GNU gcj

Tom Tromey

For the GCC-4.0.4 Version*

Published by the Free Software Foundation
59 Temple Place - Suite 330
Boston, MA 02111-1307, USA

Copyright © 2001, 2002, 2003, 2004, 2005 Free Software Foundation, Inc.

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.2 or any later version published by the Free Software Foundation; with the Invariant Sections being "GNU General Public License", the Front-Cover texts being (a) (see below), and with the Back-Cover Texts being (b) (see below). A copy of the license is included in the section entitled "GNU Free Documentation License".

(a) The FSF's Front-Cover Text is:

A GNU Manual

(b) The FSF's Back-Cover Text is:

You have freedom to copy and modify this GNU Manual, like GNU software. Copies published by the Free Software Foundation raise funds for GNU development.

Table of Contents

Introduction 1

GNU GENERAL PUBLIC LICENSE 2

Preamble 2
TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION 3
How to Apply These Terms to Your New Programs 7

GNU Free Documentation License 8

ADDENDUM: How to use this License for your documents 14

1 Invoking gcj 15

1.1 Input and output files 15
1.2 Input Options 15
1.3 Encodings 16
1.4 Warnings 17
1.5 Linking 17
1.6 Code Generation 18
1.7 Configure-time Options 20

2 Compatibility with the Java Platform 21

2.1 Standard features not yet supported 21
2.2 Extra features unique to gcj 21

3 Invoking gcjh 23

4 Invoking gjnih 25

5 Invoking jv-scan 27

6 Invoking jcf-dump 28

7 Invoking gij 29

8 Invoking gcj-dbtool 31

9 Invoking jv-convert 32

10 Invoking grmic 33

11 Invoking grmiregistry 34

12 About CNI 35

12.1 Basic concepts 35
12.1.1 Limitations 35
12.2 Packages 36
12.2.1 Leaving out package names 36
12.3 Primitive types 37
12.3.1 Reference types associated with primitive types 37
12.4 Interfaces 37
12.5 Objects and Classes 38
12.5.1 Classes 38
12.5.2 Object fields 38
12.5.3 Access specifiers 39
12.6 Class Initialization 39
12.7 Object allocation 39
12.8 Memory allocation 40
12.9 Arrays 40
12.9.1 Creating arrays 41
12.10 Methods 41
12.10.1 Overloading 41
12.10.2 Static methods 42
12.10.3 Object Constructors 42
12.10.4 Instance methods 42
12.10.5 Interface methods 42
12.11 Strings 43
12.12 Interoperating with C/C++ 43
12.12.1 RawData 44
12.12.2 RawDataManaged 45
12.12.3 Native memory allocation 45
12.12.4 Posix signals 45
12.13 Exception Handling 45
12.14 Synchronization 46
12.15 Invocation 47
12.15.1 Handling uncaught exceptions 48
12.15.2 Example 48
12.16 Reflection 49

13 System properties 50

13.1 Standard Properties 50
13.2 GNU Classpath Properties 52
13.3 libgcj Runtime Properties 53

14 Resources 55

Introduction

This manual describes how to use `gcj`, the GNU compiler for the Java programming language. `gcj` can generate both '`.class`' files and object files, and it can read both Java source code and '`.class`' files.

GNU GENERAL PUBLIC LICENSE

Version 2, June 1991

Copyright © 1989, 1991 Free Software Foundation, Inc.
59 Temple Place - Suite 330, Boston, MA 02111-1307, USA

Everyone is permitted to copy and distribute verbatim copies
of this license document, but changing it is not allowed.

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software—to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Library General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

   Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

   You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

   a. You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.

   b. You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.

   c. If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

   These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:
   a. Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
   b. Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
   c. Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

   The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

   If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.

6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.

7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

   If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

   It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

   This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

   Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software

which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.
12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

```
one line to give the program's name and a brief idea of what it does.
Copyright (C) year  name of author

This program is free software; you can redistribute it and/or modify
it under the terms of the GNU General Public License as published by
the Free Software Foundation; either version 2 of the License, or
(at your option) any later version.

This program is distributed in the hope that it will be useful,
but WITHOUT ANY WARRANTY; without even the implied warranty of
MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.  See the
GNU General Public License for more details.

You should have received a copy of the GNU General Public License
along with this program; if not, write to the Free Software Foundation,
Inc., 59 Temple Place - Suite 330, Boston, MA  02111-1307, USA.
```

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this when it starts in an interactive mode:

```
Gnomovision version 69, Copyright (C) year name of author
Gnomovision comes with ABSOLUTELY NO WARRANTY; for details
type ‘show w’.
This is free software, and you are welcome to redistribute it
under certain conditions; type ‘show c’ for details.
```

The hypothetical commands ‘`show w`’ and ‘`show c`’ should show the appropriate parts of the General Public License. Of course, the commands you use may be called something other than ‘`show w`’ and ‘`show c`’; they could even be mouse-clicks or menu items—whatever suits your program.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a "copyright disclaimer" for the program, if necessary. Here is a sample; alter the names:

```
Yoyodyne, Inc., hereby disclaims all copyright interest in the program
‘Gnomovision’ (which makes passes at compilers) written by James Hacker.

signature of Ty Coon, 1 April 1989
Ty Coon, President of Vice
```

This General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the GNU Library General Public License instead of this License.

GNU Free Documentation License

Version 1.2, November 2002

Copyright © 2000,2001,2002 Free Software Foundation, Inc.
59 Temple Place, Suite 330, Boston, MA 02111-1307, USA

Everyone is permitted to copy and distribute verbatim copies
of this license document, but changing it is not allowed.

0. PREAMBLE

The purpose of this License is to make a manual, textbook, or other functional and useful document *free* in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or non-commercially. Secondarily, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of "copyleft", which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The "Document", below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as "you". You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A "Modified Version" of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A "Secondary Section" is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document's overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The "Invariant Sections" are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released

under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The "Cover Texts" are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A "Transparent" copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not "Transparent" is called "Opaque".

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The "Title Page" means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, "Title Page" means the text near the most prominent appearance of the work's title, preceding the beginning of the body of the text.

A section "Entitled XYZ" means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as "Acknowledgements", "Dedications", "Endorsements", or "History".) To "Preserve the Title" of such a section when you modify the Document means that it remains a section "Entitled XYZ" according to this definition.

The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

2. VERBATIM COPYING

   You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and

that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

3. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.

B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.

C. State on the Title page the name of the publisher of the Modified Version, as the publisher.

D. Preserve all the copyright notices of the Document.

E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.

F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.

G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.

H. Include an unaltered copy of this License.

I. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.

J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.

K. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.

L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.

M. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.

N. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.

O. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section Entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties—for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

5. COMBINING DOCUMENTS

   You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

   The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

   In the combination, you must combine any sections Entitled "History" in the various original documents, forming one section Entitled "History"; likewise combine any sections Entitled "Acknowledgements", and any sections Entitled "Dedications". You must delete all sections Entitled "Endorsements."

6. COLLECTIONS OF DOCUMENTS

   You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

   You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

7. AGGREGATION WITH INDEPENDENT WORKS

   A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called

an "aggregate" if the copyright resulting from the compilation is not used to limit the legal rights of the compilation's users beyond what the individual works permit. When the Document is included an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document's Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

8. TRANSLATION

   Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warrany Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

   If a section in the Document is Entitled "Acknowledgements", "Dedications", or "History", the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

9. TERMINATION

   You may not copy, modify, sublicense, or distribute the Document except as expressly provided for under this License. Any other attempt to copy, modify, sublicense or distribute the Document is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

10. FUTURE REVISIONS OF THIS LICENSE

    The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See `http://www.gnu.org/copyleft/`.

    Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License "or any later version" applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation.

ADDENDUM: How to use this License for your documents

To use this License in a document you have written, include a copy of the License in the document and put the following copyright and license notices just after the title page:

```
Copyright (C)  year  your name.
Permission is granted to copy, distribute and/or modify this document
under the terms of the GNU Free Documentation License, Version 1.2
or any later version published by the Free Software Foundation;
with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts.
A copy of the license is included in the section entitled ``GNU
Free Documentation License''.
```

If you have Invariant Sections, Front-Cover Texts and Back-Cover Texts, replace the "with...Texts." line with this:

```
with the Invariant Sections being list their titles, with
the Front-Cover Texts being list, and with the Back-Cover Texts
being list.
```

If you have Invariant Sections without Cover Texts, or some other combination of the three, merge those two alternatives to suit the situation.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.

1 Invoking gcj

As `gcj` is just another front end to `gcc`, it supports many of the same options as gcc. See section "Option Summary" in *Using the GNU Compiler Collection (GCC)*. This manual only documents the options specific to `gcj`.

