The One Laptop Per Child (OLPC) Audio Subsystem

Jaya Kumar
Independent Developer,
Gurgaon, India — KL, Malaysia
jayakumar.alsa@gmail.com

Abstract

The OLPC is a Linux based laptop-like device intended as an educational tool targeted at developing countries. The audio subsystem of the OLPC faces typical challenges such as minimizing power consumption, performance/quality and component cost pressures and tradeoffs, as well as less common challenges such as the need to repurpose audio input as an oscilloscope or analog input system. This paper explains issues encountered during support and development of ALSA and low level audio support on the OLPC. It will also touch on possible future plans for the low level audio software side of the OLPC.

Keywords

OLPC audio, Analog Input, Speaker-Microphone feedback, Power Management

1 Introduction

From a hardware perspective, the OLPC is a fairly full featured embedded device. It has strong multimedia capabilities. This includes a decent audio subsystem, decent video subsystem and even a video input subsystem. Table 1 provides an overview of the hardware features. Figure 1 is a picture of the beta test (B-Test1) unit of the OLPC courtesy of the OLPC team[1].

The OLPC kernel is currently based on a 2.6.19 kernel, thus incorporating ALSA 1.0.13. The operating system for the OLPC is based on a stripped down and heavily customized version of Fedora. The userspace audio engine is the csound server. Current audio applications incorporated in the OLPC operating system are TamTam and Squeak eToys. Lots of other multimedia applications have been run on this system including Quake, mplayer, Collabra's Telepathy among others.

Figure 1: OLPC B-Test1 Unit

2 Hardware Architecture

The OLPC board currently uses an embedded x86 architecture. The audio controller for this architecture is a core within the cs5536 ASIC as shown in Figure 2. The cs5536 is the integrated southbridge for this architecture. This controller interfaces with an AD1888 AC97 codec from Analog Devices.

The communication between the Geode GX2 cpu and the cs5536 is PCI [1]. However, neither device implements a fully traditional PCI bus controller [2]. For example, the AC97 Controller (ACC) in the cs5536 is not actually a PCI core but rather communicates via GeodeLink with the GeodeLink PCI SouthBridge which then handles communication with the Geode to stream to/from host memory. The important PCI command types such as Memory Read/Write, I/O Read/Write and others are fully implemented by the GeodeLink PCI SouthBridge. But PCI Configuration Read/Write support is not implemented by either the Geode GX2 or the CS5536. This affects system software, including audio software, in a

---

[1] The image is Copyright (C) OLPC, used with permission under Creative Commons Attribution2.5 License

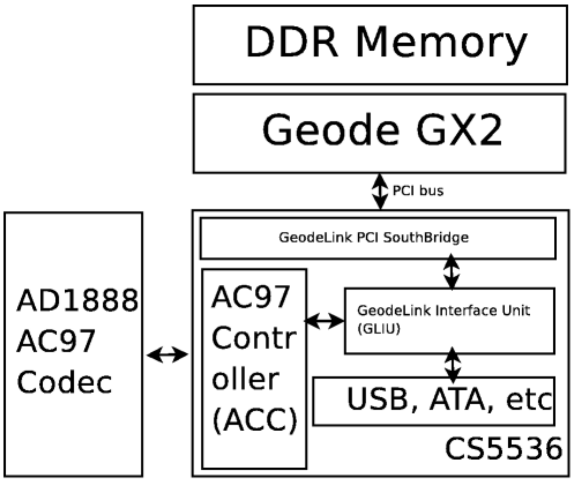

Figure 2: GX-CS-AC97 Architecture

unique way when contrasted with other x86 systems.

In typical Geode systems including the current version of the OLPC, the BIOS (LinuxBIOS on the OLPC) utilizes a legacy module called VSA (Virtual System Architecture) that identifies GeodeLink devices and virtualizes the PCI configuration space registers. This means that when a driver attempts to access a PCI configuration register, a SMI (system management interrupt) is generated. The SMI is then handled by software, in this case, the VSA code within the BIOS that performs MSR (machine specific register) based reads/writes to perform the appropriate task.

