i.MX Linux Reference Manual

Contents

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

Chapter 3
Storage

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

Chapter 4
Connectivity

Chapter 5
Graphics

Chapter 6
Video

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

Chapter 7
Audio

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

Chapter 8
Security

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

Chapter 9
Unit Tests

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

Chapter 10
Revision History

Chapter 1
Introduction

1.1 Overview

The i.MX family Linux Board Support Package (BSP) supports the Linux Operating System (OS) on the i.MX application processors.

The purpose of this software package is to support Linux OS on the i.MX family of Integrated Circuits (ICs) and their associated platforms. It provides the necessary software to interface the standard open-source Linux kernel to the i.MX hardware. The goal is to enable i.MX customers to rapidly build products based on i.MX devices that use the Linux OS.

The BSP is not a platform or product reference implementation. It does not contain all of the product-specific drivers, hardware-independent software stacks, Graphical User Interface (GUI) components, Java Virtual Machine (JVM), and applications required for a product. Some of these are made available in their original open-source form as part of the base kernel.

The BSP is not intended to be used for silicon verification. While it can play a role in this, the BSP functionality and the tests run on the BSP do not have sufficient coverage to replace traditional silicon verification test suites.

1.1.1 Software Base

The i.MX BSP is based on version 4.9.123 of the Linux kernel from the official Linux kernel website (www.kernel.org). It is enhanced with the features provided by NXP.

On Linux to change the configuration using the menu configuration with a Yocto Project environment, use bitbake like this:

```
bitbake linux-imx -c menuconfig
```

1.1.2 Features

The table below describes the features supported by the BSP for specific platforms.

Table 1-1. BSP Supported Features

Feature	Description	Chapter Source	Applicable Platform
Machine-Specific Layer			
MSL	Machine-Specific Layer (MSL) supports interrupts,Timer, Memory Map, GPIO/IOMUX, SPBA, SDMA. • Interrupts GIC: The Linux kernel contains common Arm GIC interrupts handling code. • Timer (GPT): The General Purpose Timer (GPT) is set up to generate an interrupt as programmed to provide OS ticks. Linux OS facilitates timer use through various functions for timing delays, measurement, events, alarms, high-resolution timer features, and so on. Linux OS defines the MSL timer API required for the OS-tick timer and does not expose it beyond the kernel tick implementation. • GPIO/EDIO/IOMUX: The GPIO and EDIO components in the MSL provide an abstraction layer between the various drivers and the configuration and utilization of the system, including GPIO, IOMUX, and external board I/O. The IO software module is board-specific, and resides in the MSL layer as a self-contained set of files. I/O configuration changes are centralized in the GPIO module so that changes are not required in the various drivers. • SPBA: The Shared Peripheral Bus Arbiter (SPBA) provides an arbitration mechanism among multiple masters to allow access to the shared peripherals. The SPBA implementation under MSL defines the API to allow different masters to take or release ownership of a shared peripheral.	Machine-Specific Layer (MSL)	All
General Drivers			
Thermal Driver	The thermal driver will monitor the SoC's temperature in a certain frequency to protect the SoC. It defines three trip points: critical, hot, and active.	Thermal Driver	All
OProfile	OProfile is a system-wide profiler for Linux systems, capable of profiling all running code at low overhead.	OProfile	All
Pulse Width Modulator	The pulse-width modulator (PWM) has a 16-bit counter and is optimized to generate sound from stored sample audio images and generate tones.	Pulse-Width Modulator (PWM)	All
Sensors	Sensors cover accelerometer, ambient light and magnetometer sensors.	Sensors	All

Table continues on the next page...

Table 1-1. BSP Supported Features (continued)

Feature	Description	Chapter Source	Applicable Platform
Watchdog	The Watchdog Timer module protects against system failures by providing an escape from unexpected hang or infinite loop situations or programming errors.	Watchdog	All
DMA Engine			
SDMA API	The Smart Direct Memory Access (SDMA) API driver controls the SDMA hardware and provides an API to other drivers for transferring data between MCU, DSP and peripherals.	Smart Direct Memory Access (SDMA) API	All
APBH-Bridge-DMA	Both AHB-to-APBH and AHB-to-APBX DMA support configurable DMA descript chain.	AHB-to-APBH Bridge with DMA (APBH-Bridge-DMA)	All
Power Management Drivers			
Low-level Power Management	The low-level power management driver implements hardware-specific operations to meet power requirements and conserves power. Driver implementations are often different for different platforms. It is used by the DPM layer.	Low-level Power Management (PM) Driver	All
Dynamic Bus Frequency	The bus frequency driver dynamically manages the various system frequencies to improve power consumption.	Dynamic Bus Frequency Driver	i.MX 6 and i.MX 7
CPU Freq	The CPU frequency scaling allows the clock speed of CPU to be changed.	CPUFreq	All
PMIC PF Regulator	PF regulator driver provides the low-level control of the power supply regulators, selection of voltage levels, and enabling/disabling of regulators.	PF_Regulator	All
Anatop Regulator	The Anatop regulator drive provides low-level control of power supply regulators.	Anatop Regulator	i.MX 6 and i.MX 7
Connectivity Drivers			
ENET 1588 Stack	Implementation of the Precision Time Protocol (PTP) according to IEEE standard 1588.	Fast Ethernet Controller (FEC) Driver	All
Fast Ethernet Controller	The ENET Driver performs the full set of IEEE 802.3/Ethernet CSMA/CD media access control and channel interface functions.	Fast Ethernet Controller (FEC) Driver	All
FlexCAN	The FlexCAN driver provides the interfaces to send and receive CAN messages.	FlexCAN Driver	i.MX 6Quad, i.MX 6Dual, i.MX 6DualLite, i.MX 6Solo, i.MX 6UltraLite, i.MX 6SoloX
MediaLB	MediaLB is an on-PCB or inter-chip communication bus allowing applications to access the MOST Network data or communicate with other applications.	MediaLB	i.MX 6SoloX i.MX 6Quad i.MX 6Dual
PCIe	PCI Express hardware module can either be configured to act as a Root Complex or a PCIe Endpoint.	PCIe	All
Video			
Capture	Camera Overview for Camera and capture interfaces.	Capture Overview	All

Table continues on the next page...

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

Table 1-1. BSP Supported Features (continued)

Feature	Description	Chapter Source	Applicable Platform
Display	Display Overview.	Display Overview	All
VPU	The Video Processing Unit (VPU) is a multistandard video decoder and encoder that can perform decoding and encoding of various video formats.	Video Processing Unit (VPU) Driver	i.MX 6QuadPlus/ Quad/Dual/ Solo and i.MX 8
Audio Drivers			
ALSA Sound	The Advanced Linux Sound Architecture (ALSA) is a sound driver that provides ALSA and OSS compatible applications with the means to perform audio playback and recording functions.	ALSA Sound Driver	All
ASRC	The Asynchronous Sample Rate Converter (ASRC) driver provides the interfaces to access the asynchronous sample rate converter module.	Asynchronous Sample Rate Converter (ASRC)	All
S/PDIF	The S/PDIF driver is designed under the Linux ALSA subsystem. It implements one playback device for Tx and one capture device for Rx.	The Sony/Philips Digital Interface (S/PDIF) Driver	All
Storage MTD Drivers			
SPI NOR MTD	The SPI NOR MTD driver provides the support to the Atmel data Flash using the SPI interface.	SPI NOR Flash Memory Technology Device (MTD) Driver	All
NAND MTD	The NAND MTD driver interfaces with the integrated NAND controller supporting UBIFS, CRAMFS and JFFS2UBI and UBIFSCRAMFS and JFFS2 file systems.	NAND GPMI Flash Driver	i.MX 6Quad, i.MX 6Dual, i.MX 6DualLite, i.MX 6Solo, i.MX 6UltraLite, i.MX 7Dual
SATA	The SATA AHCI driver is based on the LIBATA layer of the block device infrastructure of the Linux kernel.	SATA Driver	i.MX 6QuadPlus, i.MX 6Quad, i.MX 6Dual, i.MX 8QuadMax, i.MX 8QuadXPlus
Bus Drivers			
I2C	The Lower Power I2C bus driver interfaces with the I2C bus to transfer data over the I2C bus.	Inter-IC (I2C) Driver	All
eCSPI	The low-level Enhanced Configurable Serial Peripheral Interface (ECSPI) driver interfaces a custom, kernel-space API to both ECSPI modules.	Enhanced Configurable Serial Peripheral Interface (ECSPI) Driver	All
MMC/SD/SDIO - uSDHC	The MMC/SD/SDIO Host driver implements the standard Linux driver interface to eSDHC.	MMC/SD/SDIO Host Driver	All
Connectivity Drivers			
UART	The Universal Asynchronous Receiver/Transmitter (UART) driver interfaces the serial driver API to all UART ports.	Universal Asynchronous Receiver/ Transmitter (UART) Driver	All

Table continues on the next page...

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

Table 1-1. BSP Supported Features (continued)

Feature	Description	Chapter Source	Applicable Platform
USB	The USB driver interfaces to the ARC USB-OTG controller.	CHIPIDEA USB	All

1.2 Audience

This document is targeted to individuals who will port the i.MX Linux® OS Board Support Package (BSP) to customer-specific products.

The audience is expected to have a working knowledge of the Linux kernel internals, driver models, and i.MX processors.

1.2.1 Conventions

This document uses the following notational conventions:

- Courier monospaced type indicate commands, command parameters, code examples, and file and directory names.
- *Italic* type indicates replaceable command or function parameters.
- **Bold** type indicates function names.

1.2.2 Definitions, Acronyms, and Abbreviations

The following table defines the acronyms and abbreviations used in this document.

Table 1-2. Definitions and Acronyms

Term	Definition
ADC	Asynchronous Display Controller
address translation	Address conversion from virtual domain to physical domain
API	Application Programming Interface
Arm®	Advanced RISC Machines processor architecture
AUDMUX	Digital audio MUX-provides a programmable interconnection for voice, audio, and synchronous data routing between host serial interfaces and peripheral serial interfaces
BCD	Binary Coded Decimal
bus	A path between several devices through data lines

Table continues on the next page...

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

Table 1-2. Definitions and Acronyms (continued)

Term	Definition
bus load	The percentage of time a bus is busy
CODEC	Coder/decoder or compression/decompression algorithm-used to encode and decode (or compress and decompress) various types of data
CPU	Central Processing Unit-generic term used to describe a processing core
CRC	Cyclic Redundancy Check-Bit error protection method for data communication
CSI	Camera Sensor Interface
DFS	Dynamic Frequency Scaling
DMA	Direct Memory Access-an independent block that can initiate memory-to-memory data transfers
DPM	Dynamic Power Management
DCSS	Display controller sub system
DP	Display Port - similiar IP as HDMI
DPU	Display Processor Unit
DSI	Display Serial Interface
DRM	Display Rendering Manager or Digital Rights Manager
DRAM	Dynamic Random Access Memory
DVFS	Dynamic Voltage Frequency Scaling
EMI	External Memory Interface-controls all IC external memory accesses (read/write/erase/program) from all the masters in the system
Endian	Refers to byte ordering of data in memory. Little endian means that the least significant byte of the data is stored in a lower address than the most significant byte. In big endian, the order of the bytes is reversed.
EPDC	Electrophoretic Display Controller
EPIT	Enhanced Periodic Interrupt Timer-a 32-bit set and forget timer capable of providing precise interrupts at regular intervals with minimal processor intervention
FCS	Frame Checker Sequence
FIFO	First In First Out
FIPS	Federal Information Processing Standards-United States Government technical standards published by the National Institute of Standards and Technology (NIST). NIST develops FIPS when there are compelling Federal government requirements such as for security and interoperability but no acceptable industry standards.
FIPS-140	Security requirements for cryptographic modules-Federal Information Processing Standard 140-2(FIPS 140-2) is a standard that describes US Federal government requirements that IT products should meet for Sensitive, but Unclassified (SBU) use.
Flash	A non-volatile storage device similar to EEPROM, where erasing can be done only in blocks or the entire chip.
Flash path	Path within ROM bootstrap pointing to an executable Flash application
Flush	Procedure to reach cache coherency. Refers to removing a data line from cache. This process includes cleaning the line, invalidating its VBR and resetting the tag valid indicator. The flush is triggered by a software command.
GPIO	General Purpose Input/Output
GPU	Grapics Processor Unit
hash	Hash values are produced to access secure data. A hash value (or simply hash), also called a message digest, is a number generated from a string of text. The hash is substantially smaller than the text itself, and is generated by a formula in such a way that it is extremely unlikely that some other text produces the same hash value.

Table continues on the next page...

Table 1-2. Definitions and Acronyms (continued)

Term	Definition
HDMI	High-Definition Multimedia Interface
I/O	Input/Output
ICE	In-Circuit Emulation
IP	Intellectual Property
IPU	Image Processing Unit -supports video and graphics processing functions and provides an interface to video/ still image sensors and displays
IrDA	Infrared Data Association-a nonprofit organization whose goal is to develop globally adopted specifications for infrared wireless communication.
ISR	Interrupt Service Routine
JTAG	JTAG (IEEE® Standard 1149.1) A standard specifying how to control and monitor the pins of compliant devices on a printed circuit board
Kill	Abort a memory access
KPP	KeyPad Port-16-bit peripheral used as a keypad matrix interface or as general purpose input/output (I/O)
LDB	LVDS Display Bridge
line	Refers to a unit of information in the cache that is associated with a tag
LRU	Least Recently Used-a policy for line replacement in the cache
LVDS	Low Voltage Differential Signalling
MIPI	Mobile Industry Process Interface
MMU	Memory Management Unit-a component responsible for memory protection and address translation
MPEG	Moving Picture Experts Group-an ISO committee that generates standards for digital video compression and audio. It is also the name of the algorithms used to compress moving pictures and video.
MPEG standards	Several standards of compression for moving pictures and video: • MPEG-1 is optimized for CD-ROM and is the basis for MP3 • MPEG-2 is defined for broadcast video in applications such as digital television set-top boxes and DVD • MPEG-3 was merged into MPEG-2 • MPEG-4 is a standard for low-bandwidth video telephony and multimedia on the World-Wide Web
MQSPI	Multiple Queue Serial Peripheral Interface-used to perform serial programming operations necessary to configure radio subsystems and selected peripherals
MSHC	Memory Stick Host Controller
NAND Flash	Flash ROM technology-NAND Flash architecture is one of two flash technologies (the other being NOR) used in memory cards such as the Compact Flash cards. NAND is best suited to flash devices requiring high-capacity data storage. NAND flash devices offer storage space up to 512-Mbyte and offers faster erase, write, and read capabilities over NOR architecture
NOR Flash	See NAND Flash
PCMCIA	Personal Computer Memory Card International Association-a multicompany organization that has developed a standard for small, credit card-sized devices, called PC Cards. There are three types of PCMCIA cards that have the same rectangular size (85.6 by 54 millimeters), but different widths
physical address	The address by which the memory in the system is physically accessed
PLL	Phase Locked Loop-an electronic circuit controlling an oscillator so that it maintains a constant phase angle (a lock) on the frequency of an input, or reference, signal.
PxP	Pixel Pipeline
RAM	Random Access Memory
RAM path	Path within ROM bootstrap leading to the downloading and the execution of a RAM application

Table continues on the next page...

