

Daniel James talks to developers working on the GStreamer multimedia framework

Surfing the Pipeline

GStreamer is a free software library, written in C and released under the GNU LGPL. Designed to offer a clean and simple API for both plugin and application developers, it allows the modular construction of programs including media players, audio mixers and non-linear video editors. A multi-platform technology, GStreamer has been ported to a range of operating systems and processors, including Linux on x86, PowerPC, and ARM. It is also available on Sun Solaris, Apple OS X, Microsoft Windows and IBM OS/400.

While GStreamer is not a sound server that operates between applications in the manner of JACK, it does allow the interconnection of audio and video code. Graphs of media components can be put together quickly using the GStreamer Pipeline Editor, a graphical manipulation tool based on GNOME Canvas. The editor allows prototype applications to be run immediately, and configurations saved in an XML format. For those developers who prefer the command line, there is also the `gst-launch` tool, which is described as being similar to Ecasound.

Wim Taymans is one of the two founders of the GStreamer project, and is leading the design of the forthcoming 0.9 series releases. A native of Belgium, he has been programming since the age of eight. "When I was about 16 I got involved in the Commodore 64 demo scene. I installed Linux for the first time in 1994 on my Amiga. With my background as a demo developer, I knew some assembler, and did my first assembler patches for the RTjpeg codec. I helped on the Trinity video editor project for a while until we found out that we needed a good multimedia framework first. That's when I got involved with GStreamer, which had its first version released. I very much liked the modular

design and started adding plugins for video and audio decoding. Later on I got more involved in the core design as well."

The use of GStreamer avoids duplicated effort in the creation of audio and video applications, but according to Taymans, there's more to it than that. "First of all, most multimedia applications face the same problems. They need to pass data around at high speeds, need to handle synchronisation and must be able to handle a multitude of codecs and formats. At the core, most existing players implement these things over and over again. Still, we might consider ourselves lucky that most of the codecs are abstracted in separate libraries - including `ffmpeg`, `mad`, `lame` and `libmpeg2`."

"Secondly, most existing applications only focus on playback. Free software is at the stage where we have a few very good media player applications, but reusing the existing

The biggest challenge is finding a way to integrate patent encumbered formats

framework of those applications to make them do something else - like live capture or non-linear editing - proves to be difficult. Thirdly, multimedia is a patent minefield. A licence is required to legally decode or encode various media codecs such as MP3, DVD, WMA and WMV. The companies granting these licences sometimes require closed source implementations. The biggest challenge for multimedia in free software is finding a way to integrate patent encumbered formats in a legal way."

The work of the GStreamer project, although started independently, is now part of the GNOME desktop. GStreamer forms the basis of the sound mixer included with

GNOME 2.10, and is also the back-end to the Totem media player. However, to illustrate the point that the code is not tied to GNOME, the GStreamer website has moved to a subdomain at the `freedesktop.org` site. There has also been some discussion among KDE developers about including GStreamer as part of the forthcoming KDE 4 release, perhaps as one of several optional media back-ends.

FLUENDO AND FLUMOTION

GStreamer development and other related projects are supported by Fluendo, a company based in Barcelona. Other than employing several GStreamer developers, Fluendo has sponsored Xiph.org to work on the Ogg Vorbis and Ogg Theora codecs, and has also funded the development of Totem. As Taymans explains, Fluendo has built its business around the GStreamer framework, and has already released the first fruit - the Flumotion streaming server product. "Fluendo hired some of the core GStreamer developers to improve and extend it. We

use GStreamer in the server product, for example, but we are also investing a lot of resources in improving GStreamer as a player for the desktop. We're also investing in a non-linear video editor, PitiVi."

The Flumotion server is written in Python, and acts as the glue which holds together all the necessary GStreamer pipelines and plugins. The media processing is done within GStreamer for performance reasons, while the configuration of the various elements is done using the `gst-python` bindings. Flumotion currently supports free codecs including Ogg Vorbis and Theora, with streaming of non-free formats planned for the near future. It currently uses HTTP to transport the streams to the clients, but RTP/

RTSP support in both unicast and multicast is also expected to be added.

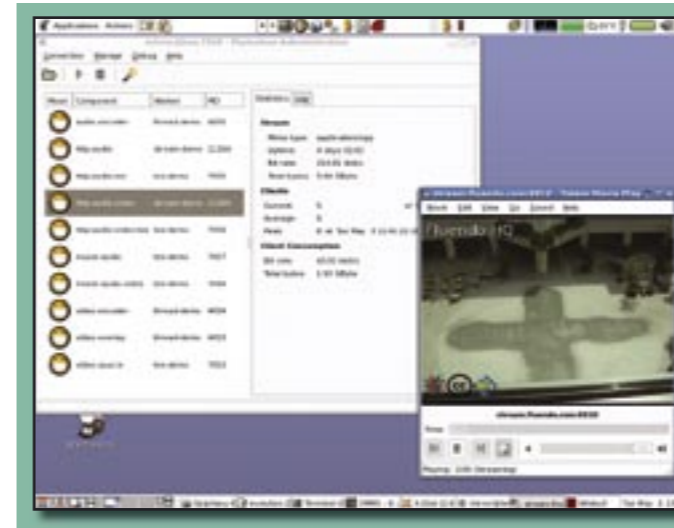
Thomas Vander Stichele is the lead developer on Flumotion, and is also the GStreamer release manager. Like Taymans, he was born in Belgium. While at university he helped start a radio station, Urgent FM. "During this radio work I got to know free software from a user perspective - our admin deployed Linux for the servers at the station. In my final year, I managed to combine my studies and my personal interests by doing a thesis on an optimised technique for mixing MP3 audio. The code written for this thesis was used at the radio

shipped first, but Real and MPEG codecs will follow quickly."