1.1 Input and output files

A `gcj` command is like a `gcc` command, in that it consists of a number of options and file names. The following kinds of input file names are supported:

file`.java`
: Java source files.

file`.class`
: Java bytecode files.

file`.zip`
file`.jar`
: An archive containing one or more `.class` files, all of which are compiled. The archive may be compressed. Files in an archive which don't end with '`.class`' are treated as resource files; they are compiled into the resulting object file as '`core:`' URLs.

`@`*file*
: A file containing a whitespace-separated list of input file names. (Currently, these must all be `.java` source files, but that may change.) Each named file is compiled, just as if it had been on the command line.

library`.a`
library`.so`
`-l`*libname*
: Libraries to use when linking. See the `gcc` manual.

You can specify more than one input file on the `gcj` command line, in which case they will all be compiled. If you specify a `-o` *FILENAME* option, all the input files will be compiled together, producing a single output file, named *FILENAME*. This is allowed even when using `-S` or `-c`, but not when using `-C` or `--resource`. (This is an extension beyond the what plain `gcc` allows.) (If more than one input file is specified, all must currently be `.java` files, though we hope to fix this.)

1.2 Input Options

`gcj` has options to control where it looks to find files it needs. For instance, `gcj` might need to load a class that is referenced by the file it has been asked to compile. Like other compilers for the Java language, `gcj` has a notion of a *class path*. There are several options and environment variables which can be used to manipulate the class path. When `gcj` looks for a given class, it searches the class path looking for matching '`.class`' or '`.java`' file. `gcj` comes with a built-in class path which points at the installed '`libgcj.jar`', a file which contains all the standard classes.

In the below, a directory or path component can refer either to an actual directory on the filesystem, or to a '`.zip`' or '`.jar`' file, which `gcj` will search as if it is a directory.

`-I`*dir* All directories specified by `-I` are kept in order and prepended to the class path constructed from all the other options. Unless compatibility with tools like `javac` is important, we recommend always using `-I` instead of the other options for manipulating the class path.

`--classpath=`*path*
This sets the class path to *path*, a colon-separated list of paths (on Windows-based systems, a semicolon-separate list of paths). This does not override the builtin ("boot") search path.

`--CLASSPATH=`*path*
Deprecated synonym for `--classpath`.

`--bootclasspath=`*path*
Where to find the standard builtin classes, such as `java.lang.String`.

`--extdirs=`*path*
For each directory in the *path*, place the contents of that directory at the end of the class path.

`CLASSPATH`
This is an environment variable which holds a list of paths.

The final class path is constructed like so:

- First come all directories specified via `-I`.
- If '`--classpath`' is specified, its value is appended. Otherwise, if the `CLASSPATH` environment variable is specified, then its value is appended. Otherwise, the current directory (`"."`) is appended.
- If `--bootclasspath` was specified, append its value. Otherwise, append the built-in system directory, '`libgcj.jar`'.
- Finally, if `--extdirs` was specified, append the contents of the specified directories at the end of the class path. Otherwise, append the contents of the built-in extdirs at `$(prefix)/share/java/ext`.

The classfile built by `gcj` for the class `java.lang.Object` (and placed in `libgcj.jar`) contains a special zero length attribute `gnu.gcj.gcj-compiled`. The compiler looks for this attribute when loading `java.lang.Object` and will report an error if it isn't found, unless it compiles to bytecode (the option `-fforce-classes-archive-check` can be used to override this behavior in this particular case.)

`-fforce-classes-archive-check`
This forces the compiler to always check for the special zero length attribute `gnu.gcj.gcj-compiled` in `java.lang.Object` and issue an error if it isn't found.

1.3 Encodings

The Java programming language uses Unicode throughout. In an effort to integrate well with other locales, `gcj` allows '`.java`' files to be written using almost any encoding. `gcj` knows how to convert these encodings into its internal encoding at compile time.

You can use the `--encoding=NAME` option to specify an encoding (of a particular character set) to use for source files. If this is not specified, the default encoding comes from your current locale. If your host system has insufficient locale support, then `gcj` assumes the default encoding to be the 'UTF-8' encoding of Unicode.

To implement `--encoding`, `gcj` simply uses the host platform's `iconv` conversion routine. This means that in practice `gcj` is limited by the capabilities of the host platform.

The names allowed for the argument `--encoding` vary from platform to platform (since they are not standardized anywhere). However, `gcj` implements the encoding named 'UTF-8' internally, so if you choose to use this for your source files you can be assured that it will work on every host.

1.4 Warnings

`gcj` implements several warnings. As with other generic `gcc` warnings, if an option of the form `-Wfoo` enables a warning, then `-Wno-foo` will disable it. Here we've chosen to document the form of the warning which will have an effect – the default being the opposite of what is listed.

`-Wredundant-modifiers`
: With this flag, `gcj` will warn about redundant modifiers. For instance, it will warn if an interface method is declared `public`.

`-Wextraneous-semicolon`
: This causes `gcj` to warn about empty statements. Empty statements have been deprecated.

`-Wno-out-of-date`
: This option will cause `gcj` not to warn when a source file is newer than its matching class file. By default `gcj` will warn about this.

`-Wno-deprecated`
: Warn if a deprecated class, method, or field is referred to.

`-Wunused` This is the same as `gcc`'s `-Wunused`.

`-Wall` This is the same as `-Wredundant-modifiers -Wextraneous-semicolon -Wunused`.

1.5 Linking

To turn a Java application into an executable program, you need to link it with the needed libraries, just as for C or C++. The linker by default looks for a global function named `main`. Since Java does not have global functions, and a collection of Java classes may have more than one class with a `main` method, you need to let the linker know which of those `main` methods it should invoke when starting the application. You can do that in any of these ways:

- Specify the class containing the desired `main` method when you link the application, using the `--main` flag, described below.
- Link the Java package(s) into a shared library (dll) rather than an executable. Then invoke the application using the `gij` program, making sure that `gij` can find the libraries it needs.

- Link the Java packages(s) with the flag `-lgij`, which links in the `main` routine from the `gij` command. This allows you to select the class whose `main` method you want to run when you run the application. You can also use other `gij` flags, such as `-D` flags to set properties. Using the `-lgij` library (rather than the `gij` program of the previous mechanism) has some advantages: it is compatible with static linking, and does not require configuring or installing libraries.

These `gij` options relate to linking an executable:

`--main=`*CLASSNAME*
: This option is used when linking to specify the name of the class whose `main` method should be invoked when the resulting executable is run.

`-D`*name*`[=`*value*`]`
: This option can only be used with `--main`. It defines a system property named *name* with value *value*. If *value* is not specified then it defaults to the empty string. These system properties are initialized at the program's startup and can be retrieved at runtime using the `java.lang.System.getProperty` method.

`-lgij`
: Create an application whose command-line processing is that of the `gij` command.

  This option is an alternative to using `--main`; you cannot use both.

1.6 Code Generation

In addition to the many `gcc` options controlling code generation, `gcj` has several options specific to itself.

`-C`
: This option is used to tell `gcj` to generate bytecode ('`.class`' files) rather than object code.

`--resource` *resource-name*
: This option is used to tell `gcj` to compile the contents of a given file to object code so it may be accessed at runtime with the core protocol handler as '`core:/`*resource-name*'. Note that *resource-name* is the name of the resource as found at runtime; for instance, it could be used in a call to `ResourceBundle.getBundle`. The actual file name to be compiled this way must be specified separately.

`-d` *directory*
: When used with `-C`, this causes all generated '`.class`' files to be put in the appropriate subdirectory of *directory*. By default they will be put in subdirectories of the current working directory.

`-fno-bounds-check`
: By default, `gcj` generates code which checks the bounds of all array indexing operations. With this option, these checks are omitted, which can improve performance for code that uses arrays extensively. Note that this can result in unpredictable behavior if the code in question actually does violate array bounds constraints. It is safe to use this option if you are sure that your code will never throw an `ArrayIndexOutOfBoundsException`.

`-fno-store-check`
: Don't generate array store checks. When storing objects into arrays, a runtime check is normally generated in order to ensure that the object is assignment compatible with the component type of the array (which may not be known at compile-time). With this option, these checks are omitted. This can improve performance for code which stores objects into arrays frequently. It is safe to use this option if you are sure your code will never throw an `ArrayStoreException`.

`-fjni`
: With `gcj` there are two options for writing native methods: CNI and JNI. By default `gcj` assumes you are using CNI. If you are compiling a class with native methods, and these methods are implemented using JNI, then you must use `-fjni`. This option causes `gcj` to generate stubs which will invoke the underlying JNI methods.

`-fno-assert`
: Don't recognize the `assert` keyword. This is for compatibility with older versions of the language specification.