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

Table 1-2. Definitions and Acronyms (continued)

Term	Definition
RGB	The RGB color model is based on the additive model in which Red, Green, and Blue light are combined to create other colors. The abbreviation RGB comes from the three primary colors in additive light models
RGBA	RGBA color space stands for Red Green Blue Alpha. The alpha channel is the transparency channel, and is unique to this color space. RGBA, like RGB, is an additive color space, so the more of a color placed, the lighter the picture gets. PNG is the best known image format that uses the RGBA color space
RNGA	Random Number Generator Accelerator-a security hardware module that produces 32-bit pseudo random numbers as part of the security module.
ROM	Read Only Memory
ROM bootstrap	Internal boot code encompassing the main boot low as well as exception vectors
RPMSG	Remote Processor Messaging
RTIC	Real-Time Integrity Checker-a security hardware module
SC	System Controller
SCC	SeCurity Controller-a security hardware module
SCFW	System Controller Firmware
SDMA	Smart Direct Memory Access
SDRAM	Synchronous Dynamic Random Access Memory
SoC	System on a Chip
SPBA	Shared Peripheral Bus Arbiter-a three-to-one IP-Bus arbiter, with a resource-locking mechanism
SPI	Serial Peripheral Interface-a full-duplex synchronous serial interface for connecting low-/medium-bandwidth external devices using four wires. SPI devices communicate using a master/slave relationship over two data lines and two control lines: *Also see SS, SCLK, MISO, and MOSI*
SRAM	Static Random Access Memory
SSI	Synchronous-Serial Interface-standardized interface for serial data transfer
TBD	To Be Determined
UART	Universal Asynchronous Receiver/Transmitter-asynchronous serial communication to external devices
UID	Unique ID-a field in the processor and CSF identifying a device or group of devices
USB	Universal Serial Bus-an external bus standard that supports high-speed data transfers. The USB 1.1 specification supports data transfer rates of up to 12 Mb/s and USB 2.0 has a maximum transfer rate of 480 Mbps. A single USB port can be used to connect up to 127 peripheral devices, such as mice, modems, and keyboards. USB also supports Plug-and-Play installation and hot plugging
USBOTG	USB On The Go-an extension of the USB 2.0 specification for connecting peripheral devices to each other. USBOTG devices, also known as dual-role peripherals, can act as limited hosts or peripherals themselves depending on how the cables are connected to the devices, and they also can connect to a host PC
VADC	Video analaog to Digital Converter
VPU	Video Processing Unit
word	A group of bits comprising 32-bits

1.3 References

i.MX has multiple families supported in software. The following are the listed families and SoCs per family. The i.MX Linux® Release Notes describes which SoC is supported in the current release. Some previously released SoCs might be buildable in the current release but not validated if they are at the previous validated level.

- i.MX 6 Family: 6QuadPlus, 6Quad, 6DualLite, 6SoloX, 6SoloLite, 6SLL, 6UltraLite, 6ULL
- i.MX 7 Family: 7Dual, 7ULP
- i.MX 8 Family: 8QuadMax
- i.MX 8M Family: 8M Quad, 8M Mini
- i.MX 8X Family: 8QuadXPlus

This release includes the following references and additional information.

- *i.MX Linux® Release Notes* (IMXLXRN) - Provides the release information.
- *i.MX Linux® User's Guide* (IMXLUG) - Contains the information on installing U-Boot and Linux OS and using i.MX-specific features.
- *i.MX Yocto Project User's Guide* (IMXLXYOCTOUG) - Contains the instructions for setting up and building Linux OS in the Yocto Project.
- *i.MX Reference Manual* (IMXLXRM) - Contains the information on Linux drivers for i.MX.
- *i.MX Graphics User's Guide* (IMXGRAPHICUG) - Describes the graphics features.
- *i.MX BSP Porting Guide* (IMXXBSPPG) - Contains the instructions on porting the BSP to a new board.
- *i.MX VPU Application Programming Interface Linux® Reference Manual* (IMXVPUAPI) - Provides the reference information on the VPU API on i.MX 6 VPU.

The quick start guides contain basic information on the board and setting it up. They are on the NXP website.

- SABRE Platform Quick Start Guide (IMX6QSDPQSG)
- SABRE Board Quick Start Guide (IMX6QSDBQSG)
- i.MX 6UltraLite EVK Quick Start Guide (IMX6ULTRALITEQSG)
- i.MX 6ULL EVK Quick Start Guide (IMX6ULLQSG)
- SABRE Automotive Infotainment Quick Start Guide (IMX6SABREINFOQSG)
- i.MX 6SoloLite Evaluation Kit Quick Start Guide (IMX6SLEVKQSG)
- i.MX 7Dual SABRE-SD Quick Start Guide (SABRESDBIMX7DUALQSG)
- i.MX 8M Quad Evaluation Kit Quick Start Guide (IMX8MQUADEVKQSG)

Documentation is available online at nxp.com.

- i.MX 6 information is at nxp.com/iMX6series

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

- i.MX SABRE information is at nxp.com/imxSABRE
- i.MX 6SoloLite EVK information is at nxp.com/6SLEVK
- i.MX 6UltraLite information is at nxp.com/iMX6UL
- i.MX 6ULL information is at nxp.com/iMX6ULL
- i.MX 7Dual information is at nxp.com/iMX7D
- i.MX 7ULP information is at nxp.com/imx7ulp
- i.MX 8 information is at nxp.com/imx8

Chapter 2
System

2.1 Machine-Specific Layer (MSL)

2.1.1 Introduction

The Machine-Specific Layer (MSL) provides the Linux kernel with the following machine-dependent components.

- Interrupts including GPIO and EDIO (only on certain platforms)
- Timer
- Memory map
- General Purpose Input/Output (GPIO) including IOMUX on certain platforms
- Clock
- Shared Peripheral Bus Arbiter (SPBA)
- Smart Direct Memory Access (SDMA)

2.1.2 Interrupts (Operation)

This section describes the hardware and software operation of interrupts on the device.

2.1.2.1 Interrupt Hardware Operation

The Interrupt Controller controls and prioritizes all internal and external interrupt sources.

The Interrupt Controller controls and prioritizes all interrupt sources, by default, all interrupts have the same priority.

Each interrupt source can be enabled or disabled by configuring the interrupt controller's registers.

There are three types of interrupts in GIC: PPI, SGI, and SPI.

- PPI is private peripheral interrupts of each CPU. It can only be handled by each CPU.
- SGI is software generated interrupts. It can be triggered by software operation, and it also can only be handled by each CPU.
- SPI is shared peripheral interrupts, which are normally external interrupt sources from SoC platform. It can be handled by all CPUs.

2.1.2.2 Interrupt Software Operation

For the Arm architecture-based processors with GIC-400 of i.MX 6 and i.MX 7 SoCs, normal interrupt and fast interrupt are two different exception types. The exception vector addresses can be configured to start at low address (0x0) or high address (0xFFFF0000) for i.MX 6 and i.MX 7 platforms. The Linux OS implementation running on the Arm architecture chooses the high-vector address model.

For Arm architecture-based processors with GIC-500 of i.MX 8 SoCs, the exception vector addresses are defined as VBAR_ELn + offset. The offset depends on which exception level the interrupt exception is taken. The file Documentation/arm/Interrupts has a description of the Arm interrupt architecture.

The software provides a processor-specific interrupt structure with callback functions defined in the irqchip structure and exports one initialization function, which is called during system startup.

Table 2-1. Interrupt Files

File	Description
drivers/irqchip/irq-gic.c	i.MX 6/7 SoCs with GIC-400
drivers/irqchip/irq-gic-v3.c	i.MX 8 SoCs with GIC-500
drivers/irqchip/irq-imx-irqsteer.c	Interrupt functions with CONFIG_IMX_IRQSTEER configuration
drivers/irqchip/irq-imx-intmux.c	Interrupt functions with CONFIG_IMX_INTMUX configuration
irq-imx-gpcv2.c	Interrupt functions with CONFIG_IMX_GPCV2 configuration

2.1.2.3 Interrupt Features

The interrupt implementation supports the following features:

- Interrupt Controller interrupt disable and enable
- Functions required by the Linux interrupt architecture as defined in the standard Arm interrupt source code

2.1.2.4   Interrupt Source Code Structure

The interrupt module is located in drivers/irqchip.

The table below lists the source files for interrupts.

Table 2-2.   Interrupt Files

File	Description
drivers/irqchip/irq-imx-irqsteer.o.c	Interrupt functions with CONFIG_IMX_IRQSTEER configuration.
drivers/irqchip/irq-imx-gpcv2.c	Interrupt functions with CONFIG_IMX_GPCV2 configuration.
drivers/irqchip/irq-imx-intmux.c	Interrupt functions for with CONFIG_IMX_INTMUX configuration.

2.1.2.5   Interrupt Programming Interface

The machine-specific interrupt implementation exports a single function.

This function initializes the Interrupt Controller hardware and registers functions for interrupt enable and disable from each interrupt source.

This is done with the global structure irq_desc of type struct irqdesc. After the initialization, the interrupt can be used by the drivers through the request_irq() function to register device-specific interrupt handlers.

In addition to the native interrupt lines supported from the Interrupt Controller, the number of interrupts is also expanded to support GPIO interrupt and (on some platforms) EDIO interrupts. This allows drivers to use the standard interrupt interface supported by Arm device running Linux OS, such as the request_irq() and free_irq() functions.

2.1.3   Timer

The Linux kernel relies on the underlying hardware to provide support for both the system timer (which generates periodic interrupts) and the dynamic timers (to schedule events).

After the system timer interrupt occurs, it does the following:

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

- Updates the system uptime
- Updates the time of day
- Reschedules a new process if the current process has exhausted its time slice
- Runs any dynamic timers that have expired
- Updates resource usage and processor time statistics

The following tables describes the different timers used.

Table 2-3.   Timers

Timer	Description
General Purpose Timer (GPT)	GPT is configured to generate a periodic interrupt at a certain interval (every 10 ms). Used by i.MX 6 to go into WFI mode. Used by i.MX 6 and i.MX 7.
Enhanced Periodic Interrupt Timer (EPIT)	Available on i.MX 6 and i.MX 7.
Arm Arch Timer	i.MX 8 usage instead of GPT
System Counter Timer	i.MX 8M and i.MX 8X usage instead of GPT

2.1.3.1   Timer Software Operation

The timer software implementation provides an initialization function that initializes the GPT with the proper clock source, interrupt mode and interrupt interval.

The timer then registers its interrupt service routine and starts timing. The interrupt service routine is required to service the OS for the purposes mentioned in the previous Section Timer. Another function provides the time elapsed as the last timer interrupt.

2.1.3.2   Timer Features

The timer implementation supports the following features:

- Functions required by Linux OS to provide the system timer and dynamic timers.
- Generates an interrupt every 10 ms for i.MX6 and i.MX 7 and every 4 ms for i.MX 8. This is based on CONFIG_HZ_XXX.

2.1.3.3   Timer Source Code Structure

Table 2-4.   Timer Files

File	Description
arch/arm/mach-imx/epit.c	Enhanced Periodic Interrupt Timer

Table continues on the next page...

Table 2-4. Timer Files (continued)

File	Description
driver/clocksource/timer-imx-sysctr.c	System Controller Timer
driver/clocksource/timer-imx-tpm.c	TPM Timer
drivers/clocksource/timer-imx-gpt.c	General Purpose Timer
drivers/clocksource/arch-arm-timer.c	Arm arch Timer

2.1.3.4   Timer Programming Interface

The timer module utilizes four hardware timers, to implement clock source and clock event objects.

This is done with the clocksource_mxc structure of struct clocksource type and clockevent_mxc structure of struct clockevent_device type. Both structures provide routines required for reading current timer values and scheduling the next timer event. The module implements a timer interrupt routine that services the Linux OS with timer events for the purposes mentioned in the beginning of this chapter.

2.1.4   Memory Map

A predefined virtual-to-physical memory map table is required for the device drivers to access to the device registers since the Linux kernel is running under the virtual address space with the Memory Management Unit (MMU) enabled.

2.1.4.1   Memory Map Hardware Operation

The MMU, as part of the Arm core, provides the virtual to physical address mapping defined by the page table. For more information, see the *Arm Technical Reference Manual* (TRM) from Arm Limited.

2.1.4.2   Memory Map Software Operation (only for i.MX 6 or i.MX 7)

A table mapping the virtual memory to physical memory is implemented for i.MX 6 and i.MX 7 platforms as defined in the arch/arm/mach-imx/pm-imx*.c file.

2.1.4.3 Memory Map Features

The Memory Map implementation programs the Memory Map module to creates the physical to virtual memory map for all the I/O modules.

2.1.4.4 Memory Map Source Code Structure

The Memory Map module implementation is used for i.MX 6 and 7 in arch/arm/mach-imx.

Table below lists the source file for the memory map for i.MX 6 and i.MX 7.

Table 2-5. Memory Map Files

File	Description
arch/arm/mach-imx/mx6.h	Header files for the i.MX 6 I/O module physical addresses
arch/arm/mach-imx/mx7.h	Header file for the i.MX 7Dual I/O module physical addresses
arch/arm/mach-imx/mx7ulp.h	Header file for the i.MX 7ULP I/O module physical addresses
arch/arm/mach-imx/hardware.h	Memory map definition file and macros for the following: • I/O module physical and virtual base addresses • Physical to virtual virtual mapping macros.
arch/arm/mach-imx/pm-imx6.c	Memory Map module for i.MX 6
arch/arm/mach-imx/pm-imx7.c	Memory Map module for i.MX 7Dual
arch/arm/mach-imx/pm-imx7ulp.c	Memory Map module for i.MX 7ULP
arch/arm/mach-imx/pm-rpmsg.c	Memory Map RPMSG

2.1.5 IOMUX

The limited number of pins of highly integrated processors can have multiple purposes.

The IOMUX module controls a pin usage so that the same pin can be configured for different purposes and can be used by different modules.

This is a common way to reduce the pin count while meeting the requirements from various customers. Platforms that do not have the IOMUX hardware module can do pin muxing through the GPIO module.

The IOMUX module provides the multiplexing control so that each pin may be configured either as a functional pin or as a GPIO pin. A functional pin can be subdivided into either a primary function or alternate functions. The pin operation is controlled by a specific hardware module. A GPIO pin, is controlled by the user through software with further configuration through the GPIO module. For example, the TXD1 pin might have the following functions:

- TXD1-internal UART1 Transmit Data. This is the primary function of this pin.
- UART2 DTR-alternate mode 3
- LCDC_CLS-alternate mode 4
- GPIO4[22]-alternate mode 5
- SLCDC_DATA[8]-alternate mode 6

If the hardware modes are chosen at the system integration level, this pin is dedicated only to that purpose and cannot be changed by software. Otherwise, the IOMUX module needs to be configured to serve a particular purpose that is dictated by the system (board) design. If the pin is connected to an external UART transceiver and therefore to be used as the UART data transmit signal, it should be configured as the primary function. If the pin is connected to an external Ethernet controller for interrupting the Arm core, then it should be configured as GPIO input pin with interrupt enabled. Again, be aware that the software does not have control over what function a pin should have. The software only configures pin usage according to the system design.

2.1.5.1 IOMUX Hardware Operation

The following discussion applies only to those processors that have an IOMUX hardware module.

The IOMUX controller registers are briefly described in this section.

For detailed information, see the pin multiplexing section of the IC Reference Manual.

- SW_MUX_CTL-Selects the primary or alternate function of a pin. Also enables loopback mode when applicable.
- SW_SELECT_INPUT-Controls pin input path. This register is only required when multiple pads drive the same internal port.
- SW_PAD_CTL-Control pad slew rate, driver strength, pull-up/down resistance, and so on.

2.1.5.2 IOMUX Software Operation

The IOMUX software implementation provides an API to set up pin functionality and pad features.

2.1.5.3 IOMUX Features

The IOMUX implementation programs the IOMUX module to configure the pins that are supported by the hardware.

2.1.5.4   IOMUX Source Code Structure

Table below lists the source files for the IOMUX module. The files are in the drivers/princtrl/freescale folder.

