <pre>F(arg)</pre>
typedef names in formal parameter lists	You can use typedef names as formal parameter names in a function declaration. "Hides" the typedef declaration.	Disallows the use of an identifier declared as a typedef name as a formal parameter.
Implementation-specific initializations of aggregates	Uses a bottom-up algorithm when parsing and processing partially elided initializers within braces: <pre>struct{ int a[3]; int b; }\ w[]={ {1}, 2};</pre> <p>yields</p> <pre>sizeof(w)=16 w[0].a=1,0,0 w[0].b=2</pre>	Uses a top-down parsing algorithm. For example: <pre>struct{int a[3];int b;}\ w[]={ {1}, 2};</pre> <p>yields</p> <pre>sizeof(w)=32 w[0].a=1,0,0 w[0].b=0 w[1].a=2,0,0 w[1].b=0</pre>
Comments spanning include files	Allows comments which start in an #include file to be terminated by the file that includes the first file.	Comments are replaced by a white-space character in the translation phase of the compilation, which occurs before the #include directive is processed.

TABLE E-1 K&R Sun C Incompatibilities With Sun ANSI/ISO C (Continued)

Topic	Sun C (K&R)	Sun ANSI/ISO C
Formal parameter substitution within a character constant	<p>Substitutes characters within a character constant when it matches the replacement list macro:</p> <pre>#define charize(c) 'c' charize(Z)</pre> <p>yields:</p> <pre>'z'</pre>	<p>The character is not replaced:</p> <pre>#define charize(c) 'c' charize(Z)</pre> <p>yields:</p> <pre>'c'</pre>
Formal parameter substitution within a string constant	<p>The preprocessor substitutes a formal parameter when enclosed within a string constant:</p> <pre>#define stringize(str) 'str' stringize(foo)</pre> <p>yields:</p> <pre>"foo"</pre>	<p>The # preprocessor operator should be used:</p> <pre>#define stringize(str) 'str' stringize(foo)</pre> <p>yields:</p> <pre>"str"</pre>
Preprocessor built into the compiler "front-end"	<p>Compiler calls <code>cpp(1)</code>.</p> <p>Components used in the compiling are:</p> <pre>cpp ccom ipropt cg inline as ld</pre> <p>Note: <code>ipropt</code> and <code>cg</code> are invoked only with the following options:</p> <pre>-O -x02 -x03 -x04 -xa -fast</pre> <p><code>inline</code> is invoked only if an inline template file (<code>file.i1</code>) is provided.</p>	<p>Preprocessor (<code>cpp</code>) is built directly into <code>acompp</code>, so <code>cpp</code> is not directly involved, except in <code>-xs</code> mode.</p> <p>Components used in the compiling are:</p> <pre>cpp (-xs mode only) acompp ipropt cg ld</pre> <p>Note: <code>ipropt</code> and <code>cg</code> are invoked only with the following options:</p> <pre>-O -x02 -x03 -x04 -xa -fast</pre>

TABLE E-1 K&R Sun C Incompatibilities With Sun ANSI/ISO C (*Continued*)

Topic	Sun C (K&R)	Sun ANSI/ISO C
Line concatenation with backslash	Does not recognize the backslash character in this context.	Requires that a newline character immediately preceded by a backslash character be spliced together.
Trigraphs in string literals	Does not support this ANSI/ISO C feature.	
asm keyword	asm is a keyword.	asm is treated as an ordinary identifier.
Linkage of identifiers	Does not treat uninitialized static declarations as tentative declarations. As a consequence, the second declaration will generate a 'redeclaration' error, as in: <pre>static int i = 1; static int i;</pre>	Treats uninitialized static declarations as tentative declarations.
Name spaces	Distinguishes only three : struct/ union/enum tags, members of struct/union/enum, and everything else.	Recognizes four distinct name spaces: label names, tags (the names that follow the keywords struct, union or enum), members of struct/union/enum, and ordinary identifiers.
long double type	Not supported.	Allows long double type declaration.
Floating point constants	The floating point suffixes f, l, F, and L, are not supported.	
Unsuffixed integer constants can have different types	The integer constant suffixes u and U are not supported.	
Wide character constants	Does not accept the ANSI/ISO C syntax for wide character constants, as in: <pre>wchar_t wc = L'x';</pre>	Supports this syntax.
'\a' and '\x'	Treats them as the characters 'a' and 'x'.	Treats '\a' and '\x' as special escape sequences.
Concatenation of string literals	Does not support the ANSI/ISO C concatenation of adjacent string literals.	