`-fno-optimize-static-class-initialization`
: When the optimization level is greater or equal to `-O2`, `gcj` will try to optimize the way calls into the runtime are made to initialize static classes upon their first use (this optimization isn't carried out if `-C` was specified.) When compiling to native code, `-fno-optimize-static-class-initialization` will turn this optimization off, regardless of the optimization level in use.

`--disable-assertions[=`*class-or-package*`]`
: Don't include code for checking assertions in the compiled code. If `=`*class-or-package* is missing disables assertion code generation for all classes, unless overridden by a more specific `--enable-assertions` flag. If *class-or-package* is a class name, only disables generating assertion checks within the named class or its inner classes. If *class-or-package* is a package name, disables generating assertion checks within the named package or a subpackage.

  By default, assertions are enabled when generating class files or when not optimizing, and disabled when generating optimized binaries.

`--enable-assertions[=`*class-or-package*`]`
: Generates code to check assertions. The option is perhaps misnamed, as you still need to turn on assertion checking at run-time, and we don't support any easy way to do that. So this flag isn't very useful yet, except to partially override `--disable-assertions`.

`-findirect-dispatch`
: `gcj` has a special binary compatibility ABI, which is enabled by the `-findirect-dispatch` option. In this mode, the code generated by `gcj` honors the binary compatibility guarantees in the Java Language Specification, and the resulting object files do not need to be directly linked against their dependencies. Instead, all dependencies are looked up at runtime. This allows free mixing of interpreted and compiled code.

  Note that, at present, `-findirect-dispatch` can only be used when compiling '`.class`' files. It will not work when compiling from source. CNI also does not

yet work with the binary compatibility ABI. These restrictions will be lifted in some future release.

However, if you compile CNI code with the standard ABI, you can call it from code built with the binary compatibility ABI.

1.7 Configure-time Options

Some `gcj` code generations options affect the resulting ABI, and so can only be meaningfully given when `libgcj`, the runtime package, is configured. `libgcj` puts the appropriate options from this group into a '`spec`' file which is read by `gcj`. These options are listed here for completeness; if you are using `libgcj` then you won't want to touch these options.

`-fuse-boehm-gc`
: This enables the use of the Boehm GC bitmap marking code. In particular this causes `gcj` to put an object marking descriptor into each vtable.

`-fhash-synchronization`
: By default, synchronization data (the data used for `synchronize`, `wait`, and `notify`) is pointed to by a word in each object. With this option `gcj` assumes that this information is stored in a hash table and not in the object itself.

`-fuse-divide-subroutine`
: On some systems, a library routine is called to perform integer division. This is required to get exception handling correct when dividing by zero.

`-fcheck-references`
: On some systems it's necessary to insert inline checks whenever accessing an object via a reference. On other systems you won't need this because null pointer accesses are caught automatically by the processor.

2 Compatibility with the Java Platform

As we believe it is important that the Java platform not be fragmented, `gcj` and `libgcj` try to conform to the relevant Java specifications. However, limited manpower and incomplete and unclear documentation work against us. So, there are caveats to using `gcj`.

2.1 Standard features not yet supported

This list of compatibility issues is by no means complete.

- `gcj` implements the JDK 1.2 language. It supports inner classes and the new 1.4 `assert` keyword. It does not yet support the Java 2 `strictfp` keyword (it recognizes the keyword but ignores it).
- `libgcj` is largely compatible with the JDK 1.2 libraries. However, `libgcj` is missing many packages, most notably `java.awt`. There are also individual missing classes and methods. We currently do not have a list showing differences between `libgcj` and the Java 2 platform.
- Sometimes the `libgcj` implementation of a method or class differs from the JDK implementation. This is not always a bug. Still, if it affects you, it probably makes sense to report it so that we can discuss the appropriate response.
- `gcj` does not currently allow for piecemeal replacement of components within `libgcj`. Unfortunately, programmers often want to use newer versions of certain packages, such as those provided by the Apache Software Foundation's Jakarta project. This has forced us to place the `org.w3c.dom` and `org.xml.sax` packages into their own libraries, separate from `libgcj`. If you intend to use these classes, you must link them explicitly with `-l-org-w3c-dom` and `-l-org-xml-sax`. Future versions of `gcj` may not have this restriction.

2.2 Extra features unique to gcj

The main feature of `gcj` is that it can compile programs written in the Java programming language to native code. Most extensions that have been added are to facilitate this functionality.

- `gcj` makes it easy and efficient to mix code written in Java and C++. See Chapter 12 [About CNI], page 35, for more info on how to use this in your programs.
- When you compile your classes into a shared library they can be automatically loaded by the `libgcj` system classloader. When trying to load a class `gnu.pkg.SomeClass` the system classloader will first try to load the shared library '`lib-gnu-pkg-SomeClass.so`', if that fails to load the class then it will try to load '`lib-gnu-pkg.so`' and finally when the class is still not loaded it will try to load '`lib-gnu.so`'. Note that all '`.`'s will be transformed into '`-`'s and that searching for inner classes starts with their outermost outer class. If the class cannot be found this way the system classloader tries to use the `libgcj` bytecode interpreter to load the class from the standard classpath. This process can be controlled to some degree via the `gnu.gcj.runtime.VMClassLoader.library_control` property; See Section 13.3 [libgcj Runtime Properties], page 53.

- `libgcj` includes a special '`gcjlib`' URL type. A URL of this form is like a `jar` URL, and looks like '`gcjlib:/path/to/shared/library.so!/path/to/resource`'. An access to one of these URLs causes the shared library to be `dlopen()`d, and then the resource is looked for in that library. These URLs are most useful when used in conjunction with `java.net.URLClassLoader`. Note that, due to implementation limitations, currently any such URL can be accessed by only one class loader, and libraries are never unloaded. This means some care must be exercised to make sure that a `gcjlib` URL is not accessed by more than one class loader at once. In a future release this limitation will be lifted, and such libraries will be mapped privately.
- A program compiled by `gcj` will examine the `GCJ_PROPERTIES` environment variable and change its behavior in some ways. In particular `GCJ_PROPERTIES` holds a list of assignments to global properties, such as would be set with the '`-D`' option to `java`. For instance, '`java.compiler=gcj`' is a valid (but currently meaningless) setting.

3 Invoking gcjh

The `gcjh` program is used to generate header files from class files. It can generate both CNI and JNI header files, as well as stub implementation files which can be used as a basis for implementing the required native methods.

`-stubs` This causes `gcjh` to generate stub files instead of header files. By default the stub file will be named after the class, with a suffix of '`.cc`'. In JNI mode, the default output file will have the suffix '`.c`'.

`-jni` This tells `gcjh` to generate a JNI header or stub. By default, CNI headers are generated.

`-force` This option forces `gcjh` to write the output file.

`-old` This option is accepted but ignored for compatibility.

`-trace` This option is accepted but ignored for compatibility.

`-J` *option*
This option is accepted but ignored for compatibility.

`-add` *text*
Inserts *text* into the class body. This is ignored in JNI mode.

`-append` *text*
Inserts *text* into the header file after the class declaration. This is ignored in JNI mode.

`-friend` *text*
Inserts *text* into the class as a `friend` declaration. This is ignored in JNI mode.

`-prepend` *text*
Inserts *text* into the header file before the class declaration. This is ignored in JNI mode.

`--classpath=`*path*
`--CLASSPATH=`*path*
`--bootclasspath=`*path*
`-I`*directory*
`-d` *directory*
`-o` *file* These options are all identical to the corresponding `gcj` options.

`-o` *file* Sets the output file name. This cannot be used if there is more than one class on the command line.

`-td` *directory*
Sets the name of the directory to use for temporary files.

`-M` Print all dependencies to stdout; suppress ordinary output.

`-MM` Print non-system dependencies to stdout; suppress ordinary output.

`-MD` Print all dependencies to stdout.

`-MMD` Print non-system dependencies to stdout.

`--help` Print help about `gcjh` and exit. No further processing is done.

`--version`
Print version information for `gcjh` and exit. No further processing is done.

`-v, --verbose`
Print extra information while running.

All remaining options are considered to be names of classes.

4 Invoking gjnih

The `gjnih` program is used to generate JNI header files from class files. Running it is equivalent to running `gcjh -jni`.

`-stubs` This causes `gjnih` to generate stub files instead of header files. By default the stub file will be named after the class, with a suffix of '`.c`'.

`-jni` This option specifies the default behavior which is to generate a JNI header or stub.

`-force` This option forces `gjnih` to write the output file.

`-old` This option is accepted but ignored for compatibility.

`-trace` This option is accepted but ignored for compatibility.

`-J` *option*
This option is accepted but ignored for compatibility.

`-add` *text*
Inserts *text* into the class body. This is ignored in by `gjnih`.

`-append` *text*
Inserts *text* into the header file after the class declaration. This is ignored in by `gjnih`.

`-friend` *text*
Inserts *text* into the class as a `friend` declaration. This is ignored by `gjnih`.