Table 2-6.   IOMUX Files

File	Description
drivers/pinctrl/freescale/pinctrl-imx.c	i.MX pinctrl core driver
drivers/pinctrl/freescale/pinctrl-im6sl.c	i.MX 6SoloLite pinctrl driver
drivers/pinctrl/freescale/pinctrl-imx6q.c	i.MX 6Quad/DualLite pinctrl driver
drivers/pinctrl/freescale/pinctrl-imx6sx.c	i.MX 6SoloX pinctrl driver
drivers/pinctrl/freescale/pinctrl-imx6sll.c	i.MX 6SLL pinctrl driver
drivers/pinctrl/freescale/pinctrl-imx6ul.c	i.MX 6UltraLite and 6ULL pinctrl driver
drivers/pinctrl/freescale/pinctrl-imx7d.c	i.MX 7Dual pinctrl driver
drivers/pinctrl/freescale/pinctrl-imx7ulp.c	i.MX 7ULP pinctrl driver
drivers/pinctrl/freescale/pinctrl-imx8qm.c	i.MX 8QuadMax pinctrl driver
drivers/pinctrl/freescale/pinctrl-imx8qxp.c	i.MX 8QuadXPlus pinctrl driver
drivers/pinctrl/freescale/pinctrl-imx8mq.c	i.MX 8M Quad pinctrl driver
drivers/pinctrl/freescale/pinctrl-imx8mm.c	i.MX 8M Mini pinctrl driver

2.1.5.5   IOMUX Programming Interface

See pinctrl binding documents Documentation/devicetree/bindings/pinctrl/fsl.

- imx-pinctrl.txt
- imx6q-pinctrl.txt
- imx6dl-pinctrl.txt
- imx6sl-pinctrl.txt
- imx6sx-pinctrl.txt
- imx6ul-pinctrl.txt
- imx7d-pinctrl.txt
- imx7ulp-pinctrl.txt
- imx8qm-pinctrl.txt
- imx8qxp-pinctrl.txt
- imx8mq-pinctrl.txt
- imx8mm-pinctrl.txt

2.1.5.6   IOMUX Control Through GPIO Module

For a multipurpose pin, the GPIO controller provides the multiplexing control so that each pin may be configured either as a functional pin, or a GPIO pin.

The operation of the functional pin, which can be subdivided into either major function or one alternate function, is controlled by a specific hardware module. If it is configured as a GPIO pin, the pin is controlled by the user through software with further configuration through the GPIO module. In addition, there are some special configurations for a GPIO pin (such as output based A_IN, B_IN, C_IN or DATA register, but input based A_OUT or B_OUT).

The following discussion applies to those platforms that control the muxing of a pin through the general purpose input/output (GPIO) module.

If the hardware modes are chosen at the system integration level, this pin is dedicated only to that purpose which cannot be changed by software. Otherwise, the GPIO module needs to be configured properly to serve a particular purpose that is dictated with the system (board) design. If this pin is connected to an external UART transceiver, it should be configured as the primary function or if this pin is connected to an external Ethernet controller for interrupting the core, then it should be configured as GPIO input pin with interrupt enabled. The software does not have control over what function a pin should have. The software only configures a pin for that usage according to the system design.

2.1.5.6.1   GPIO Hardware Operation

The GPIO controller module is divided into MUX control and PULLUP control sub modules. The following sections briefly describe the hardware operation. For detailed information, see the relevant device documentation.

2.1.5.6.1.1   Muxing Control

The GPIO In Use Registers control a multiplexer in the GPIO module.

The settings in these registers choose if a pin is utilized for a peripheral function or for its GPIO function. One 32-bit general purpose register is dedicated to each GPIO port. These registers may be used for software control of IOMUX block of the GPIO.

2.1.5.6.1.2   PULLUP Control

The GPIO module has a PULLUP control register (PUEN) for each GPIO port to control every pin of that port.

2.1.5.6.2   GPIO Software Operation (general)

The GPIO software implementation provides an API to setup pin functionality and pad features.

2.1.5.6.3   GPIO Implementation

The GPIO implementation programs the GPIO module to configure the pins that are supported by the hardware.

2.1.6   General Purpose Input/Output (GPIO)

The GPIO module provides general-purpose pins that can be configured as either inputs or outputs.

When configured as an output, the pin state (high or low) can be controlled by writing to an internal register. When configured as an input, the pin input state can be read from an internal register.

2.1.6.1   GPIO Software Operation

The general purpose input/output (GPIO) module provides an API to configure the i.MX processor external pins and a central place to control the GPIO interrupts.

The GPIO utility functions should be called to configure a pin instead of directly accessing the GPIO registers. The GPIO interrupt implementation contains functions, such as the interrupt service routine (ISR) registration/un-registration and ISR dispatching once an interrupt occurs. All driver-specific GPIO setup functions should be made during device initialization in the MSL layer to provide better portability and maintainability. This GPIO interrupt is initialized automatically during the system startup.

If a pin is configured as GPIO by the IOMUX, the state of the pin should also be set since it is not initialized by a dedicated hardware module. Setting the pad pull-up, pull-down, slew rate and so on, with the pad control function may be required as well.

2.1.6.1.1   API for GPIO

API for GPIO lists the features supported by the GPIO implementation.

The GPIO implementation supports the following features:

- An API for registering an interrupt service routine to a GPIO interrupt. This is made possible as the number of interrupts defined by NR_IRQS is expanded to accommodate all the possible GPIO pins that are capable of generating interrupts.
- Functions to request and free an IOMUX pin. If a pin is used as GPIO, another set of request/free function calls are provided. The user should check the return value of the request calls to see if the pin has already been reserved before modifying the pin state. The free function calls should be made when the pin is not needed. See the API document for more details.
- Aligned parameter passing for both IOMUX and GPIO function calls. In this implementation the same enumeration for iomux_pins is used for both IOMUX and GPIO calls and the user does not have to figure out in which bit position a pin is located in the GPIO module.
- Minimal changes required for the public drivers such as Ethernet and UART drivers as no special GPIO function call is needed for registering an interrupt.

2.1.6.2 GPIO Features

This GPIO implementation supports the following features:

- Implements the functions for accessing the GPIO hardware modules
- Provides a way to control GPIO signal direction and GPIO interrupts

2.1.6.3 GPIO Module Source Code Structure

All of the GPIO module source code is in the GPIO framework, in the following files, located in the directories indicated at the beginning of this chapter:

Table 2-7. GPIO Files

File	Description
drivers/gpio/gpio-mxc.c	Function implementation

2.1.6.4 GPIO Programming Interface 2

For more information, see the Documentation/gpio/gpio.txt under Linux source code directory for the programming interface.

2.1.7   Clock

The Linux clock framework relies on the underlying hardware to provide support for clock tree management.

The following table describes different clock hardware used.

File	Description
Clock controller module (CCM)	i.MX 6Quad/DualLite/SoloLite/SoloX/UltraLite/ULL/SLL, i.MX 7Dual, i.MX 8M Quad, and i.MX 8M Mini
Peripheral clock control (PCC) and System clock generator (SCG)	i.MX 7ULP
Distributed slave system controller (DSC)	i.MX 8QuadMax/8QuadXPlus

2.1.7.1   Clock Software Operation

The clock software implementation provides an initialization function that initializes the clock tree according to hardware clock type and settings, and then provides clock operation callbacks to operate the hardware clock module.

2.1.7.2   Clock Features

The clock implementation supports the following features according to different clock types:

- Prepare/Unprepare a clock.
- Enable/Disable a clock.
- Get/Set the clock rate.
- Get/Set the clock parent.

2.1.7.3   Source Code Structure

The source code structure is as follows.

File	Description
drivers/clk/imx/clk-imx6q.c	i.MX 6Quad/6DualLite clock driver
drivers/clk/imx/clk-imx6sl.c	i.MX 6SoloLite clock driver
drivers/clk/imx/clk-imx6sx.c	i.MX 6SoloX clock driver
drivers/clk/imx/clk-imx6ul.c	i.MX 6UltraLite and 6ULL clock driver

Table continues on the next page...

File	Description
drivers/clk/imx/clk-imx6sll.c	i.MX 6SLL clock driver
drivers/clk/imx/clk-imx7d.c	i.MX 7Dual clock driver
drivers/clk/imx/clk-imx7ulp.c	i.MX 7ULP clock driver
drivers/clk/imx/clk-imx8qm.c	i.MX 8QuadMax clock driver
drivers/clk/imx/clk-imx8qxp.c	i.MX 8QuadXPlus clock driver
drivers/clk/imx/clk-imx8mq.c	i.MX 8MQuad clock driver
drivers/clk/imx/clk-imx8mm.c	i.MX 8MMini clock driver

2.1.7.4

Different clock types provide different clock operation callbacks. Device drivers call standard clock APIs to clock framework and eventually call into platform clock driver, and the corresponding clock node's operation callback is executed.

2.2 System Controller

2.2.1 Introduction

The System Controller is provided on i.MX 8 and i.MX 8X families and provides an abstraction to many underlying features of the hardware and runs on a Cortex-M processor which executes SC firwmare (SCFW). This overview describes the features of th eSCFW and APIs exposed to other software components.

The System Controller features include:

- System Intiialization and Boot - The SC firmware runs on the SCU immediately after the SCU Read-only-memory (ROM) finishes loading code/data images from the first container. It is responsible for initializing many aspects of the system. This includes additional power and clock configuration and resource isolation hardware configuration. By default, the SC firmware configures the primary boot core to own most of the resources and launches the boot core. Additional configuration can be done by boot code.
- System Controller Communication - Other software components in the system communicate to the SC via an exposed API library. This library is implemented to make Remote Procedure Calls (RPC) via an underlying Inter-Processor Communication (IPC) mechanism. The IPC is facilitated by a hardware-based mailbox system. Software components (Linux, QNX, FreeRTOS, KSDK) delivered

for i.MX8 already include ports of the client API. Other 3rd parties will need to first port the API to their environment before the API can be used. The porting kit release includes archives of the client API for existing SW. These can be used as reference for porting the client API. All that needs to be implemented is the IPC layer which will utilize messaging units (MU) to communicate with the SCFW.

- Power Management - All aspects of power management including power control, bias control, clock control, reset control, and wake-up event monitoring are grouped within the SC Power Management service.

  - **Power Control** - The SC firmware is responsible for centralized management of power controls and external power management devices. It manages the power state and voltage of power domains as well as bias control. It also resets peripherals as required due to power state transitions. This is immmplemented with the API by communicating power state needs for individual resources.
  - **Clock Control** - The SC firmware is responsible for centralized management of clock controls. This includes clock sources such as oscillators and PLLs as well as clock dividers, muxes, and gates. This is implemented with the API by communicating clocking needs for individual resources.
  - **Reset Control** - The SC firmware is responsible for reset control. This includes booting/rebooting a partition, obtaining reset reasons, and starting/stopping of CPUs.

  Before any hardware in the SoC can be used, SW must first power up the resource and enable any clocks that it requires, otherwise access will generate a bus error.

- Resource Management - SC firmware is responsible for managing ownership and access permissions to system resources. The features of the resource management service supported by SC firmware include:

  - Management of system resources such as SoC peripherals, memory regions, and pads
  - Allows resources to be partitioned into different ownership groupings that are associated with different execution environments including multiple operating systems executing on different cores, TrustZone, and hypervisor
  - Associates ownership with requests from messaging units within a resource partition
  - Allows memory to be divided into memory regions that are then managed like other resources
  - Allows owners to configure access permissions to resources
  - Configures hardware components to provide hardware enforced isolation
  - Configures hardware components to directly control secure/nonsecure attribute driven on bus fabric

- Provides ownership and access permission information to other system controller functions (e.g. pad ownership information to the pad muxing functions)
- Protection of resources is provided in two ways. First, the SCFW itself checks resource access rights when API calls are made that affect a specific resource. Depending on the API call, this may require that the caller be the owner, parent of the owner, or an ancestor of the owner. Second, any hardware available to enforce access controls is configured based on the RM state. This includes the configuration of IP such as XRDC2, XRDC, or RDC, as well as management pages of IP like CAAM.
- Pad Configuration - Pad configuration is managed by SC firmware. The pad configuration features supported by the SC firmware include:
  - Configuring the mux, input/output connection, and low-power isolation mode.
  - Configuring the technology-specific pad setting such as drive strength, pullup/pulldown, etc.
  - Configuring compensation for pad groups with dual voltage capability.
- Timers - Many timer oriented services are grouped within the SC Timer service. This includes watchdogs, RTC, and system counter.
  - **Watchdog** - The SC firmware provides "virtual" watchdogs for all execution environments. Features include update of the watchdog timeout, start/stop of the watchdog, refresh of the watchdog, return of the watchdog status such as maximum watchdog timeout that can be set, watchdog timeout interval, and watchdog timeout interval remaining.
  - **Real-Time-Clock** - The SC firmware is responsible for providing access to the RTC. Features include setting the time, getting the time, and setting alarms.
  - **System Counter** - The SC firmware is responsible for providing access to the SYSCTR. Features incude setting an absolute alarm or a relative, periodic alarm. Reading is done directly via local hardware interfaces available for each CPU.
- Interrupts - The System Controller needs a method to inform users about asynchronous notification events. This is done via the Interrupt service. The service provides APIs to enable/disable interrupts to the user and to read the status of pending interrupts. Reading the status automatically clears any pending state.
- Miscellaneous - On previous i.MX devices, miscellaneous features were controlled using IOMUX GPR registers with signals connected to configurable hardware. This functionality is being replaced with DSC GPR signals. SC firmware is responsible for programming the GPR signals to configure these subsystem features. The SC firmware also responsible for monitoring various temperature, voltage, and clock sensors.
  - **Controls** - The SC firmware provides access to miscellaneous controls. Features include software request to set (write) miscellaneous controls and software request to get (read) miscellaneous controls.

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

- **Security** - The SC firmware provides access to several security functions including image loading and authentication.
- **DMA** - The SC firmware provides access to DMA channel grouping and priority functions.
- **Temp** - The SC firmware provides access to temperature sensors.

With this abstraction some hardware described in the SoC Reference Manual that is used by the SCFW is not directly accessible to other cores. This includes the following

- All resources in the SCU subsystem (SCU M4, SCU LPUART, SCU LPI2C, etc.).
- All resource accessed via MSI links from the SCU subsystem (inc. pads, DSC, XRDC2, eCSR)
- OCRAM controller, CAAM MP, eDMA MP and LPCG
- DB STC and LPCG, IMG GPR
- GIC/IRQSTR LPCG, IRQSTR.SCU and IRQSTR.CTI
- Any other resources reserved by the port of the SCFW to the board

The System Controll firwmware known as SCFW provided with each release works with associated i.MX reference boards and a porting kit is provided that provides a subset of source that can be customized for new boards. This porting kit is avaiable on nxp.com and includes a porting guide.

2.3 Boot Image

2.3.1 Introduction

For i.MX 6 and i.MX 7, the boot image uses only the U-Boot bootloader. For the SoC in the i.MX 8 series, the boot image is more complex and includes U-Boot as well various firmware required for a successful boot. This chapter describes the additional components for an i.MX 8 series boot loader.

For i.MX 7ULP, the boot partition requires the Arm Cortex M-4 SDK flash since the Arm Cortex M-4 boots the U-Boot boot loader, but other i.MX 6 and i.MX7 with Arm Cortex M-4 cores do not require this for succesful boot.

The i.MX 8 bootloader is created using imx-mkimage tool available on imx-mkimage on code aurora forum/ and all i.MX 8 Series require Arm trusted firmware available on imx-atf on code aurora forum/.

For details on how to use the imx-mkimage tool to create an i.MX boot partition, refer to the *i.MX Linux User's Guide*. This tool for execution requires the following components.

For i.MX 8M Quad and i.MX 8M Mini, the following firmware is needed:

- Synopys DDR frimware.
- Signed HDMI firmware - that integrates with the DCSS driver. HDMI firmware is for i.MX 8MQuad only.
- Arm Trusted firmware - bl31-*soc*.

For i.MX 8QuadMax, the following firmware is needed:

- System Controller Firmware (SCFW)
- Arm Trusted firmware - bl31-*soc*

For i.MX 8QuadXPlus, the following firmware is needed:

- System Controller Firmware (SCFW)
- Arm Trusted firmware - bl31-*soc*
- SECO firmware container image (ahab-container.img)

All the i.MX series require Arm trusted firmware and U-boot. Also i.MX SoC supporting optee (all i.MX 6, 7 and 8M families) enabled with Optee boot need the tee.bin created from building optee_ox.

Type 1 hypervisors like Xen are are part of the boot loader. However, Type 2 hypervisors such as jailhouse and kvm are not.