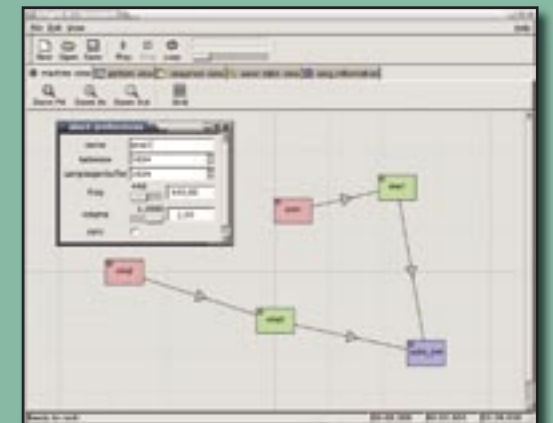
PRO AUDIO WITH GSTREAMER

Now that content creation with GStreamer is looking feasible, Taymans indicates that the framework will soon have better support for users working with audio. "I personally want to pay more attention to the pro audio environment, which should be visible in the redesign done in the 0.9 branch. In a pro audio application, the processing is driven from the audiosink. This means that the audiosink will periodically request data from the pipeline to

leading the effort to make sure the future GStreamer 0.10 release is even more pro-audio friendly. He is also the lead developer behind Buzztard, a free software replacement for the Buzz music-making environment on Windows. A PhD student based in Germany, Kost has been developing software for half of his 31 years. "My first computer was an AMIGA, for which I developed a shareware sound processing application called SoundFX. During the last four years I worked on my PhD project, which is open source and developed on the UNIX platform, and in the last year I've been working on Buzztard."



The Flumotion streaming media server was built using GStreamer



Buzztard is a GStreamer-based application

station - we were probably the first radio station streaming on the internet in Europe, and we were doing a nicely mixed MP3 stream on a lowly Pentium 200 MMX."

"After my studies, I started as an engineer at a new commercial radio station, 4FM. Various people, including myself, have created applications using GStreamer for very specific deployments (in my case, recording separate items from a newscaster reading out the news live on radio, and triggering RDS feeds based on high-frequency tones in the radio signal). After being CTO there for a year, I moved to Barcelona and helped start up Fluendo."

"The creators of GStreamer felt that a company like Apple would probably never open-source a Quicktime plugin, but they might be willing to release a binary-only one that could be used from this open source framework. Of course, all of the code implementing patented formats is pretty much impossible for distributors to ship. So users are left to get packages from 'questionable' sites. Fluendo as a company will be selling plugins for codecs like these, so that GNU/Linux users can finally play back these formats completely legally. The Windows Media codecs will probably be

write to the audio card. With the 0.8 branch it was somewhat possible to drive the pipeline from the audiosink, but not optimally. The scheduling changes in the 0.9 branch should make this much more efficient."

Taymans explains that version 0.9 is being redesigned to be threadsafe. "This was one of the biggest problems in 0.8, as most multimedia pipelines at some point do multiple tasks at the same time, like audio and video decoding. It is now also explicit that all media processing happens in separate threads to the application. This allows for a wide range of new possibilities like processing data in the callback from JACK, allowing us to do low-latency. The downside is that multithreaded programs are harder to write and debug."

"A lot will change in the way the pipeline is scheduled. We've removed one of the more scary and error prone parts of the core and replaced it with something much more simple based on threads. Elements will be more in control of how they receive and process data, making it possible to make more interesting elements."

Stefan Kost is one of the main audio developers on the GStreamer project, and is

Kost got involved with the GStreamer project after an extensive search for a good media framework to base Buzztard on. "On one hand I needed something which supported playing audio on several platforms, and on the other hand I was looking for something which would help me to control plugin graphs. I've found those needs answered by the features of the GStreamer project, and I've started to learn about the API by helping out with the docs. I've discussed the features I need with the core developers, and these extensions are now being added to GStreamer one by one. Managing a plugin graph of media generators, processors and sinks is not as easy as it sounds. I have a strong belief in the synergistic effect of having several application developers discussing their needs and designing a unified API that suits all of them."

Key Links

- GStreamer gstreamer.freedesktop.org
- Fluendo www.fluendo.com
- Urgent FM www.urgent.fm
- 4FM www.4fm.be
- Buzztard www.buzztard.org