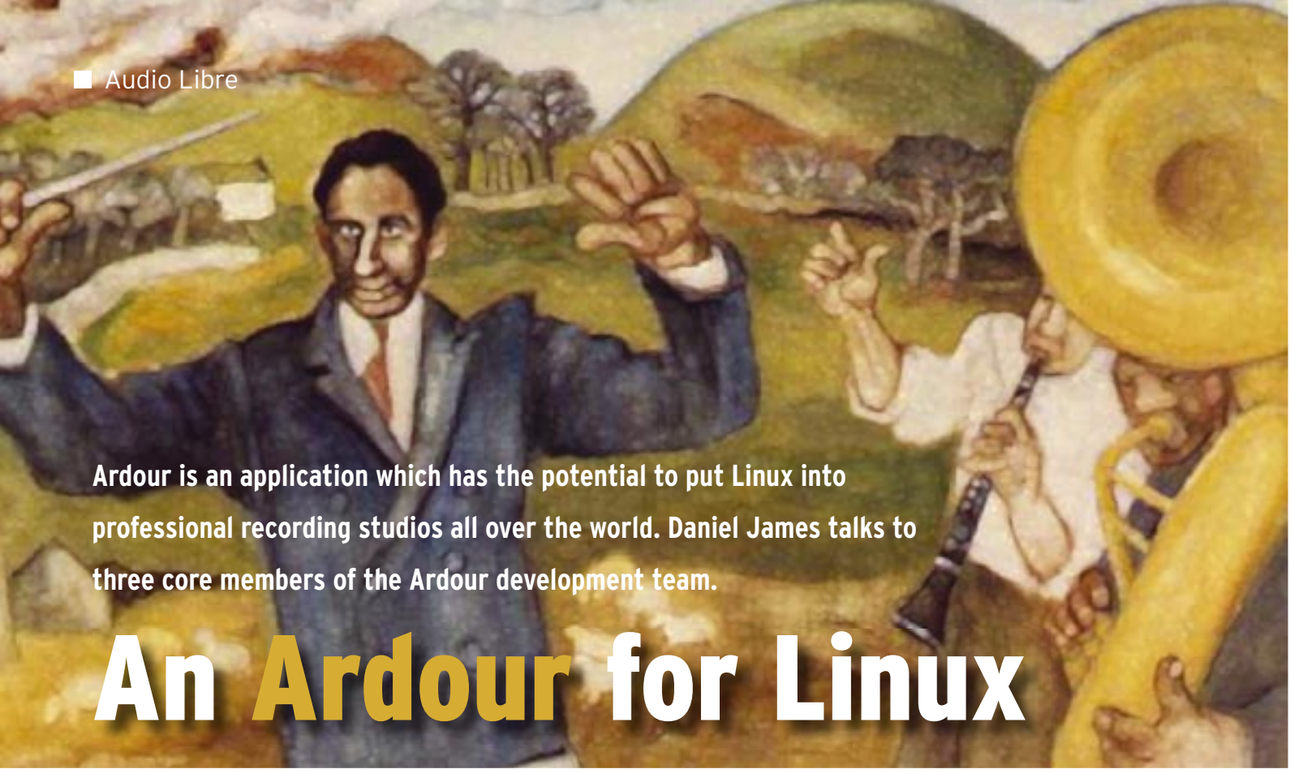

Ardour is an application which has the potential to put Linux into professional recording studios all over the world. Daniel James talks to three core members of the Ardour development team.

An Ardour for Linux

While recording studios once depended on analogue tape machines and a packet of razor blades, replaced through the 1980's and 1990's by digital tape formats, hard disc based recording and editing systems are now almost ubiquitous. These can take the form of rack-mounted units perceived as pure hardware by musicians, even though they are almost certainly running some form of embedded system. More commonly, they are found integrated with a desktop Mac or

Digidesign's ProTools enjoys considerable vendor lock-in within the audio industry

Windows machine, marketed generically as a 'digital audio workstation'. The market leader in DAWs is Digidesign's ProTools, which is based around proprietary DSP cards and audio interfaces, and therefore enjoys considerable vendor lock-in within the audio industry.

