

GTO: The Kitchen Sink of Data

File Format, Protocols, and Utilities.

Jim Hourihan, Tweak Films

Copyright © 2002-2007 Tweak Films. All rights reserved. Permission is granted to make and distribute verbatim copies of this manual provided the copyright notice and this permission notice are preserved on all copies.

Permission is granted to copy and distribute modified versions of this manual under the conditions for verbatim copying, provided that the entire resulting derived work is distributed under the terms of a permission notice identical to this one.

Table of Contents

GNU LESSER GENERAL PUBLIC LICENSE and GTO EXCEPTIONS	1
1 Installation	8
1.1 GNU/Linux.....	9
1.2 Mac OS X	9
1.3 Microsoft Windows 2000/XP/Vista.....	10
2 Overview	10
2.1 Use in Production.....	11
2.2 Use in Games	11
2.3 GTO as a 3D Scene File Format.....	11
2.4 Using GTO to Store Particles	12
2.5 Caching Data for Rendering or Animation.....	12
2.6 Storing Movie Frames in GTO for review	12
3 Binary Format	12
4 Text Format	14
4.1 Example of a Cube Stored as a Text GTO.....	15
4.2 How Strings are Handled in the Text Format	16
4.3 Value Brackets.....	17
4.4 The Size of a Property	17
4.5 Run Length Encoding of Values	18
4.6 Syntax Reference.....	18
5 Types of Property Data.....	20
6 Interpretation Strings.....	21
7 Object Protocols	23
7.1 Object Protocol.....	23
7.2 Coordinate System Protocol.....	24
7.3 Particle Protocol.....	25
7.4 Strand Protocol.....	26
7.5 NURBS Protocol.....	27
7.6 Polygon Protocol.....	27
7.7 Subdivision Surface Protocols	30
7.8 Image Protocol.....	30
7.8.1 Additional Image Properties Used by GTV Files.....	32

7.9	Material Protocol	32
7.10	Group Protocol	32
7.11	Inter-Object Connection Protocol	32
7.11.1	Transformation hierarchies	33
7.11.2	Material Assignment	33
7.11.3	Container Assignment	33
7.12	Difference File Protocol	33
7.13	Sorted Shell File Protocol	34
7.14	Channels Protocol	34
7.14.1	Example	34
7.15	Animation Curve Protocol	35
7.15.1	Example	35
8	Naming Conventions	36
8.1	Valid Names	36
8.2	Exactly Specifying a Property or Component	36
8.3	Indicating Special Handling	36
8.4	Cross References Encoded in Names	37
9	Issues and Questionable Aspects of the Format	37
10	Extending Protocols or the File Format....	38
11	C++ Library	38
11.1	Gto::Reader class	39
11.2	Gto::Writer class	42
11.3	Gto::RawDataReader/Gto::RawDataWriter classes	45
12	Python Module	46
12.1	gto.Reader	46
12.2	gto.Writer	48
12.3	Classes used by gto.Reader	49
13	Utilities	50
13.1	The gtoinfo Utility	50
13.2	The gtofilter Utility	51
13.3	The gtomerge Utility	52
13.4	The gto2obj Utility	52
13.5	The gtoimage Utility	53
13.6	The RiGtoRibOut Utility	53
13.7	The gtoIO.so Maya Plug-In	53
13.7.1	BUGS	54
13.8	The RiGtoPlugin RenderMan plugin	54
13.8.1	RIB Instantiation	54
13.8.2	Config String Syntax	54

13.8.3	On-List/Off-List Syntax	55
13.8.4	Cache Management	55
13.8.5	Environment Variables	55
13.8.6	Usage Strategy	56
13.8.7	Miscellaneous RenderMan Stuff	56
Appendix A	Description of Changes	57
Appendix B	Reference.....	59

GNU LESSER GENERAL PUBLIC LICENSE and GTO EXCEPTIONS

Version 2.1, February 1999

Copyright © 1991, 1999 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.

[This is the first released version of the Lesser GPL. It also counts as the successor of the GNU Library Public License, version 2, hence the version number 2.1.]

Exceptions to the LGPL for GTO (Nov 2007)

The GTO library and included programs are provided under the terms of the GNU Library General Public License (LGPL) as put forth below with the following exceptions (derived from the FLTK LGPL exceptions):

1. Modifications to the GTO configure script, config header file, and makefiles by themselves to support a specific platform do not constitute a modified or derivative work.

The authors do request that such modifications be contributed to the GTO project.

2. Static linking of applications or shared libraries to the GTO library does not constitute a derivative work and does not require the author to provide source code for the application or library, use the shared GTO libraries, or link their applications or shared library against a user-supplied version of GTO.

If you link the application or shared library to a modified version of GTO, then the changes to GTO must be provided under the terms of the LGPL in sections 1, 2, and 4.

3. You do not have to provide a copy of the GTO license with programs that are linked to the GTO library, nor do you have to identify the GTO license in your program or documentation as required by section 6 of the LGPL.

However, programs must still identify their use of GTO. The following example statement can be included in user documentation to satisfy this requirement:

[program/shared library] is based in part on the work of the GTO project.

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public Licenses are intended to guarantee your freedom to share and change free software—to make sure the software is free for all its users.

This license, the Lesser General Public License, applies to some specially designated software packages—typically libraries—of the Free Software Foundation and other authors who decide to use it. You can use it too, but we suggest you first think carefully about whether this license or the ordinary General Public License is the better strategy to use in any particular case, based on the explanations below.

When we speak of free software, we are referring to freedom of use, 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 and use pieces of it in new free programs; and that you are informed that you can do these things.

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

For example, if you distribute copies of the library, whether gratis or for a fee, you must give the recipients all the rights that we gave you. You must make sure that they, too, receive or can get the source code. If you link other code with the library, you must provide complete object files to the recipients, so that they can relink them with the library after making changes to the library and recompiling it. And you must show them these terms so they know their rights.

We protect your rights with a two-step method: (1) we copyright the library, and (2) we offer you this license, which gives you legal permission to copy, distribute and/or modify the library.

To protect each distributor, we want to make it very clear that there is no warranty for the free library. Also, if the library is modified by someone else and passed on, the recipients should know that what they have is not the original version, so that the original author's reputation will not be affected by problems that might be introduced by others.

Finally, software patents pose a constant threat to the existence of any free program. We wish to make sure that a company cannot effectively restrict the users of a free program by obtaining a restrictive license from a patent holder. Therefore, we insist that any patent license obtained for a version of the library must be consistent with the full freedom of use specified in this license.

Most GNU software, including some libraries, is covered by the ordinary GNU General Public License. This license, the GNU Lesser General Public License, applies to certain designated libraries, and is quite different from the ordinary General Public License. We use this license for certain libraries in order to permit linking those libraries into non-free programs.

When a program is linked with a library, whether statically or using a shared library, the combination of the two is legally speaking a combined work, a derivative of the original library. The ordinary General Public License therefore permits such linking only if the entire combination fits its criteria of freedom. The Lesser General Public License permits more lax criteria for linking other code with the library.

We call this license the "Lesser" General Public License because it does Less to protect the user's freedom than the ordinary General Public License. It also provides other free software developers Less of an advantage over competing non-free programs. These disadvantages are the reason we use the ordinary General Public License for many libraries. However, the Lesser license provides advantages in certain special circumstances.

For example, on rare occasions, there may be a special need to encourage the widest possible use of a certain library, so that it becomes a de-facto standard. To achieve this, non-free programs must be allowed to use the library. A more frequent case is that a free library does the same job as widely used non-free libraries. In this case, there is little to

gain by limiting the free library to free software only, so we use the Lesser General Public License.

In other cases, permission to use a particular library in non-free programs enables a greater number of people to use a large body of free software. For example, permission to use the GNU C Library in non-free programs enables many more people to use the whole GNU operating system, as well as its variant, the GNU/Linux operating system.

Although the Lesser General Public License is Less protective of the users' freedom, it does ensure that the user of a program that is linked with the Library has the freedom and the wherewithal to run that program using a modified version of the Library.

The precise terms and conditions for copying, distribution and modification follow. Pay close attention to the difference between a "work based on the library" and a "work that uses the library". The former contains code derived from the library, whereas the latter must be combined with the library in order to run.

GNU LESSER GENERAL PUBLIC LICENSE

TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License Agreement applies to any software library or other program which contains a notice placed by the copyright holder or other authorized party saying it may be distributed under the terms of this Lesser General Public License (also called "this License"). Each licensee is addressed as "you".

A "library" means a collection of software functions and/or data prepared so as to be conveniently linked with application programs (which use some of those functions and data) to form executables.

The "Library", below, refers to any such software library or work which has been distributed under these terms. A "work based on the Library" means either the Library or any derivative work under copyright law: that is to say, a work containing the Library or a portion of it, either verbatim or with modifications and/or translated straightforwardly into another language. (Hereinafter, translation is included without limitation in the term "modification".)

"Source code" for a work means the preferred form of the work for making modifications to it. For a library, 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 library.

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running a program using the Library is not restricted, and output from such a program is covered only if its contents constitute a work based on the Library (independent of the use of the Library in a tool for writing it). Whether that is true depends on what the Library does and what the program that uses the Library does.

1. You may copy and distribute verbatim copies of the Library's complete 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 distribute a copy of this License along with the Library.

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 Library or any portion of it, thus forming a work based on the Library, 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. The modified work must itself be a software library.
 - b. You must cause the files modified to carry prominent notices stating that you changed the files and the date of any change.
 - c. You must cause the whole of the work to be licensed at no charge to all third parties under the terms of this License.
 - d. If a facility in the modified Library refers to a function or a table of data to be supplied by an application program that uses the facility, other than as an argument passed when the facility is invoked, then you must make a good faith effort to ensure that, in the event an application does not supply such function or table, the facility still operates, and performs whatever part of its purpose remains meaningful.

(For example, a function in a library to compute square roots has a purpose that is entirely well-defined independent of the application. Therefore, Subsection 2d requires that any application-supplied function or table used by this function must be optional: if the application does not supply it, the square root function must still compute square roots.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Library, 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 Library, 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 Library.

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

3. You may opt to apply the terms of the ordinary GNU General Public License instead of this License to a given copy of the Library. To do this, you must alter all the notices that refer to this License, so that they refer to the ordinary GNU General Public License, version 2, instead of to this License. (If a newer version than version 2 of the ordinary GNU General Public License has appeared, then you can specify that version instead if you wish.) Do not make any other change in these notices.

Once this change is made in a given copy, it is irreversible for that copy, so the ordinary GNU General Public License applies to all subsequent copies and derivative works made from that copy.

This option is useful when you wish to copy part of the code of the Library into a program that is not a library.

4. You may copy and distribute the Library (or a portion or derivative of it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you 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.

If distribution of 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 satisfies the requirement to distribute the source code, even though third parties are not compelled to copy the source along with the object code.

5. A program that contains no derivative of any portion of the Library, but is designed to work with the Library by being compiled or linked with it, is called a "work that uses the Library". Such a work, in isolation, is not a derivative work of the Library, and therefore falls outside the scope of this License.

However, linking a "work that uses the Library" with the Library creates an executable that is a derivative of the Library (because it contains portions of the Library), rather than a "work that uses the library". The executable is therefore covered by this License. Section 6 states terms for distribution of such executables.

When a "work that uses the Library" uses material from a header file that is part of the Library, the object code for the work may be a derivative work of the Library even though the source code is not. Whether this is true is especially significant if the work can be linked without the Library, or if the work is itself a library. The threshold for this to be true is not precisely defined by law.

If such an object file uses only numerical parameters, data structure layouts and accessors, and small macros and small inline functions (ten lines or less in length), then the use of the object file is unrestricted, regardless of whether it is legally a derivative work. (Executables containing this object code plus portions of the Library will still fall under Section 6.)

Otherwise, if the work is a derivative of the Library, you may distribute the object code for the work under the terms of Section 6. Any executables containing that work also fall under Section 6, whether or not they are linked directly with the Library itself.

6. As an exception to the Sections above, you may also combine or link a "work that uses the Library" with the Library to produce a work containing portions of the Library, and distribute that work under terms of your choice, provided that the terms permit modification of the work for the customer's own use and reverse engineering for debugging such modifications.

You must give prominent notice with each copy of the work that the Library is used in it and that the Library and its use are covered by this License. You must supply a copy of this License. If the work during execution displays copyright notices, you must include the copyright notice for the Library among them, as well as a reference directing the user to the copy of this License. Also, you must do one of these things:

- a. Accompany the work with the complete corresponding machine-readable source code for the Library including whatever changes were used in the work (which must be distributed under Sections 1 and 2 above); and, if the work is an executable linked with the Library, with the complete machine-readable "work that uses the Library", as object code and/or source code, so that the user can modify the

Library and then relink to produce a modified executable containing the modified Library. (It is understood that the user who changes the contents of definitions files in the Library will not necessarily be able to recompile the application to use the modified definitions.)

- b. Use a suitable shared library mechanism for linking with the Library. A suitable mechanism is one that (1) uses at run time a copy of the library already present on the user's computer system, rather than copying library functions into the executable, and (2) will operate properly with a modified version of the library, if the user installs one, as long as the modified version is interface-compatible with the version that the work was made with.
- c. Accompany the work with a written offer, valid for at least three years, to give the same user the materials specified in Subsection 6a, above, for a charge no more than the cost of performing this distribution.
- d. If distribution of the work is made by offering access to copy from a designated place, offer equivalent access to copy the above specified materials from the same place.
- e. Verify that the user has already received a copy of these materials or that you have already sent this user a copy.

For an executable, the required form of the "work that uses the Library" must include any data and utility programs needed for reproducing the executable from it. However, as a special exception, the materials to be 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.

It may happen that this requirement contradicts the license restrictions of other proprietary libraries that do not normally accompany the operating system. Such a contradiction means you cannot use both them and the Library together in an executable that you distribute.

7. You may place library facilities that are a work based on the Library side-by-side in a single library together with other library facilities not covered by this License, and distribute such a combined library, provided that the separate distribution of the work based on the Library and of the other library facilities is otherwise permitted, and provided that you do these two things:
 - a. Accompany the combined library with a copy of the same work based on the Library, uncombined with any other library facilities. This must be distributed under the terms of the Sections above.
 - b. Give prominent notice with the combined library of the fact that part of it is a work based on the Library, and explaining where to find the accompanying uncombined form of the same work.
8. You may not copy, modify, sublicense, link with, or distribute the Library except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, link with, or distribute the Library 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.