TABLE E-1 K&R Sun C Incompatibilities With Sun ANSI/ISO C (*Continued*)

Topic	Sun C (K&R)	Sun ANSI/ISO C
Wide character string literal syntax	Does not support the ANSI/ISO C wide character, string literal syntax shown in this example: <pre>wchar_t *ws = L"hello";</pre>	Supports this syntax.
Pointers: void * versus char *	Supports the ANSI/ISO C void * feature.	
Unary plus operator	Does not support this ANSI/ISO C feature.	
Function prototypes—ellipses	Not supported.	ANSI/ISO C defines the use of ellipses "..." to denote a variable argument parameter list.
Type definitions	Disallows typedefs to be redeclared in an inner block by another declaration with the same type name.	Allows typedefs to be redeclared in an inner block by another declaration with the same type name.
Initialization of extern variables	Does not support the initialization of variables explicitly declared as extern.	Treats the initialization of variables explicitly declared as extern, as definitions.
Initialization of aggregates	Does not support the ANSI/ISO C initialization of unions or automatic structures.	
Prototypes	Does not support this ANSI/ISO C feature.	
Syntax of preprocessing directive	Recognizes only those directives with a # in the first column.	ANSI/ISO C allows leading white-space characters before a # directive.
The # preprocessor operator	Does not support the ANSI/ISO C # preprocessor operator.	

TABLE E-1 K&R Sun C Incompatibilities With Sun ANSI/ISO C (*Continued*)

Topic	Sun C (K&R)	Sun ANSI/ISO C
#error directive	Does not support this ANSI/ISO C feature.	
Preprocessor directives	Supports two pragmas, <code>unknown_control_flow</code> and <code>makes_regs_inconsistent</code> along with the <code>#ident</code> directive. The preprocessor issues warnings when it finds unrecognized pragmas.	Does not specify its behavior for unrecognized pragmas.
Predefined macro names	These ANSI/ISO C-defined macro names are not defined: <code>__STDC__</code> <code>__DATE__</code> <code>__TIME__</code> <code>__LINE__</code>	

The Difference Between Sun C and ANSI/ISO C as Set by `-Xs`

This section describes the differences in compiler behavior when using the `-Xs` option. The `-Xs` option tries to emulate Sun C 1.0, and Sun C 1.1 (K&R style), but in some cases it cannot emulate the previous behavior.

TABLE E-2 `-Xs` Behavior

Data Type	Sun C (K&R)	Sun ANSI/ISO C
Aggregate initialization: struct { int a[3]; int b; } w[] = {{1},2};	<code>sizeof (w) = 16</code> <code>w[0].a = 1, 0, 0</code> <code>w[0].b = 2</code>	<code>sizeof(w) = 32</code> <code>w[0].a = 1, 0, 0</code> <code>w[0].b = 0</code> <code>w[1].a = 2, 0, 0</code> <code>w[1].b = 0</code>
Incomplete struct, union, enum declaration	<code>struct fq {</code> <code>int i;</code> <code>struct unknown;</code> <code>};</code>	Does not allow incomplete struct, union, and enum declaration.
Switch expression integral type	Allows non-integral type.	Does not allow non-integral type.
Order of precedence	Allows: <code>if (rcount > count += index)</code>	Does not allow: <code>if (rcount > count += index)</code>
unsigned, short, and long typedef declarations	Allows: <code>typedef short small</code> <code>unsigned small;</code>	Does not allow (all modes).

TABLE E-2 -Xs Behavior (Continued)

Data Type	Sun C (K&R)	Sun ANSI/ISO C
struct or union tag mismatch in nested struct or union declarations	Allows tag mismatch: <pre> struct x { int i; } s1; /* K&R treats as a struct */ { union x s2; } </pre>	Does not allow tag mismatch in nested struct or union declaration.
Incomplete struct or union type	Ignores an incomplete type declaration.	<pre> struct x { int i; } s1; main() { struct x; struct y { struct x fl /* in K&R, fl refers */ /* to outer struct */ } s2; struct x { int i; }; } </pre>
Casts as lvalues	Allows: <pre> (char *) ip = &foo; </pre>	Does not allow casts as lvalues (all modes).

Keywords

The following tables list the keywords for the ANSI/ISO C Standard, the Sun ANSI/ISO C compiler, and the Sun C compiler.

The first table lists the keywords defined by the ANSI/ISO C standard.

TABLE E-3 ANSI/ISO C Standard Keywords

auto	break	case	char
const	continue	default	do
double	else	enum	extern
float	for	goto	if
int	long	register	return
short	signed	sizeof	static
struct	switch	typedef	union
unsigned	void	volatile	while

Sun ANSI/ISO defines one additional keyword, `asm`. However, `asm` is not supported in `-Xc` mode.