As the capabilities of general-purpose CPUs have increased exponentially, DAWs designed to use host-based processing have been able to challenge dedicated hardware systems. With the price/performance ratio of commodity processors now easily able to beat ProTools DSP cards that start at just under US$8000 each, a business case for open platforms in the recording studio is easy to make. On Linux, the principal project offering a direct replacement for ProTools is Ardour, free software which has been under development for several years and is now approaching beta release in the 0.9 series.

Paul Davis is the lead programmer on Ardour, which he stated hacking on after a period as one of the original developers of Amazon.com. Born in Britain but resident in the USA for the last fifteen years, Davis formed the company Linux Audio Systems to develop and provide consulting services for Ardour and other audio software. "In the autumn of 1999, the RME Hammerfall interface appeared. The idea of getting 26 channels of digital audio in and out of my computer was strangely compelling, so I bought the card and wrote a Linux device driver for it. By December of 1999, the driver was working, and I felt that I should write an application to do something with this US$600 piece of hardware. The early version of Ardour was basically an attempt to replace three Alesis M20 ADAT recorders (tape-based digital machines) that were in use at a studio of a friend of mine. The recording and playback part was surprisingly easy to get working - I think I was recording 24 tracks simultaneously within a few months of starting."

"However, within a year or so, it became abundantly clear that just being able to record and play back was close to useless when there were no Linux audio editors that could do anything useful with the data. So the focus of the Ardour project shifted to include an editor, and as we have all discovered, that was a massively more ambitious undertaking than writing a recorder. The first attempt involved using Bill Schottstaedt's 'snd' editor, which was capable of being built into Ardour."

"For various technical reasons, this didn't work very well, and a new plan emerged. Taybin Rutkin, who had already been involved with the project for a while, started this effort with some initial attempts at a new canvas-based approach to the Ardour editor's user interface. This has evolved into the incredibly complex and very powerful non-linear, non-destructive editor that Ardour represents today." Rutkin is based in New York, and has been working on Ardour since early 2001. "I mentioned to a friend of mine that I was interested in writing a simple four-track recording application, and he introduced me to Ardour and Paul. I dove in and wrote the save file system, and I've stayed with the project since."

Jesse Chappell, from Washington DC, has also been a substantial contributor to the Ardour project. "I am a programmer and a musician, and in late 2001 I decided it was time to fuse

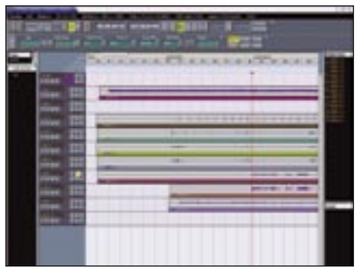

Ardour's editor window, playing back a recording session

my two interests together. I wrote SooperLooper, which emulated the Echoplex Digital Pro realtime looping sampler. In order to actually use it in live performance, I put together a rack-mount 1U computer running Linux and shoved it into my bass rig. The success of that effort, at least on a personal level, inspired me to look for other ways to do fun Linux audio development work."

"I was interested in a Linux-based DAW primarily to record my band, and I identified Ardour as the most promising software for doing serious multitrack recording and editing. Starting around June 2002 I dug into the Ardour source code and starting helping out in small ways. Over time, I grew more active on the mailing list and helped inspire and implement some important design elements. Among the things I've worked on over the last couple of years are support for multichannel tracks, live waveform updating, seamless looping, and most recently support for generic MIDI control surfaces, including motorised fader feedback."