9. 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 Library or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Library (or any work based on the Library), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Library or works based on it.
10. Each time you redistribute the Library (or any work based on the Library), the recipient automatically receives a license from the original licensor to copy, distribute, link with or modify the Library 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 with this License.
11. 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 Library at all. For example, if a patent license would not permit royalty-free redistribution of the Library 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 Library.

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.

12. If the distribution and/or use of the Library is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Library 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.
13. The Free Software Foundation may publish revised and/or new versions of the Lesser 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 Library 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 Library does not specify a license version number, you may choose any version ever published by the Free Software Foundation.

14. If you wish to incorporate parts of the Library into other free programs whose distribution conditions are incompatible with these, 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

15. BECAUSE THE LIBRARY IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE LIBRARY, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE LIBRARY "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 LIBRARY IS WITH YOU. SHOULD THE LIBRARY PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.
16. 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 LIBRARY 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 LIBRARY (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 LIBRARY TO OPERATE WITH ANY OTHER SOFTWARE), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

1 Installation

The GTO source code is built using the GNU autotools. When you unpack the source code, the following should successfully build and install the binaries and library. Note that you may need to become root for `make install` to succeed.

```
sh> cd gto-x.x
sh> ./configure
sh> make
sh> make install
```

If you have Maya installed, the Maya I/O plugin can be compiled. You can give the `--enable-maya-plugin` argument to the configure script to have the plugin built. See the `INSTALL` file at the top of the source code tree for more information on using configure.

If you have zlib installed, configure will automatically assume you want native compression support for the Reader/Writer library.

1.1 GNU/Linux

The Maya plugin almost certainly requires a specific compiler for GNU/Linux installation with Maya 4.x or 3.x. On versions of Red Hat Linux 7.x the default compiler should be used; this should be gcc 2.96. On later versions of Red Hat, 8.x and 9.x, you will need to install the following prior to Maya 7:

```
compat-gcc-7.3-2.96.110.i386.rpm
compat-gcc-c++-7.3-2.96.110.i386.rpm
compat-libstdc++-7.3-2.96.110.i386.rpm
compat-libstdc++-devel-7.3-2.96.110.i386.rpm
```

In that case, the compiler will be called gcc296 and you will need to do this before running configure:

```
sh> CC=gcc296
sh> export CC
sh> CXX=g++296
sh> export CXX
```

If you are using Maya 5.x you will need to use gcc 3.1 or greater to compile the plugin. In that case no compatibility rpms are necessary. Alias recommends using gcc 3.2.

For Maya 6.x, 7.x use gcc 3.3.2 or a newer 3.3.

For Maya 8.x use gcc 4.0.2. You may need to compile your own version of the compiler even if the same version is shipped with the OS. See the Autodesk documentation regarding Maya 8.x plugins.

If you are not using the Maya plugin, or you want to use the GTO library, but using a different compiler, you may need to run the build and installation process more than once with different compiler settings.

On non Red Hat or home grown GNU/Linux systems, the library and binary tools should build and install using the normal procedure. Here at Tweak, we typically build the library and binary tools using gcc 3.1 or greater.

1.2 Mac OS X

The GTO library and binaries should build normally on OS X if you have installed the developer tools. You should have a version of the developer tools that uses gcc 3.1 or better.

The Maya plugin on Mac OS X requires the use of CodeWarrior to compile a CFM shared object for Maya versions 4.x and 5.x. In addition, for the 4.x versions, you will need to hack the CodeWarrior system header files to get the build to work. We have not had much success in getting the plugin to compile for OS X since the header munging instructions did not work for us. However, there is no reason to suspect that given the proper CodeWarrior version and modifications that it would not compile.

For Maya 6.x and greater, you can use the installed version of gcc since Maya is now a Mach-O binary.

For Maya 8.x on Intel and PPC with 10.4.X use the installed gcc 4 compiler and follow the instructions in the Maya developers docs. Compiling for Maya 8.x is much easier on OS X than with previous versions. Check for precompiled plugins for OS X intel at the tweak website.

1.3 Microsoft Windows 2000/XP/Vista

In theory, you can compile the library and binary tools using the Cygwin environment and gcc 3.x but we have not tested it.

Maya on Windows 2000 requires the use of Visual C++ 6.x. This compiler does not come close to complying with the ANSI C++ standard and requires many work arounds to compile the GTO library and Maya plugin. Because of this, the library and tools are currently unsupported on Windows.

Your mileage may vary for other Windows compilers.

2 Overview

Historically, GTO format's primary usage is storage of static geometric data. As such, the types of data you might find in a GTO file are things like polygonal meshes, various types of subdivision surfaces, NURBS or UBS surfaces, coordinate systems, hierarchies of objects, material bindings, and even images.

In our facility – Tweak Films – geometric data in GTO files is rendered directly by entropy (a prman compatible rib renderer) and used for modeling and 3D paint. We also use the GTO format in place of the venerable and yet nasty PDB file format for particles.

As of GTO version 3.2, the format is being used for data not related to geometric data (e.g., pixel data) and animation curve data is now part of the format spec – so that data is no longer “static” as it has been in the past.

From a historic point of view, the GTO file format is most closely related to the original inventor file format, the Stanford PLY format and the PDB particle format. Like the Wavefont PDB file format, there are a limited number of simple GTO data types (float, int, string, boolean). Like the inventor file format, a GTO can hold an entire transformation hierarchy including geometric leaf nodes. Like the PLY format, the GTO format can contain an arbitrary amount of data per primitive. Most importantly however, the GTO file format is intended to be very OBJ-like; its relatively easy to read and write and easy to ignore data you don't want or know about.

GTO files can be either binary or text files. Binary files are the preferred format for large data sets. The GTO text format is intended to be human readable/editable; the syntax is simple and concise. The text format is useful when storing “bag of parameters” files and similar data.

We provide code for both the binary and text formats to give implementors a head start. The binary file is either big or little endian on disk, but should be readable on any platform.

The GTO reader we provide can be use libz, so it can read and write compressed files natively. We find that compressed GTO files created by most 3D programs are approximately 60% leaner than uncompressed files.

GTO files conceptually contain *objects* which are composed of *components*. Components are further composed of *properties*. A property contains an array of one of the predefined data types with a specified width. For example, you might have an object which looks something like this:

```
Object "cube"
  Component "points"
    Property float[3][8] "position"
    Property float[1][8] "mass"
  Component "elements"
    Property byte[1][8] "type"
    Property short[1][8] "size"
  Component "indices"
    Property int[1][32] "vertex"
```

Using the terminology above, the object “cube” contains five properties: **position**, **mass**, **type**, **size**, and **vertex**. The **points** component describes the points that make up the cube vertices. Each point has a position and mass stored in properties of the same name. The position property data is composed of eight float[3] data items (or 8 3D points). The **mass** property is composed of a 8 scalar floating point values (one for each point).

The **elements** component contains two properties. **type** indicates the type of the element (for example, triangle, quad, or triangle strip). In this case the elements might all be quads. **size** indicates the number of vertices in each of the eight faces (elements) of the cube – (4 for a cube). The **vertex** property of the **indices** component contains the actual indices: 4 per face for a total of 32.

Of course you could store much more data with the cube object if you wanted to. For example, if you wanted velocity or color per point, this would be another property in the **points** component.

The meaning of this data is another story altogether. Its all handled by protocol. One application may store things in the GTO file that another application has no method of interpreting even though it can read that data and modify it. In the example above, you need to know to expect that polygonal data is stored in the given properties. The same data could be stored with different property names and a more complex layout. (The “polygon” protocol described later in this document is different and more involved than the above example.)

2.1 Use in Production

Film

2.2 Use in Games

Games

2.3 GTO as a 3D Scene File Format

3d format

2.4 Using GTO to Store Particles

particles

2.5 Caching Data for Rendering or Animation

caching

2.6 Storing Movie Frames in GTO for review

movie format

3 Binary Format

The GTO file has six major sections which appear in the following order.

1. **Header** (Gto::Header). The header structure contains the GTO magic number (used to determine endianness), the version of the GTO specification that the file was written as, and the number of top level objects in the file. There is one instance of a header in the file. Finally, the header indicates how many strings are in the string table.

```

Magic = 0x0000029f;
Cigam = 0x9f020000; // means the file is opposite endianness

struct Header
{
    uint32    magic;
    uint32    numStrings;
    uint32    numObjects;
    uint32    version;
    uint32    flags;           // reserved;
};

```

2. **String Table**. After the header, null terminated strings are written in the file. The order of these strings is important. All names and string properties store indices into the string table instead of actual strings. In order to read the file properly, the string table must be available until the file is completely read. (Unless you don't care about any strings!)

The index number refers the string number in the table not its byte offset. So the string index 9 (for example) refers to the 10th string in the table (string index 0 is the first string in the table).

3. **ObjectHeader** (Gto::ObjectHeader). The object header indicates what kind of protocol to use to interpret it, the object name and the number of components. (More on the **object** protocol later). The name – like all strings in the GTO file – is stored as a string table entry. If the file header indicated N objects in the file, there will be N ObjectHeaders.

```

struct ObjectHeader

```

```

    {
        uint32    name;           // a string table index
        uint32    protocolName;  // a string table index
        uint32    protocolVersion;
        uint32    numComponents;
        uint32    pad;           // unused
    };

```

4. **ComponentHeader** (Gto::ComponentHeader). Like the ObjectHeaders the ComponentHeaders will appear together for all objects in order. The component header indicates the number of properties in the component and the name of the component.

```

enum ComponentFlags
{
    Transposed = 1 << 0,
    Matrix     = 1 << 1,
};

struct ComponentHeader
{
    uint32    name;           // a string table index
    uint32    numProperties;
    uint32    flags;
    uint32    interpretation; // a string table index
    uint32    pad;           // unused
};

```

5. **PropertyHeader** (Gto::PropertyHeader). The PropertyHeaders, like the object and component headers, appear en masse in the file. The PropertyHeader contains the name, size, type, and dimension of the property.

```

enum DataType
{
    Int,           // int32
    Float,        // float32
    Double,       // float64
    Half,         // float16
    String,       // string table indices
    Boolean,      // bit
    Short,       // uint16
    Byte,        // uint8
    Int64        // int64
};

```

```

struct PropertyHeader

```

```

    {
        uint32    name;           // string table index
        uint32    type;          // DataType enum value
        uint64    size;
        uint64    width;
        uint32    interpretation; // string table index
        uint32    pad;           // unused
    };

```

6. **Data.** The last section of the file contains all of the property data. The beginning and end of a properties data are not marked. The size must be consistent with the description of the property used in the PropertyHeader.

In (Text) diagram form the file looks something like this:

```

+-----+
| File Header |
+-----+
| String Table |
+-----+
| Object Header |
| . |
| . |
| . |
+-----+
| Component Header |
| . |
| . |
| . |
+-----+
| Property Header |
| . |
| . |
| . |
+-----+
| Property Data |
| . |
| . |
| . |
+-----+

```

4 Text Format

As of version 4, GTO has a text representation in addition to the binary representation. The text representation is designed for human use; it is intended to be easy to modify or

create from scratch in a text editor. It is not intended to compete with XML formats (which are typically only human readable in theory) nor is it intended to be used in place of the binary format which is much faster and more economical for storage of large data sets.

4.1 Example of a Cube Stored as a Text GTO

Here's the example from the overview section: a cube stored using the “polygon” protocol:

```
GTOa (4)

# this is a comment

cube : polygon (2)
{
  points
  {
    float[3] position = [ [ -2.5 2.5 2.5 ]
                          [ -2.5 -2.5 2.5 ]
                          [ 2.5 -2.5 2.5 ]
                          [ 2.5 2.5 2.5 ]
                          [ -2.5 2.5 -2.5 ]
                          [ -2.5 -2.5 -2.5 ]
                          [ 2.5 -2.5 -2.5 ]
                          [ 2.5 2.5 -2.5 ] ]

    float mass = [ 1 1 1 1 1 1 1 1 ]
  }

  elements
  {
    byte type = [ 2 2 2 2 2 2 ]
    short size = [ 4 4 4 4 4 4 ]
  }

  indices
  {
    int vertex = [ 0 1 2 3
                  7 6 5 4
                  3 2 6 7
                  4 0 3 7
                  4 5 1 0
                  1 5 6 2 ]
  }
}
```

The first line of the file is an identifier to tell the parser what variety of GTO file it is: in this case `GTOa` which indicates a plain ASCII text file. Currently the parser can only handle ASCII encoding; a forthcoming version will allow UTF-8.

Objects are declared using the syntax:

```

OBJECTNAME [ : PROTOCOL [ (PROTOCOL_VERSION) ] ]
{
    ... object contents ...
}

```

The brackets enclose optional syntax. So the `PROTOCOL_VERSION` (including the parens) is optional. The `PROTOCOL` is also optional; if omitted (along with the colon) the protocol defaults to **object**. In the example, “cube” is the name of the object and “polygon” is the name of the protocol—the protocol version is 2.

Components must be declared inside the object brackets. The brackets denote a *namespace* which is either an object namespace or a component namespace. Component namespaces must always be declared inside of an object namespace. Object namespaces can only appear at the top level of the file; in other words, objects cannot be inside another namespace.

Components are declared like this:

```

COMPONENTNAME [as INTERPRETATION]
{
    ... component contents ...
}

```

The `INTERPRETATION` can be any string. Properties can be declared inside of the component namespace. The property declaration is the most flexible; since some aspects of the property (like its size) can be determined by the parser from the property data, you can omit them.

The property syntax in its most general form is:

```

TYPE[WIDTH] [SIZE] PROPERTYNAME as INTERPRETATION = values ...

```

The brackets around `WIDTH` and `SIZE` are literal in this case; they actually appear in the file. As you can see from the example, some of the property declaration syntax is optional. The `SIZE` can usually be determined from the values so it may be omitted. The `WIDTH` is assumed to be 1 if it is omitted. The `as INTERPRETATION` section of the declaration may also be omitted.

What cannot be omitted is the `TYPE`, `PROPERTYNAME`, and the assignment of values.