Keywords in Sun C are listed below.

TABLE E-4 Sun C (K&R) Keywords

asm	auto	break	case
char	continue	default	do
double	else	enum	extern
float	for	fortran	goto
if	int	long	register
return	short	sizeof	static
struct	switch	typedef	union
unsigned	void	while	

Implementation-Specific Information of OpenMP

This appendix details the implementation specific details of OpenMP C and C++ Application Program Interface Version 1.0 - October 1998 (available from: <http://www.openmp.org>)

- In the absence of an explicitly defined `OMP_SCHEDULE` environment variable, this implementation uses static scheduling.
- In the absence of an explicitly defined schedule clause, the default is static scheduling.
- If you do not explicitly specify the number of threads in a team through either the `omp_set_num_threads` function or the `OMP_NUM_THREADS` environment variable, the default is 1.
- If you do not explicitly specify whether dynamic adjustment of threads is enabled through either the `omp_set_dynamic` function or the `OMP_DYNAMIC` environment variable, the default is enabled dynamic adjustment.
- Nested parallelism is not supported and is disabled by default.
- The default value of `OMP_NUM_THREADS` is 1.

Index

SYMBOLS

`#assert`, 21, 93
`#define`, 22
`#include`
 adding header files with, 89
`#pragma`, 94 to 103, 190 to 193
`.profile` filename extension, 69
`//` comment indicators
 in C99, 319
 with `-xCC`, 48
`_ _asm` keyword, 86
`__'uname -s'__'uname -r'`, 22, 104, 177
`__BUILT_IN_VA_ARG_INCR`, 22, 104, 177
`__DATE__`, 309
`__func__`, 319
`__i386`, 22, 104, 177
`__lint` predefined token, 177
`__MATHERR_ERRNO_DONTCARE`, 26
`__PRAGMA_REDEFINE_EXTNAME`, 104
`__RESTRICT`, 22, 104, 177
`__sparc`, 22, 104, 177
`__sparcv9`, 22, 172, 177
`__sun`, 22, 104, 177
`__SUNPRO_C`, 22, 104, 177
`__SVR4`, 22, 104, 177
`__TIME__`, 309
`__unix`, 22, 104, 177
`_OPENMP`, 66
`_REENTRANT-lthread`, 34
`_Restrict`, 86

A

accessing
 compilers and tools, 4
 documentation, 6
 man pages, 4
`acomp` (C compiler), 10
alias disambiguation, 189 to 205
alignment of structures, 306
ANSI/ISO 9899-1990 standard, 9
ANSI/ISO C vs. K&R C, 13, 38 to 39, 332
any level alias disambiguation, 41
arithmetic conversions, 85, 87
`_ _asm` keyword, 86
assembler, 11
assembly in source, 86
`#assert`, 21, 93
Auto-Read, 72

B

basic level alias disambiguation, 41
basic mode of lint, 157
behavior, implementation-defined, 297 to 317
binding
 static vs. dynamic, 21
bit-field
 as impacted by transition to ANSI/ISO C, 246
 portability of constants assigned to, 183
 promotion of, 217
 treating as signed or unsigned, 306

- bits, in execution character set, 299
- bitwise
 - operations on signed integers, 302
- buffering, 314

C

C compiler

- accessing, 4
- changing default dirs searched for libraries, 14
- compilation modes and dependencies, 103
- compiling a program, 13 to 14
- components, 10
- driver call to incremental linker, 138
- options passed to linker, 82
- options summary table, 14

C programming tools, 11

C99

- // comment indicators, 319
- __func__ support, 319
- inline function specifier, 319
- list of keywords, 319
- variable length arrays, 319

case statements, 307

cc compiler options, 13 to 82

- #, 18, 21
- ###, 18, 21
- A, 17, 21
- B, 19, 21
- C, 17, 22
- c, 18, 22
- D, 17, 22
- d, 19, 23
- dalign, 23
- E, 17, 23
- eroff, 19, 23
- errtags, 19, 24
- errwarn, 19, 24
- fast, 14, 25
- fd, 17, 26
- flags, 27
- fnonstd, 16, 27
- fns, 16, 27
- fprecision, 16, 28
- fround, 16, 28
- fsimple, 16, 28
- fsingle, 16, 29

