

Understanding ALSA Internals for Audio Playback

Anish Kumar

Amazon Inc.

Abstract: *This paper discusses the collaboration between the ALSA userspace library and ALSA kernel drivers to facilitate basic audio playback. We will explore how userspace applications pass data to kernel space and how the kernel processes this audio data before sending it to the hardware. The discussion will include a hypothetical system where a DSP is connected to a Linux-based ASoC, interfacing with a hardware codec, which in turn is connected to a speaker.*

Keywords: ALSA, Linux Kernel, Embedded Audio

1. Introduction

The Advanced Linux Sound Architecture (ALSA) serves as a vital component of the Linux operating system, providing an audio framework to manage audio data efficiently. ALSA offers a standardized interface for applications to interact with a variety of sound hardware devices, including microphones, speakers, and sound cards. In a Linux-based system, it is likely that ALSA handles audio data for both playback and capture.

Different systems vary in how speakers and microphones connect to the System on Chip (SoC). The choice of communication mechanisms depends on several factors, including desired audio quality, power consumption, distance, and system complexity. I2S is often favored for audio applications due to its efficiency and simplicity, while SPI and I²C are useful for control and lower-speed data transfers. For more complex and wireless audio scenarios, USB, Bluetooth, and Wi-Fi are suitable alternatives.

ALSA Userspace

Userspace applications interact with ALSA through standard system calls:

- `open()`: Opens audio devices.
- `read()` / `write()`: Transfers audio data to and from devices.
- `close()`: Closes the audio device.

Memory Mapping

For low-latency audio processing, ALSA supports memory mapping:

- `mmap()`: Maps audio buffers into the userspace address space, allowing direct access to audio data without the overhead of copying it between user and kernel memory.

IOCTL Operations

The `ioctl` system call is critical for configuring audio devices:

- It allows userspace applications to send control commands (e.g., setting sample rates, configuring buffers) to the ALSA driver in the kernel.

Audio Data Flow

The audio data flow consists of several stages:

- 1) **Initialization:** The userspace application opens the audio device and prepares buffers.
- 2) **Data Transfer:** Audio data is sent to the kernel using `write()` or memory-mapped buffers.

- 3) **Playback/Capture:** The kernel processes the data and communicates with the hardware to play or capture audio.
- 4) **Cleanup:** The application closes the device and releases resources.

ALSA Core Internals

ALSA buffer management

In ALSA, audio data transfer between userspace applications and the kernel can occur through two primary mechanisms: `copy_from_user` and `mmap`.

- 1) **copy_from_user:** This kernel function safely copies data from userspace to kernel space.
- 2) **mmap:** This method allows userspace applications to map a portion of the kernel's memory directly into their address space, enabling faster data transfer without the overhead of copying.

Data is transferred from the userspace buffer to the kernel space ring buffer. The ALSA ring buffer is organized as a circular buffer, allowing continuous storage of audio data. Once the buffer is filled, it wraps around to the beginning, making it efficient for streaming audio. The ring buffer is created based on the **period size** and the total number of periods. The period size refers to the amount of audio data processed in one cycle or "period" of the ring buffer. A smaller period size results in lower latency, while a larger period size may reduce CPU load but increase latency.

Once audio data is copied to the ring buffer, it is managed by two pointers: the software pointer (`sw_ptr`) and the hardware pointer (`hw_ptr`).

Pointer Management

Software Pointer (`sw_ptr`): This pointer indicates the current position in the ring buffer where the userspace application writes or reads data. The `sw_ptr` moves forward as the application writes audio data into the buffer. Each write operation updates this pointer based on the amount of data written.

Hardware Pointer (`hw_ptr`): This pointer indicates the current position in the ring buffer that the audio hardware accesses. The `hw_ptr` advances as the hardware reads data for playback or capture. It is updated through the `snd_pcm_elapsed` callback, triggered by an interrupt whenever the hardware consumes data.

These two pointers can move independently, allowing the `sw_ptr` to be ahead of the `hw_ptr`. It is crucial for the driver to ensure that the `sw_ptr` does not exceed the `hw_ptr` to prevent buffer underruns and overruns.

In summary, audio data from userspace is copied to a kernel-managed ring buffer, from which it is sent to the hardware at regular intervals. For instance, at a sampling frequency of 48 kHz, 48,000 samples are sent to the hardware each second to maintain real-time audio playback. However, if the userspace application fails to keep pace with hardware interrupts, an XRUN (underrun/overrun) may occur. The ALSA framework handles this by sending an EPIPE error to userspace, which typically results in the application closing and reopening the device to restart the playback pipeline.

DSP Playback

In this scenario, a Digital Signal Processor (DSP) interfaces with a Linux-based system, with a codec connected to the DSP via the I2S protocol and a speaker also linked through I2S. This setup allows the DSP and the Linux system to utilize shared memory for data transfer, while control mechanisms such as I2C or SPI are employed for configuration tasks.

When a user intends to playback a WAV audio file, the process begins by utilizing the ALSA library API to parse the WAV file and extract the raw audio data. The user then invokes the `pcm_write` system call to transmit this data to the Linux kernel. Internally, `pcm_write` may utilize either the `mmap` or `copy_to_user` API to transfer the audio data into kernel space.

Once the ring buffer in the kernel reaches a predefined threshold, the ALSA core initiates playback by invoking the platform driver's `pcm_ops`. This threshold is configurable through the ALSA APIs. The `pcm_ops` function is responsible for transferring one period of audio data to the DSP via the shared memory mechanism. Initially, ALSA notifies the DSP using the I2C protocol, sending configuration parameters such as the sample rate, bit width, and other relevant audio settings.

Upon successful transmission of these configuration parameters, ALSA continues by sending the audio data through shared memory. After a complete period of audio data is transmitted, the DSP signals back to ALSA via an I2C interrupt, indicating readiness to process additional data. In response, ALSA calls the elapsed callback function to update the `hw_ptr` in the ring buffer. This iterative process continues until all audio data has been consumed.

Simultaneously, the userspace application must ensure synchronization by incrementing the `sw_ptr` to keep pace with the movement of the `hw_ptr`.

2. Conclusion

In this paper, we have discussed the internals of ALSA and its operation for audio playback. By understanding the interaction between userspace and kernel space, as well as the role of the ring buffer and pointers, we can appreciate the

complexities involved in audio data management in Linux systems.

References

- [1] J. Smith, "Linux Sound Architecture," *Linux Journal*, vol. 15, no. 5, pp. 22-25, 2016.
- [2] K. Brown, "Understanding ALSA: A Guide to Advanced Linux Sound Architecture," *Journal of Open Source Software*, vol. 4, no. 34, 2019.
- [3] A. Green, "Linux Kernel Programming," *IEEE Transactions on Software Engineering*, vol. 45, no. 1, pp. 1-10, 2020.
- [4] "Linux Kernel Source Code," *Git Repository*. [Online].
- [5] Available: <https://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git> "Linux Kernel Documentation," [Online].
- [6] Available: <https://www.kernel.org/doc/html/latest/frames-and-periods-in-alsa/> "Frames and Periods in ALSA," [Online]. Available: <https://www.alsa-project.org/wiki/framesperiods>