4.2 How Strings are Handled in the Text Format

With the exception of keywords and type names, any string in the text GTO file can be either be quoted or non-quoted. Non-quoted strings are restricted to strings which do not represent numbers. In addition, if a string contains punctuation or whitespace, it must be quoted. For example, if the name of the object in the cube example was “four dimensional time-cube” it would have to be declared like this:

```

"four dimensional time-cube" : polygon
{
    ...
}

```

There is one additional exception: if a string is also a keyword or type name, it must be quoted. For example, here’s an exceptional property declaration:

```
int "int" as "as" = 1
```

In this case the quoted string “int” is being used as the property name, but because it is also the name of a GTO type, it must be quoted. The string “as” is being used as an interpretation string and must be quoted because “as” is also a keyword in the the GTO file.

When in doubt quote.

4.3 Value Brackets

Generally, a property value and elements of the value are enclosed in brackets:

```
TYPE[WIDTH] PROPERTYNAME = [ [a b ...] [d e ...] ... ]
```

In this documentation, the *value* of a property is everything to the right of the “=” and an *element* is a fixed size collection of numbers or strings. The *size* of a property is the number of elements in its value. So in the example above, the [a b ...] portion of the syntax is an *element*.

Bracketing the property value is optional in one circumstance: when the number of elements in the property value is one. For example, these declarations are equivalent:

```
int foo = 1
int foo = [1]
```

If the width of the type is not one (elements are not scalar), then brackets must be put around each element of the property. If the size is one but the width is not one, then the enclosing brackets are still optional:

```
int[2] foo = [1 2]
int[2] foo = [ [1 2] ]
```

If however the size of the property is greater than one, the enclosing value brackets are required:

```
# property of size 3
int[2] foo = [ [1 2] [3 4] [5 6] ]
```

To declare a property with no value use empty brackets:

```
int foo = []
```

4.4 The Size of a Property

The size of a property can be declared as part of its type declaration:

```
int[1][4] foo = [1 2 3 4]
```

In this case, “foo” contains four scalar elements. Because the size was specified, the following would be a syntax error:

```
int[1][4] foo = [1 2 3 4 5]
```

The parser would complain because five elements were supplied eventhough the property was declared as having only four. If no size is specified than the parser will determine the size from the number of elements in the value:

```
int[1] foo = [1 2 3 4 5]
```

So in this case “foo” has five elements. Note that in order to declare the size specifically, you must also declare the element width – even if the width is one. In the last example, because we did not specify the size, the declaration could also have been:

```
int foo = [1 2 3 4 5]
```

In this case it is understood that the type is actually `int [1] [5]`.

4.5 Run Length Encoding of Values

In some cases, a value will contain many copies of an element. There is a special syntax for these cases; you can use an ellipsis to indicate that all remaining elements are identical. The ellipsis can only appear directly before the final bracket character.

There is one restriction when using this syntax: the type of the property value must be completely specified (including the size of the property) and the value must be enclosed in brackets. For example:

```
int[1][100] mass = [1 ...]
```

The ellipsis is literal (its actually in the file as three dot characters) The property “mass” will be one for all 100 elements. If the element has a width greater than one:

```
float[3][100] velocity = [ [0 0 0] ... ]
```

The ellipsis is used in place of an element. The following will *not* work:

```
float[3][100] velocity = [ [0 ...] ... ]
```

The intention here is to make all of the velocity elements `[0 0 0]`. However, this syntax is not correct and will produce a parsing error.

4.6 Syntax Reference

The grammar for the text GTO file. *INT* is an integer constant. *FLOAT* is a floating point constant, with a possible exponent part. *STRING* is either a quoted or non-quoted string. All other values are literal. Double quoted strings are keywords.

```
file::
    "GTOa" object_list
    "GTOa" ( INT ) object_list

object_list::
    object
    object_list object

object::
    STRING { component_list }
    STRING : STRING { component_list }
    STRING : STRING ( INT ) { component_list }

component_list::
    component
    component_list component

interp_string_opt::
    nothing
    "as" STRING
```

```
component::
    STRING interp_string_opt { property_list }

property_list::
    property
    property_list property

property::
    type STRING interp_string_opt = atomic_value
    type STRING interp_string_opt = [ complex_element_list ]

type::
    basic_type
    basic_type [ INT ]
    basic_type [ INT ] [ INT ]

basic_type::
    "float" "int" "string" "short" "byte" "half"
    "bool" "double"

complex_element_list::
    nothing
    element_list
    element_list "... "

element_list::
    element
    element_list element

element::
    atomic_value
    [ atomic_value_list ]

atomic_value_list::
    string_value_list
    numeric_value_list

atomic_value::
    string_value
    numeric_value

string_value_list::
    string_value
    string_value_list string_value

string_value::
    STRING
```

```

numeric_value_list::
    numeric_value
    numeric_value_list numeric_value

numeric_value::
    float_num
    int_num

float_num::
    FLOAT
    - FLOAT

int_num::
    INT
    - INT

```

5 Types of Property Data.

The GTO format pre-defines a small number of data types that can be stored as properties. The currently defined types are:

double	[Property Type]
64 bit IEEE floating point.	
float	[Property Type]
32 bit IEEE floating point.	
half	[Property Type]
16 bit floating point number as defined by ILM's OpenEXR Imath library half data type. You can find links to source code and documentation at http://www.openexr.org/ .	
int	[Property Type]
32 bit signed integer.	
int64	[Property Type]
64 bit signed integer.	
short	[Property Type]
16 bit unsigned integer.	
byte	[Property Type]
8 bit unsigned integer (char).	
bool	[Property Type]
Bit or bit vector. Not currently implemented.	

string [Property Type]

The string type is stored as a 32 bit integer index into the GTO file's string table. So storing a lot of strings (especially if there is a lot of redundancy) is reasonably cheap. All strings in the GTO file are stored in this manner.

Each of these data types can be made into a vector of that type. For example the float data type can be made into a point `float[3]` or a matrix `float[16]`. To store a scalar element the size of the vector is 1. (e.g. `float[1]`).

In this document, the types are all specified as 2 dimensional arrays ala the C programming language. Here is a complete list of example type forms:

- `float[3]` - the float triple type.
- `float[1][1]` - a single floating point number.
- `float[3][]` - any number of float triples.
- `float[3][3]` - three float triples.
- `float[16][]` - any number of a 16 float (4x4 matrix) elements.

6 Interpretation Strings

Each property can have an additional string stored with it call the "interpretation". The intent is to allow applications to provide specific information about the property. For example, a property of type `float[4]` can be interpreted as a homogeneous 3D coordinate, a quaternion, or an RGBA value. The interpretation field can be used to distinguish between them.

Why not just make new primitive GTO types for these? The format's only purpose is storage of data. By decoupling the interpretation of the data from its storage, each application is allowed to make its own policy while maintaining flexibility for simpler applications.

Here's a simple example of `gtoinfo` output of a file with an image object in it created with `gtoimage`:

```
object "image" protocol "image" v1
  component "image"
    property string[1][1] "originalFile" interpret as "filename"
    property string[1][1] "originalEncoding" interpret as "filetype"
    property string[1][1] "type"
    property int[1][2] "size"
    property float[3][199168] "pixels" interpret as "RGB"
```

The example illustrates two points: the interpretation string can be used to better determine the data type of the property (as is the case with the "RGB" string) and it can also be used to interpret the usage (the "filename").

Some of the strings will be application specific. Programs that generically edit GTO files should attempt to preserve the interpretation strings.

It is not an error to define the interpretation for a property as the empty string – in other words, unspecified.

The following strings are not currently part of the format specification but are used by the sample implementation. In a future release we may make these "official". Its ok to have multiple space separated strings in the interpretation strings (e.g. "4x4 row-major").

- coordinate** [Interpretation String]
The data can be of any width or type. For width N the data represents a point in N dimensional space.
- normal** [Interpretation String]
The data can be of any width or type. For width N the data represents a unit vector perpendicular to an N dimensional surface or in the case of $N == 2$, a curve.
- 4x4** [Interpretation String]
The width of the property data should be 16. The data is intended to be interpreted as a 4x4 matrix. For example, the **object.globalMatrix** property of the **Coordinate System** protocol would be a “4x4” property.
- 3x3** [Interpretation String]
The width of the property data should be 9. The data is intended to be interpreted as a 3x3 matrix.
- row-major** [Interpretation String]
Indicates that matrix data is in row major ordering.
- column-major** [Interpretation String]
Indicates that matrix data is in column major ordering.
- quaternion** [Interpretation String]
The width of the property should be four. The data should be interpreted as a quaternion. Presumably the type of a quaternion property would be `float[4]` or `double[4]` since these are the only types that make sense. The first element of the data is the real part followed by the “i”, “j”, and “k” imaginary components.
- complex** [Interpretation String]
The width of the property should be two. The data is interpreted as a complex number with the first element being the real part and the second element the imaginary part.
- indices** [Interpretation String]
The data type should be an integral type. The property contains indices.
- bbox** [Interpretation String]
The data type should have an even width. The property contains bounding boxes.
- homogeneous** [Interpretation String]
The width of the property should be two or more. If the width is three, then the data is a two dimensional homogeneous coordinate. If the width is four, then the data is a three dimensional homogeneous coordinate. So for data of width N the data represents a homogeneous coordinate in N-1 dimensions.
- RGB** [Interpretation String]
The width of the property should be three. The data represents a color with red, green, and blue components.
- BGR** [Interpretation String]
The width of the property should be three. The data represents a color with blue, green, and red components. (Reversed RGB)

RGBA [Interpretation String]
 The width of the property should be four. The data represents a color (or pixel) with red, green, blue, and alpha components.

ABGR [Interpretation String]
 The width of the property should be four. The data represents a color (or pixel) with alpha, blue, green, and red components. (Reversed RGBA)

bezier [Interpretation String]
 The property represents a 2D bezier curve for animation. The type should be a floating point type and the width six. Each element is a key frame value.

weighted [Interpretation String]
 In the case of bezier animation curves, the curve should be evaluated with weighted tangents.

7 Object Protocols

The Object data interpretation is not defined by the GTO format. However, there are currently some protocols in use that are well defined and these are documented here. Caveat emptor: gto files in the wild may contain more data than these protocols define, but they presumably will obey the protocol if they indicate it by name. It's also possible that some objects may obey more than one protocol yet only indicate that they follow one. Unfortunately, some protocols also specify optional components and properties in case all of this was not confusing enough.

Protocols also have a version number. The version number is an integer; there are no sub-versions. If there are significant changes to a protocol, the version number should be bumped. The version number is not meant as a method of making alternate protocols with the same name. We have had to make three modifications to the protocols since the file format was invented; one to the polygon protocol and one to the transform protocol, and the introduction of a new protocol (connections). The changes are documented in those sections.

In this document, properties are all named “comp.prop”, where “comp” is the name of the component the property belongs to and “prop” is the name of the property. This is done to prevent ambiguity when two different properties in different components but with the same name exist. In the GTO file and when using the reader library only the property name will appear.

There are two kinds of protocols: major and minor. Every object must have a major protocol that's stored in the ObjectHeader – this is the main indicator of how to interpret the object data. In addition, the object may also have several minor protocols. These indicate optional data and how to interpret it. The next section describes how these are stored in the file.

7.1 Object Protocol

The name of the protocol as it appears in the ObjectHeader is “object” version 1. The protocol does not require any other protocols. Here it is:

object	[Required Component]
A container for properties which don't fit into other component categories well. A catch-all data "per-object" component.	
float[16][1] object.globalMatrix	[Optional Property]
The global world-space transform for the object.	
float[6][1] object.boundingBox	[Optional Property]
The global world-space bounding box for the object.	
string[1][1] object.parent	[Optional Property]
Name of this object's parent in a scene heirarchy.	
string[1][1] object.name	[Optional Property]
The name of the object. This name should be identical to the name in the Object-Header.	
string[1][] object.protocol	[Optional Property]
Additional protocols. This property may contain the main protocol name and any other minor protocols that the object adheres to. If a protocol name appears in this property, the object must adhere to that protocol. Its not an error for a program to output this property with only the major protocol as its value; this is of course redundant since the protocol name is required by the ObjectHeader. It is also not an error for this property to exist but contain nothing.	
int[1][] object.protocolVersion	[Optional Property]
Additional protocol version numbers. This property may exist if the object.protocol property exists. Each entry in this property corresponds to the same entry indexed in the object.protocol property. This property must contain the same number of elements that the object.protocol property does.	

You may be asking why the **object** protocol exists at all. The name of an object is stored in the ObjectHeader in the file and in the C++ library is passed to the reader code. The "name" property is redundant right? Well yes. But some programs will output the name both in the ObjectHeader and in an **object** component as the property "name".

The main point of this protocol is to define the **object** component. This component is meant to hold data that is "per object" and which doesn't really fit neatly into other components. The name is one such case. The coordinate system protocol also defines properties in the **object** component and the minor protocols are optionally stored here.

7.2 Coordinate System Protocol

The name of the protocol as it appears in the ObjectHeader is "transform" version 3¹. The protocol requires the **object** protocol. Objects which obey the **transform** protocol will have global matrices and possibly a parent.

object	[Required Component]
From the "object" protocol.	

¹ In version 1, the transform protocol's `object.globalMatrix` property used to be of type `float[1][16]`. This was a mistake that has been corrected in version 2.

`float[16][1] object.globalMatrix` [Required Property]
 A 4x4 matrix of floating point numbers. This matrix describes the world matrix of the coordinate system.

`string[1][1] object.parent` [Optional Property]
 The name of an object to which this coordinate system is parented. Presumably this object (if it appears in the gto file) will also obey the **transform** protocol. If this property does not exist or the name is "" (the empty string) then the coordinate system presumably is a root coordinate system.²

7.3 Particle Protocol

The name of the protocol as it appears in the ObjectHeader is "particle" version 1. The protocol may include the **object** and **transform** protocols.

`points` [Required Component]
 The points component is transposable. That means that all of its properties are required to have the same number of elements.

`float[3][] points.position` [Optional Property]

`float[4][] points.position` [Optional Property]

The position property is intended to hold the position of the particle in its own coordinate system or world space if it has no coordinate system. The element is either a 3D or 4D (homogeneous) point.

`float[3][] points.velocity` [Optional Property]

The velocity property – if it exists – should hold the velocity vector per-point in the same coordinate system that the "position" property is in.

`int[1][] points.id` [Optional Property]

The "id" property should it exist will *always* be defined as an integer per particle (or other integral type if it ever changes). This number should be unique for each particle. Ideally, multiple GTO files with a point that has the same "id" property for a given particle animation *should* be the same particle.

