Not only, but ALSA

Between the program and the machine, come the drivers. Daniel James talks to three core members of the Advanced Linux Sound Architecture project

Audio applications for Linux would be of little use without a solid driver infrastructure, and yet it's only been in the last few years that free software drivers have become permanently established on the platform. For some people, the pivotal moment came in 2002 when Linus Torvalds accepted the work of the Advanced Linux Sound Architecture project, or ALSA, into the main Linux kernel tree. Before then, the kernel included Open Sound System drivers, known as OSS/Free - a parallel project to the original, proprietary code developed to make sound hardware work under UNIX and Linux.

This major replacement of kernel code wasn't considered lightly, and ALSA had to mature for several years before it was considered acceptable by core Linux hackers. The origins of ALSA go back over a decade, and a large group of developers have contributed to make the code robust enough for merging into the mainstream kernel. A grep for MODULE_AUTHOR in the ALSA source returns 57 names, but project members point out that many other people contributed patches, or provided test hardware.

Thanks to the work of the ALSA project, Linux users can now access a broad range of audio interfaces with both analogue and digital I/O, including budget chipsets built into motherboards, PCI cards, PCMCIA adaptors for laptops, and external USB devices. FireWire device support remains at the experimental stage for the time being; but in most categories, there is a solution available. The usual process of natural selection in free software driver development means that proprietary or poorly performing hardware is not so well supported; nevertheless, drivers are likely to be maintained long after support for any given version of Windows has been dropped by the hardware manufacturer.

ALSA had to mature for several years before it was considered acceptable

AN ANCIENT HISTORY

The development of what would become ALSA began when Jaroslav Kysela, a computer science student at the University of South Bohemia in the Czech Republic, discovered Linux. In 1994, he started working on The Linux Ultra Sound Project, based around a low-level driver for the Gravis UltraSound series of audio cards. "From this time, I systematically evaluated the Linux sound APIs and drivers, and tried to push them forward. In 1999 I founded the ALSA project. Playing with soundcards was fun in the first years and now, it is both fun and a full time job for me." Kysela was recruited by SUSE to work on ALSA, and since the ALSA merge in Linux 2.5.5 he has been the sound system maintainer on the Linux kernel team.

Also working for SUSE on the ALSA code is Takashi Iwai. "I studied Material Engineering at the University of Tokyo in Japan, and used Linux during research work on computer simulation. I started coding the audio-related stuff when I casually received a SoundBlaster AWE board that didn't work with Linux. Having non-working hardware is a great temptation to turn a good man to the dark side of programming!" Iwai got involved in the ALSA project around 1999, contributing the code for MIDI WaveTable support and hacking on the sequencer. "Since then, I've got involved more and more deeply with ALSA. It's become my job now."

Other members of the ALSA team contribute to the project on a voluntary basis. Clemens Ladisch is currently studying computer science at the Martin Luther University in Halle, Germany, and maintains the drivers for USB sound devices. "I had been using Linux for many years but became interested in ALSA when my new SC-8820 USB MIDI device didn't work with Linux.†Now my contributions to ALSA are mostly MIDI related, but the amount depends on how much free time I can spare." Other significant contributors to ALSA code include Abramo Bagnara, who designed the initial PCM API for ALSA 0.9, and Frank van de Pol who designed the sequencer API. Ardour project founder Paul Davis developed several drivers, including two for the pro-audio RME interfaces, and James Courtier-Dutton is working on drivers for the newer Sound Blaster cards.

STARTING FROM SCRATCH

When asked why they think the Open Sound System model failed to retain the interest of Linux developers, the ALSA team agrees that it was both for technical reasons, and the damaging split between free and proprietary OSS versions. (Kysela, Iwai and Ladisch answer as one, on behalf of the ALSA project.) "We think that both reasons are correct. Technically seen, OSS is (or was) a good system: it is small and simple. It was built, in the initial phase, using information about only the cheap consumer cards. However, it lacks extensibility. This is a critical problem for modern systems, and OSS could not be improved in this regard due to the proprietary version. Also, the split was not very helpful, because there are not many active kernel sound developers, and they would much rather work on the fully open source project by nature."