2.4 Anatop Regulator Driver

2.4.1 Introduction

The Anatop regulator driver provides the low-level control of the power supply regulators, and selection of voltage levels.

This device driver makes use of the regulator core driver to access the Anatop hardware control registers and is only supported on i.MX 6 and i.MX 7.

2.4.2 Hardware Operation

The Power Management Unit on the die is built to simplify the external power interface and allow the die to be configured in a power appropriate manner. The power system consists of the input power sources and their characteristics, the integrated power transforming and controlling elements, and the final load interconnection and requirements.

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

Utilizing 7 LDO regulators, the number of external supplies is greatly reduced. If the backup coin and USB inputs are neglected, then the number of external supplies is reduced to two. Missing from this external supply total are the necessary external supplies to power the desired memory interface. This will change depending on the type of external memory selected. Other supplies might also be necessary to supply the voltage to the different I/O power segments if their I/O voltage needs to be different than what is provided above.

Some internal regulator can be bypassed, so that the external PMIC can supply power directly to decrease power number, such as VDD_SOC and VDD_ARM.

2.4.3   Software Operation

The Anatop regulator client driver performs operations by reconfiguring the Anatop hardware control registers. This is done by calling regulator core APIs with the required register settings.

2.4.4   Driver Features

The Anatop regulator driver is based on regulator core driver. A list of services provided for regulator control can be found here.

- Switch ON/OFF all voltage regulators.
- Set the value for all voltage regulators.
- Get the current value for all voltage regulators.

2.4.5   Driver Interface Details

Access to the Anatop regulator is provided through the API of the regulator core driver. The Anatop regulator driver provides the following regulator controls:

- Seven LDO regulators
- All of the regulator functions are handled by setting the appropriate Anatop hardware register values. This is done by calling the regulator core APIs to access the Anatop hardware registers.

2.4.6   Regulator APIs

The regulator power architecture is designed to provide a generic interface to voltage and current regulators within the Linux kernel. It is intended to provide voltage and current control to client or consumer drivers and also provide status information to user space applications through a sysfs interface. The intention is to allow systems to dynamically control regulator output to save power and prolong battery life. This applies to both voltage regulators (where voltage output is controllable) and current sinks (where current output is controllable).

For more details visit opensource.wolfsonmicro.com/node/15

Under this framework, most power operations can be done by the following unified API calls:

- `regulator_get` used to lookup and obtain a reference to a regulator:
  - `struct regulator *regulator_get(struct device *dev, const char *id);`

- `regulator_put` used to free the regulator source:
  - `void regulator_put(struct regulator *regulator, struct device *dev);`

- `regulator_enable` used to enable regulator output:
  - `int regulator_enable(struct regulator *regulator);`

- `regulator_disable` used to disable regulator output:
  - `int regulator_disable(struct regulator *regulator);`

- `regulator_is_enabled` is the regulator output enabled:
  - `int regulator_is_enabled(struct regulator *regulator);`

- `regulator_set_voltage` used to set regulator output voltage:
  - `int regulator_set_voltage(struct regulator *regulator, int uV);`

- `regulator_get_voltage` used to get regulator output voltage:
  - `int regulator_get_voltage(struct regulator *regulator);`

For more APIs and details in the regulator core source code inside the Linux kernel see: drivers/regulator/core.c.

2.4.7   Source Code Structure

The Anatop regulator driver is located in the drivers/regulator directory:

Table 2-8.   Anatop Power Management Driver Files

File	Description
drivers/regulator/core.c	Regulator interface
drivers/regulator/anatop-regulator.c	Anatop regulator client driver

The Anatop regulators are registered in each SoC-specific dts file in arch/arm/boot/dts.

2.4.8 Menu Configuration Options

In menu configuration enable the following module:

- Device Drivers > Voltage and Current regulator support > Anatop Regulator Support.
- System Type > Freescale i.MX on-chip ANATOP LDO regulators.

2.5 Power Management

2.5.1 Low Level Power Management (PM)

2.5.1.1 Introduction

Information found here describes the low-level Power Management (PM) driver which controls the low-power modes.

The following describes the differences between how power management is handled for each supported i.MX family.

Table 2-9. Power Management Modes

i.MX Family	Supported Low Power Modes
i.MX 6	RUN, WAIT, STOP, and DORMANT
i.MX 7	RUN, WAIT, STOP, DORMANT, and LPSR
i.MX 8M	RUN, IDLE, SUSPEND, and SNVS
i.MX 8 and i.MX 8X	None - handled by the System Controller

Table below lists the detailed clock information for the different low power modes.

Table 2-10. Low Power Modes

Mode	Core	Modules	PLL	CKIH/FPM	CKIL
RUN	Active	Active, Idle or Disable	On	On	On
WAIT	Disable	Active, Idle or Disable	On	On	On
STOP	Disable	Disable	Off	On	On
LPSR	Power off	Disable	Off	Off	On

Table continues on the next page...

Table 2-10. Low Power Modes (continued)

Mode	Core	Modules	PLL	CKIH/FPM	CKIL
DORMANT	Power off	Disable	Off	Off	On
SNVS	Power off	Disable	Off	Off	On

For detailed information about low power modes, see the Applications Processor Reference Manual associated with SoC.

2.5.1.2 Software Operation

The i.MX 6 and i.MX 7 power management driver maps the low-power modes to the kernel power management states as listed below:

- Standby-maps to STOP mode, which offers significant power saving, as all blocks in the system are put into a low-power state, except for Arm® core, which is still powered on, and memory is placed in self-refresh mode to retain its contents.
- Mem (suspend to RAM) maps to DORMANT mode, which offers most significant power saving, as all blocks in the system are put into a low-power state, except for memory, which is placed in self-refresh mode to retain its contents. If there is "fsl,enable-lpsr" defined in DTB ocrams node, mem is mapped to LPSR mode instead of DORMANT, and all the blocks in the system are put into power off state, except the LPSR, SNVS, and DRAM power domains.
- System idle maps to WAIT mode.
- If Arm Cortex®-M4 processor is alive together with Arm Cortex-A processor before the kernel enters standby/mem mode, and if Arm Cortex-M4 processor is not in its low-power idle mode, Arm Cortex-A processor triggers the SOC to enter WAIT mode instead of STOP mode to make sure that Arm Cortex-M4 processor can continue running.

The i.MX 6 and i.MX 7 power management driver performs the following steps to enter and exit low power mode:

1. Allow the Cortex-A platform to issue a deep sleep mode request.
2. If STOP or DORMANT mode:
   - Program i.MX 6 CCM_CLPCR or i.MX 7 GPC_LPCR_A7_BSC and GPC_SLPCR registers to set low-power control register.
   - If DORMANT mode, request switching off CPU power when pdn_req is asserted.
   - Request switching off embedded memory peripheral power when pdn_req is asserted.
   - Program GPC mask register to unmask wakeup interrupts.

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

3. Call cpu_do_idle to execute WFI pending instructions for wait mode.
4. Execute imx6_suspend or imx7_suspend in IRAM.
5. In DORMANT mode, save Arm context, and change the drive strength of DDR PADs as "low" to minimize the power leakage in DDR PADs. Execute WFI pending instructions for stop mode.
6. Generate a wakeup interrupt and exit low-power mode. In DORMANT mode, restore Arm core and DDR drive strength.

In DORMANT mode, the i.MX 6 and i.MX 7 can assert the PMIC_STBY_REQ pin to the PMIC and request a voltage change. The U-Boot or Machine-Specific Layer (MSL) usually sets the standby voltage in STOP mode according to i.MX 6 and i.MX 7 data sheet.

On i.MX 8M Family the power management driver uses the following modes.

- RUN Mode: In this mode, the Quad-A53 CPU core is active and running. Some portions can be shut off for power saving.
- IDLE Mode: This mode is defined as a mode which CPU can automatically enter when there is no thread running and all high-speed devices are not active. The CPU can be put into power gated state but with L2 data retained, DRAM and bus clock are reduced, and most of the internal logics are clock gated but still remain powered.
- SUSPEND Mode: This mode is defined as the most power saving mode where all the clocks are off and all the unnecessary power supplies are off. Cortex-A53 CPU platform is fully power gated. All the internal digital logics and analog circuits that can be powered down will be off.
- SNVS Mode: This mode is also called RTC mode. In this mode, only the power for the SNVS domain remains on to keep RTC and SNVS logic alive.

On i.MX 8 and i.MX 8X:

- No hardware low-power mode is available.
- All low-power modes are implemented in system controller firmware (SCFW) using software method.
- SCFW powers off clusters/CPUs when the system is suspended.

2.5.1.3 Source Code Structure

Table below shows Power Management driver source files.

Table 2-11. Power Management Driver Files

File	Description
• arch/arm/mach-imx/pm-imx6.c	Supports i.MX 6 QuadPlus/Quad/Dual/Solo power management operation

Table continues on the next page...

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

Table 2-11. Power Management Driver Files (continued)

File	Description
• arch/arm/mach-imx/suspend-imx6.S • arch/arm/mach-imx/cpuidle-imx6q.c	
• arch/arm/mach-imx/pm-imx6.c • arch/arm/mach-imx/suspend-imx6.S • arch/arm/mach-imx/cpuidle-imx6sl.c • arch/arm/mach-imx/imx6sl_low_power_idle.S	Supports i.MX 6 SoloLite power management operation
• arch/arm/mach-imx/pm-imx6.c • arch/arm/mach-imx/suspend-imx6.S • arch/arm/mach-imx/cpuidle-imx6sll.c • arch/arm/mach-imx/imx6sll_low_power_idle.S	Supports i.MX 6 SLL power management operation
• arch/arm/mach-imx/pm-imx6.c • arch/arm/mach-imx/suspend-imx6.S • arch/arm/mach-imx/cpuidle-imx6ul.c • arch/arm/mach-imx/imx6ul_low_power_idle.S	Supports i.MX 6 UltraLite power management operation
• arch/arm/mach-imx/pm-imx6.c • arch/arm/mach-imx/suspend-imx6.S • arch/arm/mach-imx/cpuidle-imx6ul.c • arch/arm/mach-imx/imx6ull_low_power_idle.S	Supports i.MX 6 ULL power management operation
• arch/arm/mach-imx/pm-imx6.c • arch/arm/mach-imx/suspend-imx6.S • arch/arm/mach-imx/cpuidle-imx6sx.c • arch/arm/mach-imx/imx6sx_low_power_idle.S	Supports i.MX 6 SoloX power management operation
• arch/arm/mach-imx/pm-imx7.c • arch/arm/mach-imx/suspend-imx7.S • arch/arm/mach-imx/cpuidle-imx7d.c • arch/arm/mach-imx/imx7d_low_power_idle.S	Supports i.MX 7Dual power management operation
• arch/arm/mach-imx/pm-imx7ulp.c • arch/arm/mach-imx/suspend-imx7ulp.S • arch/arm/mach-imx/cpuidle-imx7.c	Supports i.MX 7ULP power management operation
• drivers/soc/pm-domain-imx8.h	Supports i.MX 8, 8X and 8M power domains
ARM Trusted firmware exists in imx-atf on code aurora forum/.	Supports i.MX 8, 8X, and 8M use arm trusted firmware for power management operation

2.5.1.4 Menu Configuration Options

In menu configuration enable the CONFIG_PM: CONFIG_PM builds support for power management. By default, this option is Y In menuconfig, this option is available under: Power management options > Power Management support.

In menu configuration enable the CONFIG_SUSPEND. CONFIG_SUSPEND builds support for suspend. In menuconfig, this option is available under: Power management options > Suspend to RAM and standby

2.5.1.5 Programming Interface

Look in the cpu_idle for each SoC as shown in the source code structure table and search for lpm. This will be the API for lower power mode. This implements all the steps required to put the system into WAIT and STOP modes.

2.5.2 PMIC PF Regulator

2.5.2.1 Introduction

PF100/200/300 is a PMIC chip.

PF200/PF3000 is based on PF100 with little change, since they share the same PF100 driver. PF100 regulator driver provides the low-level control of the power supply regulators, selection of voltage levels, and enabling/disabling of regulators. This device driver makes use of the PF100 regulator driver to access the PF100 hardware control registers. PF100 regulator driver is based on regulator core driver and it is attached to kernel I2C bus.

PF8100/8200 PMIC is designed for i.MX 8 and i.MX 8X families and is controlled by system controller firmware (SCFW) since it is a system-level device. SCFW creates some specific power resource for the Linux touch, such as "SC_R_BOARD_R0".

2.5.2.2 Hardware Operation

PMIC PF regulator provides reference and supply voltages for the application processor and peripheral devices.

Four buck (step down) converters (up to 6 independent output) and one boost (step up) converter are included. The buck converters provide the power supply to processor cores and to other low voltage circuits such as memory. Dynamic voltage scaling is provided to allow controlled supply rail adjustments for the processor cores and/or other circuitry.

Linear regulators are directly supplied from the battery or from the switchers and include supplies for I/O and peripherals, audio, camera, BT, WLAN, and so on. Naming conventions are suggestive of typical or possible use case applications, but the switchers and regulators may be utilized for other system power requirements within the guidelines of specified capabilities.

The only power on event of PF100 is PWRON is high, and the only power off event of PF100 is PWRON is low. PMIC_ON_REQ pin of i.MX 6, which is controlled by SNVS block of i.MX 6, will connect with PWRON pin of PF100 to control PF100 on/off, so that system can power off.

2.5.2.3 Software Operation

PMIC PF regulator client driver performs operations by reconfiguring the PMIC hardware control registers.

Some of the PMIC power management operations depend on the system design and configuration. For example, if the system is powered by a power source other than the PMIC, then turning off or adjusting the PMIC voltage regulators has no effect. Conversely, if the system is powered by the PMIC, then any changes that use the power management driver and the regulator client driver can affect the operation or stability of the entire system.

2.5.2.4 Driver Features

PMIC PF regulator driver is based on regulator core driver. It provides the following services for regulator control of the PMIC component:

- Switch ON/OFF all voltage regulators.
- Set the value for all voltage regulators.
- Get the current value for all voltage regulators.

2.5.2.5 Regulator APIs

The regulator power architecture is designed to provide a generic interface to voltage and current regulators within the Linux kernel.

It is intended to provide voltage and current control to client or consumer drivers and to provide status information to user space applications through a sysfs interface. The intention is to allow systems to dynamically control regulator output to save power and prolong battery life. This applies to both voltage regulators (where voltage output is controllable) and current sinks (where current output is controllable).

For more details, see opensource.wolfsonmicro.com/node/15

Under this framework, most power operations can be done by the following unified API calls:

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

- `regulator_get` is an unified API call to lookup and obtain a reference to a regulator:

  ```
  struct regulator *regulator_get(struct device *dev, const char *id);
  ```
- `regulator_put` is an unified API call to free the regulator source:

  ```
  void regulator_put(struct regulator *regulator, struct device *dev);
  ```
- `regulator_enable` is an unified API call to enable regulator output:

  ```
  int regulator_enable(struct regulator *regulator);
  ```
- `regulator_disable` is an unified API call to disable regulator output:

  ```
  int regulator_disable(struct regulator *regulator);
  ```
- `regulator_is_enabled` is the regulator output enabled:

  ```
  int regulator_is_enabled(struct regulator *regulator);
  ```
- `regulator_set_voltage` is an unified API call to set regulator output voltage:

  ```
  int regulator_set_voltage(struct regulator *regulator, int uV);
  ```
- `regulator_get_voltage` is an unified API call to get regulator output voltage:

  ```
  int regulator_get_voltage(struct regulator *regulator);
  ```

You can find more APIs and details in the regulator core source code inside the Linux kernel at:

```
drivers/regulator/core.c
```

2.5.2.6 Driver Architecture

The following figure shows the basic architecture of the PMIC PF regulator driver.