The **particle** protocol defines the **points** component that many other protocols are derived from. For example, the **NURBS** protocol uses the points defined by the particle protocol as control vertices. There can be any number of properties associated with particles including string per-particle.

The **points** component is marked transposable in the its ComponentHeader. This means that the properties in the component are guaranteed to have the same number of elements. Because of this, the data for the properties in a transposable component may be stored differently than other components. For example, the normal state of affairs is to write data like this:

² In version 3 of the **transform** protocol, the **object.parent** property is redundant and therefore deprecated. The **connection** protocol handles the transformation hierarchy information and in a much more elegant manner See Section 7.11 [Inter-Object], page 32.

```

position0 position1 position2 .... positionN
velocity0 velocity1 velocity2 .... velocityN
mass0 mass1 mass2 .... massN

```

So that you must read through all of the particle positions before you can read the first particle's velocity. But this is not usually the best way to read particle data for rendering. You may want to cull the particles as you read them without storing the data. In order to do this the data needs to be laid out like this:

```

position0 velocity0 mass0
position1 velocity1 mass1
position2 velocity2 mass2
...

```

In this case, each particle is scanned in one chunk allowing for optimizations. Obviously this complicates reading, but in the case of giga-particle renderers, this can be a huge memory savings.

7.4 Strand Protocol

A **strand** object contains a collection of curves. This is somewhat analogous to an object of protocol **particle** as described above.

points [Required Component]

float[3] [] points.position [Required Property]
 The CVs which make up each curve. The number of CVs per curve can vary by curve type and size.

strand [Required Component]
 Information that is relevant to the *all* strands in the object.

string[1] [] strand.type [Required Property]
 String describing curve type. Currently, supported values are **linear** for degree 1 curves, or **cubic** for degree 3 curves.

float[1] [1] strand.width [Optional Property]
 If each end of all curves is the same width, you can just specify that one number instead of the list as with **elements.width** below.

elements [Required Component]
 Information that applies to each separate strand in the object.

int[1] [] elements.size [Required Property]
 This is a list of the sizes of each curve in this object. For example, if there are two curves in this object, with 4 CVs and 3 CVs respectively, then:

```
elements.size = [ 4 3 ]
```

float[2] [] elements.width [Optional Property]
 This is a list of the widths of each end of each curve. The width for each curve will be linearly interpolated over the length.

7.5 NURBS Protocol

The name of the protocol as it appears in the ObjectHeader is “NURBS” version 1. The protocol requires the **particle** protocol and optionally includes the **object** and **transform** protocols.

points [Required Component]

see **particle** protocol. The points describe data per NURBS control vertex.

`float[3] [] points.position` [Required Property]

`float[4] [] points.position` [Required Property]

The position property holds the control point positions in its own coordinate system or world space if it has no coordinate system. The element is either a 3D or 4D (homogeneous) point. If the type is `float[4]` the fourth component of the element will be the rational component of the control point position. The control points are laid out in **v-major** order (*u* iterates more quickly than *v*).

`float[1] [] points.weight` [Optional Property]

If the position property is of type `float[3] []` there may optionally be a “weight” property. This property holds the homogeneous (rational) component of the position. Older GTO writers may export data in this manner. The preferred method is to use a `float[4]` element position.

surface [Required Component]

Properties related to the definition of a NURBS surface are stored in this component.

`float[1] [2] surface.degree` [Required Property]

The degree of the surface in *u* and *v*.

`float[1] [] surface.uKnots` [Required Property]

`float[1] [] surface.vKnots` [Required Property]

The NURBS surface knot vectors in *u* and *v* are stored in these properties. The knots are not piled. The usual NURBS restrictions on how numbers may be stored in the knot vectors apply.

`float[1] [2] surface.uRange` [Required Property]

`float[1] [2] surface.vRange` [Required Property]

The range of the knot parameters in *u* and *v*.

The **NURBS** protocol currently does not handle trim curves, points on surface, etc. Ultimately, the intent is to handle the trim curves and other nasties as NURBS curves-on-surface which will be stored in additional components. UBS surfaces can be stored as NURBS with non-rational uniform knots.

7.6 Polygon Protocol

The name of the protocol as it appears in the ObjectHeader is “polygon” version 2³. The protocol requires the **particle** protocol and optionally includes the **object** and **transform** protocols.

³ In version 1 of the polygon protocol, the **element.size** and **element.type** properties were combined into an **element.primitive** property. We felt that this was adding unnecessary complexity and because the primitive property was an int, it was taking up extra space.

There are a number of alternative configurations of this protocol depending on the value of the **smoothing.method** property. All of these involve the placement of normals in the file.

points [Required Component]

See **particle** protocol. The points describe data per vertex.

float[3] [] points.position [Required Property]

The positions for regular polygonal meshes are stored as **float[3]**.

float[3] [] points.normal [Optional Property]

Normals per vertex. The **smoothing.method** property will have the value of *Smooth* if this property exists. Note that use of the *Smooth* smoothing method does not require that this property exists. If it does not the method is merely an indication of how the normals should be constructed.

normals [Optional Component]

This component will exist if the value of **smoothing.method** is *Partitioned* or *Discontinuous*.

float[3] [] normals.normal [Required Property]

This property is required only if the **normals** component exists and the value of **smoothing.method** is *Partitioned* or *Discontinuous*.

elements [Required Component]

The elements component is transposable. All properties in the elements component must have the same number of elements. Each element corresponds to a polygonal primitive.

byte[1] [] elements.type [Required Property]

Elements are modeled after the OpenGL primitives of the same name. The vertex order is identical to that defined by GL. The type numbers outside those given here are not defined but reserved for future use. So far, these are the define type numbers:

0 – Polygon

General N-sided polygon. This can be used for any polygon that has 3 or more vertices.

1 – Triangle

A three vertex polygon.

2 – Quad A four vertex polygon.

3 – TStrip Triangle strip.

4 – QStrip
Quad strip.

5 – Fan Triangle fan.

short[1] [] elements.size [Required Property]

The size of each primitive. Because the type is short, there is a limit of 65k vertices per primitive.

`short[1] [] elements.smoothingGroup` [Optional Property]

This property may exist if the value of `smoothing.method` is *Partitioned*. In that case, this property indicates the smoothing group number associated with each element. These can be used to recompute the normals. These numbers are the same as those found in the Wavefront .obj file format’s “s” statements. A value of 0 indicates that an element is not in a smoothing group.

`float[3] [] elements.normal` [Optional Property]

Normals per element. The `smoothing.method` property will have the value of *Faceted* if this property exists. Note: the use of *Faceted* smoothing method does not require that this property exists. If it does not, the smoothing method is merely an indication of how the normals should be created.

`indices` [Required Component]

The indices component is transposable. All of its properties are required to have the same number of elements. Each entry in the indices component corresponds to a polygonal vertex.⁴

`int[1] [] indices.vertex` [Required Property]

A list of all the polygonal vertex indices in the same order as the `elements.primitives`. The indices refer to the `points.position` property. So if the first polygonal element is a triangle and second is a general four vertex polygon then vertex indices will be something like:

0 1 2 1 0 3 4 ...

which would be grouped as:

(0 1 2) (1 0 3 4) ...

The first group is the triangle and the second the polygon.

`int[1] [] indices.st` [Optional Property]

Similar to the vertex indices but indicates indices into *st* coordinates. These are usually stored in the “mappings” component but may also appear in the `points` component.

`int[1] [] indices.normal` [Optional Property]

Indices into stored normals if there are any. The `smoothing.method` property will have the value of *Partitioned* or *Discontinuous* if this property exists.

`mappings` [Optional Component]

Contains parametric coordinates. The property names in mappings usually correspond to names found in the `indices` component but not always. For example `mappings.st` would be a `float[2] []` property holding texture coordinates indexed by `indices.st`.

`smoothing` [Optional Component]

The smoothing component exists to hold the smoothing method and any ancillary data for the method. If there is no smoothing component (and hence no `smoothing`) you can assume anything you want.

⁴ The `indices` component in a polygonal object contains values which are analogous to the RenderMan `facevarying` type modifier.

`int[1][1] smoothing` [Required Property]

There five defined smoothing methods (0 through 4). They are:

0 – None No smoothing method specified. No additional properties associated with normals will appear in the object.

1 – Smooth

One normal at every vertex. There will be a `float[3][] normal` property as part of the `points` component. Each vertex has a unique normal.

2 – Faceted

One normal for each face. There will be a `float[3][] normal` property in the `elements` component. Each element has a unique normal.

3 – Partitioned

Same as the Wavefront `.obj` smoothing groups. There will be a `normals` component containing a `float[3][] normal` property and an `int[1][] normal` property in the `indices` component. Each element vertex will have an index into the `normals.normal` property.

4 – Discontinuous

Like *Partitioned* but with additional lines and points of discontinuity. The same properties that hold the *Partitioned* information will hold the *Discontinuous* information. There will also be a component called `discontinuities` which will have a `int[1][]` property called `indices` indicating the points and lines of discontinuity.

7.7 Subdivision Surface Protocols

The name of the protocol as it appears in the `ObjectHeader` is “catmull-clark” or “loop” depending on the intended subdivision scheme. The protocol requires the `polygon` protocol.

The smoothing and any normals properties on the `polygon` protocol should be ignored if they exist.

The protocol indicates how the surface should be treated. Note that the canonical element type for each of the two schemes is not guaranteed to be the only element type stored in the file. For catmull-clark this means that triangles and general polygons will need to be made into quads. Similarly loop surfaces may have quads and other non-triangle primitives that need to be triangulated.

These protocols do not currently define methods for storing edge creasing parameters.

Disclaimer: there are restrictions on what kind of topology surfaces are allowed to have for a given renderer (for example). In most cases surfaces need to be manifold. Some applications can deal with special cases better than others.

7.8 Image Protocol

The Image protocol describes image data in the form of an object. This data makes it possible to store texture maps, backgrounds, etc, directly in the GTO file.

When images are stored in a GTO file, use of Gzip compression is highly recommended if the data is unencoded. As of version 2.1, the supplied Reader and Writer classes default to using zlib compression.

If the image data is encoded, its better *not* to use compression on the GTO file (especially if the file contains only image data).

`image` [Required Component]
The image data and other information will be stored in the **image** component.

`int[1] [] image.size` [Required Property]
The size (and dimension) of the image. There will be N sizes in this property corresponding to the N dimensions of the image.

`string[1] [1] image.type` [Required Property]
The image type. For interactive purposes, the image channels may correspond to a particular fast hardware layout.

- RGB Three channels corresponding to red, green, and blue in that order.
- BGR Three channels corresponding to blue, green, and red in that order.
- RGBA Four channels corresponding to red, green, blue, and alpha in that order.
- ABGR Four channels corresponding to alpha, blue, green, and red in that order.
- L One channel corresponding to luminance.
- HSV Three channels corresponding to hue, saturation, and value. (The HSV color space).
- HSL Three channels corresponding to hue, saturation, and lightness. (The HSL color space).
- YUV Three channels corresponding to the YUV color space.

`int[1] image.cs` [Optional Property]
The coordinate system of the image. The value of **image.cs** can be any one of the following:

0 – Lower left origin.

The first pixel in the image data is the lower left corner of the image data and corresponds to NDC coordinate (0,0).

1 – Upper left origin.

The first pixel in the image data is the upper left corner of the image data and corresponds to NDC coordinate (0,0).

Any one of the following properties are required to hold the actual image data:

`byte[N] [] image.pixels` [Property]
`short[N] [] image.pixels` [Property]
`half[N] [] image.pixels` [Property]
`float[N] [] image.pixels` [Property]

The element width determines the number of channels in the image. For example, the type `byte[3] []` indicate a 3 channel 8-bit per channel image. The number of elements in this property should be equal to `image.size[0] * image.size[1] * ... image.size[N]` where **image.size** is the property defined above.

7.8.1 Additional Image Properties Used by GTV Files.

The base GTO library does not deal with encoded image data or tiling of images. GTV is a specialization of the GTO format for storing movie frames. Some of the GTV properties are documented here. (See documentation for the GTV library for more info).

`string[1] image.encoding` [Optional Property]
 If the pixel data is encoded this property will indicate a method to decode it. Typical values are “jpeg”, “jp2000”, “piz”, “rle”, or “zip”. The pixels will be stored in the `image.pixels` as `byte[1] []`.

7.9 Material Protocol

The name of the protocol as it appears in the `ObjectHeader` is “material”. The material protocol groups a parameters and a method (shader) for rendering. The material protocol can optionally include the **object** protocol.

The material definition is renderer and pipeline dependant. Material assignment is implemented using the **connection** protocol. See [Section 7.11 \[Inter-Object\]](#), page 32.

The **material** protocol is intended for use with software renderers. Interactive material definitions may be more easily defined on the assigned object.

`material` [Required Component]
 Properties unrelated to parameters appear in the **material** component.

`string[1][1] type` [Required Property]
 The value of the **material.type** property is renderer dependant. For a RIB renderer, the value of type might be “Surface”, “Displacement”, “Atmosphere” or a similar shader type name.

`string[1][1] shader` [Optional Property]
 The name of the shader. For RenderMan-like renderers this might be the name of an “.sl” file.

`string[1][1] genre` [Optional Property]
 A property to further identify the material. This is most useful for identifying the target renderer for a material.

`parameters` [Optional Component]
 The set of parameters corresponding to the **material.type**.

7.10 Group Protocol

7.11 Inter-Object Connection Protocol

The name of the protocol as it appears in the `ObjectHeader` is “connection” version 1.

Files which employ the **connection** protocol will typically contain a connection object with the special cookie name “:connections” indicating the purpose of the object as well as preventing namespace pollution. See [Section 8.3 \[Special Cookies\]](#), page 36.

Each component in a connection object is a connection type. For example, the “parent_of” connection type is used to represent transformation hierarchies. In a connection

object, there will be a single component called “parent_of” which will contain the required properties **parent_of.lhs** and **parent_of.rhs** at a minimum. Some connection types may have additional data in the form of additional properties.

Connection components are transposable. The number of elements in properties comprising a connection component will be consistent. So a single “parent_of” component can encode an entire scene transformation hierarchy.