2.1 Mechanical Architecture

The OLPC exposes on-board audio via left and right speakers grill-mounted on both sides of the front fascia as pictured in Figure 3. On-board audio input is exposed via a ported microphone mounted on the front fascia. External audio output and input is enabled via two standard red and green 3.5mm jacks on the left side of the display head. The audio input jack is dual purposed for analog input as well. A general picture of the OLPC B-Test1 unit is provided in Figure 1.

3 ALSA on OLPC

The lowest level of the ALSA subsystem on the OLPC is the two hardware drivers. These are the cs5535audio driver and the AD1888 AC97 driver. Both have been part of the ALSA tree prior to initiation of the OLPC project. The cs5535audio driver prior to OLPC was a fairly standard ALSA PCI driver. The same is true of the AD1888 AC97 driver.

Hardware	Features
CPU	AMD Geode GX2 (2 Watts @ 366 MHz)
Memory	128 MB DDR400 SDRAM
Storage	512 MB NAND Flash
Audio	CS5536 ACC AD1888 AC97
Camera	VGA CMOS
Video	200dpi 7" LCD 1200x900 (BW reflective) 640x480 (Color transmissive) 250 nits
Wifi	Marvell 88W8388 802.11b/g/s
USB	3 x USB2.0 ports
SD	1 x SD slot
Battery	2 hour NiMH
Input	Keyboard, touchpad

Table 1: Hardware Features

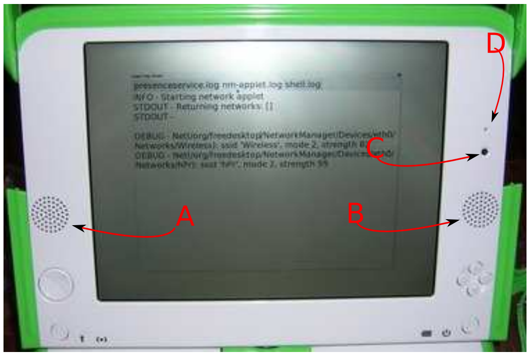

Figure 3: OLPC Front Fascia
A - Left Speaker
B - Right Speaker
C - CMOS Video Camera
D - Microphone

3.1 cs5535audio OLPC support

Prior to OLPC, the cs5535audio driver was typically utilized in always-on devices in industrial control systems. In those environments, power management was not an attribute that was sought after. The introduction of the OLPC project gave rise to the motivation to add power management functionality to cs5535audio. This support was fairly straightforward in combination with the AC97 suspend/resume support that is part of ALSA. From the ACC perspective, the task to be done by the driver is very typical. It utilizes ALSA to suspend PCMs, to take the AC97 codec through the suspend sequence and finally to turn off the AC Link and to take the ACC to D3 state. Further additions to cs5535audio to support OLPC were related to the analog input which is covered in Section 4.5.

3.2 AD1888 OLPC support

There is broad support for almost all of the AD1xxx series AC97 codecs from Analog Devices in ALSA. Testing with the OLPC found a minor issue related to duplicate controls. Further additions to support OLPC were bug fixes associated with power management, new mixer controls and analog input support.

3.3 AC97 Power Management

One of the important changes to ALSA that is very relevant to the OLPC is the addition of aggressive power-saving support of AC97 codecs. This addition was done by Takashi Iwai. The way that it works is that when all PCMs are closed and a reasonable delay has passed suggesting that audio activity has ceased, ALSA proceeds with switching capable AC97 codecs into suspend mode. Upon opening of a PCM, ALSA resumes the AC97 codec. This allows for reduced power consumption without noticeable loss of functionality.

4 Problems

Several interesting problems were identified as part of the bringup of audio on the OLPC. Some of these issues remain open.

4.1 AC97 read failure

One of the first problems identified with ALSA on the OLPC was two failed reads of the AC97 codec. Specifically, ALSA's AC97 support code attempts to read the AC97_VENDOR_ID2 on the AD1888 codec. This AC97 read fails even with a large timeout. The failure is derived from the ACC not asserting a bit in its status register called STS_NEW. However, the value that is read back from the AC97 register despite the lack of status bit assertion is 0x7E805368. Thus, the 0x5368 which is returned is correct, as per the AD1888 datasheet thereby allowing everything else to proceed as normal despite the apparent read failure. This issue was observed only with the OLPC board. Other systems such as cs5535/realtek combinations and cs5536/wolfson combinations have not exhibited this symptom. The cause of this read failure has not yet been identified and remains an open issue.