Device drivers

PF100 driver

```
┌─────────────────────────┐
│  Regulator core driver  │
└─────────────────────────┘
             │
             ▼
┌─────────────────────────┐
│  PF100 regulator driver │
└─────────────────────────┘
             │
             ▼
┌─────────────────────────┐
│    I2C or SPI driver    │
└─────────────────────────┘
```
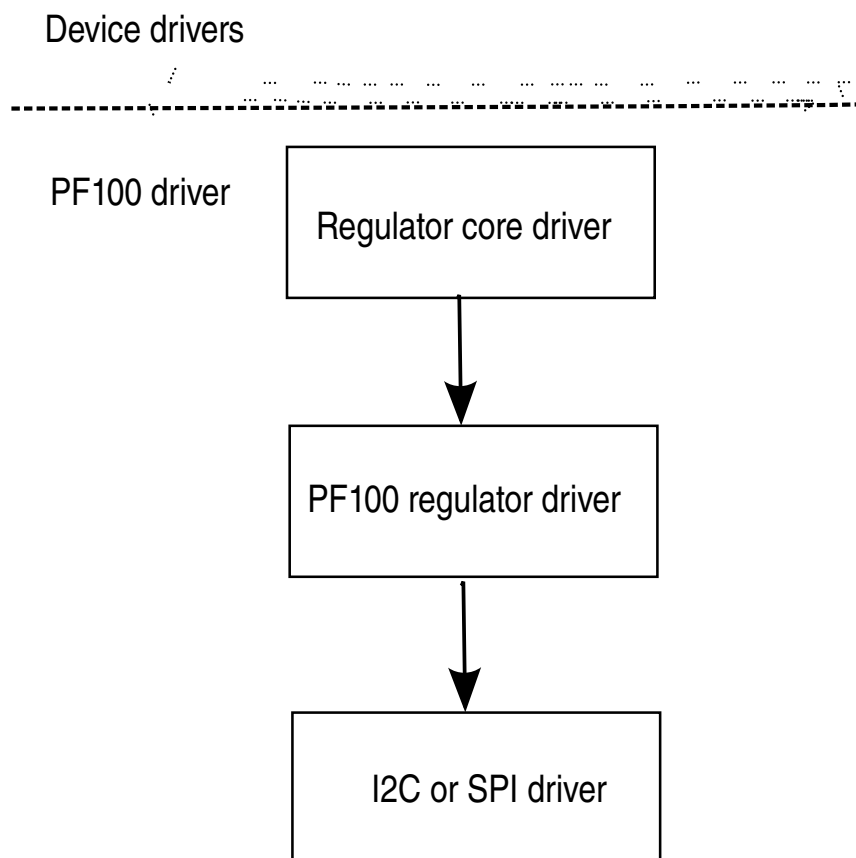
2.5.2.7 Driver Interface Details

Access to PFUZE100 regulator is provided through the API of the regulator core driver.

PFUZE100 regulator driver provides the following regulator controls:

- 4 buck switch regulators on normal mode (up to 6 independent rails): SW1AB, SW1C, SW2, SW3A, SW3B, and SW4.
- Buck switch can be programmed to a state of standby with specific register (PFUZE100_SWxSTANDBY) in advance.
- 6 Linear Regulators: VGEN1, VGEN2, VGEN3, VGEN4, VGEN5, and VGEN6.
- 1 LDO/Switch supply for VSNVS support on i.MX processors.
- 1 Low current, high accuracy, voltage reference for DDR Memory reference voltage.
- 1 Boost regulator with USB OTG support.
- Most power rails from PFUZE100 have been programmed properly according to the hardware design. Therefore, you can't find the kernel using PFUZE100 regulators. PFUZE100 regulator driver has implemented these regulators so that customers can use it freely if default PFUZE100 value can't meet their hardware design.

2.5.2.8 Source Code Structure

The PFUZE regulator driver is located in the drivers/regulator directory:

Table 2-12. PFUZE Driver Files

File	Description
drivers/regulator/pfuze100-regulator.c	Implementation of the PFUZE100 regulator client driver.
drivers/regulator/pf1550.c	Implementation of the PFUZE1550 regulator client driver.
drivers/regulator/pf1550-regulator-rpmsg.c	Implementation of the PFUZE150 regulator RPMSG code.

There is no board file related to PMIC. Some PFUZE driver code was moved to U-Boot, such as standby voltage setting. Some code is implemented by DTS file. Search for PFUZE100 in Uboot source and pfuze in device trees dtsi files in i.MX 6 and i.MX7 in arch/arm/boot/dts and for i.MX 8M in arch/arm64/boot/dts.

2.5.2.9 Menu Configuration Options

In menu configuration enable the following module:

Device Drivers > Voltage and Current regulator support > Freescale PFUZE100/200/3000 regulator driver.

2.5.3   CPU Frequency Scaling (CPUFREQ)

2.5.3.1   Introduction

The CPU frequency scaling device driver allows the clock speed of the CPU to be changed on the fly. Once the CPU frequency is changed, the voltage of the necessary power supplies are changed to the voltage value defined in device tree scripts (DTS). This method can reduce power consumption (thus saving battery power), because the CPU uses less power as the clock speed is reduced.

2.5.3.2   Software Operation

The CPUFREQ device driver is designed to change the CPU frequency and voltage on the fly.

If the frequency is not defined in DTS, the CPUFREQ driver changes the CPU frequency to the nearest higher frequency in the array. The frequencies are manipulated using the clock framework API, while the voltage is set using the regulators API. The CPU frequencies in the array are based on the boot CPU frequency. Interactive CPU frequency governor is used which cannot be changed manually. To change CPU frequency manually, the userspace CPU frequency governor can be used. By default, the conservative CPU frequency governor is used.

See the API document for more information on the functions implemented in the driver.

To view what values the CPU frequency can be changed to in KHz (the values in the first column are the frequency values), use this command:

```
cat /sys/devices/system/cpu/cpu0/cpufreq/stats/time_in_state
```

To change the CPU frequency to a value that is given by using the command above (for example, to 792 MHz) use this command:

```
echo 792000 > /sys/devices/system/cpu/cpu0/cpufreq/scaling_setspeed
```

The frequency 792000 is in KHz, which is 792 MHz.

The maximum frequency can be checked using this command:

```
cat /sys/devices/system/cpu/cpu0/cpufreq/scaling_max_freq
```

Use the following command to view the current CPU frequency in KHz:

```
cat /sys/devices/system/cpu/cpu0/cpufreq/cpuinfo_cur_freq
```

Use the following command to view available governors:

```
cat /sys/devices/system/cpu/cpu0/cpufreq/scaling_available_governors
```

Use the following command to change to interactive CPU frequency governor:

```
echo interactive > /sys/devices/system/cpu/cpu0/cpufreq/scaling_governor
```

2.5.3.3   Source Code Structure

Table below shows the source files and headers available in the following directory:

Table 2-13.   CPUFREQ Driver Files

File	Description
drivers/cpufreq/imx6q-cpufreq.c	i.MX 6 CPUFREQ functions
drivers/cpufreq/imx7-cpufreq.c	i.MX 7 CPUFREQ functions
drivers/cpufreq/imx7ulp-cpufreq.c	i.MX 7ULP CPUFREQ functions
drivers/cpufreq/imx8-cpufreq.c	i.MX 8 and i.MX 8X CPUFREQ functions
drivers/cpufreq/imx8mq-cpufreq.c	i.MX 8M CPUFREQ functions

For CPU frequency working point settings, see the SOC corresponding dtsi file in arch/arm/boot/dts for i.MX 6 and i.MX7 and arch/arm64/boot/dts for i.MX 8, i.MX 8X and i.MX 8M.

2.5.3.4   Menu Configuration Options

The following Linux kernel configuration is provided for this module:

- CONFIG_CPU_FREQ; In menuconfig, this option is located under:
    - CPU Power Management > CPU Frequency scaling
- The following options can be selected:

- CPU Frequency scaling
- CPU frequency translation statistics
- Default CPU frequency governor (conservative)(interactive)
- Performance governor
- Powersave governor
- Userspace governor for userspace frequency scaling
- Interactive CPU frequency policy governor
- Conservative CPU frequency governor
- Schedutil CPU frequency governor
- CPU frequency driver for i.MX CPUs

2.5.4 Dynamic Bus Frequency

2.5.4.1 Introduction

To improve power consumption, the Bus Frequency driver dynamically manages the various system frequencies for i.MX 6, i.MX 7, and i.MX 8M families.

The frequency changes are transparent to the higher layers and require no intervention from the drivers or middleware. Depending on activity of the peripheral devices and CPU loading, the bus frequency driver varies the DDR frequency between 24 MHz and its maximum frequency. Similarly, the AHB frequency is varied between 24 MHz and its maximum frequency.

2.5.4.2 Operation

The Bus Frequency driver is part of the power management module in the Linux BSP. The main purpose of this driver is to scale the various operating frequency of the system clocks (like AHB, DDR, AXI etc.) based on peripheral activity and CPU loading.

2.5.4.3 Software Operation

The bus frequency depends on the request and release of device drivers for its operation. Drivers will call bus frequency APIs to request or release the bus setpoint they want. The bus frequency will set the system frequency to highest frequency setpoint based on the peripherals that are currently requesting.

To enable the bus frequency driver, use the following command:

```
echo 1 > /sys/bus/platform/drivers/imx_busfreq/soc\:busfreq/enable
```

To disable the bus frequency driver, use the following command:

```
echo 0 > /sys/bus/platform/drivers/imx_busfreq/soc\:busfreq/enable
```

If Arm Cortex-M4 processor is alive with Arm Cortex-A processor together, Arm Cortex-M4 processor also requests or releases bus frequency high setpoint for its operation. This means that Arm Cortex-A processor treats Arm Cortex-M4 processor as one of its high-speed devices.

The setpoint modes do the following:

- High Frequency Setpoint mode is used when most peripherals that need higher frequency for good performance are active. For example, video playback and graphics processing.
- Audio Playback setpoints mode is used in audio playback mode.
- Low Frequency setpoint mode is used when the system is idle waiting for user input (display is off). For 8M, this mode is used when no peripheral device request high mode or audio mode.

The following table explains the software setpoints for each Family.

Table 2-14.   BusFrequency Set Points

SoC	Setpoints
MX6	- High Frequency Setpoint: AHB is at 132 MHz, AXI is at 264 MHz. - Audio Playback setpoints: On i.MX 6, AHB is at 25 MHz, AXI is at 50 MHz, and DDR is at 50 MHz for DDR3 and 100 MHz for LPDDR2.. - Low Frequency setpoint: AHB is at 24 MHz, AXI is at 24 MHz, and DDR is at 24 MHz.
MX 7Dual	- High Frequency Setpoint: AHB is at 135 MHz, AXI is at 332 MHz, and DDR is at the maximum frequency. - Audio Playback setpoints: AHB is at 24 MHz, AXI is at 24 MHz, and DDR is at 100 MHz. - Low Frequency setpoint: AHB is at 24 MHz, AXI is at 24 MHz, and DDR is at 24 MHz.
MX8M	- High bus frequency mode: The DDRC core clock is set to 800 MHz. The DDRC APB clock is set to 200 MHz. The NOC clock is set to 800 MHz. The main AXI cock is set to 333 MHz, and the AHB clock is set to 133 MHz. - Audio bus frequency mode: The DDRC core clock is set to 25 MHz, The DDRC APB clock is set to 20 MHz, the NOC clock is set to 100 MHz. Tthe main AXI clock is set to 25 MHz, and the AHB clock is set to 20 MHz. The DDR PLL is powered down for power saving. - Low bus frequency mode: The DDRC core clock is set to 25 MHz. The DDRC APB clock is set to 20 MHz. The NOC clock is set to 100 MHz. The main AXI clock is set to 25 MHz. The AHB clock is set to 20MHz. The DDR PLL is powered down for power saving.

2.5.4.4   Source Code Structure

The following table lists the source files and headers:

Table 2-15.   BusFrequency Driver Files

File	Description
arch/arm/mach-imx/busfreq-imx.c	i.MX 6 and i.MX 7 Bus Frequency functions
include/linux/busfreq-imx.h	i.MX Bus Frequency API Definitions
arch/arm/mach-imx/busfreq_ddr3.c	i.MX 6 and i.MX 7 DDR3 Bus Frequency functions
arch/arm/mach-imx/busfreq_lpddr2.c	i.MX 6 and i.MX 7 LPDDR2 Bus Frequency functions
arch/arm/mach-imx/lpddr2_freq_imx6.S	i.MX 6 LPDDR2 Bus Frequency functions
arch/arm/mach-imx/lpddr2_freq_imx6q.S	i.MX 6 QuadPlus/Quad/Dual/Solo/SoloLite LPDDR2 Bus Frequency functions
arch/arm/mach-imx/lpddr2_freq_imx6sll.S	i.MX 6 SLL LPDDR2 Bus Frequency functions
arch/arm/mach-imx/lpddr2_freq_imx6sx.S	i.MX 6 SoloX LPDDR2 Bus Frequency functions
arch/arm/mach-imx/lpddr3_freq_imx.S	i.MX 6 and i.MX 7 LPDDR3 Bus Frequency functions
arch/arm/mach-imx/ddr3_freq_imx6.S	i.MX 6 Bus Frequency functions
arch/arm/mach-imx/ddr3_freq_imx6sx.S	i.MX 6 SoloX Bus Frequency functions
arch/arm/mach-imx/ddr3_freq_imx7d.S	i.MX 7 Dual DDR3 Bus Frequency functions
drivers/doc/imx/busfreq-imx8mq.c	i.MX 8M Bus Frequency functions

Bus frequency modes are defined in the SoC dtsi files in arch/arm/boot/dts for i.MX 6 and i.MX 7 and arch/arm64/boot/dts for i.MX 8M.

2.5.4.5   Menu Configuration Options

There are no menu configuration options for this driver. The Bus Frequency drivers are included and enabled by default for the SoC that support bus frequency drivers.

2.5.5   Battery Charging

2.5.5.1   Introduction

Battery Charging is supported by the max8903-charger for the i.MX 6 SABRE SD boards.

2.5.5.2   Software Operation

2.5.5.3 Source Code Structure

The battery charging source is based in drivers/power/supply/sabresd_battery.c

2.5.5.4 Menu Configuration Options

In menu configuration enable the following module:

Device Drivers > Power supply class support > Sabresd Board Battery DC-DC Charger for USB and Adapter Power.

2.6 OProfile

2.6.1 Introduction

OProfile is a system-wide profiler capable of profiling all running code at low overhead.

OProfile consists of a kernel driver, a daemon for collecting sample data, and several post-profiling tools for turning data into information.

2.6.1.1 Overview

OProfile leverages the hardware performance counters of the CPU to enable profiling of a wide variety of interesting statistics, which can also be used for basic time-spent profiling.

All code is profiled: hardware and software interrupt handlers, kernel modules, the kernel, shared libraries, and applications.

2.6.1.2 Features

OProfile has the following features.

- Unobtrusive-No special recompilations or wrapper libraries are necessary. Even debug symbols (-g option to gcc) are not necessary unless users want to produce annotated source. No kernel patch is needed; just insert the module.

- System-wide profiling-All code running on the system is profiled, enabling analysis of system performance.
- Performance counter support-Enables collection of various low-level data and association for particular sections of code.
- Call-graph support-OProfile can provide gprof-style call-graph profiling data.
- Low overhead-OProfile has a typical overhead of 1-8% depending on the sampling frequency and workload.
- Post-profile analysis-Profile data can be produced on the function-level or instruction-level detail. Source trees, annotated with profile information, can be created. A hit list of applications and functions that utilize the most CPU time across the whole system can be produced.
- System support-Works with any i.MX supported kernel.

2.6.1.3 Hardware Operation

OProfile is a statistical continuous profiler.

Profiles are generated by regularly sampling the current registers on each CPU (from an interrupt handler, the saved PC value at the time of interrupt is stored), and converting that runtime PC value into something meaningful to the programmer.

OProfile achieves this by taking the stream of sampled PC values, along with the detail of which task was running at the time of the interrupt, and converting the values into a file offset against a particular binary file. Each PC value is thus converted into a tuple (group or set) of binary-image offset. The userspace tools can use this data to reconstruct where the code came from, including the particular assembly instructions, symbol, and source line (through the binary debug information if present).

