Pluggin' away

Steve Harris is one of the most prolific DSP programmers working on Linux audio software, with a large collection of effects plugins and other software to his name. He talks about his work to Daniel James

DSP, or digital signal processing, has had a profound effect on the audio industry. Recording studios were once dependent on large racks of analogue equipment, because their engineering tradition mirrored the UNIX philosophy of 'one box to do one job well'. A hardware effects unit might be plugged in only to add reverberation, while another would exist solely for limiting the output signal to tape; complex routing of analogue signals meant a spaghetti-like tangle of patch cables was required. These factors meant that audio production became highly capital and maintenance intensive, and not very portable either.

Once audio is broken into digital chunks, it becomes possible to perform these signal processing functions using general-purpose computing hardware. If the processor is sufficiently powerful that the manipulation of the audio can take place in real time - which in this context means with latency acceptable to the human ear - then digital hardware can replace specialised analogue equipment. It's only been in the last few years that typical desktop PCs have started to become capable of bearing serious real-time DSP loads on the CPU alone.

Now it's possible to emulate a full room of audio hardware with a single laptop, or an embedded system in a 2U case. While purists may claim that the digital emulations might not have all the qualities of the original equipment, certain facts are undeniable. Not only can audio production be highly portable, and high-quality recordings be made at a much lower cost, but the technology can be made available to a far greater number of people than ever had access to traditional studio equipment.

These pure software DSPs are inserted into the signal chain within or between applications, and are therefore known as 'plugins', following the analogue hardware tradition. Various proprietary formats for these plugins exist, but the world of free software has its own, open standard: LADSPA. The Linux Audio Developers' Simple Plugin API now has well over 100 plugins available, covering most of the typical audio processing tasks and more besides.

HOSTS BUT NO PLUGINS

Steve Harris has probably written more of the LADSPA plugins than any other single developer, with 84 released to date. An academic programmer working on artificial intelligence projects at the UK's University of Southampton, he is also a musician. Part of his time is spent on his DSP research interests, the fruits of which have all been released as free software.

Harris started exploring the possibilities of using Linux for music production in the mid-90's, but like most people at the time still had to rely on dedicated studio hardware. He'd also used SGI and Solaris machines, the Amiga for MIDI and sample processing, and later tried Windows PCs. "I didn't really feel that Windows was a very stable platform for audio work, and I have never enjoyed developing under Windows, so I never wrote my own software for that platform. I wrote some very primitive DSP software for the Amiga around '92, but I had little understanding of the theory, and there was limited realtime processing I could do with the Amiga's CPU. So there wasn't as much satisfaction as with realtime plugins."

I didn't really feel that Windows was a very stable platform for audio work

"I decided to write LADSPA plugins when I discovered the Linux music scene, read up on the LADSPA API, and discovered that there were several hosts but not many plugins. I took a liking to the API, as it is very clean and simple, and regular enough that I could use code generators to handle most of the mundane housekeeping. Most of my plugins so far have been software implementations of known techniques, and they vary greatly in complexity. Some are very simple, and took a few hours to write, while some took weeks."

"I read up on DSP theory when I started coding for LADSPA. I had some experience with modular synths already so I knew the high level stuff, but the inner workings of filters, for example, were a mystery to me at that point. It's a very interesting subject, and I still have a lot to learn." Other than the advantage of being a free software standard, LADSPA has some technical benefits over competing proprietary plugins, such as Steinberg's widely-used VST format. "The API is much cleaner than VST as it hasn't been evolved to the same extent. The authors learnt from the mistakes made in VST specification development, so there is less cruft."

Plugin automation, which allows DSP settings to be recorded, saved with a project and played back along a timeline, is a killer feature of digital audio mixers. In the analogue days, the best sound engineers could do was mark the position of all the knobs in the rack on a sheet of paper, an approach which was time consuming, and hardly dynamic. Fortunately, LADSPA was created with automation in mind: "By design, it is possible for the host to automate the controls of any plugin, without any explicit support from the plugin author. VST has automation support, but it is somewhat complicated, and not widely supported. On the other hand VST has some features that LADSPA does not, such as custom graphical user interface support, which is not really standard in LADSPA yet."