- fstore, 16, 29
- ftrap, 16, 30
- G, 19, 30
- g, 19, 30
- H, 17, 31
- h, 19, 31
- I, 17, 32
- i, 19, 32
- keeptmp, 18, 33
- KPIC, 32
- Kpic, 32
- L, 19, 33
- l, 19, 33
- mc, 19, 33
- misalign, 33
- misalign2, 33
- mr, 19, 34
- mt, 16, 34
- native, 34
- nofstore, 16, 34
- noqueue, 14, 34
- O, 35
- o, 18, 35
- P, 17, 35
- p, 14, 35
- Q, 20, 35
- qp, 35
- R, 20, 36
- S, 18, 36
- s, 19, 36
- U, 17, 36
- V, 18, 37
- v, 19, 37
- W, 18, 37
- w, 19, 38
- X, 18, 38
- x386, 15, 39
- x486, 15, 39
- xa, 15, 39
- xalias_level
 - examples, 196 to 205
 - explanation, 189
 - grouped by functionality, 15
 - syntax for cc, 40
- xarch, 20, 42
- xautopar, 16, 47
- xbuiltin, 15, 48
- xc99, 17
- xcache, 20, 49

- xCC, 17, 48
- xcg, 20, 50
- xchar_byte_order, 15, 51
- xchip, 20, 51
- xcode, 20, 52
- xcrossfile, 15, 54
- xcsi, 17, 54
- xdepend, 15, 55
- xe, 19, 55
- xexplicitpar, 16, 55
- xF, 15, 56
- xhelp, 18, 56
- xildoff, 20, 56
- xildon, 20, 57
- xinline, 15, 57
- xipo, 15, 58
- xlibmieee, 16, 59
- xlibmil, 15, 59
- xlic_lib, 15, 59
- xlicinfo, 14, 60
- xloopinfo, 16, 60
- xM, 17, 60
- xM1, 17, 61
- xmaxopt, 15, 61
- xmemalign, 15, 62
- xMerge, 20, 61
- xnolib, 20, 63
- xnolibmil, 16, 20, 63
- xO, 15, 64
- xopenmp, 16, 66
- xP, 17, 67
- xparallel, 17, 67
- xpentium, 15, 67
- xpg, 18, 68
- xprefetch, 15, 68
- xprofile, 15, 68
- xreduction, 17, 70
- xregs, 20, 70
- xrestrict, 15, 71
- Xs, 323, 332
- xs, 19, 72
- xsafe, 15, 72
- xsb, 18, 72
- xsbfast, 18, 73
- xsfpcnst, 16, 73
- xspace, 15, 73
- xstrconst, 20, 73
- xtarget, 20, 73
- xtemp, 18, 79
- xtime, 18, 79
- xtransition, 19, 79, 85
- xtrigraphs, 18, 80
- xunroll, 15, 80
- xvector, 16, 81
- xvpara, 17, 19, 81
- Y, 18, 81
- YA, 18, 81
- YI, 18, 82
- YP, 18, 82
- YS, 18, 82
- Z11, 17, 82
- cg (code generator), 11
- cg386 (intermediate language translator), 10
- char
 - storage allocation for, 285
- character
 - bits in set, 299
 - decimal point, 316
 - mapping set, 299
 - multibyte, shift status, 299
 - set, collation sequence, 317
 - single-character character constant, 308
 - source and execution of set, 299
 - space, 314
 - testing of sets, 310
- clock function, 316
- code generator, 11
- code optimization
 - by using `-fast`, 25
 - optimizer, 10
 - with `-xO`, 64
- codegen (code generator), 11
- comments
 - preventing removal by preprocessor, 22
 - using `//` by issuing `-xCC`, 48
 - using `//` in C99, 319
- compatibility options, 13, 38
- consistency checks by `lint`, 182
- const, 224 to 226, 245
- constants
 - promotion of integral, 218
 - specific to Sun ANSI/ISO C, 88 to 89
- conversion
 - integers, 302
- conversions, 85, 87
- cpp (C preprocessor), 10

- cscope, 265 to 284
 - command-line use, 267, 275 to 277
 - editing source files, 266 to 267, 274 to 275, 283 to 284
 - environment setup, 266 to 267, 284
 - environment variables, 278
 - searching source files, 265 to 266, 267, 268 to 274
 - See also* Source Browser
 - usage examples, 266 to 275, 279 to 283

D

- data types
 - long long, 86
 - unsigned long long, 86
- `__DATE__`, 309
- date and time formats, 317
- dbx tool
 - disable Auto-Read for, 72
 - initializes faster, 72
 - symbol table information for, 30
- debugging information, removing, 36
- decimal-point character, 316
- declarators, 307
- default
 - compiler behavior, 38
 - handling and SIGILL, 313
 - locale, 300
- default dirs searched for libraries, 14
- deferred-link messages, 142
- `#define`, 22
- diagnostics, format, 298
- directives, 93
- documentation index, 6
- documentation, accessing, 6
- domain errors, math functions, 311
- double
 - storage allocation for, 286
- dynamic linking, 23

