

Where Are We Going And Why Aren't We There Yet ?

A Presentation Proposal for LAC 2005, Karlsruhe

Dave Phillips

linux-sound.org

400 Glessner Avenue

Findlay OH USA 45840

dlphillips@woh.rr.com

Abstract

A survey of Linux audio development since LAC 2004. Commentary on trends and unusual development tracks, seen from an experienced user's perspective. Magic predictions and forecasts based on the author's experience as the maintainer of the Linux Sound & Music Applications website, as a professional journalist specializing in Linux audio, and as a Linux-based practicing musician.

Keywords

history, survey, forecast, user experience, magic

1 Introduction

Linux sound and music software developers have created a unique world populated by some remarkable programs, tools, and utilities. ALSA has been integrated with the kernel sources, the Rosegarden audio/MIDI sequencer has reached its 1.0 milestone, and Ardour and JACK will soon attain their own 1.0 releases. Sophisticated audio and GUI toolkits provide the means to create more attractive and better-performing sound and music programs, and users are succeeding in actually using them.

2 A Brief Status Report

The Linux Sound & Music Applications site is the online world's most popular website devoted to Linux audio software. Maintaining the site is an interesting task, one in which we watch the philosophy of "Let 10,000 flowers blossom!" become a reality. It can be difficult to distinguish between a trend and the merely trendy, but after a decade of development there are definite strong currents of activity.

The past year has been a year of maturities for Linux audio software at both the system and application development levels. To the interested

user, the adoption of the ALSA sound system into the Linux kernel means that Linux can start to provide sound services to whatever degree required. Desktop audio/video aficionados can enjoy better support for the capabilities of their soundcards. Users seeking support for more professional needs can find drivers for some pro-audio hardware.

In addition to this advanced basic support there are patches for the Linux kernel that can dramatically reduce performance latency, bringing Linux into serious consideration as a viable professional-grade platform for digital audio production needs, at least at the hardware level. It is important to note that these patches are not merely technically interesting, that they are being used on production-grade systems now. Furthermore, there is a continuing effort to reduce or eliminate the need for patching at all, giving Linux superior audio capabilities out-of-the-box.

ALSA has passed its 1.0 release, as has Erik de Castro Lopo's necessary libsndfile. JACK is currently at 0.99, and the low-latency kernel patches have been well-tested in real-world application. The combined significance of these development tracks indicates that Linux is well on its way to becoming a viable contender in the sound and MIDI software arenas.

Support for the LADSPA plugin API has been an expected aspect of Linux audio applications for a few years. LADSPA limits are clear and self-imposed, but users want services more like those provided by VST/VSTi plugins on their host platforms. Support for running VST/VSTi plugins under Linux has also inspired users to ask for a more flexible audio/MIDI plugin API. At this time the most likely candidate is the DSSI (Disposable SoftSynth Interface) from the Rosegarden developers. The DSSI has much to recommend it, including support for LADSPA and an interface for plugin instruments (a la VSTi plugins).

In this author's opinion the union of ALSA,

JACK, and LADSPA should be regarded as the base system for serious audio under Linux. However, the world of Linux audio is not defined only by the AJL alliance. Other interesting and useful projects are going on with broader intentions that include Linux as a target platform.

The PortAudio/MIDI libraries have been adopted as the cross-platform solution to Csound5's audio/MIDI needs. Support for PortAudio has appeared in Hydrogen CVS sources, and it is already a nominal driver choice for JACK.

GRAME's MidiShare is not a newcomer to the Linux sound software world, but it is beginning to see some wider implementation. Among its virtues, MidiShare provides a flexible MIDI multiplexing system similar to the ALSA sequencer (it can even be an ALSA sequencer client). The system has been most recently adopted by Rick Taube's Common Music and the fluidsynth project.

Sound support in Java has been useful for a few years. All too often more attention has been paid to Java's licensing issues than to its audio capabilities. Many excellent Java-based applications run quite nicely on Linux, including the jMusic software, JSynthEdit, and Phil Burk's excellent jSyn plugin synthesizer.

At the level of the normal user the applications development track of Linux audio is simply amazing. Most of the major categories for music software have been filled or are being filled soon by mature applications. Ardour is designed for high-end digital audio production, Rosegarden covers the popular all-in-one Cubase-style mode, Audacity, Snd, and ReZound provide excellent editing software, Hydrogen takes care of the drum machine/rhythm programmer category, and MusE and Rosegarden cover the standard MIDI sequencer environment. Denemo and Rosegarden can be used as front-ends for LilyPond, providing a workpath for very high-quality music notation.

Notably missing from that list are samplers and universal editor/librarian software for hardware synthesizers. However, the LinuxSampler project is rapidly approaching general usability, and while the JSynthEdit project's pace is slow it does remain in development. Some similar projects have appeared in the past year, but none have advanced as far as JSynthEdit.