`-prepend` *text*
Inserts *text* into the header file before the class declaration. This is ignored in by `gjnih`.

`--classpath=`*path*
`--CLASSPATH=`*path*
`--bootclasspath=`*path*
`-I`*directory*
`-d` *directory*
`-o` *file* These options are all identical to the corresponding `gcj` options.

`-o` *file* Sets the output file name. This cannot be used if there is more than one class on the command line.

`-td` *directory*
Sets the name of the directory to use for temporary files.

`-M` Print all dependencies to stdout; suppress ordinary output.

`-MM` Print non-system dependencies to stdout; suppress ordinary output.

`-MD` Print all dependencies to stdout.

`-MMD` Print non-system dependencies to stdout.

`--help` Print help about `gjnih` and exit. No further processing is done.

`--version`
: Print version information for `gjnih` and exit. No further processing is done.

`-v, --verbose`
: Print extra information while running.

All remaining options are considered to be names of classes.

5 Invoking jv-scan

The `jv-scan` program can be used to print information about a Java source file (‘`.java`’ file).

`--no-assert`
: Don’t recognize the `assert` keyword, for backwards compatibility with older versions of the language specification.

`--complexity`
: This prints a complexity measure, related to cyclomatic complexity, for each input file.

`--encoding=`*`name`*
: This works like the corresponding `gcj` option.

`--print-main`
: This prints the name of the class in this file containing a `main` method.

`--list-class`
: This lists the names of all classes defined in the input files.

`--list-filename`
: If `--list-class` is given, this option causes `jv-scan` to also print the name of the file in which each class was found.

`-o` *`file`*
: Print output to the named file.

`--help`
: Print help, then exit.

`--version`
: Print version number, then exit.

6 Invoking jcf-dump

This is a class file examiner, similar to `javap`. It will print information about a number of classes, which are specified by class name or file name.

`-c` Disassemble method bodies. By default method bodies are not printed.

`--print-constants`
Print the constant pool. When printing a reference to a constant also print its index in the constant pool.

`--javap` Generate output in `javap` format. The implementation of this feature is very incomplete.

`--classpath=`*`path`*
`--CLASSPATH=`*`path`*
`-I`*`directory`*
`-o` *`file`* These options as the same as the corresponding `gcj` options.

`--help` Print help, then exit.

`--version`
Print version number, then exit.

`-v, --verbose`
Print extra information while running. Implies `--print-constants`.

7 Invoking gij

`gij` is a Java bytecode interpreter included with `libgcj`. `gij` is not available on every platform; porting it requires a small amount of assembly programming which has not been done for all the targets supported by `gcj`.

The primary argument to `gij` is the name of a class or, with `-jar`, a jar file. Options before this argument are interpreted by `gij`; remaining options are passed to the interpreted program.

If a class name is specified and this class does not have a `main` method with the appropriate signature (a `static void` method with a `String[]` as its sole argument), then `gij` will print an error and exit.

If a jar file is specified then `gij` will use information in it to determine which class' `main` method will be invoked.

`gij` will invoke the `main` method with all the remaining command-line options.

Note that `gij` is not limited to interpreting code. Because `libgcj` includes a class loader which can dynamically load shared objects, it is possible to give `gij` the name of a class which has been compiled and put into a shared library on the class path.

`-cp` *path*
`-classpath` *path*
: Set the initial class path. The class path is used for finding class and resource files. If specified, this option overrides the `CLASSPATH` environment variable. Note that this option is ignored if `-jar` is used.

`-D`*name*`[=`*value*`]`
: This defines a system property named *name* with value *value*. If *value* is not specified then it defaults to the empty string. These system properties are initialized at the program's startup and can be retrieved at runtime using the `java.lang.System.getProperty` method.

`-ms=`*number*
: Equivalent to `-Xms`.

`-mx=`*number*
: Equivalent to `-Xmx`.

`-noverify`
: Do not verify compliance of bytecode with the VM specification. In addition, this option disables type verification which is otherwise performed on BC-ABI compiled code.

`-X`
`-X`*argument*
: Supplying `-X` by itself will cause `gij` to list all the supported `-X` options. Currently these options are supported:

  `-Xms`*size* Set the initial heap size.

  `-Xmx`*size* Set the maximum heap size.

  `-Xss`*size* Set the thread stack size.

Unrecognized `-X` options are ignored, for compatibility with other runtimes.

`-jar` This indicates that the name passed to `gij` should be interpreted as the name of a jar file, not a class.

`--help`
`-?` Print help, then exit.

`--showversion`
Print version number and continue.

`--fullversion`
Print detailed version information, then exit.

`--version`
Print version number, then exit.

`-verbose`
`-verbose:class`
Each time a class is initialized, print a short message on standard error.

`gij` also recognizes and ignores the following options, for compatibility with existing application launch scripts: `-client`, `-server`, `-hotspot`, `-jrockit`, `-agentlib`, `-agentpath`, `-debug`, `-d32`, `-d64`, `-javaagent`, `-noclassgc`, `-verify`, and `-verifyremote`.

8 Invoking gcj-dbtool.

`gcj-dbtool` is a tool for creating and manipulating class file mapping databases. `libgcj` can use these databases to find a shared library corresponding to the bytecode representation of a class. This functionality is useful for ahead-of-time compilation of a program that has no knowledge of `gcj`.

`gcj-dbtool` works best if all the jar files added to it are compiled using `-findirect-dispatch`.

Note that `gcj-dbtool` is currently available as "preview technology". We believe it is a reasonable way to allow application-transparent ahead-of-time compilation, but this is an unexplored area. We welcome your comments.

`-n` *DBFILE* [*SIZE*]
: This creates a new database. Currently, databases cannot be resized; you can choose a larger initial size if desired. The default size is 32,749.

`-a` *DBFILE JARFILE LIB*
`-f` *DBFILE JARFILE LIB*
: This adds a jar file to the database. For each class file in the jar, a cryptographic signature of the bytecode representation of the class is recorded in the database. At runtime, a class is looked up by its signature and the compiled form of the class is looked for in the corresponding shared library. The '`-a`' option will verify that *LIB* exists before adding it to the database; '`-f`' skips this check.

['-']['-0'] `-m` *DBFILE DBFILE*,[*DBFILE*]
: Merge a number of databases. The output database overwrites any existing database. To add databases into an existing database, include the destination in the list of sources.

  If '`-`' or '`-0`' are used, the list of files to read is taken from standard input instead of the command line. For '`-0`', Input filenames are terminated by a null character instead of by whitespace. Useful when arguments might contain white space. The GNU find -print0 option produces input suitable for this mode.

`-t` *DBFILE*
: Test a database.

`-l` *DBFILE*
: List the contents of a database.

`-p`
: Print the name of the default database. If there is no default database, this prints a blank line. If *LIBDIR* is specified, use it instead of the default library directory component of the database name.

`--help`
: Print a help message, then exit.

`--version`
`-v`
: Print version information, then exit.

9 Invoking jv-convert

`jv-convert` [‘`OPTION`’] ... [*INPUTFILE* [*OUTPUTFILE*]]

`jv-convert` is a utility included with `libgcj` which converts a file from one encoding to another. It is similar to the Unix `iconv` utility.

The encodings supported by `jv-convert` are platform-dependent. Currently there is no way to get a list of all supported encodings.

`--encoding` *name*
`--from` *name*
: Use *name* as the input encoding. The default is the current locale’s encoding.

`--to` *name*
: Use *name* as the output encoding. The default is the `JavaSrc` encoding; this is ASCII with ‘`\u`’ escapes for non-ASCII characters.

`-i` *file*
: Read from *file*. The default is to read from standard input.

`-o` *file*
: Write to *file*. The default is to write to standard output.

`--reverse`
: Swap the input and output encodings.

`--help`
: Print a help message, then exit.

`--version`
: Print version information, then exit.

10 Invoking grmic

`grmic` [‘`OPTION`’] . . . *class* . . .

`grmic` is a utility included with `libgcj` which generates stubs for remote objects.

Note that this program isn't yet fully compatible with the JDK `grmic`. Some options, such as ‘`-classpath`’, are recognized but currently ignored. We have left these options undocumented for now.

Long options can also be given with a GNU-style leading ‘`--`’. For instance, ‘`--help`’ is accepted.

`-keep`
`-keepgenerated`
: By default, `grmic` deletes intermediate files. Either of these options causes it not to delete such files.

`-v1.1` Cause `grmic` to create stubs and skeletons for the 1.1 protocol version.

`-vcompat` Cause `grmic` to create stubs and skeletons compatible with both the 1.1 and 1.2 protocol versions. This is the default.

`-v1.2` Cause `grmic` to create stubs and skeletons for the 1.2 protocol version.

`-nocompile`
: Don't compile the generated files.