Plugin automation is a killer feature of digital audio mixers

Apple's part-proprietary UNIX, or OS X, has its own, incompatible plugin standard, known as Audio Units. Due to factors including the legacy of Apple hardware in recording studios and the acquisition of niche software house Emagic, which produces the Logic sequencing application, Apple remains significant in the studio market. "LADSPA is very different to the Audio Units specification. AU has many features that LADSPA does not, but the API is much more complicated. It's really quite hard to compare them - from the users point of view, the main difference is probably that AU supports MIDI inputs, which is not handled by LADSPA, though it is supported by DSSI, the Disposable Soft Synth Interface, which is an extension to LADSPA" (see the Audio Libre column in last month's LinuxUser & Developer, issue 44). "DSSI is not as widely supported as LADSPA yet though."

The LADSPA SC4 compressor plugin working in Ardour

JAMin provides several mastering tools in one GUI

JAMIN 'TIL THE JAM IS THROUGH

Harris has written several applications for Linux, including the JACK Meterbridge, an emulator of traditional audio meters, and the Time Machine recorder that buffers audio output in order to capture sounds that conventional recorders would miss. However, the most prominent project he's initiated is JAMin, a tool for audio 'mastering', or 'post-production' as it is known in the broadcast industries. The original meaning of the term mastering is somewhat obsolete, as it referred to the process of making master discs for record stamping, but today it usually means making audio recordings suitable for a potentially diverse range of target devices, be they huge club sound systems or portable radios with tiny speakers. To oversimplify what is undoubtedly an art form in its own right, audio mastering involves raising and balancing subjective volume level for the listener, without exceeding the limits of the distribution technology so that audible distortion occurs on playback.

JAMin, or the Jack Audio Mastering Interface, is a 'missing link' application for free software, given that tools for both recording audio and burning CDs on Linux already existed. Without a mastering tool, music created on Linux systems wouldn't sound as loud, or consistent, as that created with proprietary software or dedicated hardware. "The motivation for JAMin was when Jörn Nettingsmeier pointed out that some of my recordings sounded bad on headphones due to stereo positioning problems (he was being polite). I did some reading around the subject, and discovered that Linux had no real mastering toolchain. So that led to some developments in JACK Meterbridge, and inspired me to start work on what would become JAMin."

"I never really got it working well enough to be useful, and the UI that I built was terrible, so the DSP codebase just sat around for a year or so until I put out a call for developers. Then Jan Depner, Jack O'Quin, Ron Parker, Patrick Shirkey and I formed the JAMin project. I didn't have the time, experience or motivation to complete it on my own, and without them it would never have got anywhere."

JAMin is based on DSP plugins, but they are gathered together into a unique interface which provides the standard features needed by mastering engineers. There is a 1024-band equaliser with a hand-drawn curve control, in addition to a more conventional 30-band graphic equaliser, with a spectrum analyser which can highlight problem areas of the frequency range. A three-band compressor with adjustable crossover points and a look-ahead limiter allow average level to be raised without exceeding the maximum peak possible with digital audio playback. JAMin does not have the ability to record - instead, it sits between JACK-enabled applications, processing the audio streams that pass through.

"At the time I started I hadn't heard any Linux-made recordings that could be compared to commercial recordings, so my motivation was mostly for the challenge. JAMin as it stands now is massively overspecified to fix the problems I had in my recordings - which I still haven't found time to fix!"

Key Links

The LADSPA standard
www.ladspa.org

Steve Harris's plugin collection
www.plugin.org.uk

Meterbridge
www.plugin.org.uk/meterbridge/

Time Machine
www.plugin.org.uk/timemachine/

JAMin mastering tool
jamin.sourceforge.net