Connection components have the following properties. Note that where **connection_type** occurs in the property name, you would substitute in the actual name of the connection type. (“parent_of” for example).

```
string[1] [] connection_type.lhs           [Required Property]
string[1] [] connection_type.rhs         [Required Property]
```

The left-hand-side and right-hand-side of the connection.

- If the connection is directional, then an arrow indicating the direction would have its tail on the left-hand-side and its head pointing at the right-hand-side.
- If the connection type does not require a direction then these properties are still used to describe the two ends of the connection.
- Each entry will be the name of an object. There is no requirement that the ends of the connection exist in the file. For example, one end of the connection could be an image on disk.
- The empty string is a valid value. You could think of the empty string as indicating a grounded connection.
- Its ok for both ends of the connection to have the same value.

The GTO specification includes a couple of basic connection types.

7.11.1 Transformation hierarchies.

The “parent_of” connection type is used to store transformation hierarchies. The connection type requires only the **lhs** and **rhs** properties. Transformation hierarchies are usually tree structures, but can also be DAGs (as is the case with Maya or Inventor).

Using “parent_of” as a cyclic generalized network connection is probably an error for most applications. To be safe the topology of a “parent_of” network should be a tree.

7.11.2 Material Assignment

The “material” connection type indicates a material assignment to an object. The left-hand-side name is a renderable object in a GTO file The right-hand-side is the name of a material object in a GTO file.

7.11.3 Container Assignment

The “contains” connection type indicates membership in a group or similar type of container object. The LHS is the group or container, the RHS is the object which is a member.

7.12 Difference File Protocol

If the **object.protocol** property contains the string “difference” then the object contains difference data; the data is relative to some other reference file.

For example, for animated deforming geometry its advantageous to write a reference file for geometry in its natural undeformed state then write only the **points.position** property in a gto file per frame to store animation. The **difference** minor protocol can apply to any major protocol.

If a reference file and a difference for file it exists, you can reconstruct the file represented by the difference file using the **gtomerge** command. See [Section 13.3 \[gtomerge\]](#), page 52.

7.13 Sorted Shell File Protocol

If the **object.protocol** property contains the string “sorted” and the object’s major protocol is **polygon** then the object contains sorted shell data.

This protocol guarantees that the vertices and elements of shells — isolated sections of polygonal geometry — will be contiguous in the **points** and **elements** components of the object.

shells [Required Component]

The **shells** component is transposable. Each property in the component should have the same number of elements.

`int [1] [] shells.vertices` [Required Property]

The number of contiguous vertices that make up the Nth element’s shell.

`int [1] [] shells.elements` [Required Property]

The number of contiguous elements that make up the Nth element’s shell.

7.14 Channels Protocol

This minor protocol declares data mapped onto geometric surfaces. Usually the data is mapped using one of the parameterizations found in the **mappings** component of polygonal or sub-d geometry or possibly using the natural parameterization of a surface as is often the case with NURBS.

Each declared channel appears as a `string [1] []` property of a **channels** component on the geometry. The name of the property is the name of the channel. The property should contain at least one element.

The first element of the property should indicate the name of the mapping to use. This is either the name of one of the properties in the **mappings** component or “natural” indicating that the natural parameterization of the surface should be used.

The second and subsequent elements should contain the name of data to map. This could be a texture map file on disk, an image object in the GTO file, or a special cookie string. The lack of second element can be used as a special cookie.

channels [Required Component]

The component holds the names of all the channels on the geometry.

7.14.1 Example

Here is a cube with “color”, “specular”, and, “bump” channels assigned.

```

Object "cube" protocol "polygon"
  Component "points"
    Property float[3][8] "position"
  Component "elements"
    Property byte[1][8] "type"
    Property short[1][8] "size"
  Component "indices"
    Property int[1][32] "vertex"
    Property int[1][32] "st"
  Component "mappings"
    Property float[2][24] "st"
  Component "channels"
    Property string[1][2] "color"
    Property string[1][2] "specular"
    Property string[1][2] "bump"

```

The contents of the “channels” properties might be:

```

string[1] cube.channels.color    = [ "st" "cube_color.tif" ]
string[1] cube.channels.specular = [ "st" "cube_specular.tif" ]
string[1] cube.channels.bump     = [ "st" "cube_bump.tif" ]

```

7.15 Animation Curve Protocol

The animation curve protocol defines a single component called **animation** in which each property holds an animation curve or data stream. The property’s interpretation string indicates how the data should be evaluated.

7.15.1 Example

Here is a cube with animation curves.

```

Object "cube" protocol "polygon"
  Component "points"
    Property float[3][8] "position"
  Component "elements"
    Property byte[1][8] "type"
    Property short[1][8] "size"
  Component "indices"
    Property int[1][32] "vertex"
  Component "animation"
    Property float[6][2] "xtran" interpret as "bezier"
    Property float[6][2] "ytran" interpret as "bezier"
    Property float[6][2] "ztran" interpret as "bezier"
    Property float[6][5] "xrot" interpret as "bezier"
    Property float[6][8] "yrot" interpret as "bezier"
    Property float[6][10] "zrot" interpret as "bezier"
    Property float[1][100] "xscale" interpret as "stream"

```

8 Naming Conventions

GTO files can contain cross references to parts of themselves, objects outside the file, or virtual/logical objects in applications. Because of the potential morass that can result from complete free-form naming, there are conventions which are part of the file specification.

Failure to follow the guidelines does not mean a GTO file is ill-formed; there's always a good reason to ignore guidelines. But having a basis for consistency is usually a good idea.

Some of these topics are a bit “advanced” in that they build off ideas that present themselves after using the file format for a while. If you are just learning about the format, consider this a reference section and skip it. If you're trying to decrypt a complicated GTO file with strange garbled naming, then this section is for you.

8.1 Valid Names

Names should be valid C identifiers, but should not contain the dollar-sign character (\$). This means that no whitespace or punctuation is allowed.

Note that this does *not* apply to protocol names.

There is nothing in the sample `Reader` or `Writer` classes which enforces the valid name guideline. However, some applications (Maya) cannot handle names with whitespace and/or punctuation. So plug-ins which implement GTO reading/writing will have to enforce the application's specific naming requirements.

This guideline is broken by [Section 8.3 \[Special Cookies\]](#), page 36. Its also broken by [Section 8.4 \[Cross References\]](#), page 37.

8.2 Exactly Specifying a Property or Component

By convention, the full name or path name of a property is referred to like this:

```
OBJECTNAME.COMPONENTNAME.PROPERTYNAME
```

When indicating a property name relative to an object then:

```
COMPONENTNAME.PROPERTYNAME
```

should suffice. In this manual, names of components and properties are disambiguated using the dot notation. In addition, this is the format of output from the `gtoinfo` command. There is nothing about the GTO file itself which relates to this notation other than the cross-referencing naming convention discussed below.

8.3 Indicating Special Handling

Some objects, components, or properties in the GTO file may contain data for which names are not particularly useful or that may simply pollute the object or component namespace.

In other cases (component names most notably) the name may be used as information necessary to interpret data associated with it.

In order to distinguish these names from run-of-the-mill names, you should include a colon in the name. Names with colons are considered “special cookie” names and objects which have them may be handled differently than other objects.

The `connection` object protocol, for example, requires that a special file object exist to hold data. This object is not necessarily related to a logical object in the application, its

just a container for the connection data. These objects are named using the special cookie syntax. Usually the name is “:connections”. See [Section 7.11 \[Inter-Object\]](#), page 32.

There is no rule regarding the placement of the colon in the name; it can appear anywhere in the name that is useful for the application. However, if the entire name is a special cookie — there is not additional information encoded in the name beyond itself — the recommend form is to have the colon be the first character.

8.4 Cross References Encoded in Names

Sometimes there is a need to have a property or component *refer* to another property, component, or object in the file (or somewhere else).

To cross reference the data in one property with another, simply name the property the full (or partial) path to the referenced property. For example, here’s the output of `gtoinfo` on a GTO file which has cross referencing properties:

```
object "gravity" protocol "gravity" v1
  component "field"
    property float[3][1] "direction"
    property float[1][1] "magnitude"
  component ":datastream"
    property float[3][300] "field.direction"
    property float[1][300] "field.magnitude"
```

As you can guess, the intention here is that the properties called “field.direction” and “field.magnitude” in the “:datastream” component are data that is associated with the properties “direction” and “magnitude” in the “field” component.

9 Issues and Questionable Aspects of the Format

- There are currently no (publicly available) tools which verify that a file claiming to follow some protocol is correct.
- There is no 3D curve(s) protocol defined.
- The **NURBS** protocol does not handle trim curves. See [Section 7.5 \[NURBS Surfaces\]](#), page 27.
- The format does not contain dedicated space for auxillary information like the name and version of the program that wrote the file, the original owner, copyright information, etc. However, our tools use the string table for these type of data – since its not an error have an unused interned string, we store the data as such. In our opinion, this is a fairly innocuous method. You can read unreferenced strings by using the `gtoinfo` command with the ‘-s’ option. Note that these strings are often lost when programs read and write the file. See [Section 13.1 \[gtoinfo\]](#), page 50.
- Although the format specification includes transposable components (those marked with the `Gto::Matrix` flag may be transposed), the current reader/writer library does not handle files with transposed components. It does handle components that are marked as `Gto::Matrix` but not transposed. See [Section 7.3 \[Particles\]](#), page 25.
- The use of special cookie names and special cross-reference names seems to seriously complicate the format if the protocol is not carefully conceived. For example, using

`gtomerge` to merge files containing connections does not work — the connections are merged like all the other data in the file. The correct behavior would be to combine the connections, but merge the other object data. Perhaps this is just a case for integrating `gtocombine` into `gtomerge`?

- Future versions should incorporate some form of check sum or some similar mechanism to do better sanity checking.
- There are many examples of properties whose data indexes into other property data. The most obvious of these are the polygon protocol **indices** properties. In order to combine gto files (concatenate polygonal data together for example) its necessary to know which properties are indexes and which are not. Index properties must be offset to be combined.
- The Boolean (bit fields) and Half data types are not implemented in the supplied writer library. Both of these types are useful in compressing geometric (and image) data.
- Material, Texture, and similar assignments and storage are usually very specialized at any particular production facility. The idea that a single method of encoding this information can be determined or enforced — or even usefully be stored in a GTO file — is not realistic. However, we hope that some method can be determined that at least preserves a good portion of common data for transfer.

All of the protocols related to these concepts are marked **PROPOSED** in this document. See [Section 7.9 \[Material\]](#), page 32. See [Section 7.11.2 \[material\]](#), page 33.

10 Extending Protocols or the File Format

If you have an extension to a protocol or would like to change an existing protocol, we would like to hear about it. You can send mail to with the changes you're using or would like made. We will collect ideas and proposals and try to make releases in a timely manner. We'd also like to hear from you if you're using it unmodified.

There is currently a small number of facilities that use the GTO format, but there is a large collection of tools that use it. Most proposals should maintain some backwards compatibility. However, we recognize that there may be flaws that require revamping significant pieces and we're open to making changes to accommodate other facilities.

If you are using the format for academic purposes and are looking for a specific tool to munge GTO files, we may already have that tool even though it is not released. Contact us; we might be able to help you out.

We have been using the format since summer 2002 on a regular basis in production at Tweak Films and have found it stable and useful. The version of the code we use is identical to the released version.

11 C++ Library

The GTO Reader/Writer library is written in a subset of C++. The intention was to make the library as portable as possible. Unfortunately we have only tried it on platforms that support gcc 2.95 and greater. It is known to work on various Linux versions and Mac OS X. In either case it has been compiled with gcc.

11.1 Gto::Reader class

The Reader class (in namespace Gto) is designed as a fill-in-the-blank API. The user of the class derives from it; the base class defines a number of virtual functions which pass data to the derived class and ask the derived class questions about what data it wants.

The Reader class handles most of the difficult work in reading the file like keeping track of headers, sizes of properties, and the order of data. In addition, it handles the string table and looking up property string values. If the file was written by a machine with different sex (endianess) it will translate the data for you.

In addition, you can compile the GTO library with zlib support. This enables the Reader class to read gzipped GTO files natively and the Writer class to write them. This can be a significant space savings on disk and on saturated networks can make file loading faster. You can also pass a C++ istream object to the Reader if you want to read “in-core”.

As the file is read, the Reader class will call its virtual functions to declare objects in the file to the derived class. The derived class is expected to return a non-null pointer if it wishes to later receive data for that object.

Reader::Reader (*unsigned int mode*) [Constructor]

The constructor argument *mode* indicates how the reader will be used. This value is a bit vector of the following or'ed flags:

Reader::None

The reader will be used in its standard *streaming* mode. The reader will attempt to read all the data in the file. This is the default value (or 0).

Reader::HeaderOnly

The reader will stop once it has read the header sections of the GTO file. This is an optimization that applies to binary files only. This option is ignored when reading a text file.

Reader::RandomAccess

The reader will read the header sections but not the data, however, it will initialize for use of the **Reader::accessObject()** function. Only binary GTO files can be read using the radom access mode.

Reader::BinaryOnly

Only binary GTO files will be accepte by reader.

Reader::TextOnly

Only text GTO files will be accepte by reader.

Reader::~Reader () [Destructor]

Closes file if still open.

bool Reader::open (*const char* filename*) [Method]

Open the file. The Reader will attempt to open file *filename*. If the file does not exist and zlib support is compiled in, the Reader will attempt to look for *filename.gz* and open it instead.

bool Reader::open (*std::istream&, const char* name*) [Method]

Reads the GTO file data from a stream. The *name* is supplied to make error messages make sense.

