

The monthly GNU Column

BRAVE GNU WORLD



GNU grep

ard Rosenkraenzer is the new maintainer of GNU grep, so I'd like to take opportunity to write a little bit about project. Most GNU/Linux users probably already know GNU grep and use it on a daily basis, but there may be some who are not yet familiar with it, so I'll give a introduction to its functionality. GNU grep searches files or standard input streams – usually text strings – and

outputs lines matching or containing this pattern. It is also possible to search multiple files in order to determine which files contain the pattern. Typical uses go from processing texts and finding passages, to narrowing the output of other programs down to relevant information.

grep is definitely one of the standard commands of any Unix-like system and GNU grep contains not only the standard features but also options like recursive search of directories or word-based output.

Since it is already being used for many years on a multitude of systems, there can be no doubt about its readiness for daily use. I'm not quite sure about the exact age of GNU grep, but the ChangeLog goes back to 1993 and the first copyright notice by the FSF originates in 1988, which is about three years before the creation of the Linux kernel. This feature is interesting not despite but because of the age of the project, because it shows how lively Free Software is – even when it reaches such a high age.

During those years, over 80 people contributed to GNU grep and development still continues. On the ToDo list are some small changes for complete POSIX.2 conformance and a few new options. Among them are for instance a `--max-count` option, a switch for PCRE (Perl Compatible Regular Expressions) or a highlighting option for colour-marked output.

Even such an old, stable and widespread package needs help. In particular users willing to test development versions for portability beyond GNU/Linux and FreeBSD and bug-hunting are very welcome. Also the current maintainer only speaks

English, German and French, so he cannot truly determine whether the multibyte-support is actually working. So users in multibyte-areas should feel especially encouraged to participate in the testing.

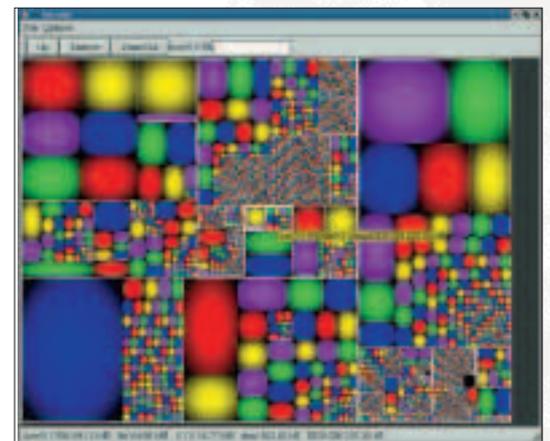
QTreeMap

The QTreeMap project by Alexander Rawass implements Treemaps in a Qt-Widget under the GNU Lesser General Public License. Deep hierarchies or trees are usually displayed as a kind of structure that can be (un)folded via mouse click. Treemaps offer the capability to display these complete hierarchies at a single glance.

The principle is understood more easily when giving an example, so I'll explain it with the operation of KDirStat, which uses QTreeMaps in order to visualise the hard disk usage.

QTreeMaps are displayed in rectangular areas. The total area of the rectangle represents the total size of the partition to be visualized. Directories and files are displayed as areas proportional to their size. A directory structure using a third of the available space on a partition would also get one third of the available display area. A subdirectory of this structure containing half the size of the whole structure would get half of its display size and so on.

Treemaps are especially useful in situations where size-dependent hierarchies have to be displayed, like



Squarified Treemaps with Fast Bump shading

Welcome to another issue of Georg CF Greve's Brave GNU World. After introducing a classic, this issue will soon become relatively technical, but it should be interesting even to not-so-technical readers



filesystems, network traffic or content/organisation management.

QTreeMap supports classic Treemaps, quadratic Treemaps and different colouring schemes (also based on regular expressions). The generated Treemaps can be loaded and saved as XML and saving them as bitmap is also possible.

According to the author, the special problems of QTreeMap are the usual KDE-specific problems as well as lack of the Cushion algorithm, which is present in Sequoia View, the proprietary program inspiring QTreeMap, but which could not be implemented because Alexander wasn't able to figure out the mathematical concepts for it.