Regularly sampling the PC value like this approximates what actually was executed and how often and, more often than not, this statistical approximation is good enough to reflect reality. In common operation, the time between each sample interrupt is regulated by a fixed number of clock cycles. This implies that the results reflect where the CPU is spending the most time. This is a very useful information source for performance analysis.

The ARM CPU provides hardware performance counters capable of measuring these events at the hardware level. Typically, these counters increment once per each event and generate an interrupt on reaching some pre-defined number of events. OProfile can use these interrupts to generate samples and the profile results are a statistical approximation of which code caused how many instances of the given event.

2.6.1.4   Architecture-specific Components

OProfile supports the hardware performance counters available on a particular architecture. Code for managing the details of setting up and managing these counters can be located in the kernel source tree in the relevant arch/arm/oprofile directory. The architecture-specific implementation operates through filling in the oprofile_operations structure at initialization. This provides a set of operations, such as setup(), start(), stop(), and so on, that manage the hardware-specific details the performance counter registers.

The other important facility available to the architecture code is oprofile_add_sample(). This is where a particular sample taken at interrupt time is fed into the generic OProfile driver code.

2.6.1.5   oprofilefs Pseudo Filesystem

OProfile implements a pseudo-filesystem known as oprofilefs, which is mounted from userspace at /dev/oprofile. This consists of small files for reporting and receiving configuration from userspace, as well as the actual character device that the OProfile userspace receives samples from. At setup() time, the architecture-specific code may add further configuration files related to the details of the performance counters. The filesystem also contains a stats directory with a number of useful counters for various OProfile events.

2.6.1.6   Generic Kernel Driver

The generic kernel driver resides in drivers/oprofile, and forms the core of how OProfile operates in the kernel. The generic kernel driver takes samples delivered from the architecture-specific code (through oprofile_add_sample()), and buffers this data (in a transformed configuration) until releasing the data to the userspace daemon through the /dev/oprofile/buffer character device.

2.6.1.7   OProfile Daemon

The OProfile userspace daemon takes the raw data provided by the kernel and writes it to the disk. It takes the single data stream from the kernel and logs sample data against a number of sample files (available in /var/lib/oprofile/samples/current/). For the benefit of the separate functionality, the names and paths of these sample files are changed to reflect where the samples were from. This can include thread IDs, the binary file path, the event type used, and more.

After this final step from interrupt to disk file, the data is now persistent (that is, changes in the running of the system do not invalidate stored data). This enables the post-profiling tools to run on this data at any time (assuming the original binary files are still available and unchanged).

2.6.1.8 Post Profiling Tools

The collected data must be presented to the user in a useful form. This is the job of the post-profiling tools. In general, they collate a subset of the available sample files, load and process each one correlated against the relevant binary file, and produce user readable information.

2.6.1.9 Interrupt Requirements

The number of interrupts generated with respect to the OProfile driver are numerous. The latency requirements are not needed.

The rate at which interrupts are generated depends on the event.

2.7 Pulse-Width Modulator (PWM)

2.7.1 Introduction

The pulse-width modulator (PWM) has a 16-bit counter and is optimized to generate sound from stored sample audio images and generate tones. The PWM also provides control for the back light.

The PWM has 16-bit resolution and uses a 4x16 data FIFO to generate sound. The software module is composed of a Linux driver that allows privileged users to control the backlight by the appropriate duty cycle of the PWM Output (PWMO) signal.

2.7.2 Hardware Operation

The figure below shows the PWM block diagram.
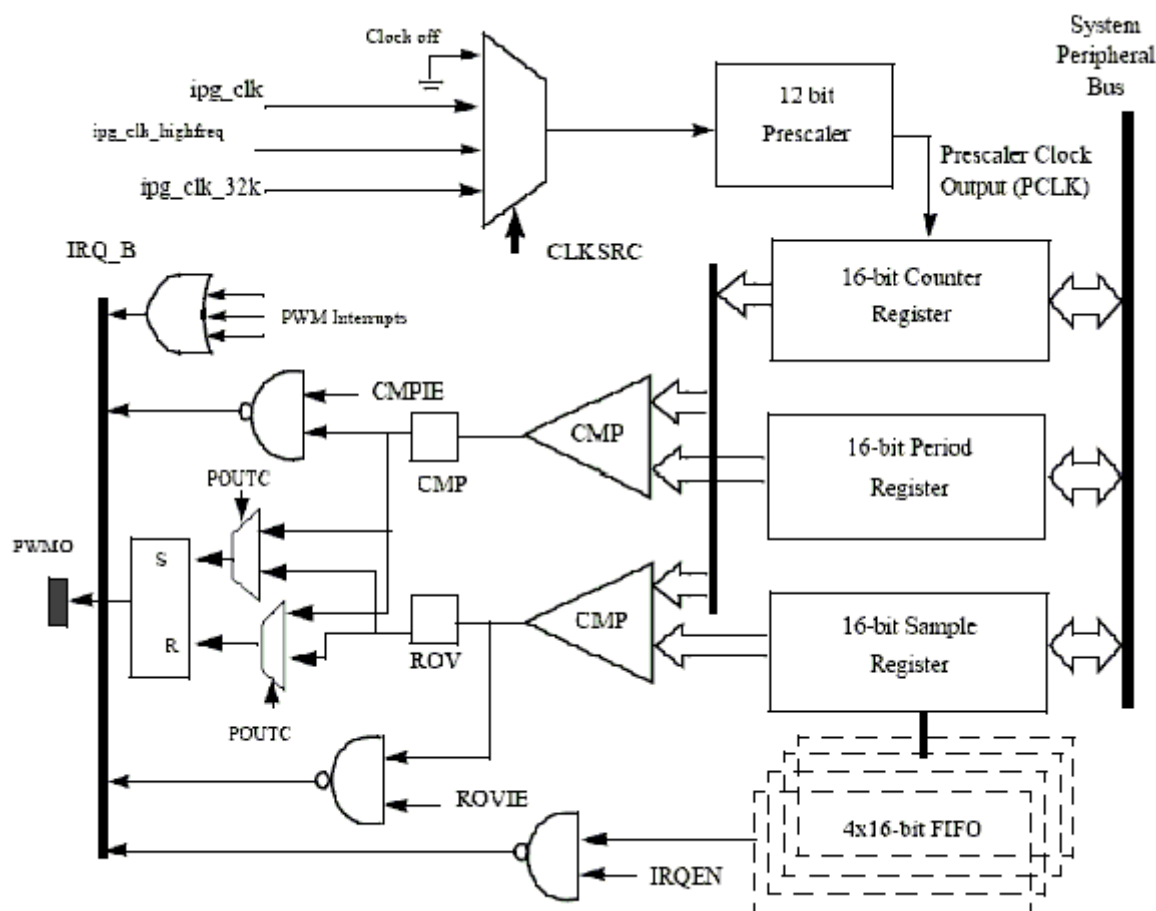
Figure 2-2. PWM Block Diagram

The PWM follows IP Bus protocol for interfacing with the processor core. It does not interface with any other modules inside the device except for the clock and reset inputs from the Clock Control Module (CCM) and interrupt signals to the processor interrupt handler. The PWM includes a single external output signal, PMWO. The PWM includes the following internal signals:

- Three clock inputs
- Four interrupt lines
- One hardware reset line
- Four low power and debug mode signals
- Four scan signals
- Standard IP slave bus signals

2.7.3 Clocks

The clock that feeds the prescaler can be selected from:

- High frequency clock-provided by the CCM. The PWM can be run on this clock in low power mode.
- Low reference clock - 32 KHz low reference clock provided by the CCM. The PWM can be run on this clock in the low power mode.
- Global functional clock - for normal operations. In low power modes this clock can be switched off.

The clock input source is determined by the CLKSRC field of the PWM control register. The CLKSRC value should only be changed when the PWM is disabled.

2.7.4   Software Operation

The PWM device driver reduces the amount of power sent to a load by varying the width of a series of pulses to the power source. One common and effective use of the PWM is controlling the backlight of a QVGA panel with a variable duty cycle.

Table below provides a summary of the interface functions in source code.

Table 2-16.   PWM Driver Summary

Function	Description
struct pwm_device *pwm_get(struct device *dev, const char *con_id)	Look up and request a PWM device
void pwm_put(struct pwm_device *pwm)	Release a PWM device
int pwm_config(struct pwm_device *pwm, int duty_ns, int period_ns)	Change a PWM device configuration
int pwm_enable(struct pwm_device *pwm)	Start a PWM output toggling
int pwm_disable(struct pwm_device *pwm)	Stop a PWM output toggling

The function pwm_config() includes most of the configuration tasks for the PWM module, including the clock source option, period and duty cycle of the PWM output signal. It is recommended to select the peripheral clock of the PWM module, rather than the local functional clock, as the local functional clock can change.

2.7.5   Driver Features

The PWM driver includes the following software and hardware support:

- Duty cycle modulation
- Varying output intervals
- Two power management modes - full on and full off

2.7.6 Source Code Structure

Table 2-17. PWM Driver Files

File	Description
drivers/pwm/pwm.h	Functions declaration
drivers/pwm/pwm-imx.c	i.MX Pulse Width modulation Functions

2.7.7 Menu Configuration Options

In menu configuration enable the following module:

- Device Drivers > Pulse-Width Modulation (PWM) Support > i.MX PWM support
- Select the following option to enable the Backlight driver:

  Device Drivers > Graphics support > Backlight & LCD device support > Generic
  PWM based Backlight Driver

2.8 Remote Processor Messaging

2.8.1 Introduction

With the newest multicore architecture designed by using the Arm Cortex®-A series
processors and the ArmCortex-M series processors, industrial applications can achieve
greater power efficiency for a reduced carbon footprint. This reduces power consumption
without performance deterioration.

A homogeneous SoC would traditionally run a single operating system (OS) that controls
all the memory. The OS or a hypervisor would handle task management among available
cores to maximize system utilization. Such a system is called Symmetric MultiProcessing
(SMP).

A heterogeneous multicore chip where different processing cores running different
instruction sets and different OSs. Each processing core handles a specific task as
required. Such a system is called Asymmetric Multiprocessing (AMP). To understand the
distinction between the SMP and AMP systems, it is possible for a homogeneous
multicore SoC to be an AMP system but a heterogeneous multicore SoC cannot be an
SMP system.

A multicore architecture brings new challenges to the system design, because the software must be rewritten to distribute tasks across the available cores. In addition, all the peripheral resources need to be properly allocated to avoid resource contention and achieve efficient sharing of the data spaces between the cores. A multicore SoC also needs mechanisms for reliable communication and synchronization among tasks running on different processing cores.

RPMsg is a virtio-based messaging bus, which allows kernel drivers to communicate with remote processors available on the system. In turn, drivers could then expose appropriate user space interfaces if needed. Every RPMsg device is a communication channel with a remote processor (so the RPMsg devices are called channels). Channels are identified by a textual name and have a local ("source") RPMsg address, and remote ("destination") RPMsg address. For more information, see www.kernel.org/doc/ Documentation/rpmsg.txt.

As shown in the following figure, the messages pass between endpoints through bidirectional connection-less communication channels.
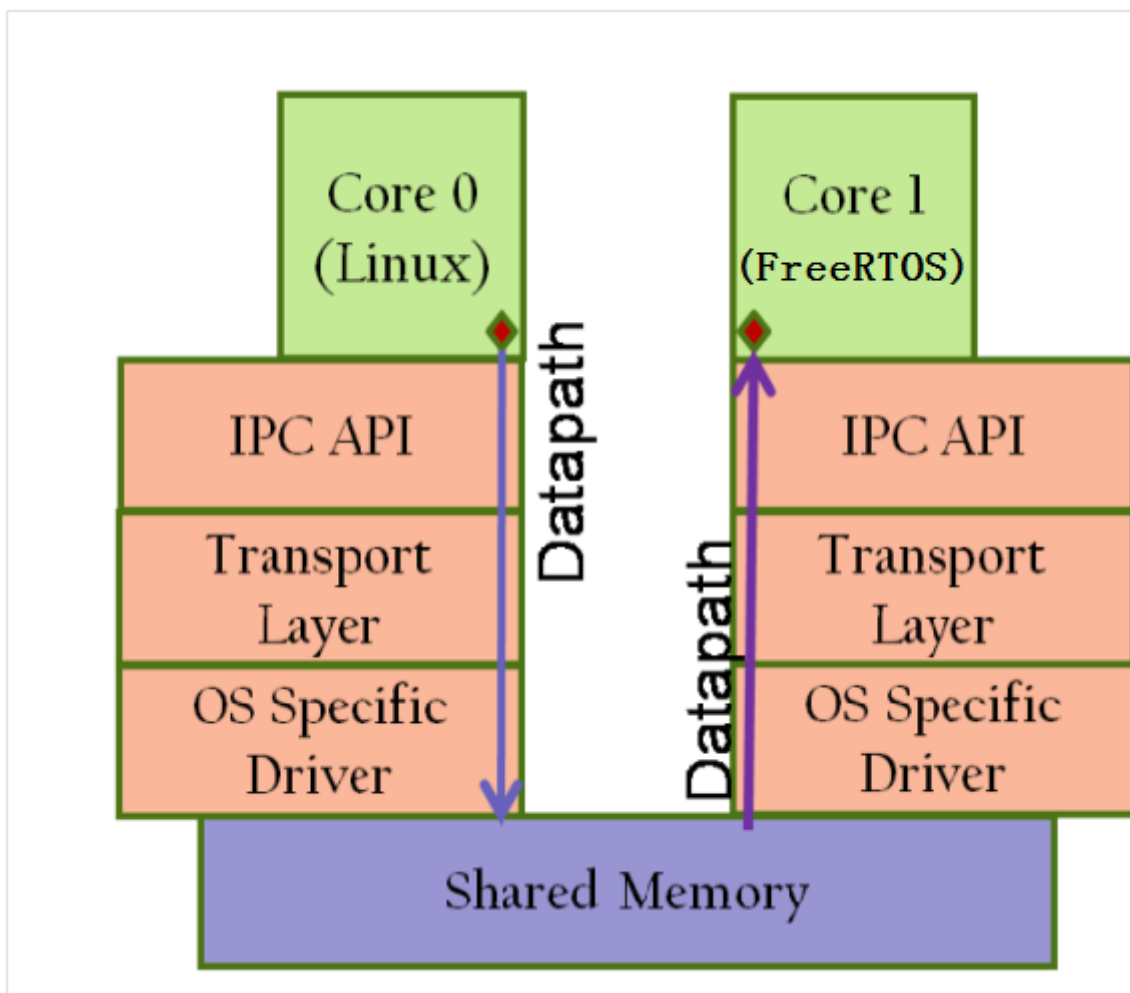
Figure 2-3. New multicore, multiOS architecture

2.8.2 Features

- Designed for low-latency and low overhead operation, and compliant with the Linux RPMsg framework.
- Optimized for embedded environments with constrained CPU and memory resources.
- Implementation by using shared memory without data translation or message headers.
- Application communication by using a client-server methodology.
- Dynamic allocation of the RPMsg channels.