A host of smaller, more focused applications continues to thrive. Programs such as Jesse Chappell's FreqTweak and SooperLooper, Rui Capela's QJackCtl, and holborn's midirgui indicate that useful Linux audio software is becoming more easily written and that there is still a need for small

focused applications. Of course the on-going development of graphics toolkits such as GTK, QT, and FLTK has had a profound effect on the usability of Linux applications.

Csound represents yet another significant class of sound software for Linux, that of the traditional language-based sound synthesis environment. The currently cutting-edge Csound is Csound5, basically a complete reorganization and rewrite (where necessary) of the Csound code base. Improvements include extensive modularization, internal support for Python scripting, and an enhanced cross-platform build system. The Linux version of Csound5 is already remarkable, with excellent realtime audio and MIDI performance capability.

One downside to the increasing capabilities of the Linux sound system is the increasing complexity of Linux itself. For most users it is decidedly uncomfortable and uninteresting to perform the necessary system modifications themselves, but happily the AGNULA/Demudi and Planet CCRMA systems have brought near-painless Linux audio system installation to the masses. However, given the resistance of said masses, we have seen the rise of the "live" Linux multimedia-optimized CD. These systems allow provide a safe and very effective means of introducing not only Linux audio capabilities but Linux in general, without alteration of the host system. The Fervent Software company has taken advantage of this trend and released their Studio To Go! commercially. I believe that these live CDs have enormous potential for Linux evangelization generally, and they may be a particular blessing for the expansion of interest in Linux audio capabilities.

3 Visibility Is Clear

Linux audio software is becoming a serious alternative for serious users. Composers of all sorts, pro-audio recordists, sound synthesis mavens, audio/video DJs and performance artists, all these and many other sound & music people are using this software on a productive daily basis. More "music made with Linux" has appeared in the past year than in the entire previous decade, and coverage of Linux audio regularly appears in major Linux journals. Articles on Linux audio software have appeared in serious audio journals such as Sound On Sound and the Computer Music Journal.

Some of the significant events acknowledging Linux audio software included Ron Parker's demonstrations of the viability of Ardour and JAMin in a commercial recording environment, Criscabello's announcement that he'd recorded Gilberto Gil with software libre, and the awards received for Hydrogen and JACK. Small steps perhaps, but they mark the steady progress of the development in this domain.

4 Some Problems

There is no perfection here. Lack of documentation continues to be a primary issue for many new users. Hardware manufacturers still refuse to massively embrace Linux audio development. Many features common in Win/Mac music software are still missing in their Linux counterparts. Many application types are still poorly represented or not represented at all.

Community efforts towards addressing documentation issues include various wikis (Ardour, Pd) and a few focus groups (Hydrogen, Csound5), and while many applications do have excellent docs, the lack of system-comprehensive documentation still plagues the new user, particularly when troubleshooting or attempting to optimize some aspect of an increasingly complex system. The problem is familiar and remains with us: writing good documentation is difficult and there are too few good writers with the time and energy to spare for the work required. Nevertheless, the impact of the documentation wikis is yet to be felt, they may yet prove to be a salvation for the befuddled user.

Hardware support still remains problematic. ALSA support expanded to the Echo cards, and the AudioScience company announced native Linux driver support for their high-end audio boards, but no train of manufacturers hopped on the Linux sound support bandwagon. I'm not sure what needs to happen to convince soundcard and audio hardware manufacturers that they need to support Linux, and I believe that this issue needs some more focused discussion in the community.

Limited hardware support is often worse than none at all. ALSA developers are working to provide sound services as complete as their Win/Mac counterparts, but there are still problems with regard to surround sound systems (3D, 5.1) and access to on-board DSP chipsets.

A glance through any popular music publication clearly shows that Linux audio software, wonderful as it is, definitely lacks the variety of the Win/Mac worlds. Users new to the Linux audio

world often lament the absence of programs such as Acid, Fruity Loops, or Ableton Live, and I have already mentioned the dearth of editor/librarian software for hardware MIDI synthesizers. The situation is surely improving, but there are still several application types awaiting project involvement.

5 Summary Conclusions

The good news far outweighs the bad, and the bad news itself can be dealt with in productive ways. The development and user communities continue to thrive, long-term projects proceed, more people are coming into our world, and more music is being made. Coordinated efforts need to be made to bring about greater program documentation and manufacturer participation, but whatever difficulties we encounter, the history of Linux software development advises us to never say never.

6 Acknowledgements

The author thanks the entire community of Linux audio developers for their enormous contribution to music and sound artists the world over. The author also thanks the community of Linux audio software users for their experiences, advice, and suggestions regarding a thousand sound- and music-related topics. Finally, the author extends great gratitude to the faculty and staff at ZKM for their continued support for this conference.