ALSA developers had mixed feelings about the existence of the proprietary OSS drivers. "It is a good question whether the commercial version was good or evil. It's true that Linux did have the 'least' support for many devices thanks to commercial OSS. On the other hand, the existence of binary-only drivers surely demotivated the hardware manufacturers to open their specs and co-operate with open-source projects like ALSA.†Hopefully, the situation is getting better nowadays."

The embryonic ALSA team quickly realised that it was necessary to start from scratch to create a new sound architecture, rather than attempt to improve on the existing OSS/Free base. "It was a very early moment. We soon faced the limitation of OSS, and needed a new infrastructure to achieve better support of sound cards. That was the starting point for the development of ALSA. We tried to define better APIs, suitable for both consumer and professional audio, and MIDI. We had false starts in the beginning, and the current (third or fourth) version of the APIs is sufficient for most developers."

"OSS depends highly on the abstraction level defined within its simple APIs. We created many extensions, inventing methods which are now reused by the OSS developers to enhance the original OSS APIs. The OSS APIs are strictly between user space and kernel. Our 'public' APIs are defined with a new layer, alsa-lib. This library sits in user space and provides backwards compatibility. It hides all changes, in case drivers define new interfaces for user space. It also means we can do user space drivers - for example, for Bluetooth and Firewire devices."

OPEN CHIPSETS

Some audio equipment vendors have helped the ALSA project by providing chipset specifications; a minority refuse to provide any information on their products at all. The creation of drivers for some of these 'closed' products has relied on trial and error to divulge the operation of interfaces. "Receiving the chipset specification was a hard job in the early stage of ALSA development. We had to do reverse engineering more or less. As we recall, Trident was the first company to be co-operative among the hardware manufacturers. They contributed driver code and gave us an open datasheet. After that, some companies producing high-end audio cards like RME, M-Audio and TerraTec showed interest and provided us the datasheet and test boards for the support of many, but not all, of their products."

Non-working hardware is a great temptation to turn a good man to the dark side of programming!

ALSA has mature support for multichannel interfaces, including the M-Audio Delta 1010

"Nowadays we have a good relationship with many hardware companies in the wide range from the on-motherboard sound chip manufacturers like Intel, VIA and ATI to the high-end audio manufacturers. The reverse engineering is still sometimes necessary, but we're trying to avoid it as much as possible. It would be much better to have contact with the manufacturer than be shooting in the dark."

Barriers for free software driver development remain. FireWire audio interfaces are a particular difficulty, since the technology is controlled by the 1394 Trade Association, and the manufacturers of these devices rarely follow published standards. "The most annoying barrier is the presence of closed or restricted specifications. In our opinion, the policies should only cover hardware vendors, and these associations should give out all necessary information to write drivers for free."

THE FUTURE OF ALSA

The purchase of SUSE by Novell has made little difference to the ALSA project, according to the team members who now work for the US company. So far there has not been much integration with Novell's desktop Linux initiatives. "Our work has not changed, at least for the time being. Note that only Jaroslav Kysela and Takashi Iwai are SUSE employees. Other ALSA developers are working without SUSE sponsorship. Actually, there is no direct communication between the SUSE ALSA team and other projects in Novell. We directly communicate mostly with the YaST team, to provide good sound configuration for SUSE distributions. Takashi also maintains several audio related packages for SUSE.

The ALSA project includes scope for audio-over-network devices, a current example of which is Yamaha's proprietary 'mLAN' system. "We would like to cover all hardware, and audio-over-network seems to be a very interesting area. Unfortunately, mLAN is a closed specification. The question is if the current network layer in the Linux kernel is capable of doing lowlatency operations for realtime audio. Fortunately, we might be able to improve the situation, because thanks to open source, we can work with the Linux network code."

Key links

ALSA project homepage
www.alsa-project.org

SUSE Linux
www.suse.com

RME
www.rme-audio.com

M-Audio
www.m-audio.com

Terratec
en.terratec.net

1394 Trade Association
www.1394ta.org

mLAN Alliance
www.yamaha.co.jp/tech/1394mLAN/english/