4.2 AD1888 duplicate controls

Marcelo Tosatti was the first to identify this issue on the OLPC. This problem had to do with the fact that ALSA creates the majority of controls in *snd_ac97_mixer_build*. For example, one of those controls was the "Surround Playback" control. On the AD1888, this mixer creation conflicted with the same control being provided by the *snd_ac97_controls_ad18xx_surround* table thereby resulting in a fail out. The solution was trivial which was to precondition the initial mixer creation with the AC97_AD_MULTI flag that represents that set of Analog Device codecs.

4.3 AD1888 magic resume register

As with many hardware designs, the AD1888 contains its share of magic registers that are not formally documented[3]. An unusual suspend/resume bug triggered the discovery of one such register. The observed symptom is best described by the test sequence:

```
insmod snd-cs5535audio.ko
amixer set PCM 100 on
amixer set Master 100 on
aplay /tree/test.wav
# playback works fine
# now suspend to D3
echo -n 3 > /sys/devices/pci0000
\:00/0000\:00\:0f.3/power/state
# try to play
aplay /tree/test.wav
# playback is blocked as expected
# now resume to D0
echo -n 0 > /sys/devices/pci0000
\:00/0000\:00\:0f.3/power/state
# now try to playback
aplay /tree/test.wav
# playback proceeds at normal rate
but no sound is heard on the output
```

After some debug, a register was found that made things work after resume which is AC97 register 0x60, or AC97_CODEC_CLASS_REV. The AD1888 datasheet does not refer to this register. The AC'97 2.3 spec says:

> Extended Codec Registers Page Structure Definition Page 00 of the

> Extended Codec Registers is reserved for vendor specific use. Driver writers should not access these registers unless the Vendor ID register has been checked first to ensure that the vendor of the AC '97 component has been identified and the usage of the vendor defined registers is understood.

Tweaking that register on resume allowed things to work correctly. The reason why this is the case is still unknown.

4.4 Uniqueness of PCI IDs

As was mentioned in Section 2, the BIOS currently provides the virtualized PCI configuration registers. This implies that the BIOS software loaded on to the SPI flash (PROM equivalent) is what determines the PCI Vendor/Subvendor and Device/Subdevice ID for the system. Almost all Geode systems use either the BIOS from Insyde Systems (a BIOS vendor) or LinuxBIOS. In the latter case, LinuxBIOS on the OLPC currently has a legacy build of the VSA included within it. This perpetuates the situation where almost all cs5536 based boards have exactly the same PCI Vendor/Subvendor and Device/Subdevice IDs. This makes it rather hard for the ALSA driver to identify which board it is being used on in order to apply board specific quirks.

In the case of the OLPC, it is sufficiently different that it warranted having a CONFIG_OLPC config entry. This enables the various associated drivers to apply OLPC specific quirks using that flag.

4.5 Analog Input support

An interesting feature on the OLPC is that it can repurpose the microphone input jack as an analog input. This capability can be used to interface the OLPC with analog input devices such as a photodiode. This was demonstrated by Barry Vercoe at WorldComp '06, where a spoon with a photodiode connected to the analog input was used as a conductor's baton to control audio.

The method by which the analog input capability is implemented is by having the trace from the input jack split to two separate paths. One path is a standard resistive-capacitive path to the microphone input pin on the AD1888. The other path is an unbiased direct path to the same pin. The selection of these paths is done via a standard analog switch controlled by an 8051 based embedded controller (EC). The EC is then interfaced to the host via GPIO pins. From an audio software perspective, communication with the EC is currently via x86 port IO. The B2 model of the OLPC will move this analog switch control over to a Geode GPIO, thereby avoiding the communication with the EC.

There are therefore several aspects to analog input support. The first set are features of the AC97 codec. These are the ability to disable the V_{Ref} bias that is normally internally applied to the microphone input and the ability to disable the high pass filter that is internally applied to limit input to typical audio frequency ranges. The second set is the need for the host to instruct the EC to toggle the analog switch. The current implementation for analog input support on OLPC is done by adding an OLPC specific quirk to cs5535audio. This quirk adds a new mixer control named "Analog Input". This function performs the task of toggling both the EC bits and the AC97 bits in order to achieve the desired functionality. This therefore exposes the analog input mode decision to userspace through the standard ALSA mixer API and can be controlled via amixer and other regular tools as well.