`-verbose` Print information about what `grmic` is doing.

`-d` *directory*
: Put output files in *directory*. By default the files are put in the current working directory.

`-help` Print a help message, then exit.

`-version` Print version information, then exit.

11 Invoking grmiregistry

`grmic` [‘`OPTION`’] . . . [*port*]

`grmiregistry` starts a remote object registry on the current host. If no port number is specified, then port 1099 is used.

`--help` Print a help message, then exit.

`--version`
Print version information, then exit.

12 About CNI

This documents CNI, the Compiled Native Interface, which is is a convenient way to write Java native methods using C++. This is a more efficient, more convenient, but less portable alternative to the standard JNI (Java Native Interface).

12.1 Basic concepts

In terms of languages features, Java is mostly a subset of C++. Java has a few important extensions, plus a powerful standard class library, but on the whole that does not change the basic similarity. Java is a hybrid object-oriented language, with a few native types, in addition to class types. It is class-based, where a class may have static as well as per-object fields, and static as well as instance methods. Non-static methods may be virtual, and may be overloaded. Overloading is resolved at compile time by matching the actual argument types against the parameter types. Virtual methods are implemented using indirect calls through a dispatch table (virtual function table). Objects are allocated on the heap, and initialized using a constructor method. Classes are organized in a package hierarchy.

All of the listed attributes are also true of C++, though C++ has extra features (for example in C++ objects may be allocated not just on the heap, but also statically or in a local stack frame). Because `gcj` uses the same compiler technology as G++ (the GNU C++ compiler), it is possible to make the intersection of the two languages use the same ABI (object representation and calling conventions). The key idea in CNI is that Java objects are C++ objects, and all Java classes are C++ classes (but not the other way around). So the most important task in integrating Java and C++ is to remove gratuitous incompatibilities.

You write CNI code as a regular C++ source file. (You do have to use a Java/CNI-aware C++ compiler, specifically a recent version of G++.)

A CNI C++ source file must have:

```
#include <gcj/cni.h>
```

and then must include one header file for each Java class it uses, e.g.:

```
#include <java/lang/Character.h>
#include <java/util/Date.h>
#include <java/lang/IndexOutOfBoundsException.h>
```

These header files are automatically generated by `gcjh`.

CNI provides some functions and macros to make using Java objects and primitive types from C++ easier. In general, these CNI functions and macros start with the `Jv` prefix, for example the function `JvNewObjectArray`. This convention is used to avoid conflicts with other libraries. Internal functions in CNI start with the prefix `_Jv_`. You should not call these; if you find a need to, let us know and we will try to come up with an alternate solution.

12.1.1 Limitations

Whilst a Java class is just a C++ class that doesn't mean that you are freed from the shackles of Java, a CNI C++ class must adhere to the rules of the Java programming language.

For example: it is not possible to declare a method in a CNI class that will take a C string (`char*`) as an argument, or to declare a member variable of some non-Java datatype.

12.2 Packages

The only global names in Java are class names, and packages. A *package* can contain zero or more classes, and also zero or more sub-packages. Every class belongs to either an unnamed package or a package that has a hierarchical and globally unique name.

A Java package is mapped to a C++ *namespace*. The Java class `java.lang.String` is in the package `java.lang`, which is a sub-package of `java`. The C++ equivalent is the class `java::lang::String`, which is in the namespace `java::lang` which is in the namespace `java`.

Here is how you could express this:

```
(// Declare the class(es), possibly in a header file:
namespace java {
  namespace lang {
    class Object;
    class String;
    ...
  }
}

class java::lang::String : public java::lang::Object
{
  ...
};
```

The `gcjh` tool automatically generates the necessary namespace declarations.

12.2.1 Leaving out package names

Always using the fully-qualified name of a java class can be tiresomely verbose. Using the full qualified name also ties the code to a single package making code changes necessary should the class move from one package to another. The Java `package` declaration specifies that the following class declarations are in the named package, without having to explicitly name the full package qualifiers. The `package` declaration can be followed by zero or more `import` declarations, which allows either a single class or all the classes in a package to be named by a simple identifier. C++ provides something similar with the `using` declaration and directive.

In Java:

```
import package-name.class-name;
```

allows the program text to refer to *class-name* as a shorthand for the fully qualified name: *package-name.class-name*.

To achieve the same effect C++, you have to do this:

```
using package-name::class-name;
```

Java can also cause imports on demand, like this:

```
import package-name.*;
```

Doing this allows any class from the package *package-name* to be referred to only by its class-name within the program text.

The same effect can be achieved in C++ like this:

```
using namespace package-name;
```

12.3 Primitive types

Java provides 8 *primitives* types which represent integers, floats, characters and booleans (and also the void type). C++ has its own very similar concrete types. Such types in C++ however are not always implemented in the same way (an int might be 16, 32 or 64 bits for example) so CNI provides a special C++ type for each primitive Java type:

Java type	C/C++ typename	Description
`char`	`jchar`	16 bit Unicode character
`boolean`	`jboolean`	logical (true or false) values
`byte`	`jbyte`	8-bit signed integer
`short`	`jshort`	16 bit signed integer
`int`	`jint`	32 bit signed integer
`long`	`jlong`	64 bit signed integer
`float`	`jfloat`	32 bit IEEE floating point number
`double`	`jdouble`	64 bit IEEE floating point number
`void`	`void`	no value

When referring to a Java type You should always use these C++ typenames (e.g.: `jint`) to avoid disappointment.

12.3.1 Reference types associated with primitive types

In Java each primitive type has an associated reference type, e.g.: `boolean` has an associated `java.lang.Boolean` class. In order to make working with such classes easier GCJ provides the macro `JvPrimClass`:

`JvPrimClass` *type* [macro]

Return a pointer to the `Class` object corresponding to the type supplied.

`JvPrimClass(void)` ⇒ `java.lang.Void.TYPE`

12.4 Interfaces

A Java class can *implement* zero or more *interfaces*, in addition to inheriting from a single base class.

CNI allows CNI code to implement methods of interfaces. You can also call methods through interface references, with some limitations.

CNI doesn't understand interface inheritance at all yet. So, you can only call an interface method when the declared type of the field being called matches the interface which declares that method. The workaround is to cast the interface reference to the right superinterface.

For example if you have:

```
interface A
{
  void a();
}

interface B extends A
```

```
{
  void b();
}
```

and declare a variable of type `B` in C++, you can't call `a()` unless you cast it to an `A` first.

12.5 Objects and Classes

12.5.1 Classes

All Java classes are derived from `java.lang.Object`. C++ does not have a unique root class, but we use the C++ class `java::lang::Object` as the C++ version of the `java.lang.Object` Java class. All other Java classes are mapped into corresponding C++ classes derived from `java::lang::Object`.

Interface inheritance (the `implements` keyword) is currently not reflected in the C++ mapping.

12.5.2 Object fields

Each object contains an object header, followed by the instance fields of the class, in order. The object header consists of a single pointer to a dispatch or virtual function table. (There may be extra fields *in front of* the object, for example for memory management, but this is invisible to the application, and the reference to the object points to the dispatch table pointer.)

The fields are laid out in the same order, alignment, and size as in C++. Specifically, 8-bite and 16-bit native types (`byte`, `short`, `char`, and `boolean`) are *not* widened to 32 bits. Note that the Java VM does extend 8-bit and 16-bit types to 32 bits when on the VM stack or temporary registers.

If you include the `gcjh`-generated header for a class, you can access fields of Java classes in the *natural* way. For example, given the following Java class:

```
public class Int
{
  public int i;
  public Integer (int i) { this.i = i; }
  public static zero = new Integer(0);
}
```

you can write:

```
#include <gcj/cni.h>;
#include <Int>;

Int*
mult (Int *p, jint k)
{
  if (k == 0)
    return Int::zero;  // Static member access.
  return new Int(p->i * k);
}
```

12.5.3 Access specifiers

CNI does not strictly enforce the Java access specifiers, because Java permissions cannot be directly mapped into C++ permission. Private Java fields and methods are mapped to private C++ fields and methods, but other fields and methods are mapped to public fields and methods.

12.6 Class Initialization

Java requires that each class be automatically initialized at the time of the first active use. Initializing a class involves initializing the static fields, running code in class initializer methods, and initializing base classes. There may also be some implementation specific actions, such as allocating `String` objects corresponding to string literals in the code.

The GCJ compiler inserts calls to `JvInitClass` at appropriate places to ensure that a class is initialized when required. The C++ compiler does not insert these calls automatically—it is the programmer's responsibility to make sure classes are initialized. However, this is fairly painless because of the conventions assumed by the Java system.