2.8.3 Source Code

RPMSG driver software is in drivers/rpmsg

Table 2-18. RPMSG Source

File	Description
drivers/rpmsg/virtio_rpmsg_bus.c	Common code
drivers/rpmsg/imx_rpmsg.c	i.MX platform-related code
drivers/rpmsg/imx_rpmsg_pingpong.c	i.MX RPMsg ping-pong tests
drivers/rpmsg/imx_rpmsg_tty.c	i.MX RPMsg TTY driver

2.8.4 Menu Configuration Options

In menu configuration enable the following module:

- Device Drivers > IMX RPMSG pingpong driver -- loadable modules only
- Device Drivers > IMX RPMSG tty driver -- loadable modules only

2.8.5 Running i.MX RPMsg Test Programs

To run the i.MX RPMsg test program, perform the following operations:
1. Make sure that the proper Cortex-M4 processor RTOS and Linux images are used.
   For example on the i.MX 7Dual platforms:
   - rpmsg_pingpong_sdk_7dsdb.bin -> ping-pong test used on the i.MX 7Dual SDB board
   - rpmsg_str_echo_sdk_7dsdb.bin -> tty string echo test used on the i.MX 7Dual SDB board
   - rpmsg_pingpong_sdk_7dval.bin -> ping-pong test used on the i.MX 7Dual 12x12 LPDDR3 Arm2 board
   - rpmsg_str_echo_sdk_7dval.bin -> tty string echo test used on the i.MX 7Dual 12x12 LPDDR3 Arm2 board
2. Load the Cortex-M4 processor RTOS image, and kick it off in U-Boot.
   Load the Cortex-M4 processor RTOS image by the TFTP server or by the bootable SD card in U-Boot.
   - Load the Cortex-M4 processor RTOS image by the TFTP server. For example,
     1. Boot into U-Boot and stop.
     2. Use the following command to TFTP the responding Cortex-M4 processor RTOS image and boot it.

```
dhcp 0x7e0000 10.192.242.53:rpmsg_pingpong_sdk_7dval.bin;  bootaux 0x7e0000
```

- Load the Cortex-M4 processor RTOS image by the SD card. For example,
    1. Created A bootable SD card by the MFGtools. Then, copy the Cortex-M4 processor RTOS files to the first partition formatted by the VFAT file system.
    2. Change the default Cortex-M4 processor RTOS name of the U-Boot.

    ```
    setenv m4image '<The name of the M4/RTOS image>';save
    ```
    3. Set up a boot args used by the Cortex-M4 processor.

    ```
    setenv run_m4_tcm 'if run loadm4image; then cp.b ${loadaddr} 0x7e0000 0x8000;
    bootaux 0x7e0000; fi'; save
    ```
    4. Modify the original bootcmd by adding `run run_m4_tcm`".

    ```
    setenv bootcmd "run run_m4_tcm; <original contents of the bootcmd>"; save
    ```

    ### NOTE
    "uart_from_osc" is mandatory required by i.MX 6SoloX when the Cortex-M4 processor RTOS image is running. Therefore, the mmcargs of U-Boot should be modified on i.MX 6SoloX.

    ```
    setenv mmcargs 'setenv bootargs console=${console},$
    {baudrate} root=${mmcroot}, uart_from_osc';save
    ```

3. Run the RPMsg test program.
    a. Make sure that `imx_rpmsg_pingpong.ko` and `imx_rpmsg_tty.ko` are built out.
    b. Use `insmod imx_rpmsg_pingpong.ko` or `insmod imx_rpmsg_tty.ko` to run the test program.

    ### NOTE
    Do not run different test programs at the same time.

    c. Run the following command and ensure that the RPMsg TTY receiving program is running at backend when starting RPMsg TTY tests.

    ```
    /unit_tests/mxc_mcc_tty_test.out /dev/ttyRPMSG30 115200 R 100 1000 &
    ```

    Logs at the Linux OS side:

    ```
    insmod imx_rpmsg_tty.ko
    imx_rpmsg_tty rpmsg0: new channel: 0x400 -> 0x1!
    Install rpmsg tty driver!
    echo deadbeaf > /dev/ttyRPMSG30
    imx_rpmsg_tty rpmsg0: msg(<- src 0x1) deadbeaf len 8
    ```

2.9 Thermal

2.9.1 Introduction

Thermal driver is a necessary driver for monitoring and protecting the SoC. The thermal driver monitors the SoC temperature in a certain frequency from an internal thermal sensor.

It defines two trip points: critical and passive. Cooling device will take actions to protect the SoC according to the different trip points that SoC has reached:

- When reaching critical point, cooling device will shut down the system.
- When reaching passive point, cooling device will lower CPU frequency and notify GPU/VPU to run at a lower frequency.
- When the temperature drops to 10 °C below passive point, cooling device will release all the cooling actions.

Thermal driver has two parts:

- Thermal zone defines trip points and monitors the SoC's temperature.
- Cooling device takes the actions according to the different trip points.

The critical and passive points threshold are confiugured in the following files.

- i.MX 6 and i.MX 7 SoCs configure this in drivers/thermal/imx_thermal.c
- i.MX 8M SoCs configure this in their dtsi file and specify CONFIG_IMX8M_THERMAL in defconfig.
- i.MX 8 and i.MX 8X SoCsconfigure this in their dtsi file and specify CONFIG_IMX_SC_THERMAL in defconfig.

2.9.2 Software Operation

The thermal driver registers a thermal zone and a cooling device. A structure,`thermal_zone_device_ops`, describes the necessary interface that the thermal framework needs. The framework will call the related thermal zone interface to monitor the SoC temperature and do the cooling protection.

The thermal driver can be accessed through the following interface:

- /sys/bus/platform/drivers/imx_thermal for i.MX 6 and i.MX 7.
- /sys/bus/platform/drivers/imx_sc_thermal for i.MX 8 and i.MX 8X.
- /sys/bus/platform/drivers/qoriq_thermal for i.MX 8M Quad.
- /sys/class/thermal/thermal_zone0/temp for i.MX 8M Mini.

2.9.3 Source Code Structure

Table below shows the driver source files available in drivers/thermal:

Table 2-19. Thermal Driver Files

File	Description
imx_thermal.c, device_cooling.c	Thermal zone driver source file for i.MX 6 or i.MX 7.
qoriq_thermal.c, device_cooling.c	Thermal zone driver source files for i.MX 8M.
imx_sc_thermal.c, device_cooling.c	Thermal zone driver source files for i.MX 8 and i.MX 8X.

2.9.4 Menu Configuration Options

In menu configuration enable the following module:

- For i.MX6 and i.MX7: Device Drivers > Generic Thermal sysfs driver > Temperature sensor driver for i.MX SoCs.

- For i.MX 8QuadMax and i.MX 8QuadXPlus: Device Drivers > Generic Thermal sysfs driver > thermal sensor driver for NXP i.MX8 SoCs

2.10 Sensors

2.10.1 Introduction

Sensors include a group of drivers for Accelerometer, Pressure, Gyroscope, Ambient Light, and Magnetometer.

Sensors are configured in the device trees for each board.

i.MX supports accelerometers for the following SoC:

- i.MX 6SABRE-SD, 6SABRE-AI, and 6SoloX use the MMA8451 sensor
- i.MX 6 SoloLite uses the MMA8450 sensor.
- i.MX 6UltraLite and 6ULL EVK use the FXLS8571Q sensor.
- i.MX 7Dual SABRE-SD and i.MX 8QuadMax and i.MX 8QuadXPlus use the FX0S8700 sensor.

i.MX Supports pressure sensor MPL3115 for the following SoC:

- i.MX 7 Dual SABRE-SD

- i.MX 8 QuadMax
- i.MX 8 QuadXPlus

i.MX Supports gyroscope sensor FXAS2100 for the following SoC:

- i.MX 7 Dual SABRE-SD

i.MX Supports ambient light sensor ISL29023 for the following SoC:

- i.MX 6 SABRE-SD and 6 SABRE-AI
- i.MX 6 SoloX
- i.MX 8 QuadMax
- i.MX 8 QuadXPlus.

i.MX supports magnetometer sensors MAG3110 for the following SoC:

- i.MX 6 SABRE-SD
- i.MX 6, SoloLite,
- i.MX 6 UL EVK
- i.MX 6 ULL EVK
- i.MX 6SoloX

2.10.2 Sensor Driver Software Operation

2.10.3 Source Code Structure

Table below shows the driver source files available in the directory:

Table 2-20. Sensor Driver Files

File	Description
drivers/hwmon/mxc_mma8451.c	Acceleromater Sensor
drivers/misc/fxos8700.c	Acceleromater and Magnetometer Sensor
drivers/misc/fxls8471.c	Acceleromater and MagnetometerLight Sensor
drivers/input/misc/isl29023.c	Ambient Light Sensor
drivers/input/misc/fxas2100x.c	MagnetometerLight Sensor
drivers/hwmon/mag3110.c	Magnetometer Sensor

2.10.4 Menu Configuration Options

In menu configuration enable the following module:

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

- Device Drivers > Misc devices > Intersil ISL29020 ambient light sensor
- Device Drivers > Misc devices > Freescale FXOS8700 M+G combo sensor
- Device Drivers > Misc devices > Freescale FXAS2100X gyroscope sensor
- Device Drivers > Hardware Monitoring support > Freescale MAG3110 e-compass sensor
- Device Drivers > Hardware Monitoring support > MMA8451 device driver

2.11 Watchdog (WDOG)

2.11.1 Introduction

The Watchdog Timer module protects against system failures by providing an escape from unexpected hang or infinite loop situations or programming errors.

Some platforms may have two WDOG modules with one of them having interrupt capability. i.MX 6 and 7Dual share the same watch dog driver with i.MX 8M. i.MX 7ULP has a separate watchdog driver. i.MX 8 and i.MX 8X share a virtual watchdog driver interface through system controller firmware.

2.11.2 Hardware Operation

After the WDOG timer is activated, it must be serviced by software on a periodic basis.

If servicing does not take place in time, the WDOG times out. Upon a time-out, the WDOG either asserts the wdog_b signal or a wdog_rst_b system reset signal, depending on software configuration. The watchdog module cannot be deactivated after it is activated.

2.11.3 Software Operation

The Linux OS has a standard WDOG interface that allows support of a WDOG driver for a specific platform.

WDOG can be suspended/resumed in STOP/DOZE and WAIT modes independently. Since some bits of the WDOG registers are only one-time programmable after booting, ensure these registers are written correctly.

2.11.4   Generic WDOG

The generic WGOD driver is implemented in the drivers/watchdog/imx2_wdt.c file.

It provides functions for various IOCTLs and read/write calls from the user level program to control the WDOG.

2.11.5   Driver Features

This WDOG implementation includes the following features:

- Generates the reset signal if it is enabled but not serviced within a predefined timeout value (defined in milliseconds in one of the WDOG source files)
- Does not generate the reset signal if it is serviced within a predefined timeout value
- Provides IOCTL/read/write required by the standard WDOG subsystem

2.11.6   Source Code Structure

Table below shows the source files for watchdog WDOG drivers that are in drivers/watchdog.

Table 2-21.   Watchdog Driver Files

File	Description
driveers/watchdog/imx2_wdt.c	i.MX 6, i.MX 7Dual and i.MX 8M watchdog function implementations. For i.MX 6 and i.MX 7, the watchdog system reset function is located under arch/arm/mach-imx/system.c.
drivers/watchdog/imx7ulp_wdt.c	i.MX 7ULP watchdog function implementations
drivers/watchdog/imx8_wdt.c	On i.MX 8 and i.MX 8X, the software watchdog used in system controller firmware (SCFW) and kernel call those interfaces by virtual watchdog driver imx8_wdt.c. This is not used for i.MX 8M.

2.11.7   Menu Configuration Options

In menu configuration enable the following module:

Device Drivers > Watchdog Timer Support > IMX2+ Watchdog

Device Drivers > Watchdog Timer Support > IMX7ULP Watchdog

Device Drivers > Watchdog Timer Support > IMX8 Watchdog

2.11.8 Programming Interface

The following IOCTLs are supported in the WDOG driver:

- WDIOC_GETSUPPORT
- WDIOC_GETSTATUS
- WDIOC_GETBOOTSTATUS
- WDIOC_KEEPALIVE
- WDIOC_SETTIMEOUT
- WDIOC_GETTIMEOUT

For detailed descriptions about these IOCTLs, see `Documentation/watchdog`.

Chapter 3
Storage

3.1 AHB-to-APBH Bridge with DMA (APBH-Bridge-DMA)

3.1.1 Overview

The AHB-to-APBH bridge provides the processor with an inexpensive peripheral attachment bus running on the AHB's HCLK. The H in APBH denotes that the APBH is synchronous to HCLK.

The AHB-to-APBH bridge includes the AHB-to-APB PIO bridge for a memory-mapped I/O to the APB devices, as well as a central DMA facility for devices on this bus and a vectored interrupt controller for the Arm core. Each one of the APB peripherals, including the vectored interrupt controller, is documented in their own chapters elsewhere in this document.

There is no separate DMA bus for these devices. Contention between the DMA's use of the APBH bus and the AHB-to-APB bridge functions' use of the APBH is mediated by an internal arbitration logic. For contention between these two units, the DMA is favored and the AHB slave will report "not ready" through its HREADY output until the bridge transfer can complete. The arbiter tracks repeated lockouts and inverts the priority, guaranteeing the Arm platform every fourth transfer on the APB.

3.1.1.1 Hardware Operation

The SDMA controller is responsible for transferring data between the MCU memory space and peripherals and includes the following features.

- Multichannel DMA supporting up to 32 time-division multiplexed DMA channels
- Powered by a 16-bit Instruction-Set micro-RISC engine
- Each channel executes a specific script
- Very fast context-switching with two-level priority based preemptive multitasking

- 4 Kbytes ROM containing startup scripts (that is, boot code) and other common utilities that can be referenced by RAM-located scripts
- 8 Kbyte RAM area is divided into a processor context area and a code space area used to store channel scripts that are downloaded from the system memory.

3.1.1.2 Software Operation

The DMA supports sixteen channels of DMA services, as shown in the following table. The shared DMA resource allows each independent channel to follow a simple chained command list. Command chains are built up using the general structure.

Table 3-1. APBH DMA Channel Assignments

APBH DMA CHANNEL #	USAGE
0	GPMI0
1	GPMI1
2	GPMI2
3	GPMI3
4	GPMI4
5	GPMI5
6	GPMI6
7	GPMI7
8	EMPTY
9	EMPTY
10	EMPTY
11	EMPTY
12	EMPTY
13	EMPTY
14	EMPTY
15	EMPTY

3.1.1.3 Source Code Structure

The table below shows the source files available in drivers/dma/

Table 3-2. APBH DMA Source Files

File	Description
mxs-dma.c	APBH DMA implement driver

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

3.1.1.4 Menu Configuration Options

In menu configuration enable the following module:

- Device Drivers > DMA Engine support > MXS DMA support.

3.1.1.5 Programming Interface

The module implements standard DMA API. See the API documents, which are located in the Linux documentation package, for more information on the functions implemented in the driver such as GPMI NAND driver.

3.2 EIM NOR

3.2.1 Introduction

The External Interface Module (EIM) NOR driver supports the Parallel NOR flash.

3.2.2 Hardware Operation

By default, there is a parallel NOR in the i.MX 6Quad/6Dual SABRE-AI boards. The parallel NOR has more pins than the SPI NOR. On some boards, the M29W256GL7AN6E is equipped. Refer to the datasheet for details on the parallel NOR.

3.2.3 Software Operation

Similar to the SPI NOR, the parallel NOR uses the MTD subsystem. Because the parallel NOR is very small, you may only use the jffs2 but cannot use the UBIFS for it.

3.2.4 Source Code

Table 3-3. WEIM-NOR Driver Files

File	Description
drivers/bus/imx-weim.c	Timing only changes for Parallel NOR WEIM-NOR source

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

3.2.5   Enabling the EIM NOR

Refer to the DTS file to enable the EIM NOR: imx6q-sabreauto-gpmi-weim.dts or imx6dl-sabreauto-gpmi-weim.dts.

3.3   MMC/SD/SDIO Host

3.3.1   Introduction

The MultiMediaCard (MMC)/ Secure Digital (SD)/ Secure Digital Input Output (SDIO) Host driver implements a standard Linux driver interface to the ultra MMC/SD host controller (uSDHC).

The host driver is part of the Linux kernel MMC framework.

The MMC driver has the following features:

* 1-bit or 4-bit operation for SD3.0 and SDIO 2.0 cards (so far we support SDIO v2.0 (AR6003 is verified)).
* Supports card insertion and removal detections.
* Supports the standard MMC commands.
* PIO and DMA data transfers.
* Supports power management.
* Supports 1/4 8-bit operations for MMC cards.
* For i.MX 6, USDHC supports eMMC4.4 SDR and DDR modes.
* For i.MX 7Dual, USDHC supports eMMC5.0, which includes HS400 and HS200.
* Supports SD3.0 SDR50 and SDR104 modes.