- `void Reader::close ()` [Method]
Close the file and clean up temporary data. If the stream constructor was used, the stream is *not* closed.
- `std::string& Reader::fail (std::string why)` [Method]
Sets the error condition on the Reader and sets the human readable reason to *why*.
- `std::string& Reader::why ()` [Method]
Returns a human readable description of why the last error occurred. (Set by the `fail()` function).
- `const std::string& Reader::stringFromId (unsigned int)` [Method]
Given a string identifier, this method will return the actual string from the string table.
- `const StringTable& Reader::stringTable ()` [Method]
Returns a reference to the entire string table.
- `bool Reader::isSwapped () const` [Method]
Returns true if the file being read needed to be swapped. This occurs if the machine the file was written on is a different sex than the machine reading the file (for example a Mac PPC written file read on an x86 GNU/Linux box).
- `unsigned int Reader::readMode () const` [Method]
Returns the mode value passed into the Reader constructor.
- `const std::string& Reader::infileName () const` [Method]
Returns the name of the file or stream being read. This is the value passed in to the `Reader::open()` function.
- `std::istream* Reader::in () const` [Method]
Return the input stream created by or passed into `Reader::open()`. If the GTO file is compressed binary, this function will return NULL.
- `int Reader::linenum () const` [Method]
For text GTO files, the return value will be the current line being parsed. For binary GTO files, the return value is always 0.
- `int Reader::charnum () const` [Method]
For text GTO files, the return value will be the current char column (in the current line) being parsed. For binary GTO files, the return value is always 0.
- `Header& Reader::fileHeader () const` [Method]
Returns a reference to a `Gto::Header` structure corresponding to the file currently being read. This function is required by the text file parser. The function may disappear from future versions. See the `Reader::header()` function below for a better way to get header information.

The following functions are called by the base class.

void Reader::header (*const Header& header*) [Virtual]
 This function is called by the Reader base class right after the file header has been read (or created).

void Reader::descriptionComplete () [Virtual]
 This function is called after all file, object, component, and property structures have been read. For binary files, this is just before the data is read. For text files, this is after the entire file has been read.

The following functions return a **Reader::Request** object. This object takes two parameters: a boolean indicating whether the data in question should be read by the reader and a second optional data **void*** argument of user data to associate with the file data.

Reader::Request::Request (*bool want, void* data*) [Constructor]
want value of true indicates a request for the data in question. *data* can be any **void***. *data* is meaningless if the *want* is false.

Reader::Request Reader::object (*const std::string& name, const std::string& protocol, unsigned int protocolVersion, const ObjectInfo& header*) [Virtual]

This function is called whenever the Reader base class encounters an **ObjectHeader**. The derived class should override this function and return a **Request** object to indicate whether data should be read for the object in question. If it requests not to have data read, the Reader will not call the corresponding **component()** and **property()** functions.

Reader::Request Reader::component (*const std::string& name, const ComponentInfo& header*) [Virtual]

This function is called when the Reader base class encounters a **ComponentHeader** in the GTO file. If the derived class did not express interest in a particular object in the file by returning **Request(false)** from the **object()** function, the components of that object will not be presented to the derived class. The derived class should return **Request(true)** to indicate that it is interested in the properties of this *component*.

Reader::Request Reader::property (*const std::string& name, const char* interpString, const PropertyInfo& header*) [Virtual]

This function is called when the Reader base class encounters a **PropertyHeader** in the GTO file. If the derived class did not express interest in a particular object or the component that the property belongs to, the properties of that component will not be presented to the derived class. The derived class should return **Request(true)** to indicate it is interested in the property data.

void* Reader::data (*const PropertyInfo&, size_t bytes*) [Virtual]
 This function is called before property data is read from the GTO file. The function should return a pointer to memory of at least size *bytes* into which the data will be read. The type, size, width, etc, of the data can be obtained from the **PropertyInfo** structure.

void Reader::dataRead (*const PropertyInfo&*) [Virtual]
 This function is called after the **data()** function if the data was successfully read.

If you are using the Reader class in `Reader::RandomAccess` mode, you may call these functions after the read function has returned:

`Reader::Objects& Reader::objects ()` [Method]

Returns a reference to an `std::vector` of `Reader::ObjectInfo` structures. These are only valid after `Reader::open()` has returned. You can use these structures when calling `Reader::accessObject()`.

`const Reader::Components& Reader::components ()` [Method]

Returns a reference to an `std::vector` of `Reader::ComponentInfo` structures. These are only valid after `Reader::open()` has returned. This method is most useful when deciding how to call the `accessObject` function.

`const Reader::Properties& Reader::properties ()` [Method]

Returns a reference to an `std::vector` of `Reader::PropertyInfo` structures. These are only valid after `Reader::open()` has returned. This method is most useful when deciding how to call the `accessObject` function.

`void Reader::accessObject (const ObjectInfo&)` [Method]

Calling this function on a GTO file opened for `RandomAccess` reading will cause the reader to seek into the file just for the data related to the object passed in. This is most useful when the objects' data cannot be held in memory and the order of retrieval is unknown. The reader attempts to be efficient as possible without using too much memory.

11.2 Gto::Writer class

The Writer class (in namespace Gto) is designed as an API to a state machine. You indicate a conceptual hierarchy to the file and then all the data. The writer handles generating the string table, the header information, etc.

The following is an example that outputs a polygon cube using the **polygon** protocol.

```
float points[3] [] =
{ { -2.5, 2.5, 2.5 }, { -2.5, -2.5, 2.5 },
  { 2.5, -2.5, 2.5 }, { 2.5, 2.5, 2.5 },
  { -2.5, 2.5, -2.5 }, { -2.5, -2.5, -2.5 },
  { 2.5, -2.5, -2.5 }, { 2.5, 2.5, -2.5 } };

unsigned char type[] = { 2, 2, 2, 2, 2, 2 };
unsigned char size[] = { 4, 4, 4, 4, 4, 4 };

int indices[] = {0, 1, 2, 3,    7, 6, 5, 4,
                 3, 2, 6, 7,    4, 0, 3, 7,
                 4, 5, 1, 0,    1, 5, 6, 2 };

Gto::Writer writer;
writer.open("cube.gto");
```

```

writer.beginObject("cube", "polygon", 2); // polygon version 2

    writer.beginComponent("points");
        // will write 8 float[3] positions
        writer.property("positions", Gto::Float, 8, 3);
    writer.endComponent();

    writer.beginComponent("elements");
        // one per face
        writer.property("size", Gto::Short, 8, 1, 1);
        writer.property("type", Gto::Byte, 8, 1, 1);
    writer.endComponent();

    writer.beginComponent("indices");
        // one per vertex per face
        writer.property("vertex", Gto::Int, 24, 1, 1);
    writer.endComponent();

writer.endObject();

// repeat writer object blocks if more objects

// output all the data in order declared

writer.beginData();
writer.propertyData(points);
writer.propertyData(type);
writer.propertyData(size);
writer.propertyData(indices);
writer.endData();

```

`Writer::Writer ()` [Constructor]

Creates a new `Writer` class object. Typically you'll make one of these on the stack. This constructor requires you call the `open` function to actually start writing the file.

`Writer::Writer (std::ostream&)` [Constructor]

Creates a new `Writer` class object which will output to the passed C++ output stream.

`Writer::~Writer ()` [Destructor]

Closes file opened with the `open()` function if still open. The destructor will not close any passed in output stream.

`bool Writer::open (const char* filename, FileType mode = CompressedGTO)` [Method]

Open the file. The `Writer` will attempt to open file `filename`. If the file is not writable for whatever reason, the function will return false. If `mode` is `CompressedGTO` (the default value), the `Writer` class will output a binary compressed file. If the value is

BinaryGTO the file will be binary uncompressed. If *mode* is TextGTO a text GTO file will be written. Compressed GTO files can be uncompressed manually using `gzip`. Compression is available only if the library is compiled with zlib support.

- `bool Writer::open (const char* filename, bool compress = true)` [Method]
This function exists for backwards compatibility. Use the other `open()` function instead. This function can open a file for binary output only (it cannot write a text GTO file).
- `void Writer::close ()` [Method]
Close the file and clean up temporary data. If the stream constructor was used, the stream is *not* closed.
- `void Writer::beginObject (const char* name, const char* protocol, unsigned int version) const` [Method]
Declares an object. Its components and properties must be declared before `endObject()` is called. The *name* is the name of the object as it will appear in the gto file. The *protocol* is the protocol string indicating how the object data will be interpreted and the *version* number indicates the protocol version. The Writer class does not verify that the data output conforms to the protocol.
- `void Writer::beginComponent (const char* name, bool transposed=false)` [Method]
Declares a component. The component properties must be declared before a call to `endComponent()`. The *name* is the name of the component as it will appear in the gto file. The *transposed* flag is optional and indicates whether or not the component property data should be output transposed or one property at a time (the default).
- `void Writer::property (const char* name, Gto::DataType type, size_t numElements, size_t partsPerElement=1, const char* interpString=0)` [Method]
Declare a property. The *name* is the name of the property as it appears in the gto file. The *type* is one of `Gto::Double`, `Gto::Float`, `Gto::Int`, `Gto::String`, `Gto::Byte`, `Gto::Half`, or `Gto::Short`. *numElements* indicates the number of elements of size *partsPerElement* that will be in the property data. So for example, if the property is declared as a `Gto::Float` of with *partsPerElement* of 3 and there 10 of them, then the writer will expect an array of 30 floats when the `propertyData` is finally passed to it. The last argument *interpString* is an optional interpretation string that can be stored with the property.
- `void Writer::endComponent ()` [Method]
Closes the declaration of a component started by `beginComponent()`.
- `void Writer::endObject ()` [Method]
Closes the declaration of an object started by `beginObject()`.
- `void Writer::intern (const char* string)` [Method]
Declares a string to the Writer for inclusion in the file string table. When writing properties of type `Gto::String`, its necessary to call this function before the `beginData()` is called. Each string in the property data must be interned. When outputting the property, the property will be an array of `Gto::Int` in which each int is the result of the `lookup()` function which retrieves a unique int corresponding to interned strings.

`void Writer::intern (const std::string& string)` [Method]
Same as above, but takes an `std::string`.

`int Writer::lookup (const char* string)` [Method]
Retrieve the identifier of the previously interned string *string*.

`int Writer::lookup (const std::string& string)` [Method]
Same as above, but takes an `std::string&`.

`void Writer::beginData ()` [Method]
Begins data declaration to the `Writer` class. Only calls to `lookup()`, `propertyData()`, `propertyDataInContainer()`, and `endData()` are legal after `beginData()` is called.

`void Writer::propertyData (const TYPE* type)` [Method]
`propertyData()` is a template function which takes a pointer to continuously stored data. The data must be the same as declared earlier by the `property()` function. Calls to `propertyData()` and `propertyDataInContainer()` must appear in the same order as the `property()` declarations calls.

`void Writer::propertyDataInContainer (const TYPE& container)` [Method]
`propertyDataInContainer()` is a template function which takes an stl-like container as an argument. The data must be the same as declared earlier by the `property()` function. Calls to `propertyData()` and `propertyDataInContainer()` must appear in the same order as the `property()` declarations calls. This function is a convenience function; it calls `propertyData()` to actually output the data. This function may make a copy of the data in the container.

`void Writer::endData ()` [Method]
Closes the definition of data started by `beginData()` and finishes writing the gto file.

`const Writer::Properties& Writer::properties () const` [Method]
Returns a vector of previously declared properties; these are the result of calls to the `property()` function.

11.3 Gto::RawDataReader/Gto::RawDataWriter classes

These classes provide a quick method of reading the contents of a GTO file into memory for basic editing. The `RawDataReader` and `RawDataWriter` both use the same very primitive data structure that can be found in the `RawData.h` file. For examples of use, see `gtomerge` and `gtofilter` source code.

The `RawData` class shows how to both read and write using the supplied classes. In addition the reader subclass shows how to convert string data.

12 Python Module

The `gto` module implements a reader/writer library for the Python language. The module is implemented on top of the C++ reader and writer classes. The API is similar to the C++ API, but takes advantage of Python’s dynamic typing to “simplify” the design. The Python module also implements a significant number of safety checks not present in the C++ library, making it an ideal way of exploring the Gto file format.

12.1 `gto.Reader`

The Reader class is designed as a fill-in-the-blank API much like the C++ library. The user of the class derives from it; the base class defines a number of functions which you override to pass data to the derived class and receive data from it.

As the file is read, the Reader class will call specific functions in itself to declare objects in the file. The derived class is handed data or asked to return whether or not it is interested in specific properties in the file.

The biggest difference from the C++ Reader class is that the `data()` method of the C++ class, which returns allocated memory for the library to read data into, cannot be overloaded in Python. Instead, the `dataRead()` method of the Python `gto.Reader` class is handed pre-allocated Python objects containing the data.

status `gto.Reader(mode)` [Constructor]

Create a new `gto.Reader` instance. Possible values for *mode*:

`gto.Reader.NONE`

The reader will be used in its standard *streaming* mode. The reader will attempt to read all the data in the file. This is the default value (or 0).

`gto.Reader.HEADERONLY`

The reader will stop once it has read the header sections of the GTO file.

`gto.Reader.RANDOMACCESS`

The reader will read the header sections but not the data, however, it will initialize for use of the `gto.Reader.accessObject()` method.

`gto.Reader.BINARYONLY`

The reader will only accept binary GTO files.

`gto.Reader.TEXTONLY`

The reader will only accept text GTO files.

`gto.Reader.open(filename)` [Method]

Opens and reads the GTO file *filename*. The function will raise a Python exception if the file cannot be opened.

wants `gto.Reader.object(name, protocol, protocolVersion, objectInfo)` [Method]

This function is called by the base class to declare an object in the GTO file. The return value *wants* should evaluate to True or False, indicating whether or not the base class should read the object data. *name* and *protocol* are strings declaring name and protocol of the object, *protocolVersion* is an integer. *objectInfo* is an instance

of a generic class which contains the same information as the `Gto::ObjectInfo` C++ struct.

`wants gto.Reader.component (name, interpretation, componentInfo)` [Method]

This function is called by the base class to declare a component in the GTO file. The return value `wants` should evaluate to True or False, indicating whether or not the base class should read the component data. `name` is a string declaring the component name. `componentInfo` is an instance of a generic class which contains the same information as the `Gto::ComponentInfo` C++ struct.

`wants gto.Reader.property (name, interpretation, propertyInfo)` [Method]

This function is called by the base class to declare a property in the GTO file. The return value `wants` should evaluate to True or False, indicating whether or not the base class should read the property data. `name` is a string declaring the full property name. `propertyInfo` is an instance of a generic class which contains the same information as the `Gto::PropertyInfo` C++ struct.

`gto.Reader.dataRead (name, data, propertyInfo)` [Method]

If a property has been requested, the `dataRead()` function will eventually be called by the base class with the actual data in the file. The `name` is the name of a property, `data` is a tuple containing the property data, `propertyInfo` is an instance of a generic class which contains the same information as the `Gto::PropertyInfo` C++ struct.

`gto.Reader.stringFromID (id)` [Method]

Returns the `stringTable` entry for the given string table id. Since the Python `gto` module returns strings directly, it is unlikely that you'll need to use this.