First, `libgcj` will make sure a class is initialized before an instance of that object is created. This is one of the responsibilities of the `new` operation. This is taken care of both in Java code, and in C++ code. When G++ sees a `new` of a Java class, it will call a routine in `libgcj` to allocate the object, and that routine will take care of initializing the class. Note however that this does not happen for Java arrays; you must allocate those using the appropriate CNI function. It follows that you can access an instance field, or call an instance (non-static) method and be safe in the knowledge that the class and all of its base classes have been initialized.

Invoking a static method is also safe. This is because the Java compiler adds code to the start of a static method to make sure the class is initialized. However, the C++ compiler does not add this extra code. Hence, if you write a native static method using CNI, you are responsible for calling `JvInitClass` before doing anything else in the method (unless you are sure it is safe to leave it out).

Accessing a static field also requires the class of the field to be initialized. The Java compiler will generate code to call `Jv_InitClass` before getting or setting the field. However, the C++ compiler will not generate this extra code, so it is your responsibility to make sure the class is initialized before you access a static field from C++.

12.7 Object allocation

New Java objects are allocated using a *class instance creation expression*, e.g.:

```
new Type ( ... )
```

The same syntax is used in C++. The main difference is that C++ objects have to be explicitly deleted; in Java they are automatically deleted by the garbage collector. Using CNI, you can allocate a new Java object using standard C++ syntax and the C++ compiler will allocate memory from the garbage collector. If you have overloaded constructors, the compiler will choose the correct one using standard C++ overload resolution rules.

For example:

```
java::util::Hashtable *ht = new java::util::Hashtable(120);
```

12.8 Memory allocation

When allocting memory in CNI methods it is best to handle out-of-memory conditions by throwing a Java exception. These functions are provided for that purpose:

`void* JvMalloc (`*jsize* ***size***`)` [Function]
Calls malloc. Throws `java.lang.OutOfMemoryError` if allocation fails.

`void* JvRealloc (`*void** ***ptr***, *jsize* ***size***`)` [Function]
Calls realloc. Throws `java.lang.OutOfMemoryError` if reallocation fails.

`void JvFree (`*void** ***ptr***`)` [Function]
Calls free.

12.9 Arrays

While in many ways Java is similar to C and C++, it is quite different in its treatment of arrays. C arrays are based on the idea of pointer arithmetic, which would be incompatible with Java's security requirements. Java arrays are true objects (array types inherit from `java.lang.Object`). An array-valued variable is one that contains a reference (pointer) to an array object.

Referencing a Java array in C++ code is done using the `JArray` template, which as defined as follows:

```
class __JArray : public java::lang::Object
{
public:
  int length;
};

template<class T>
class JArray : public __JArray
{
  T data[0];
public:
  T& operator[](jint i) { return data[i]; }
};
```

There are a number of `typedef`s which correspond to `typedef`s from the JNI. Each is the type of an array holding objects of the relevant type:

```
typedef __JArray *jarray;
typedef JArray<jobject> *jobjectArray;
typedef JArray<jboolean> *jbooleanArray;
typedef JArray<jbyte> *jbyteArray;
typedef JArray<jchar> *jcharArray;
typedef JArray<jshort> *jshortArray;
typedef JArray<jint> *jintArray;
typedef JArray<jlong> *jlongArray;
typedef JArray<jfloat> *jfloatArray;
typedef JArray<jdouble> *jdoubleArray;
```

`T* elements` (*JArray*<*T*> *array*) [Method on `template<class T>`]
: This template function can be used to get a pointer to the elements of the `array`. For instance, you can fetch a pointer to the integers that make up an `int[]` like so:

```
extern jintArray foo;
jint *intp = elements (foo);
```

The name of this function may change in the future.

`jobjectArray JvNewObjectArray` (*jsize length*, *jclass klass*, *jobject init*) [Function]
: Here `klass` is the type of elements of the array and `init` is the initial value put into every slot in the array.

12.9.1 Creating arrays

For each primitive type there is a function which can be used to create a new array of that type. The name of the function is of the form:

`JvNew`*Type*`Array`

For example:

`JvNewBooleanArray`

can be used to create an array of Java primitive boolean types.

The following function definition is the template for all such functions:

`jbooleanArray JvNewBooleanArray` (*jint length*) [Function]
: Create's an array *length* indices long.

`jsize JvGetArrayLength` (*jarray array*) [Function]
: Returns the length of the *array*.

12.10 Methods

Java methods are mapped directly into C++ methods. The header files generated by `gcjh` include the appropriate method definitions. Basically, the generated methods have the same names and *corresponding* types as the Java methods, and are called in the natural manner.

12.10.1 Overloading

Both Java and C++ provide method overloading, where multiple methods in a class have the same name, and the correct one is chosen (at compile time) depending on the argument types. The rules for choosing the correct method are (as expected) more complicated in C++ than in Java, but given a set of overloaded methods generated by `gcjh` the C++ compiler will choose the expected one.

Common assemblers and linkers are not aware of C++ overloading, so the standard implementation strategy is to encode the parameter types of a method into its assembly-level name. This encoding is called *mangling*, and the encoded name is the *mangled name*. The same mechanism is used to implement Java overloading. For C++/Java interoperability, it is important that both the Java and C++ compilers use the *same* encoding scheme.

12.10.2 Static methods

Static Java methods are invoked in CNI using the standard C++ syntax, using the `::` operator rather than the `.` operator.

For example:

```
jint i = java::lang::Math::round((jfloat) 2.3);
```

C++ method definition syntax is used to define a static native method. For example:

```
#include <java/lang/Integer>
java::lang::Integer*
java::lang::Integer::getInteger(jstring str)
{
  ...
}
```

12.10.3 Object Constructors

Constructors are called implicitly as part of object allocation using the `new` operator.

For example:

```
java::lang::Integer *x = new java::lang::Integer(234);
```

Java does not allow a constructor to be a native method. This limitation can be coded round however because a constructor can *call* a native method.

12.10.4 Instance methods

Calling a Java instance method from a C++ CNI method is done using the standard C++ syntax, e.g.:

```
// First create the Java object.
java::lang::Integer *x = new java::lang::Integer(234);
// Now call a method.
jint prim_value = x->intValue();
if (x->longValue == 0)
  ...
```

Defining a Java native instance method is also done the natural way:

```
#include <java/lang/Integer.h>

jdouble
java::lang:Integer::doubleValue()
{
  return (jdouble) value;
}
```

12.10.5 Interface methods

In Java you can call a method using an interface reference. This is supported, but not completely. See Section 12.4 [Interfaces], page 37.

12.11 Strings

CNI provides a number of utility functions for working with Java Java `String` objects. The names and interfaces are analogous to those of JNI.

`jstring JvNewString` (*const char** **`chars`**, *jsize* **`len`**) [Function]
Returns a Java `String` object with characters from the C string *chars* up to the index *len* in that array.

`jstring JvNewStringLatin1` (*const char** **`bytes`**, *jsize* **`len`**) [Function]
Returns a Java `String` made up of *len* bytes from *bytes*.

`jstring JvNewStringLatin1` (*const char** **`bytes`**) [Function]
As above but the length of the `String` is `strlen(bytes)`.

`jstring JvNewStringUTF` (*const char** **`bytes`**) [Function]
Returns a `String` which is made up of the UTF encoded characters present in the C string *bytes*.

`jchar* JvGetStringChars` (*jstring* **`str`**) [Function]
Returns a pointer to an array of characters making up the `String` *str*.

`int JvGetStringUTFLength` (*jstring* **`str`**) [Function]
Returns the number of bytes required to encode the contents of the `String` *str* in UTF-8.

`jsize JvGetStringUTFRegion` (*jstring* **`str`**, *jsize* **`start`**, *jsize* **`len`**, *char** **`buf`**) [Function]
Puts the UTF-8 encoding of a region of the `String` *str* into the buffer `buf`. The region to fetch is marked by *start* and *len*.

Note that *buf* is a buffer, not a C string. It is *not* null terminated.

12.12 Interoperating with C/C++

Because CNI is designed to represent Java classes and methods it cannot be mixed readily with C/C++ types.

One important restriction is that Java classes cannot have non-Java type instance or static variables and cannot have methods which take non-Java types as arguments or return non-Java types.

None of the following is possible with CNI:

```
class ::MyClass : public java::lang::Object
{
   char* variable;  // char* is not a valid Java type.
}


uint
::SomeClass::someMethod (char *arg)
```

```
{
  .
  .
  .
}   // uint is not a valid Java type, neither is char*
```

Of course, it is ok to use C/C++ types within the scope of a method:

```
jint
::SomeClass::otherMethod (jstring str)
{
   char *arg = ...
   .
   .
   .
}
```

12.12.1 RawData

The above restriction can be problematic, so CNI includes the `gnu.gcj.RawData` class. The `RawData` class is a *non-scanned reference* type. In other words variables declared of type `RawData` can contain any data and are not checked by the compiler or memory manager in any way.