E

- edit, source files, *See* cscope
- EDITOR, 266, 283

- ellipsis notation, 210, 212, 246
- enhanced mode of lint, 157
- environment variable
 - EDITOR as used by cscope, 266, 283
 - LANG, 300
 - LC_ALL, 300
 - LC_CTYPE, 300
 - LD_DEBUG, 153
 - LD_LIBRARY_PATH, 147
 - LD_LIBRARY_PATH as used by incremental linker, 152
 - LD_LIBRARY_PATH_64, 153
 - LD_OPTIONS, 153
 - LD_PRELOAD, 153
 - LD_RUN_PATH, 153
 - list of Sun specific variables, 83
 - OMP_DYNAMIC, 335
 - OMP_NUM_THREADS, 335
 - OMP_SCHEDULE, 335
 - PARALLEL, 48, 107
 - SUN_PROFDATA, 69, 84
 - SUN_PROFDATA_DIR, 69, 84
 - SUNPRO_SB_INIT_FILE_NAME, 84
 - SUNW_MP_THR_IDLE, 107
 - SUNW_MP_WARN, 108
 - TCOVDIR, 70
 - TERM as used by cscope, 266
 - TZ, 316
 - VPATH as used by cscope, 267
- environment variables
 - MANPATH, 5
 - PATH, 5
- ERANGE, 311
- errno, 311
- error messages, 298
 - `c_lint_messages` file, 173
 - suppressing in lint, 165
- example messages
 - full relink*, 144
 - ild version*, 144
 - new working directory*, 145
 - out of room*, 143
 - running strip*, 143
 - too many files changed*, 144
- executable, modifying, 141
- expressions, grouping and evaluation in, 238 to 241

F

- faster linking and initializing, 72
- fbe (assembler), 11
- file
 - padding by incremental linker, 138
- filename, .profile extension for, 69
- files used by `ild`, 156
- files, temporary, 83
- filters for `lint`, 187 to 188
- final production code and incremental linker, 141
- FIPS 160 standard, 9
- Fix and Continue*
 - and `ild`, 137
 - linking, 137
- float
 - storage allocation for, 286
- float expressions as single precision, 29
- floating point, 303
 - gradual underflows, 92
 - nonstandard, 27
 - nonstop, 92
 - representations, 303
 - truncation, 304
 - values, 303
- `fprintf` function, 315
- `fscanf` function, 315
- full relink, reasons for, 142
- function
 - `clock`, 316
 - `fmod`, 312
 - `fprintf`, 315
 - `fscanf`, 315
 - `omp_get_num_threads`, 335
 - `omp_set_dynamic`, 335
 - prototypes, 182
 - prototypes, `lint` checks for, 186
 - `remove`, 315
 - `rename`, 315
 - `sunw_mp_register`, 106
- function prototypes, 208 to 212
- functions with varying argument lists, 212 to 214

G

- `-g`
 - example 1, 143
 - example 2, 143
- global symbols and the incremental linker, 138
- gradual underflows, 92

H

- hardware architecture, 42
- header files
 - format for `#include` directives, 89
 - how to include, 89 to 90
 - `math.h`, 26
 - standard place, 89 to 90
 - with `lint`, 159 to 160
- how to use the incremental linker, 138

I

- `i386` predefined token, 22, 104, 177
- implementation-defined behavior, 297 to 317
- incomplete types, 241 to 243
- incremental linker (ILD)
 - as compared to linker, 140
 - as part of C compilation system, 11
 - as used by compiler driver, 138
 - bypassing with `-G`, 30
- commands
 - `-a`, 145
 - `-B`, 146
 - `-d`, 146
 - `-e`, 146
 - `-g`, 146
 - `-I`, 146
 - `-i`, 147
 - `-L`, 147
 - `-l`, 147
 - `-m`, 147
 - `-o`, 147
 - `-Q`, 148
 - `-R`, 148
 - `-s`, 148
 - `-t`, 148
 - `-u`, 148
 - `-V`, 148