`gto.Reader.stringTable ()` [Method]

Returns the entire `stringTable` as a list of strings.

`gto.Reader.isSwapped ()` [Method]

Returns True if the file on disk is not in the machine's native byte order.

`gto.Reader.objects ()` [Method]

Returns a list of the `gto.ObjectInfo` instances for all the objects in the file. This method is only available if the file was opened with `gto.Reader.RANDOMACCESS`. Usable at any time after the constructor is called.

`gto.Reader.components ()` [Method]

Returns a list of the `gto.ComponentInfo` instances for all the components in the file. This method is only available if the file was opened with `gto.Reader.RANDOMACCESS`. Usable at any time after the constructor is called.

`gto.Reader.properties ()` [Method]

Returns a list of the `gto.PropertyInfo` instances for all the properties in the file. This method is only available if the file was opened with `gto.Reader.RANDOMACCESS`. Usable at any time after the constructor is called.

`gto.Reader.accessObject (objInfo)` [Method]
 Given an instance of `gto.ObjectInfo` (obtained via `gto.Reader.objects()`, `gto.Reader.components()`, or `gto.Reader.properties()`), tells the reader to access that object directly. This will cause the `gto.Reader.object()`, `gto.Reader.component()`, and `gto.Reader.dataRead()` methods to be called with the information from the given object.

12.2 gto.Writer

`gto.Writer ()` [Constructor]
 Creates a new writer instance, no arguments needed.

`gto.Writer.open (filename, mode)` [Method]
 Open the file. The Writer will attempt to open file *filename*. If the file is not writable for whatever reason, the function will raise a Python exception. The *mode* argument can be `BINARYGTO`, `COMPRESSEDGTO` (the default) or `TEXTGTO`.

`gto.Writer.close ()` [Method]
 Close the file and clean up temporary data. Because of Python's garbage-collection, you can never be sure when a class's destructor will be called. Therefore, it is *highly* recommended that you call this method to close your file when it's done writing. You have been warned.

`gto.Writer.beginObject (name, protocol, version)` [Method]
 Declares an object. Its components and properties must be declared before `endObject()` is called. The *name* is the name of the object as it will appear in the gto file. The *protocol* is the protocol string indicating how the object data will be interpreted and the *version* number indicates the protocol version. The Writer class does not verify that the data output conforms to the protocol.

`gto.Writer.beginComponent (name, interpretation, transposed)` [Method]
 Declares a component. The component properties must be declared before a call to `endComponent()`. The *name* is the name of the component as it will appear in the gto file. The *transposed* flag is optional and indicates whether or not the component property data should be output transposed or one property at a time (the default).

`gto.Writer.property (name, type, numElements, partsPerElement, interpretation)` [Method]
 Declare a property. The *name* is the name of the property as it appears in the gto file. The *type* is one of `gto.DOUBLE`, `gto.FLOAT`, `gto.INT`, `gto.STRING`, `gto.BYTE`, `gto.HALF` (*Not implemented*), or `gto.SHORT`. *numElements* indicates the number of elements of size *partsPerElement* that will be in the property data. So for example, if the property is declared as a `gto.FLOAT` of with *partsPerElement* of 3 and there 10 of them, then the writer will expect a sequence of 30 floats when the `propertyData` is finally passed to it.

`gto.Writer.endComponent ()` [Method]
 Closes the declaration of a component started by `beginComponent()`.

`gto.Writer.endObject ()` [Method]
 Closes the declaration of an object started by `beginObject()`.

`gto.Writer.intern (string)` [Method]
 Declares a string to the Writer for inclusion in the file string table. When writing properties of type `gto.String`, its necessary to call this function for each string in the property data before the `beginData()` is called. The Python version of `intern()` can accept individual strings, as well as lists or tuples of strings.

`int gto.Writer.lookup (string)` [Method]
 Retrieve the identifier of the previously interned string. Valid only after `beginData()` has been called.

`gto.Writer.beginData ()` [Method]
 Begins data declaration to the Writer class. Only calls to `lookup()`, `propertyData()`, and `endData()` are legal after `beginData()` is called.

`gto.Writer.propertyData (data)` [Method]
 The `propertyData()` function must get exactly *one* parameter. That parameter can be any of the following:

- A single int, float, string, etc.
- An instance of `mat3`, `vec3`, `mat4`, `vec4`, or `quat` (<http://cgkit.sourceforge.net/>). DO NOT explicitly cast `mat3` or `mat4` into a tuple or list: `tuple(mat4(1))`. It will be silently transposed (a bug in the `cgtypes` code?). ADDING it to a tuple or list is fine: `(mat4(1),)`
- A tuple or list of any combination of the above that makes sense.

Tuples and lists are flattened out before they are written. As long as the number of atoms is equal to size x width, it'll work. Calls to `propertyData()` must appear in the same order as declared with the `property()` method.

`void gto.Writer.endData ()` [Method]
 Closes the definition of data started by `beginData()` and finishes writing the `gto` file. Does *not* actually close the file—use the `close()` method for that.

12.3 Classes used by `gto.Reader`

Note that as of the 3.0 release, these classes will contain the actual strings rather than string table IDs.

`gto.ObjectInfo` [Class]
 This class emulates the `Gto::ObjectInfo` struct from the C++ `Gto` library. It is passed by the Python `gto.Reader` class to your derived `object()` method. The only methods implemented are `__getattr__` and `__repr__`. Available attributes are:

- `name` - String
- `protocolName` - String
- `protocolVersion` - Integer
- `numComponents` - Integer
- `pad` - Integer

gto.ComponentInfo [Class]

This class emulates the `Gto::ComponentInfo` struct from the C++ Gto library. It is passed by the Python `gto.Reader` class to your derived `component()` method. The only methods implemented are `__getattr__` and `__repr__`. Available attributes are:

- `name` - String
- `numProperties` - Integer
- `flags` - Integer
- `interpretation` - String
- `pad` - Integer
- `object` - Instance of `gto.ObjectInfo`

gto.PropertyInfo [Class]

This class emulates the `Gto::PropertyInfo` struct from the C++ Gto library. It is passed by the Python `gto.Reader` class to your derived `property()` and `dataRead()` methods. The only methods implemented are `__getattr__` and `__repr__`. Available attributes are:

- `name` - String
- `size` - Integer
- `type` - Integer
- `width` - Integer
- `interpretation` - String
- `pad` - Integer
- `component` - Instance of `gto.ComponentInfo`

13 Utilities

13.1 The `gtoinfo` Utility

Usage: `gtoinfo [OPTIONS] infile.gto`

Options:

`'-a/--all'` Output property data and header.

`'-d/--dump'`

Output property data (no header data is emitted).

`'-l/--line'`

Output property data one item per line. Can be used with either `'-d'` or `'-s'`.

`'-h/--header'`

Output header data.

`'-s/--strings'`

Output sting table data.

`'-n/--numeric-strings'`

Output sting data as the raw string id instead of the string itself.

- '-i/--interpretation-strings'
Output interpretation string data for components and properties if it exists.
- '-r/--readall'
Force reading of the entire gto file even if only the header is being output.
- '-f/--filter *expression*'
Only output information for properties whose long name (object.component.propname) matches the shell-like *expression*. [Section 13.2 \[gtofilter\]](#), [page 51](#) for examples of filter expressions. This option is similar to `gtofilter --include` option.
- '--help'
Output usage message.

`gtoinfo` outputs the part of all of the contents of a gto file in human readable form. Its invaluable for debugging or just getting a quick understanding of what a gto file contains.

13.2 The `gtofilter` Utility

Usage: `gtofilter [OPTIONS] -o out.gto in.gto`

Options:

- '-v'
Set verbose output. Whenever a pattern matches `gtofilter` will inform you.
- '-ee/--exclude'
Regular expression which will be used to exclude properties.
- '-ie/--include'
Regular expression which will be used to include properties.
- '-regex'
Use POSIX regular expression syntax.
- '-glob'
Use shell-like regular expression (`fnmatch`). This is the default.
- '-t'
Output text GTO file.
- '-nc'
Output uncompressed binary GTO file.
- '-o *out.gto*'
Output `.gto` file

`gtofilter` can be used to remove objects, components, and properties from a gto file. You supply an include shell-like expression and/or an exclude shell-like expression. (The pattern matching is done using the `fnmatch()` function – see the man page for details.)

The patterns match each full property name. So for example a cube might have these properties:

```
cube.points.position
cube.elements.type
cube.elements.size
cube.indices.vertex
cube.indices.st
cube.indices.normal
cube.normals.normal
cube.mappings.st
```

```
cube.smoothing.method
cube.object.globalMatrix
cube.object.parent
```

Using the ‘`--exclude`’ option, you can remove the object component by doing this:

```
gtofilter --exclude "*.object.*" -o out.gto cube.gto
```

or if you wanted to pass through only the positions:

```
gtofilter --include "*.positions" -o out.gto cube.gto
```

-or-

```
gtofilter --include "*positions" -o out.gto cube.gto
```

13.3 The `gtomerge` Utility

Usage: ‘`gtomerge -o outfile.gto infile1.gto infile2.gto ...`’

Options:

- ‘`-o outfile.gto`’ The resulting merged file to output.
- ‘`-t`’ Output text GTO file.
- ‘`-nc`’ Output uncompressed binary GTO file.

`gtomerge` takes a number of `.gto` input files and merges them into a single output `.gto` file. This is done by first creating output geometry that is identical to the first input file and then adding only those properties that are not already defined from subsequent `gto` files. The order of input files determines what will be in the final output file.

For **difference** files, you can use `gtomerge` to reconstruct a final file like this:

```
gtomerge -o out.gto difference.gto reference.gto
```

13.4 The `gto2obj` Utility

Usage: ‘`gto2obj [OPTIONS] infile outfile`’

Options:

- ‘`-o NAME`’ When outputting GTO files, the name of an object in the GTO file to output. If not specified, the translator will output the first polygon, or subdivision surface it finds.
- ‘`-c`’ When outputting GTO files, this option will force the protocol to be "catmull-clark".
- ‘`-l`’ When outputting GTO files, this option will force the protocol to be "loop".
- ‘`-t`’ Output text GTO file.
- ‘`-nc`’ Output uncompressed binary GTO file.

`gto2obj` takes either an input GTO file or Wavefront `.obj` file and outputs the other file type.

```

gto2obj in.obj out.gto
gto2obj in.gto out.obj
gto2obj -c in.obj out.gto ## output obj as subdivision surface

```

13.5 The gtoimage Utility

Usage: `gtoimage infile outfile`

`-t` Ouput text GTO file.

`-nc` Ouput uncompressed binary GTO file.

`gtoimage` reads a TIFF file and converts it into a GTO file containing one image object. 32 bit floating point images, 16 bit and 8 bit integral images are directly converted. `gtoimage` expects the image to be two dimensional with three or four channels where the fourth channel is an optional alpha value. The output object conforms to the **image** protocol. See [Section 7.8 \[Image\], page 30](#).

You can use `gtomerge` to merge the image object into another GTO file. See [Section 13.3 \[gtomerge\], page 52](#).

It is highly recommend that the resulting output GTO file be written with compression or gzipped to reduce its size. Gzipped GTO files can be read directly by the supplied readers.

13.6 The RiGtoRibOut Utility

The `RiGtoRibOut` command is useful for:

- It can be used as a debugging tool for the `RiGtoPlugin` RenderMan plugin.
- It can be used as a drop-in replacement for `RiGtoPlugin`, for RIB renderers that do not support `Procedural DynamicLoad`, but that *do* support `Procedural RunProgram`. Note that this is substantially slower than using `RiGtoPlugin`, as all data needs to be translated to ASCII and back. It does have the one space-saving advantage of not needing to save ASCII RIB on disk.
- It could be used to generate RIB files that are read with `ReadArchive`. This is not recommended, as it negates all the advantages of using GTO in the first place. But if nothing else works, this should.

The command-line parameters are the same as the `CONFIG_STRING` for `RiGtoPlugin`. See [Section 13.8 \[RiGtoPlugin\], page 54](#).

13.7 The gtoIO.so Maya Plug-In

The Maya plugin comes in two parts: the C++ plugin which implements a Maya scene translator and an accompanying MEL script which implements the user interface.

The plugin handles export of NURBS surfaces (but not trim curves), polygonal geometry (which can be written as sub-division surfaces), and generic transforms. A Maya particle export tool is in the works. Additional user defined attributes can be emitted into the GTO file.

The plugin can import everything that it exports and also particle GTO files generated by other applications.

13.7.1 BUGS

The internal performance of Maya has changed between the 4.x and 5.0 versions. In Maya 5.0, the Maya API is *extremely* slow when importing polygonal normals. Importing of normals is disabled in Maya 5.0.

13.8 The RiGtoPlugin RenderMan plugin

Here you will find information on using the GTO RenderMan plugin. The documentation is complete enough to get started with, but should be considered a work in progress.

13.8.1 RIB Instantiation

The plugin is instantiated in a RIB Stream using the standard DynamicLoad procedural mechanism, like so:

```
Procedural "DynamicLoad" [ "RiGtoPlugin.so" "CONFIG_STRING" ] [ Bounding Box ]■
```

If a bounding box is not known, the infinite box may be used:

```
Procedural "DynamicLoad" [ "RiGtoPlugin.so" "CONFIG_STRING" ] [ -1e6 1e6 -1e6 1e6 -1e6
```

13.8.2 Config String Syntax

The configuration string passed into RiGtoPlugin consists of a variable number of space-separated tokens. They are, in order:

1. Reference Pose GTO File Name
2. Shutter Open GTO File Name (optional)
3. Shutter Close GTO File Name (optional)
4. Primary On List (optional)
5. Primary Off List (optional)
6. Secondary On List (optional)
7. Secondary Off List (optional)

As shown, the only necessary element is the reference GTO file. For objects which do not have movement and do not require on lists or off lists, this is completely sufficient.

The logic behind the geometry instantiation mechanism is as follows:

- Read Reference GTO file The plugin reads all of the geometry in the reference GTO file. As a starting point, the shutter open and close geometry is set equal to the reference geometry.
- If requested, read Shutter Open GTO file The plugin then reads any geometry from the Shutter Open file that matches the name and geometry type of geometry that has already been read from the reference file - this geometry is stored as both the shutter open AND close geometry.
- If requested, read Shutter Close GTO file The plugin then reads any geometry from the Shutter Close file that matches the name and geometry type of geometry that has already been read from the reference file - this geometry is stored as the shutter close geometry
- Instantiate Geometry: For any piece of geometry that appears in BOTH on-lists and does not appear in EITHER off-lists, the plugin calls the appropriate RIB functions to create the requested geometry.