QTreeMap was written in C++ after Alexander experimented with the algorithms in Python. According to him, there are two similar projects, but both of them are written in Java. As the projects KDirStat and KProf prove, QTreeMap is usable, so interested developers should take a look.

Dap

Another new member of the GNU Project is Dap by Susan Bassein. Dap, which stands for Data Analysis and Presentation, is released under the GNU General Public License.

Dap provides basic functions for data management, analysis and graphical visualisation as they are commonly used in statistical consulting and education. Also it is useful for managing sets of data; Susan herself uses Dap to do her taxes and prepare the payroll for her employees.

The program is written in C and users with C experience should have no problem using Dap after studying the examples provided with the package.

Since the GNU Project already had one statistical package called R, it now provides two alternatives. Whilst R is object-oriented, Dap follows the procedural approach. Users of the proprietary programs S or S-plus will probably prefer R. Former users of the non-Free SAS package will quickly feel at home with Dap. Dap is also more memory-friendly: while R reads the whole file into memory first, Dap works line-oriented which makes it suited for very big sets of data.

Problems mentioned by Susan are that Dap has fewer statistical tests than R and was never optimised for speed. Fixing these weaknesses as well as expansion and improvement of its functionality is the goal of further development.

Despite these problems Dap has been used for about three years now, so it has been thoroughly tested and can be recommended to interested users.

GNU ccRTP

The Goal of the ccRTP project is to implement the RFC standards for the Realtime Transport Protocol (RTP) which allows transport of time-dependent data

GNU UnRTF

GNU UnRTF is a recent addition to the GNU Project by Zachary T Smith. It enables people to transfer documents from the Rich Text Format (RTF) into other formats. RTF is often used as a transfer format by Windows users, but other text processors also use it for saving text with formatting information.

Thanks to this project it is now possible to convert these documents into pure text, HTML, LaTeX and PostScript. So anyone who uses RTF themselves or is dependent on people who do should profit from this project.

The license of the project was always the GNU General Public License, but it is possible that some people already got in touch with the project under its old name, rtf2htm.

The current focus of the development process is two-fold: the character conversation routines and the output to LaTeX. Additionally it is planned to support more target formats in the future. Regardless of the ongoing development, the project is definitely ready for use.

(like audio or video) over a network.

The streams are transmitted in packets containing time information to enable correct assignments on the receiving end. Typically this is done by UDP packets because network problems do not block the transmission, which would destroy the synchronisation.

ccRTP implements this object-oriented in a C++ class library under the GNU General Public License. The authors, David Sugar and Frederico Montesino Pouzols, want ccRTP to become the most versatile, efficient and standards-compliant RTP implementation and they have taken several steps in this direction so far.

GNU ccRTP already supports multicasting as well

ccScript

GNU ccScript is, just like Common C++, Bayonne, ccRTP and cc Audio, a project under maintenance by David Sugar, who gladly stated this makes him a Penta-Maintainer.

The ccScript C++ class library under the GPL provides a virtual machine (VM) for real-time applications in state-transition event driven systems. This assembler scripting language is being used by GNU Bayonne and other parts of the GNUComm project to script the user interaction.

As can easily be understood, defined execution speed is very important for real-time environments, which is exactly what ccScript has been written for. Any operation in ccScript is deterministic; the only exceptions to this rule are operations where this is impossible, like database lookups and the like. Because of this it cannot deal with complex or open-ended expressions.

Even if it provides general functionality and macros, ccScript should not be confused with projects like Guile or Tcl, because they are more versatile but do not have matching real-time capabilities.

As the next steps in ccScript development, the syntax will be restructured and clarified a bit; also more parts of the language should be available in loadable modules. This project also needs help with its documentation – so if you feel you want to do this, please do.

as point-to-point transmission, multiple input-streams and the Real-Time Control Protocol (RTCP). Additionally the transition to IPv6 has been prepared and even real-time packet filtering and mixed-mode data streams are possible, which allows features like RFC 2833 signalling inside a data stream, for instance.

The high transmission rates required for video data as well as partial packet reconstruction and Class Of Service routing are possible, so the library can already be used for clients and servers.