- xildoff, 149
- xildon, 149
- YP, 149
- z, 149
- z defs, 149
- z i_dryrun, 150
- z i_full, 150
- z i_noincr, 150
- z i_quiet, 150
- z i_verbose, 150
- z nodefs, 150

commands accepted from compiler

- a, 151
- e, 151
- I, 151
- m, 151
- t, 151
- u, 152

environment variables

- LD_DEBUG, 153
- LD_LIBRARY_PATH, 152
- LD_LIBRARY_PATH_64, 153
- LD_OPTIONS, 153
- LD_PRELOAD, 153
- LD_RUN_PATH, 153

examples, 143 to 145

figure explaining, 139

global symbols, 138

how it works, 140

how to use, 138

impact of changing object files, 141

impact on final production code, 141

introduction, 137

invalidating object files, 138

invoking with `-g`, 30

limitations, 141

linker commands passed to, 138

overview, 138

performance improvement with `-xs`, 72

relink messages, 142

relocation records, 138

saved files, 138

See linker

symbol references, 140

timestamps, 138

turning off with `-xildoff`, 56

turning on with `-xildon`, 57

unsupported commands

- D, 155

- F, 155
- M, 156
- r, 156

unsupported commands from linker

- B, 154
- b, 154
- G, 154
- h, 155
- z muldefs, 155
- z text, 155

inline expansion templates, 59, 63

inline function specifier for C99, 319

inlining, 59

int

- storage allocation for, 285

integers, 301 to 302

integral constants, promotion of, 218

interactive device, 299

internationalization, 228 to 231, 234 to 237

interprocedural analysis pass, 58

invoking the incremental linker, 138

iropt (code optimizer), 10

isalnum, 310

isalpha, 310

isctrl, 310

islower, 310

ISO/IEC 9899:1990 Programming Language C, 9

ISO/IEC 9899:1999 Programming Language C, 9, 319

isprint, 310

isupper, 310

K

K&R C vs. ANSI/ISO C, 13, 38 to 39, 332

keywords, 86

L

LANG, 300

layout level alias disambiguation, 41

LC_ALL, 300

LC_CTYPE, 300

LD_DEBUG, 153

- LD_LIBRARY_PATH, 147, 152
- LD_LIBRARY_PATH_64, 153
- LD_OPTIONS, 153
- LD_PRELOAD, 153
- LD_RUN_PATH, 153
- libfast.a, 322
- libraries
 - default dirs searched by cc, 14
 - intrinsic name, 32
 - libfast.a, 322
 - lint, 186 to 187
 - llib-lx.ln, 186
 - renaming shared, 31
 - shared or non shared, 21
 - specifying dynamic or static links, 21
- library bindings, 21
- limit of memory allocation on stack, 285
- link, static vs. dynamic, 23
- linker
 - invoking with -G, 30
 - options not supported by incremental linker, 154
 - options received from compiler, 82
 - See incremental linker (ILD)
 - specifying dynamic or static linking in, 23
 - suppressing linking with, 22
 - using in place of incremental linker, 56
- lint
 - basic mode
 - introduced, 157
 - invoking, 158
 - commands
 - , 160
 - ###, 161
 - a, 161
 - b, 161
 - C, 161
 - c, 161
 - dirout, 161
 - err=warn, 161
 - errchk, 162
 - errfmt, 164
 - errhdr, 164
 - erroff, 165
 - errtags, 166
 - errwarn, 166
 - F, 167
 - fd, 167
 - flagsrc, 167
 - h, 167
 - I, 167
 - k, 167
 - L, 167
 - l, 168
 - m, 168
 - n, 170
 - Ncheck, 168
 - Nlevel, 169
 - o, 170
 - p, 170
 - R, 170
 - s, 170
 - u, 171
 - v, 171
 - v, 171
 - W, 171
 - x, 171
 - Xalias_level, 171
 - Xarch=v9, 172
 - XCC, 171
 - Xexplicitpar, 172
 - Xkeeptmp, 172
 - Xtemp, 172
 - Xtime, 172
 - Xtransition, 173
 - y, 173
 - consistency checks, 182
 - diagnostics, 181 to 186
 - directives, 177 to 181
 - enhanced mode
 - introduced, 157
 - invoking, 159
 - filters, 187 to 188
 - header files, finding, 159
 - how lint examines code, 159
 - introduction to, 157
 - libraries, 186 to 187
 - messages
 - c_lint_messages file, 173
 - formats of, 174 to 176
 - message ID (tag), identifying, 166, 173
 - suppressing, 173
 - portability checks, 182 to 184
 - predefined tokens, 177
 - predefinition, 93
 - questionable constructs, 184 to 186
 - recognized cc commands, 160

- llib-lx.ln library, 186
- local time zone, 316
- locale, 234, 235, 237
 - behavior, 316
 - default, 300
 - use of non-conforming, 54
- long
 - storage allocation for, 286
- long double, 295
 - storage allocation for, 286
- long int, 87
- long long, 86 to 87
 - arithmetic promotions, 87
 - passing, 295
 - representation of, 288
 - returning, 295
 - storage allocation for, 286
 - suffix, 88
 - value preserving, 88
- loops, 55