This means that you can put C/C++ data structures (including classes) in your CNI classes, as long as you use the appropriate cast.

Here are some examples:

```
class ::MyClass : public java::lang::Object
{
   gnu.gcj.RawData string;

   MyClass ();
   gnu.gcj.RawData getText ();
   void printText ();
}

::MyClass::MyClass ()
{
   char* text = ...
   string = text;
}

gnu.gcj.RawData
::MyClass::getText ()
{
   return string;
}

void
```

```
::MyClass::printText ()
{
  printf("%s\n", (char*) string);
}
```

12.12.2 RawDataManaged

`gnu.gcj.RawDataManaged` is another type used to indicate special data used by native code. Unlike the `RawData` type, fields declared as `RawDataManaged` will be "marked" by the memory manager and considered for garbage collection.

Native data which is allocated using CNI's `JvAllocBytes()` function and stored in a `RawDataManaged` will be automatically freed when the Java object it is associated with becomes unreachable.

12.12.3 Native memory allocation

`void* JvAllocBytes` (*jsize* `size`) [Function]

Allocates *size* bytes from the heap. The memory returned is zeroed. This memory is not scanned for pointers by the garbage collector, but will be freed if no references to it are discovered.

This function can be useful if you need to associate some native data with a Java object. Using a CNI's special `RawDataManaged` type, native data allocated with `JvAllocBytes` will be automatically freed when the Java object itself becomes unreachable.

12.12.4 Posix signals

On Posix based systems the `libgcj` library uses several signals internally. CNI code should not attempt to use the same signals as doing so may cause `libgcj` and/or the CNI code to fail.

SIGSEGV is used on many systems to generate `NullPointerExceptions`. SIGCHLD is used internally by `Runtime.exec()`. Several other signals (that vary from platform to platform) can be used by the memory manager and by `Thread.interrupt()`.

12.13 Exception Handling

While C++ and Java share a common exception handling framework, things are not yet perfectly integrated. The main issue is that the run-time type information facilities of the two languages are not integrated.

Still, things work fairly well. You can throw a Java exception from C++ using the ordinary `throw` construct, and this exception can be caught by Java code. Similarly, you can catch an exception thrown from Java using the C++ `catch` construct.

Here is an example:

```
if (i >= count)
   throw new java::lang::IndexOutOfBoundsException();
```

Normally, G++ will automatically detect when you are writing C++ code that uses Java exceptions, and handle them appropriately. However, if C++ code only needs to execute

destructors when Java exceptions are thrown through it, GCC will guess incorrectly. Sample problematic code:

```
struct S { ~S(); };

extern void bar();     // Is implemented in Java and may throw exceptions.

void foo()
{
  S s;
  bar();
}
```

The usual effect of an incorrect guess is a link failure, complaining of a missing routine called `__gxx_personality_v0`.

You can inform the compiler that Java exceptions are to be used in a translation unit, irrespective of what it might think, by writing `#pragma GCC java_exceptions` at the head of the file. This `#pragma` must appear before any functions that throw or catch exceptions, or run destructors when exceptions are thrown through them.

12.14 Synchronization

Each Java object has an implicit monitor. The Java VM uses the instruction `monitorenter` to acquire and lock a monitor, and `monitorexit` to release it.

The corresponding CNI macros are `JvMonitorEnter` and `JvMonitorExit` (JNI has similar methods `MonitorEnter` and `MonitorExit`).

The Java source language does not provide direct access to these primitives. Instead, there is a `synchronized` statement that does an implicit `monitorenter` before entry to the block, and does a `monitorexit` on exit from the block. Note that the lock has to be released even when the block is abnormally terminated by an exception, which means there is an implicit `try finally` surrounding synchronization locks.

From C++, it makes sense to use a destructor to release a lock. CNI defines the following utility class:

```
class JvSynchronize() {
  jobject obj;
  JvSynchronize(jobject o) { obj = o; JvMonitorEnter(o); }
  ~JvSynchronize() { JvMonitorExit(obj); }
};
```

So this Java code:

```
synchronized (OBJ)
{
   CODE
}
```

might become this C++ code:

```
{
   JvSynchronize dummy (OBJ);
   CODE;
```

```
}
```

Java also has methods with the `synchronized` attribute. This is equivalent to wrapping the entire method body in a `synchronized` statement. (Alternatively, an implementation could require the caller to do the synchronization. This is not practical for a compiler, because each virtual method call would have to test at run-time if synchronization is needed.) Since in `gcj` the `synchronized` attribute is handled by the method implementation, it is up to the programmer of a synchronized native method to handle the synchronization (in the C++ implementation of the method). In other words, you need to manually add `JvSynchronize` in a `native synchronized` method.

12.15 Invocation

CNI permits C++ applications to make calls into Java classes, in addition to allowing Java code to call into C++. Several functions, known as the *invocation API*, are provided to support this.

`jint JvCreateJavaVM` (*JvVMInitArgs** `vm_args`) [Function]

Initializes the Java runtime. This function performs essential initialization of the threads interface, garbage collector, exception handling and other key aspects of the runtime. It must be called once by an application with a non-Java `main()` function, before any other Java or CNI calls are made. It is safe, but not recommended, to call `JvCreateJavaVM()` more than once provided it is only called from a single thread. The *vmargs* parameter can be used to specify initialization parameters for the Java runtime. It may be `NULL`.

JvVMInitArgs represents a list of virtual machine initialization arguments. `JvCreateJavaVM()` ignores the version field.

```
typedef struct JvVMOption
{
  // a VM initialization option
  char* optionString;
  // extra information associated with this option
  void* extraInfo;
} JvVMOption;

typedef struct JvVMInitArgs
{
  // for compatibility with JavaVMInitArgs
  jint version;

  // number of VM initialization options
  jint nOptions;

  // an array of VM initialization options
  JvVMOption* options;

  // true if the option parser should ignore unrecognized options
  jboolean ignoreUnrecognized;
```

```
} JvVMInitArgs;
```

`JvCreateJavaVM()` returns `0` upon success, or `-1` if the runtime is already initialized.

Note: In GCJ 3.1, the `vm_args` parameter is ignored. It is recognized and used as of release 4.0.

`java::lang::Thread* JvAttachCurrentThread` (*jstring* ***name***, *java::lang::ThreadGroup** ***group***) [Function]

Registers an existing thread with the Java runtime. This must be called once from each thread, before that thread makes any other Java or CNI calls. It must be called after `JvCreateJavaVM`. *name* specifies a name for the thread. It may be `NULL`, in which case a name will be generated. *group* is the ThreadGroup in which this thread will be a member. If it is `NULL`, the thread will be a member of the main thread group. The return value is the Java `Thread` object that represents the thread. It is safe to call `JvAttachCurrentThread()` more than once from the same thread. If the thread is already attached, the call is ignored and the current thread object is returned.

`jint JvDetachCurrentThread` () [Function]

Unregisters a thread from the Java runtime. This should be called by threads that were attached using `JvAttachCurrentThread()`, after they have finished making calls to Java code. This ensures that any resources associated with the thread become eligible for garbage collection. This function returns `0` upon success, or `-1` if the current thread is not attached.

12.15.1 Handling uncaught exceptions

If an exception is thrown from Java code called using the invocation API, and no handler for the exception can be found, the runtime will abort the application. In order to make the application more robust, it is recommended that code which uses the invocation API be wrapped by a top-level try/catch block that catches all Java exceptions.

12.15.2 Example

The following code demonstrates the use of the invocation API. In this example, the C++ application initializes the Java runtime and attaches itself. The `java.lang.System` class is initialized in order to access its `out` field, and a Java string is printed. Finally, the thread is detached from the runtime once it has finished making Java calls. Everything is wrapped with a try/catch block to provide a default handler for any uncaught exceptions.

The example can be compiled with `c++ test.cc -lgcj`.

```
// test.cc
#include <gcj/cni.h>
#include <java/lang/System.h>
#include <java/io/PrintStream.h>
#include <java/lang/Throwable.h>

int main(int argc, char *argv)
{
  using namespace java::lang;
```

```
  try
  {
    JvCreateJavaVM(NULL);
    JvAttachCurrentThread(NULL, NULL);

    String *message = JvNewStringLatin1("Hello from C++");
    JvInitClass(&System::class$);
    System::out->println(message);

    JvDetachCurrentThread();
  }
  catch (Throwable *t)
  {
    System::err->println(JvNewStringLatin1("Unhandled Java exception:"));
    t->printStackTrace();
  }
}
```

12.16 Reflection

Reflection is possible with CNI code, it functions similarly to how it functions with JNI.

The types `jfieldID` and `jmethodID` are as in JNI.

The functions:

- `JvFromReflectedField`,
- `JvFromReflectedMethod`,
- `JvToReflectedField`
- `JvToFromReflectedMethod`

will be added shortly, as will other functions corresponding to JNI.