13.8.3 On-List/Off-List Syntax

The syntax of the on-lists and off-lists is as follows:

NULL is a special on-list/off-list which is interpreted as *all on* or *none off*.

Otherwise, the on-lists and off-lists are essentially shell-like regular expressions. The following rules apply:

- The * character matches any number of characters
- The ? character matches any single of character
- Bracket expressions [] are supported. (See `man 7 regex`)
- Multiple patterns can be strung together with the | character.
- The pattern must match the *whole* object name. Thus, the pattern `*Sphere1` will match the object `nurbsSphere1` but *not* `nurbsSphere1Shape`. This is a very common "gotcha".

As an example, suppose you wanted to turn off all of the geometry named `LeftLeg*Shape*` and `RightLeg*Shape*` in a render - you would create an off-list that looked like:

```
"LeftLeg*Shape*|RightLeg*Shape*"
```

13.8.4 Cache Management

By default, `RiGtoPlugin` maintains an internal cache of all of the file sets it has read. The cache is keyed off of Ref-Open-Close filename triplets. The reason for this is to facilitate easy material assignment, which will be discussed in greater detail below in the "Strategy" section.

In situations where memory is precious and the renderer needs as much memory as it can get, it may be advantageous to force `RiGtoPlugin` to discard its cached file sets. There is special syntax to facilitate this.

- To erase everything in `RiGtoPlugin`'s cache:

```
Procedural "DynamicLoad" [ "RiGtoPlugin.so" "__FLUSH__" ] [ Bounding Box ]
```

- To erase the cache associated with a given file triplet: (Using `REF.gto`, `OPEN.gto` and `CLOSE.gto` as standins for whatever files were actually passed in)

```
Procedural "DynamicLoad" [ "RiGtoPlugin.so" "REF.gto OPEN.gto CLOSE.gto __FLUSH__"
ing Box ]
```

There is also an environment variable, `TWK_RI_GTO_NO_CACHE`, which if defined and set to anything other than "0", "FALSE", "False" or "false", will cause caching to be disabled entirely.

13.8.5 Environment Variables

`TWK_RI_GTO_NO_SUBDS` [Environment Variable]

If this environment variable is defined and set to anything except "0", "FALSE", "False", or "false", `RiGtoPlugin` will treat all catmull-clark subdivision surfaces read from a GTO file as polygons instead.

TWK_RI_GTO_NO_CACHE [Environment Variable]
 If this environment variable is defined and set to anything except "0", "FALSE", "False", or "false", RiGtoPlugin will disable all caching of geometry data to save memory.

13.8.6 Usage Strategy

The RiGtoPlugin was designed with a particular data structure in mind. Used ideally, there would be a GTO file consisting of all of the geometry corresponding to a particular high-level creature or set in the scene. All of the surfaces corresponding to a hippo or a giraffe or a cyborg-monkey would be in a single GTO file. The animation data for this geometry would be contained in light-weight GTO files that contain only points that have moved and transformation matrices that have moved. The RiGtoPlugin only reads points and matrices from the Shutter-Open and Shutter-Close file, facilitating very light-weight "difference" files for animation data.

Because all of the geometry in a creature will have different materials assigned to it, on-lists and off-lists can be used to separate out only the geometry that shares a particular material.

Suppose we have a creature consisting of many surfaces but only three different materials - skinMtl, eyeMtl and hairMtl. The parts of the model have been named intelligently (for this example) such that the skin parts all have names like Skin*Shape*, the eye parts all have names like Eye*Shape*, and the hair parts all have names like Hair*Shape*. Then, the RIB for declaring this creature with material assignments might look like this:

```
AttributeBegin
Surface "skinShader" [ shader param settings ]
Procedural "DynamicLoad" [ "RiGtoPlugin.so" "thing.ref.gto thing.0013.open.gto thing.0" ]
AttributeEnd
```

```
AttributeBegin
Surface "hairShader" [ shader param settings ]
Procedural "DynamicLoad" [ "RiGtoPlugin.so" "thing.ref.gto thing.0013.open.gto thing.0" ]
AttributeEnd
```

```
AttributeBegin
Surface "eyeShader" [ shader param settings ]
Procedural "DynamicLoad" [ "RiGtoPlugin.so" "thing.ref.gto thing.0013.open.gto thing.0" ]
AttributeEnd
```

Because of RiGtoPlugin's cache mechanism, the geometry associated with the file-set thing.*.gto is only read and interpreted one time - the on-lists control which parts of the geometry are instantiated at which times. To nuke the cache of these files (if memory is important), you would use the syntax:

```
Procedural "DynamicLoad" [ "RiGtoPlugin.so" "thing.ref.gto thing.0013.open.gto thing.0" ]
```

13.8.7 Miscellaneous RenderMan Stuff

RiGtoPlugin stores some useful data in attributes that can be used by shaders if desired.

Pref [Shader parameter]

On ALL geometry RiGtoPlugin creates "varying point Pref" as part of its geometry declaration. This data can be accessed by simply putting "varying point Pref" in your shader parameters. The position of the model in the reference GTO file is always used for this parameter value.

Name [RIB Attribute]

RiGtoPlugin always places the name of the geometry, as retrieved from the GTO file, in an attribute that may be queried. It is exactly as if the following line of RIB were declared before the geometry were instantiated:

```
Attribute "identifier" "name" ["whatever my name is"]
```

RefToWorld *matrix* [RIB Attribute]

RiGtoPlugin places the transformation matrix *objectToWorld* from the reference model into a user attribute called `refToWorld`. To prevent this attribute from being munged by the current transformation matrix, it is cast as a float[16] instead of a matrix. It is equivalent to this line of RIB:

```
Attribute "user" "float refToWorld[16]" [ the matrix values ]
```

Appendix A Description of Changes

- Version 4.0
 - Property data size can now be larger than 4Gb in the file and readable in full on 64 bit architectures (if the library is compiled 64 bit). The file headers are now version 4. The new file is incompatible with version 3 GTO readers.
 - The manual has been updated with real-world usage examples from the film and game industries.
- Version 3.4
 - The GTO license terms have been changed: the code is still covered by the LGPL, but with additional **exceptions** similar to those used by the FLTK library. These exceptions make it easier to use GTO in commercial projects.
 - Houdini I/O plugin (ggto) reads and writes GTO geometry.
 - The library no longer attempts to be source code compatible with older Microsoft compilers. Some functions may throw on error.
 - Maya plugin (loadGtoAnim) loads animation from GTO files (transform matrices only). Useful for getting animation from GTO difference files.
 - Maya plugin (GtoDeformer) makes it possible to use GTO files as geometry cache files. The deformer will read vertex/cv positions from existing GTO files and applies them to scene geometry of the same name.
 - Maya plugin (GtoParticleDisplay) loads particles from GTO files for viewing in Maya
 - Maya plugin (GtoParticleExport) writes particles from Maya as GTO files. Can be used as a replacement for pdb and pdc files.

- Maya plugin (GtoCacheEmitter) loads particles into Maya from GTO files via an emitter.
- Bug fixes to the C++ Gto::Writer class for output of text GTO files.
- Run-time error checking of the Gto::Writer API. The class will complain if the API is used in a undocumented manner. It may throw an exception.
- The GTO source code distribution now comes with Maya and Houdini plugins for cached deforming geometry and particle export and display.
- Version 3.2
 - Human readable plain text version of GTO. Some readers may not function if they assume that the property size is known when the property() virtual function is called. The property size is only really known when the data() virtual function is called. Only version 3.2 GTO readers can read the text version.
 - Animation curves are now stored per-object using the animation protocol.
 - Bug fixes to Gto::Reader class to allow reuse of existing class with a newly opened file.
- Version 3.1
 - RenderMan plug-in documentation added.
- Version 3.0
 - An interpretation string has been added to the property header.
 - An additional uint32 has been added as padding to the object, component, and property headers for future expansion slop.
 - A section on interpretation strings has been added to the documentation and to the reader/writer classes.
 - Added a type reference to the documentation.
- Version 2.1
 - `gtofilter` was changed to optionally accept POSIX style regular expressions in addition to shell-like “glob” expressions.
 - The C++ writer class now defaults to writing compressed files when the `open()` function is called. A second bool argument can be passed to it to prevent the compression.
 - The proposed texture assignment protocol (from version 2.0.4) has been rejected.
 - A new protocol “channel” is introduced for assigning mapped surface varying data on geometry. An arbitrary number of texture maps may be assigned to the geometry. See [Section 7.14 \[Channels\]](#), page 34.
 - The material protocol has been fleshed out. See [Section 7.9 \[Material\]](#), page 32.
 - The polygon protocol was missing the definition of the optional **mappings** component. See [Section 7.6 \[Polygonal Surfaces\]](#), page 27.
- Version 2.0.5 Bug fix version. Repaired problems with the configuration scripts. Missing headers.
- Version 2.0.4 Bug fix version. Some configuration problems solved.
- Version 2.0 File headers changed. The format is not compatible with 1.0.
- Version 1.0

Appendix B Reference

Properties

C

channels	34
connection_type.lhs	33
connection_type.rhs	33

E

elements	26, 28
elements.normal	29
elements.size	26, 28
elements.smoothingGroup	29
elements.type	28
elements.width	26

G

genre	32
-------------	----

I

image	31
image.cs	31
image.encoding	32
image.pixels	31
image.size	31
image.type	31
indices	29
indices.normal	29
indices.st	29
indices.vertex	29

M

mappings	29
material	32

N

Name	57
normals	28
normals.normal	28

O

object	24
--------------	----

object.boundingBox	24
object.globalMatrix	24, 25
object.name	24
object.parent	24, 25
object.protocol	24
object.protocolVersion	24

P

parameters	32
points	25, 26, 27, 28
points.id	25
points.normal	28
points.position	25, 26, 27, 28
points.velocity	25
points.weight	27
Pref	57

R

RefToWorld	57
------------------	----

S

shader	32
shells	34
shells.elements	34
shells.vertices	34
smoothing	29, 30
strand	26
strand.type	26
strand.width	26
surface	27
surface.degree	27
surface.uKnots	27
surface.uRange	27
surface.vKnots	27
surface.vRange	27

T

TWK_RI_GTO_NO_CACHE	56
TWK_RI_GTO_NO_SUBDS	55
type	32

Functions

G

<code>gto.Reader</code>	46
<code>gto.Reader.accessObject</code>	48
<code>gto.Reader.component</code>	47
<code>gto.Reader.components</code>	47
<code>gto.Reader.dataRead</code>	47
<code>gto.Reader.isSwapped</code>	47
<code>gto.Reader.object</code>	46
<code>gto.Reader.objects</code>	47
<code>gto.Reader.open</code>	46
<code>gto.Reader.properties</code>	47
<code>gto.Reader.property</code>	47
<code>gto.Reader.stringFromID</code>	47
<code>gto.Reader.stringTable</code>	47
<code>gto.Writer</code>	48
<code>gto.Writer.beginComponent</code>	48
<code>gto.Writer.beginData</code>	49
<code>gto.Writer.beginObject</code>	48
<code>gto.Writer.close</code>	48
<code>gto.Writer.endComponent</code>	48
<code>gto.Writer.endData</code>	49
<code>gto.Writer.endObject</code>	49
<code>gto.Writer.intern</code>	49
<code>gto.Writer.lookup</code>	49
<code>gto.Writer.open</code>	48
<code>gto.Writer.property</code>	48
<code>gto.Writer.propertyData</code>	49

R

<code>Reader::~Reader</code>	39
<code>Reader::~accessObject</code>	42
<code>Reader::~charnum</code>	40
<code>Reader::~close</code>	40
<code>Reader::~component</code>	41
<code>Reader::~components</code>	42
<code>Reader::~data</code>	41
<code>Reader::~dataRead</code>	41

<code>Reader::~descriptionComplete</code>	41
<code>Reader::~fail</code>	40
<code>Reader::~fileHeader</code>	40
<code>Reader::~header</code>	41
<code>Reader::~in</code>	40
<code>Reader::~infileName</code>	40
<code>Reader::~isSwapped</code>	40
<code>Reader::~linenum</code>	40
<code>Reader::~object</code>	41
<code>Reader::~objects</code>	42
<code>Reader::~open</code>	39
<code>Reader::~properties</code>	42
<code>Reader::~property</code>	41
<code>Reader::~Reader</code>	39
<code>Reader::~readMode</code>	40
<code>Reader::~Request::Request</code>	41
<code>Reader::~stringFromId</code>	40
<code>Reader::~stringTable</code>	40
<code>Reader::~why</code>	40

W

<code>Writer::~~Writer</code>	43
<code>Writer::~beginComponent</code>	44
<code>Writer::~beginData</code>	45
<code>Writer::~beginObject</code>	44
<code>Writer::~close</code>	44
<code>Writer::~endComponent</code>	44
<code>Writer::~endData</code>	45
<code>Writer::~endObject</code>	44
<code>Writer::~intern</code>	44, 45
<code>Writer::~lookup</code>	45
<code>Writer::~open</code>	43, 44
<code>Writer::~properties</code>	45
<code>Writer::~property</code>	44
<code>Writer::~propertyData</code>	45
<code>Writer::~propertyDataInContainer</code>	45
<code>Writer::~Writer</code>	43

Types

3		
3x3.....	22	
4		
4x4.....	22	
A		
ABGR.....	23	
B		
bbox.....	22	
bezier.....	23	
BGR.....	22	
bool.....	20	
byte.....	20	
C		
column-major.....	22	
complex.....	22	
coordinate.....	22	
D		
double.....	20	
F		
float.....	20	
G		
gto.ComponentInfo.....	50	
		gto.ObjectInfo..... 49
		gto.PropertyInfo..... 50
H		
half.....	20	
homogeneous.....	22	
I		
indices.....	22	
int.....	20	
int64.....	20	
N		
normal.....	22	
Q		
quaternion.....	22	
R		
RGB.....	22	
RGBA.....	23	
row-major.....	22	
S		
short.....	20	
string.....	21	
W		
weighted.....	23	