M

- macro expansion, 221
- macros
 - __DATE__, 309
 - __MATHERR_ERRNO_DONTCARE, 26
 - __RESTRICT, 104
 - __TIME__, 309
- main
 - semantics of args, 298
- man pages, accessing, 4
- MANPATH environment variable,
 - setting, 5
- math functions, domain errors, 311
- mcs and strip, 143
- memory allocation on the stack, 285
- message ID (tag), 23, 24
- messages
 - deferred link, 142
 - error, 298
 - error file for lint, 173
 - ild relink, 142
- mode, compiler, 39
- MP C, 105 to 136

- multibyte characters and wide characters, 228 to 231
- multiprocessing, 105 to 136
- mwinline, 10

N

- newline, terminating, 314
- nonstop
 - floating-point arithmetic, 27, 92
- null characters not appended to data, 314
- NULL, value of, 309

O

- object file
 - impact of changes on incremental linker, 141
 - invalidated by incremental linker, 138
 - linking with ld, 22
 - producing object file for each source file, 22
 - suppressing removal of, 22
- OMP_DYNAMIC environment variable, 335
- omp_get_num_threads, 335
- OMP_NUM_THREADS environment variable, 335
- OMP_SCHEDULE environment variable, 335
- omp_set_dynamic, 335
- OpenMP
 - defining the _OPENMP preprocessor token, 66
 - how to compile for, 106
 - implementation specific information, 335
 - OMP_DYNAMIC environment variable, 335
 - omp_get_num_threads, 335
 - OMP_NUM_THREADS environment variable, 335
 - OMP_SCHEDULE environment variable, 335
 - omp_set_dynamic, 335
 - sunw_mp_register, 106
 - supported version information, 335
 - xopenmp command, 66
- optimization
 - by using -xipo, 58
 - for SPARC, 321
 - optimizer, 10
 - specify hardware architecture, 42
 - with -fast, 25
 - with -xO, 64

options

- compiler, 13 to 82
- Incremental Linke Editor `ild`, 145
- lint, 160 to 173

overview of incremental linker, 138

P

padding added to files by incremental linker, 138

padding of structures, 306

PARALLEL, 48, 107

parallelization, 105 to 136

- checking for properly parallelized loops with `-xvpara`, 81
- creating a program database with `-zll`, 82
- finding parallelized loops with `-xloopinfo`, 60
- list of compiler commands for, 16
- See also* OpenMP
- specifying multithread coding with `-mt`, 34
- specifying OpenMP pragmas with `-xopenmp`, 66
- turning on reduction recognition with `-xreduction`, 70
- turning on with `-xautopar` for multiple processors, 47
- with `-xexplicitpar`, 55
- `-xparallel` macro, 67

pass, name and version of each, 37

PATH environment variable, setting, 5

Pentium, 79

performance

- comparison of incremental linker and linker, 140
- impact of incremental linker on final production code, 141
- improving linker's with `-xs`, 72
- optimizing for SPARC, 321
- optimizing with `-fast`, 25
- optimizing with `-xO`, 64

portability checks performed by lint, 182 to 184

portability, of code, 182 to 184

`#pragma alias`, 191

`#pragma alias_level`, 191

`#pragma align`, 94

`#pragma does_not_read_global_data`, 95

`#pragma does_not_return`, 95

`#pragma does_not_write_global_data`, 96

`#pragma error_messages`, 96

`#pragma fini`, 96

`#pragma ident`, 97

`#pragma init`, 97

`#pragma inline`, 97

`#pragma int_to_unsigned`, 97

`#pragma may_not_point_to`, 193

`#pragma may_piont_to`, 192

`#pragma MP serial_loop`, 98, 129

`#pragma MP serial_loop-nested`, 98, 129

`#pragma MP taskloop`, 98, 129

`#pragma no_inline`, 97

`#pragma no_side_effect`, 98, 99

`#pragma noalias`, 192

`#pragma nomemorydepend`, 98

`#pragma opt`, 99

`#pragma pack`, 99

`#pragma pipelooop`, 99

`#pragma rarely_called`, 100

`#pragma redefine_extname`, 100

`#pragma returns_new_memory`, 101

`#pragma unknown_control_flow`, 102

`#pragma unroll`, 102

`#pragma weak`, 103

preassertions for `-Aname`, 21

predefined tokens

- `__'uname -s'_'uname -r'`, 22, 104, 177
- `__BUILTIN_VA_ARG_INCR`, 22, 104, 177
- `__i386`, 22, 104, 177
- `__lint`, 177
- `__RESTRICT`, 22, 104, 177
- `__sparc`, 22, 104, 177
- `__sparcv9`, 22, 172, 177
- `__sun`, 22, 104, 177
- `__SUNPRO_C`, 22, 104, 177
- `__SVR4`, 22, 104, 177
- `__unix`, 22, 104, 177
- `i386`, 22, 104, 177
- lint, 177
- `sparc`, 22, 104, 177
- `sun`, 22, 104, 177
- `unix`, 22, 104, 177