13 System properties

The runtime behavior of the `libgcj` library can be modified by setting certain system properties. These properties can be compiled into the program using the `-Dname[=value]` option to `gcj` or by setting them explicitly in the program by calling the `java.lang.System.setProperty()` method. Some system properties are only used for informational purposes (like giving a version number or a user name). A program can inspect the current value of a property by calling the `java.lang.System.getProperty()` method.

13.1 Standard Properties

The following properties are normally found in all implementations of the core libraries for the Java language.

`java.version`
: The `libgcj` version number.

`java.vendor`
: Set to '`The Free Software Foundation, Inc.`'

`java.vendor.url`
: Set to `http://gcc.gnu.org/java/`.

`java.home`
: The directory where `gcj` was installed. Taken from the `--prefix` option given to `configure`.

`java.class.version`
: The class format version number supported by the libgcj byte code interpreter. (Currently '`46.0`')

`java.vm.specification.version`
: The Virtual Machine Specification version implemented by `libgcj`. (Currently '`1.0`')

`java.vm.specification.vendor`
: The name of the Virtual Machine specification designer.

`java.vm.specification.name`
: The name of the Virtual Machine specification (Set to '`Java Virtual Machine Specification`').

`java.vm.version`
: The `gcj` version number.

`java.vm.vendor`
: Set to '`The Free Software Foundation, Inc.`'

`java.vm.name`
: Set to '`GNU libgcj`'.

`java.specification.version`
: The Runtime Environment specification version implemented by `libgcj`. (Currently set to '`1.3`')

`java.specification.vendor`
: The Runtime Environment specification designer.

`java.specification.name`
: The name of the Runtime Environment specification (Set to '`Java Platform API Specification`').

`java.class.path`
: The paths (jar files, zip files and directories) used for finding class files.

`java.library.path`
: Directory path used for finding native libraries.

`java.io.tmpdir`
: The directory used to put temporary files in.

`java.compiler`
: Name of the Just In Time compiler to use by the byte code interpreter. Currently not used in `libgcj`.

`java.ext.dirs`
: Directories containing jar files with extra libraries. Will be used when resolving classes.

`java.protocol.handler.pkgs`
: A '`|`' separated list of package names that is used to find classes that implement handlers for `java.net.URL`.

`java.rmi.server.codebase`
: A list of URLs that is used by the `java.rmi.server.RMIClassLoader` to load classes from.

`jdbc.drivers`
: A list of class names that will be loaded by the `java.sql.DriverManager` when it starts up.

`file.separator`
: The separator used in when directories are included in a filename (normally '`/`' or '`\`').

`file.encoding`
: The default character encoding used when converting platform native files to Unicode (usually set to '`8859_1`').

`path.separator`
: The standard separator used when a string contains multiple paths (normally '`:`' or '`;`'), the string is usually not a valid character to use in normal directory names.)

`line.separator`
: The default line separator used on the platform (normally '`\n`', '`\r`' or a combination of those two characters).

`policy.provider`
: The class name used for the default policy provider returned by `java.security.Policy.getPolicy`.

`user.name`
: The name of the user running the program. Can be the full name, the login name or empty if unknown.

`user.home`
: The default directory to put user specific files in.

`user.dir` The current working directory from which the program was started.

`user.language`
: The default language as used by the `java.util.Locale` class.

`user.region`
: The default region as used by the `java.util.Local` class.

`user.variant`
: The default variant of the language and region local used.

`user.timezone`
: The default timezone as used by the `java.util.TimeZone` class.

`os.name` The operating system/kernel name that the program runs on.

`os.arch` The hardware that we are running on.

`os.version`
: The version number of the operating system/kernel.

`awt.appletWarning`
: The string to display when an untrusted applet is displayed. Returned by `java.awt.Window.getWarningString()` when the window is "insecure".

`awt.toolkit`
: The class name used for initializing the default `java.awt.Toolkit`. Defaults to `gnu.awt.gtk.GtkToolkit`.

`http.proxyHost`
: Name of proxy host for http connections.

`http.proxyPort`
: Port number to use when a proxy host is in use.

13.2 GNU Classpath Properties

`libgcj` is based on the GNU Classpath (Essential Libraries for Java) a GNU project to create free core class libraries for use with virtual machines and compilers for the Java language. The following properties are common to libraries based on GNU Classpath.

`gcj.dumpobject`
: Enables printing serialization debugging by the `java.io.ObjectInput` and `java.io.ObjectOutput` classes when set to something else then the empty string. Only used when running a debug build of the library.

`gnu.classpath.vm.shortname`
: This is a succint name of the virtual machine. For `libgcj`, this will always be '`libgcj`'.

`gnu.classpath.home.url`
: A base URL used for finding system property files (e.g., '`classpath.security`'). By default this is a '`file:`' URL pointing to the '`lib`' directory under '`java.home`'.

13.3 libgcj Runtime Properties

The following properties are specific to the `libgcj` runtime and will normally not be found in other core libraries for the java language.

`java.fullversion`
: The combination of `java.vm.name` and `java.vm.version`.

`java.vm.info`
: Same as `java.fullversion`.

`impl.prefix`
: Used by the `java.net.DatagramSocket` class when set to something else then the empty string. When set all newly created `DatagramSocket`s will try to load a class `java.net.[impl.prefix]DatagramSocketImpl` instead of the normal `java.net.PlainDatagramSocketImpl`.

`gnu.gcj.progname`
: The name that was used to invoked the program.

`gnu.gcj.runtime.NameFinder.demangle`
: Whether names in a stack trace should be demangled. Defaults to `true`.

`gnu.gcj.runtime.NameFinder.sanitize`
: Whether calls to initialize exceptions and starting the runtime system should be removed from the stack trace. Only done when names are demangled. Defaults to `true`.

`gnu.gcj.runtime.NameFinder.remove_unknown`
: Whether calls to unknown functions (class and method names are unknown) should be removed from the stack trace. Only done when the stack is sanitized. Ignored if this means no stack trace information would be available anymore. Defaults to `true`.

`gnu.gcj.runtime.NameFinder.remove_interpreter`
: Whether runtime interpreter calls (methods in the `_Jv_InterpMethod` class and functions starting with '`ffi_`') should be removed from the stack trace. Only done when the stack is sanitized. Defaults to `true`.

`gnu.gcj.runtime.NameFinder.use_addr2line`
: Whether an external process (`addr2line` or `addr2name.awk`) should be used as fallback to convert the addresses to function names when the runtime is unable to do it through `dladdr`.

`gnu.gcj.runtime.VMClassLoader.library_control`
: This controls how shared libraries are automatically loaded by the built-in class loader. If this property is set to '`full`', a full search is done for each requested class. If this property is set to '`cache`' (the default), then any failed lookups

are cached and not tried again. If this property is set to '`never`', then lookups are never done. For more information, See Section 2.2 [Extensions], page 21.

`gnu.gcj.runtime.endorsed.dirs`
: This is like the standard `java.endorsed.dirs`, property, but specifies some extra directories which are searched after the standard endorsed directories. This is primarily useful for telling `libgcj` about additional libraries which are ordinarily incorporated into the JDK, and which should be loaded by the bootstrap class loader, but which are not yet part of `libgcj` itself for some reason.

`gnu.gcj.jit.compiler`
: This is the full path to `gcj` executable which should be used to compile classes just-in-time when `ClassLoader.defineClass` is called. If not set, `gcj` will not be invoked by the runtime; this can also be controlled via `Compiler.disable`.

`gnu.gcj.jit.options`
: This is a space-separated string of options which should be passed to `gcj` when in JIT mode. If not set, a sensible default is chosen.

`gnu.gcj.jit.cachedir`
: This is the directory where cached shared library files are stored. If not set, JIT compilation is disabled. This should never be set to a directory that is writable by any other user.

`gnu.gcj.precompiled.db.path`
: This is a sequence of file names, each referring to a file created by `gcj-dbtool`. These files will be used by `libgcj` to find shared libraries corresponding to classes that are loaded from bytecode. `libgcj` often has a built-in default database; it can be queried using `gcj-dbtool -p`.

14 Resources

While writing `gcj` and `libgcj` we have, of course, relied heavily on documentation from Sun Microsystems. In particular we have used The Java Language Specification (both first and second editions), the Java Class Libraries (volumes one and two), and the Java Virtual Machine Specification. In addition we've used the online documentation at `http://java.sun.com/`.

The current `gcj` home page is `http://gcc.gnu.org/java/`.

For more information on gcc, see `http://gcc.gnu.org/`.

Some `libgcj` testing is done using the Mauve test suite. This is a free software Java class library test suite which is being written because the JCK is not free. See `http://sources.redhat.com/mauve/` for more information.