4.6 Speaker-Microphone Feedback

As can be seen from Figure 3, the right speaker and the microphone are physically very close. One of the effects of this industrial design decision is that there is significant amounts of feedback when both the speaker and microphone are active at the same time. This has a negative effect on telephony (voice/video over IP) type applications. This problem is not yet solved. Several proposals have been put forth including one suggesting that the right speaker be disabled when the microphone is active.

5 Future Enhancements/Issues

5.1 ASoC/DAPM

One of the most interesting future issues is continued improvement of power management. With this in mind, the ALSA System On Chip and Dynamic Audio Power Management support contributed by Liam Girdwood and others is very relevant to the OLPC. In particular, the ability to dynamically switch on and off DAI capabilities, tune AC97 clock rates, and performing switch-level power management on the AC97 codec while completely transparent to applications is very attractive for minimized power consumption on portable devices like the OLPC. Work is in progress to add ASoC/DAPM support for OLPC.

5.2 Speaker-Microphone Feedback

As was mentioned before in Section 4.6, this feedback issue is a challenging problem. One of the approaches that may be of interest from the ALSA software perspective is to use a dmix type plugin to downmix stereo streams and redirect to only the left channel when it is detected that the on-board microphone is enabled.

5.3 VSA elimination

The virtualization of PCI configuration registers by the VSA can be removed. This is possible if the kernel's generic PCI configuration dependent code was replaced by OLPC specific rd/wr-msr code. That code would handle the GeodeLink support. This would also aid boot performance. This work is in progress and is being done by Mitch Bradley of the OLPC team.

5.4 Preempt and Tickless

There are two kernel features being actively tested and developed on the OLPC. These are the tickless kernel and realtime preemption support. Realtime preemption is clearly of benefit to the audio subsystem [4]. Tickless should not be detrimental to audio. Tickless will serve to reduce power consumption on OLPC and also improve timer resolution.

6 Conclusions

The OLPC provides good audio quality and performance for a device of it's type. It provides a fully featured environment for application developers, as can be seen from the wide range of audio-enabled applications that are running on it such as TamTam, eToys, and others. It presents an interesting environment for ALSA and Linux Audio in general. It has aspects of a typical embedded system. That is, it has very tightly constrained and hard-wired resources in terms of cpu, audio controller, audio codec, memory and storage. But it also has aspects of a high end system. It has a wide set of multimedia capabilities thus requiring full broad ALSA and other audio functionality. This combination will be likely to continue to provide unusual and challenging problems to be solved by the Linux Audio community.

In the event that the OLPC project is successful, ALSA and Linux Audio in general will be contributing something directly positive to the daily lives of a large number of human beings through out the planet. Further, if the commercial aspects of the OLPC project are similarly successful, then this will encourage more hardware manufacturers to become more involved in the ALSA and Linux Audio community in general. This will help ALSA and Linux Audio in general to thrive.

7 Acknowledgements

My thanks go especially to Takashi Iwai and Jaroslav Kysela for their advice on many of the OLPC audio issues and reviewing all the code for the OLPC audio changes. Special thanks to Liam Girdwood of Wolfson Micro for coming to FOSS.IN in Bangalore to present ASoC/DAPM and improve ALSA development here. Thanks to the OLPC organization for kindly supplying the board for driver testing and development. Also my sincere thanks to all the members of the ALSA community for much help over many years.

References

[1] AMD. *AMD Geode CS5536 Companion Device Data Book*. AMD, USA, 3/14/2006.

[2] AMD. Amd geode gx and lx processor based systems virtualized pci configuration space. *AMD*, 1(32663C):3–20, 2006.

[3] Analog Devices. *AC '97 SoundMAX AD1888 Codec Data Book*. Analog Devices, USA, 2005.

[4] Lee Revell. Realtime audio vs. linux 2.6. *Proceedings of Linux Audio Conference 2006*, 2006(1):21–24, 2006.