- preprocessing, 219 to 224
 - directives, 22, 89 to 90, 103, 308
 - how to preserve comments, 22
 - predefined names, 103
 - stringizing, 222
 - token pasting, 223
- preserving
 - unsigned, 85
 - value, 85
- printing, 87, 316
- profiling with `tcov`, 39
- programming tools for C, 11
- promotion, 215 to 219
 - bit-fields, 217
 - default arguments, 210
 - integral constants, 218
 - value preserving, 215
- promotions, 85

Q

- qualifiers, 307

R

- reasons for full relinks, 142
- relink messages, 142
- remove function, 315
- removing symbolic debugging information, 36
- rename function, 315
- renaming shared libraries, 31
- representation
 - floating point, 303
 - integers, 301
- reserved names, 231 to 234
 - for expansion, 233
 - for implementation use, 232
 - guidelines for choosing, 234
- `__RESTRICT` macro, 104
- right shift, 302
- rounding behavior, 92

S

- saved files and the incremental linker, 138
- search, source files, *See* `cscope`
- `setlocale(3C)`, 235, 237
- shared libraries, 141
- shared libraries, naming, 31
- shared objects as created by incremental linker, 141
- shell prompts, 3
- short
 - storage allocation for, 285
- signal, 312 to 313
- signed, 85, 300
- Solaris versions supported, 3
- source files
 - checking with `lint`, 157 to 188
 - compatibility of K&R C and ANSI/ISO C, 13
 - editing, *See* `cscope`
 - locating, 308
 - searching, *See* `cscope`
- space characters, 314
- `sparc` predefined token, 22, 104, 177
- stack, memory allocation on, 285
- standards conformance, 9, 83
- static linking, 23
- static scheduling, 335
- `std` level alias disambiguation, 42
- storage allocation for types, 285
- streams, 314
- `strict` level alias disambiguation, 42
- string literals in text segment, 73
- `strip` and `mcs`, 143
- `strong` level alias disambiguation, 42
- structure
 - alignment, 306
 - padding, 306
- `sun` predefined token, 22, 104, 177
- `SUN_PROFDATA`, 69, 84
- `SUN_PROFDATA_DIR`, 69, 84
- `SUNPRO_SB_INIT_FILE_NAME`, 84
- `SUNW_MP_THR_IDLE`, 107
- symbol references of incremental linker, 140
- symbol table for `dbx`, 72
- symbolic debugging information, removing, 36

T

- tcov
 - new style with `-xprofile`, 69
- tcov tool, 39
- TCOVDIR, 70
- Temporary files, 83
- TERM, 266
- text
 - segment and string literals, 73
 - stream, 314
- `__TIME__`, 309
- time and date formats, 317
- time to link, 137
- timestamps as used by incremental linker, 138
 - `/tmp`, 83
- TMPDIR environment variable, 83
- tokens, 219 to 224
- tools for programming with C, 11
- trigraph sequences, 220
- type conversions, 85
- type qualifiers, 224 to 228
- type-based alias-disambiguation, 189 to 205
- types
 - storage allocation for, 285
- types, compatible and composite, 244 to 246
- types, incomplete, 241 to 243
- typographic conventions, 3
- TZ, 316

U

- underflow, gradual, 27
- unix predefined token, 22, 104, 177
- unsigned, 85, 300
- using assembly in source, 86

V

- value
 - floating point, 303
 - integers, 301
 - preserving, 85
- `varargs(5)`, 210
- variable length arrays in C99, 319
- volatile, 224 to 226, 227 to 228, 245
- volatile, 307
- VPATH, 267

W

- warning messages, 298
- warning messages file for `lint`, 173
- weak level alias disambiguation, 41
- whole-program optimizations, 58
- wide character constants, 230 to 231
- wide characters, 229 to 231
- wide string literals, 230 to 231
- write on text stream, 314

Z

- `-z i_verbose` option, 142
- zero-length file, 315