Patches now seem capable of getting many people's systems to be more or less hard real-time capable

PATCHING THE KERNEL

In Ardour's early days, Linux was becoming established in the server space, but the particular demands of an audio workstation on the Linux kernel had not yet been addressed, as Davis explains. "I wanted a platform whose stability would itself be an asset, as well as capabilities for networking and data handling, which Windows and MacOS simply did not have in 1999. It turned out that although Linux was indeed a very stable platform, it was incapable of supporting the low-latency performance that was becoming expected of DAWs. Fortunately, various kernel developers including Ingo Molnar, Andrew Morton and Andrea Arcangeli came to rapidly understand what the problems were, and also to appreciate the worthiness of correcting them. Several months of work and testing led to Andrew Morton's low-latency patches for the 2.4 kernel series, and although they were never accepted into the mainstream kernel, several Linux distributions made kernels available with these patches built in."

"That situation has recently repeated itself with the 2.6 series kernels, with Ingo again doing a fine job. His patches now seem capable of getting many people's systems to be more or less hard real-time capable, implying fixed response times to hardware interrupts, and typically reducing delays in scheduling audio software caused by the kernel to the tens of microseconds range. It's an amazing accomplishment, but it is not part of the mainstream kernel and probably never will be. We will once again have to wait for Linux distributors to make this work available to users."

THE GUI DILEMMA

Creating a DAW for Linux from scratch, a decision has to be made to provide either a GUI which will be familiar to studios using other platforms, or an interface that existing Linux users without DAW experience will find comfortable. Davis chose the former: "Ardour is primarily designed for people with knowledge of the common core functionality found in all proprietary DAWs. There has been a massive convergence of features and UI design over the last few years that simultaneously makes it easier and harder for Ardour - on the one hand, there's no moving target anymore; on the other, it's harder to differentiate any given DAW from all the rest."

"However, we have been very aware of the issue of new users who are completely unfamiliar with the terminology, work flow patterns and typical UI designs of DAWs. DAWs are not like word processors, spreadsheets or email clients. They perform a totally different function working with a completely different kind of data. It is a challenge to come up with UI designs that cater well to both audiences. It doesn't take very long when working with a DAW to start wanting fast and convenient ways of doing things that until very recently you had never even thought of doing. This means we need designs that make the common 'new user' operations easily discoverable, but at the same time retain the kind of work flow model that ProTools has become appreciated for - very few people actually seem to like ProTools, but lots of people comment on how fast it is to use."

"Although we are aware of usability issues at this point in time, our focus for the 1.0 release is primarily functionality and reliability. We plan to focus much more on usability once we port the GUI to the GTK2 toolkit, immediately after the 1.0 release. GTK2 will offer us a number of improvements that will in turn let us improve Ardour in some quite dramatic ways." Rutkin adds: "I, and everyone else, I'm sure, is excited about the move to GTK2. It'll give us the widgets we need to make the interface we want Ardour to have. I come from a Macintosh background, and I think that it is possible to find the balance between the two groups of users. The main problem is that it requires a very detail-oriented approach to the UI which is a lot of work, and is also rather tedious. I think the switch to GTK2 will re-energise us, and hopefully attract some other programmers as well."

Very few people actually seem to like ProTools, but lots of people comment on how fast it is to use

Rutkin is also working on a complementary application called Ardour Session Exchange, which allows distributed studios to work on large multitrack projects, merging differences between sessions on a peer to peer basis. "Session Exchange accomplishes it major goal of letting people collaborate on sessions and only transferring the changes. Unfortunately, for large sessions, it's still a little impractical to use over the Internet - it works fine for LANs. I'm investigating ideas such as compressing the audio on the fly before transmitting it, to see if that will reduce the transmission time."

As Ardour approaches a 1.0 release, prospective users are invited to test the application and provide feedback to the project. The Ardour mailing lists, IRC channels, and the bug and feature tracking system are open to all. Ardour source code tarballs are provided, but the simplest way to get a full DAW environment up and running is to use one of the existing audio package collections, such as Planet CCRMA for Red Hat's Fedora Core distribution, the Debian-based AGNULA/DeMuDi, or Thac's RPMs for Mandrake.

Key Links

The Ardour homepage:
ardour.org

Binary packages of Ardour can be downloaded from:
Planet CCRMA ccrma.stanford.edu/planetccrma/software/

AGNULA/DeMuDi
agnula.org

Thac's RPMs
rpm.nyvalls.se/sound10.0.html