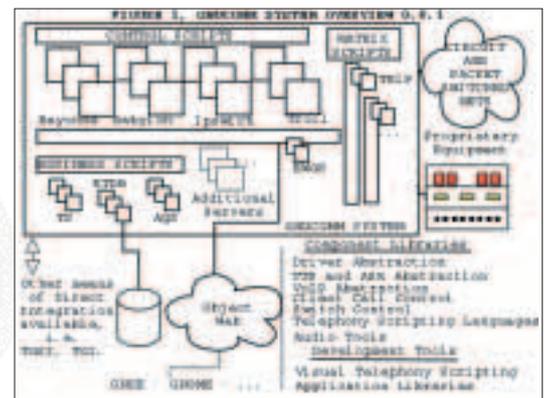
Work is currently being done on the Resource reSerVation Protocol (RSVP) and completion of the RTCP-support. The next step will be proper project documentation, which is lacking right now.

Help by real-time specialists and good authors of documentation would be very welcome. The latter in particular are harder to find than most people would assume.

GNUComm

Other projects pursue the goal of furthering the GNUComm meta project, which is about creating a complete and flexible communication environment based on interactive components.

One part of the GNUComm meta project is Bayonne, the telephony server. One of the focal aspects of GNUComm development is interoperability and integration with the GNU Enterprise project. GNU Enterprise works on a complete solution in the field of the so-called Enterprise Resource Planning (ERP) and will bring Free Software into this largely



Overview of GNUComm

proprietary area; more information can be found in issue seven of Linux Magazine.

GNU ccAudio

The ccAudio project was also started by David Sugar. As the name suggests, its goal is the creation of a general purpose library for manipulation of audio data on hard disk and in memory.

Like ccRTP, ccAudio is implemented as a C++ class library under the GNU General Public License and is, since it originates in the GNU Bayonne project, a part of the GNUComm meta project.

Currently ccAudio supports accessing audio data from the hard disk through libsndfile and other libraries and provides basic signal/audio processing facilities. It treats audio data as discrete samples and can deal with RIFF headers and such.

David considers the treatment of audio data as sets of samples rather than binary buffers in byte format to be especially notable. It is also aware of different sample encoding formats, endian ordering, and multiple channels.

Possible platforms for ccAudio are Unix-like systems as well as Win32, so developers in this area should feel encouraged to take a look at it.

Further development will aim towards better support for dynamically loadable software-codecs and also making more built-in audio-codecs available. David is also considering including fourier-transforms (FFT) and different forms of audio mixing and transforming into ccAudio. Someone with more experience in the digital audio/signal processing would be very helpful for this.



Enough

Enough said for this month's issue. As a last thing I'd like to point out that the GNU Transport Layer Security Library (GNUTLS) is looking for a new logo and has started the GNUTLS Logo Contest. If you'd like your artworks to become part of the GNU Project, this is a good opportunity.

Of course I'm asking for lots of impulses, ideas, comments, criticism and new projects by mail. ■

Info

Send ideas, comments and questions to

Brave GNU World	column@brave-gnu-world.org
Homepage of the GNU Project	http://www.gnu.org/
Homepage of Georg's Brave GNU World	http://brave-gnu-world.org
"We run GNU" initiative	http://www.gnu.org/brave-gnu-world/rungnu/rungnu.en.html
GNU grep homepage	http://www.gnu.org/software/grep/
GNU UnRTF homepage	http://www.geocities.com/tuorfa/unrtf.html
QTreeMap homepage	http://qtreetmap.sourceforge.net
KDirStat homepage	http://kdirstat.sourceforge.net/kdirstat/
KProf homepage	http://kprof.sourceforge.net
Dap homepage	http://home.earthlink.net/~sbassein/public
GNUComm project homepage	http://www.gnu.org/software/gnucomm
GNU Bayonne homepage	http://www.gnu.org/software/bayonne
GNU Enterprise homepage	http://www.gnuenterprise.org
GNU ccRTP homepage	http://www.gnu.org/software/ccrtp
GNU ccAudio homepage	http://www.gnu.org/software/ccaudio
GNU ccScript homepage	http://www.gnu.org/software/ccscript
GNUTLS Logo Contest	http://www.gnu.org/software/gnutls/logo-contest