3.3.2   Hardware Operation

The MMC communication is based on an advanced 11-pin serial bus designed to operate in a low voltage range. The uSDHC module supports MMC along with SD memory and I/O functions. The uSDHC controls the MMC, SD memory, and I/O cards by sending commands to cards and performing data accesses to and from the cards. The SD memory card system defines two alternative communication protocols: SD and SPI. The uSDHC only supports the SD bus protocol.

The uSDHC command transfer type and uSDHC command argument registers allow a command to be issued to the card. The uSDHC command, system control, and protocol control registers allow the users to specify the format of the data and response and to control the read wait cycle.

There are four 32-bit registers used to store the response from the card in the uSDHC. The uSDHC reads these four registers to get the command response directly. The uSDHC uses a fully configurable 128x32-bit FIFO for read and write. The buffer is used as temporary storage for data being transferred between the host system and the card, and vice versa. The uSDHC data buffer access register bits hold 32-bit data upon a read or write transfer.

For receiving data, the steps are as follows:

1. The uSDHC controller generates a DMA request when there are more words received in the buffer than the amount set in the RD_WML register
2. Upon receiving this request, DMA engine starts transferring data from the uSDHC FIFO to system memory by reading the data buffer access register.

For transmitting data, the steps are as follows:

1. The uSDHC controller generates a DMA request whenever the amount of the buffer space exceeds the value set in the WR_WML register.
2. Upon receiving this request, the DMA engine starts moving data from the system memory to the uSDHC FIFO by writing to the Data Buffer Access Register for a number of pre-defined bytes.

The read-only uSDHC Present State and Interrupt Status Registers provide uSDHC operations status, application FIFO status, error conditions, and interrupt status.

When certain events occur, the module has the ability to generate interrupts as well as set the corresponding Status Register bits. The uSDHC interrupt status enable and signal-enable registers allow the user to control if these interrupts occur.

3.3.3 Driver Features

The MMC driver supports the following features:

- Supports multiple uSDHC modules.
- Provides all the entry points to interface with the Linux MMC core driver.
- MMC and SD cards.
- SDIO cards.
- SD3.0 cards.
- Recognizes data transfer errors such as command time outs and CRC errors.

- Power management.
- It supports to be built as loadable or builtin module

3.3.4 Source Code Structure

Table below shows the uSDHC source files available in drivers/mmc/host.

Table 3-4. uSDHC Driver Files MMC/SD Driver Files

File	Description
drivers/mmc/host/sdhci.c	sdhci standard stack code
driers/mmc/host/sdhci-pltfm.c	sdhci platform layer
drivers/mmc/host/sdhci-esdhc-imx.c	uSDHC driver
drivers/mmc/host/sdhci-esdhc.h	uSDHC driver header file

3.3.5 Menu Configuration Options

The following Linux kernel configuration options are provided for this module.

- CONFIG_MMC builds support for the MMC bus protocol. In menuconfig, this option is available under:
  - Device Drivers > MMC/SD/SDIO Card support
  - By default, this option is Y.
- CONFIG_MMC_BLOCK builds support for MMC block device driver which can be used to mount the file system. In menuconfig, this option is available under:
  - Device Drivers > MMC/SD Card Support > MMC block device driver
  - By default, this option is Y.
- CONFIG_MMC_SDHCI_ESDHC_IMX is used for the i.MX USDHC ports. In menuconfig, this option is found under:
  - Device Drivers > MMC/SD Card Support > Secure Digital Host Controller Interface support > SDHCI support on the platform-specific bus > SDHCI platform support for the eSDHC i.MX controller

To compile SDHCI driver as a loadable module, several options should be selected as indicated below:
  - CONFIG_MMC_SDHCI=m, it can be found at Device Drivers > MMC/SD Card Support > Secure Digital Host Controller Interface support

- CONFIG_MMC_SDHCI_PLTFM=m, it can be found at Device Drivers > MMC/SD Card Support > Secure Digital Host Controller Interface support > SDHCI support on the platform-specific bus.
- CONFIG_MMC_SDHCI_ESDHC_IMX=y, it can be found at Device Drivers > MMC/SD Card Support > Secure Digital Host Controller Interface support > SDHCI support on the platform-specific bus > SDHCI platform support for the Freescale eSDHC i.MX controller

To compile SDHCI driver as a builttin module, several options should be selected as indicated below:

- CONFIG_MMC_SDHCI=y, it can be found at Device Drivers > MMC/SD Card Support > Secure Digital Host Controller Interface support
- CONFIG_MMC_SDHCI_PLTFM=y, it can be found at Device Drivers > MMC/SD Card Support > Secure Digital Host Controller Interface support > SDHCI support on the platform-specific bus.
- CONFIG_MMC_SDHCI_ESDHC_IMX=y, it can be found at Device Drivers > MMC/SD Card Support > Secure Digital Host Controller Interface support > SDHCI support on the platform-specific bus > SDHCI platform support for the Freescale eSDHC i.MX controller
- CONFIG_MMC_UNSAFE_RESUME is used for embedded systems which use a MMC/SD/SDIO card for rootfs. In menuconfig, this option is found under:

3.3.6 Device Tree Binding

Required properties:

- compatible : Should be "fsl,<chip>-esdhc"
- reg : Should contain eSDHC registers location and
- interrupts : Should contain eSDHC interrupt

Optional properties:

- non-removable : Indicate the card is wired to host permanently
- fsl,cd-internal : Indicate to use controller internal card detection
- fsl,wp-internal : Indicate to use controller internal write protection
- cd-gpios : Specify GPIOs for card detection
- wp-gpios : Specify GPIOs for write protection
- fsl,delay-line : Specify delay line value for emmc ddr mode

```
Example:usdhc@02194000 { /* uSDHC2 */
        compatible = "fsl,imx6q-usdhc";
        reg = <0x02194000 0x4000>;
        interrupts = <0 23 0x04>;
        clocks = <&clks 164>, <&clks 164>, <&clks 164>;
```

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

```
        clock-names = "ipg", "ahb", "per";
        pinctrl-names = "default";
        pinctrl-0 = <&pinctrl_usdhc2_1>;
        cd-gpios = <&gpio2 2 0>;
        wp-gpios = <&gpio2 3 0>;
        bus-width = <8>;
        no-1-8-v;
        keep-power-in-suspend;
        enable-sdio-wakeup;
        status = "okay";
};
```

Reference:

- Documentation/devicetree/bindings/mmc/fsl-imx-esdhc.txt
- arch/arm/boot/dts/imx6*.dtsi

3.3.7 Programming Interface

This driver implements the functions required by the MMC bus protocol to interface with the i.MX uSDHC module. See the Linux document generated from build: make htmldocs.

3.3.8 Loadable Module Operations

The SDHCI driver can be built as loadable or builtin module.

1. How to build SDHCI driver as loadable module.
   - CONFIG_MMC_SDHCI=m, it can be found at Device Drivers > MMC/SD Card Support > Secure Digital Host Controller Interface support
   - CONFIG_MMC_SDHCI_PLTFM=m, it can be found at Device Drivers > MMC/SD Card Support > Secure Digital Host Controller Interface support > SDHCI support on the platform-specific bus.
   - CONFIG_MMC_SDHCI_ESDHC_IMX=y, it can be found at Device Drivers > MMC/SD Card Support > Secure Digital Host Controller Interface support > SDHCI support on the platform-specific bus > SDHCI platform support for the i.MX eSDHC i.MX controller
2. How to load and unload SDHCI module.

   Due to dependency, load or unload the module following the module sequence shown below.

   run the following commands to load module:
   - load modules via insmod command, assuming the files of sdhci.ko and sdhci-platform.ko exist in current directory.

```
$> insmod sdhci.ko
$> insmod sdhci-platform.ko
```
• load modules via modprobe command, make sure the files of sdhci.ko and sdhci-platform.ko exist in corresponding kernel module lib directory.

```
$> modprobe sdhci.ko
$> modprobe sdhci-platform.ko
```

run the following commands to unload module.:
• unload modules via insmod command.

```
$> rmsmod sdhci-platform
$> rmsmod sdhci
```
• unload modules via modprobe command.

```
$> modprobe -r sdhci-platform
$> modprobe -r sdhci
```

3.4 NAND GPMI Flash

3.4.1 Introduction

The NAND Flash Memory Technology Devices (MTD) driver is used in the Generic-Purpose Media Interface (GPMI) controller on the i.MX 6 series and i.MX 7Dual.

Only the hardware-specific layer has to be implemented for the NAND MTD driver to operate.

The rest of the functionality such as Flash read/write/erase is automatically handled by the generic layer provided by the Linux MTD subsystem for NAND devices.

The NAND MTD driver interfaces with the integrated NAND controller supporting file systems, such as UBIFS, CRAMFS and JFFS2UBI and UBIFSCRAMFS and JFFS2. The driver implementation supports the lowest level operations on the external NAND Flash chip, such as block read, block write and block erase as the NAND Flash technology only supports block access. Because blocks in a NAND Flash are not guaranteed to be good, the NAND MTD driver is also able to detect bad blocks and feed that information to the upper layer to handle bad block management.

3.4.2 Hardware Operation

NAND Flash is a nonvolatile storage device used for embedded systems.

Driver does not support random accesses of memory as in the case of RAM or NOR Flash. Reading or writing to NAND Flash must be done through the GPMI. NAND Flash is a sequential access device appropriate for mass storage applications. Code stored on NAND Flash cannot be executed from there. Code must be loaded into RAM memory and executed from there. The i.MX 6 contains a hardware error-correcting block.

3.4.3 Software Operation

MTDs in Linux covers all memory devices such as RAM, ROM, and different kinds of NOR/NAND Flashes.

The MTD subsystem provides uniform access to all such devices. Above the MTD devices there could be either MTD block device emulation with a Flash file system (JFFS2) or a UBI layer. The UBI layer in turn, can have either UBIFS above the volumes or a Flash Translation Layer (FTL) with a regular file system (FAT, Ext2/3) above it. The hardware-specific driver interfaces with the GPMI module on the i.MX 6. It implements the lowest level operations such as read, write and erase. If enabled, it also provides information about partitions on the NAND device-this information has to be provided by platform code.

The NAND driver is the point where read/write errors can be recovered if possible. Hardware error correction is performed by BCH blocks and is driven by NAND drivers code.

Detailed information about NAND driver interfaces can be found at www.linux-mtd.infradead.org.

3.4.4 Basic Operations: Read/Write

The NAND driver exports the following callbacks:

```
gpmi_ecc_read_page (with ECC)
gpmi_ecc_write_page (with ECC)
gpmi_read_byte (without ECC)
gpmi_read_buf (without ECC)
gpmi_write_buf (without ECC)
gpmi_ecc_read_oob (with ECC)
gpmi_ecc_write_oob (with ECC)
```

Since Kernel 4.1, the GPMI driver provides raw read/write modes, which exports these callbacks:
- gpmi_ecc_read_page_raw (without ECC)
- gpmi_ecc_write_page_raw (without ECC)

- gpmi_ecc_read_oob_raw (without ECC)
- gpmi_ecc_write_oob_raw (without ECC)

These functions read the requested amount of data, with or without error correction. In the case of read, the gpmi_read_page() function is called, which creates the DMA chain, submits it to execute, and waits for completion. The write case is a bit more complex: the data to be written is mapped and flushed out by calling gpmi_send_page().

3.4.5 Backward Compatibility

Users should know several major GPMI NAND driver changes in kernel 4.1, which may cause incompatibility in Kernel upgrade.

- Exported necessary information to user space application (kobs-ng) through debugfs
- New BCH layout algorithm
- New raw read/write mode

In Kernel 4.1, the NAND GPMI driver exports necessary information to the upper layer through debugfs. The most common case is for the NAND burning tool, kobs-ng. Without enabling debugfs, kobs-ng may not fully use the new feature or may use inappropriate parameters. The user needs to provide the correct BCH geometry information and raw access mode to kobs-ng, if debugfs is not enabled in the customized kernel.

BCH layout in the previous kernel may not meet the NAND chip minimum ECC requirement. Since Kernel 4.1, the BCH layout algorithm, by default, uses the NAND required ECC strength and step size, which are acquired from ONFI parameters, if it is accessible. The change may not be compatible with the BCH layout settings in the previous kernel. For backward compatibility, Kernel and U-boot provide switches to use legacy BCH layout.
- For Kernel, add "fsl,legacy-bch-geometry" in the device tree file.
- For U-Boot, add "CONFIG_NAND_MXS_BCH_LEGACY_GEO" in the board configuration file.

BCH legacy layout setting must be turned on/off simultaneously in both Kernel and U-boot for alignment.

Kobs-ng checks either the Kernel version or raw mode flag in debugfs to determine whether to use new raw mode to access the NAND chip. New kobs-ng fully backward is compatible with the previous Kernel, while the old version kobs-ng cannot work on Kernel 4.1.

3.4.6   Error Correction

When reading or writing data to Flash, some bits can be flipped. This is normal behavior, and NAND drivers utilize various error correcting schemes to correct this. It could be resolved with software or hardware error correction. The GPMI driver uses only a hardware correction scheme with the help of an hardware accelerator-BCH.

For BCH, the page laylout of 2K page is (2k + 64), the page layout of 4K page is (4k + 218) the page layout of 8K page is (8K + 448).

3.4.7   Boot Control Block Management

During startup, the NAND driver scans the first block for the presence of a NAND Control Block (NCB). Its presence is detected by magic signatures. When a signature is found, the boot block candidate is checked for errors using Hamming code. If errors are found, they are fixed, if possible. If the NCB is found, it is parsed to retrieve timings for the NAND chip.

All boot control blocks are created when formatting the medium using the user space application kobs-ng .

3.4.8   Bad Block Handling

When the driver begins, by default, it builds the bad block table. It is possible to determine if a block is bad, dynamically, but to improve performance it is done at boot time. The badness of the erase block is determined by checking a pattern in the beginning of the spare area on each page of the block. However, if the chip uses hardware error correction, the bad marks falls into the ECC bytes area. Therefore, if hardware error correction is used, the bad block mark should be moved.

3.4.9   Source Code Structure

The NAND driver is located in drivers/mtd/nand/gpmi-nand.

Table below lists the source files for the NAND Driver.

Table 3-5.   NAND Driver Files

File	Description
• drivers/mtd/nand/gpmi-nand/bch-regs.h	Functions declaration

Table continues on the next page...

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

Table 3-5. NAND Driver Files (continued)

File	Description
• drivers/mtd/nand/gpmi-nand/gpmi-nand.h • drivers/mtd/nand/gpmi-nand/gpmi-regs.h	
• drivers/mtd/nand/gpmi-nand/gpmi-lib.c • drivers/mtd/nand/gpmi/nand/gpmi-nand.c	GPMI NAND Functions

3.4.10 Menu Configuration Options

To enable the NAND driver, the following options must be set:

• Device Drivers > Memory Technology Device (MTD) support > GPMI NAND Flash Controller driver

In addition, these MTD options must be enabled:

• CONFIG_MTD_NAND = [y | m]
• CONFIG_MTD = y
• CONFIG_MTD_BLOCK = y

In addition, these UBI options must be enabled:

• CONFIG_MTD_UBI=y
• CONFIG_UBIFS_FS=y

3.5 Quad Serial Peripheral Interface (QuadSPI)

3.5.1 Introduction

The Quad Serial Peripheral Interface (QuadSPI) block acts as an interface to one single or two external serial flash devices, each with up to four bidirectional data lines.

It supports the following features:

• Flexible sequence engine to support various flash vendor devices.
• Single, dual, quad and octal mode of operation.
• DDR/DTR mode wherein the data is generated on every edge of the serial flash clock.
• Support for flash data strobe signal for data sampling in DDR and SDR mode.
• DMA support to read RX Buffer data via AMBA AHB bus (64-bit width interface) or IP registers space (32-bit access).

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

3.5.2 Hardware Operation
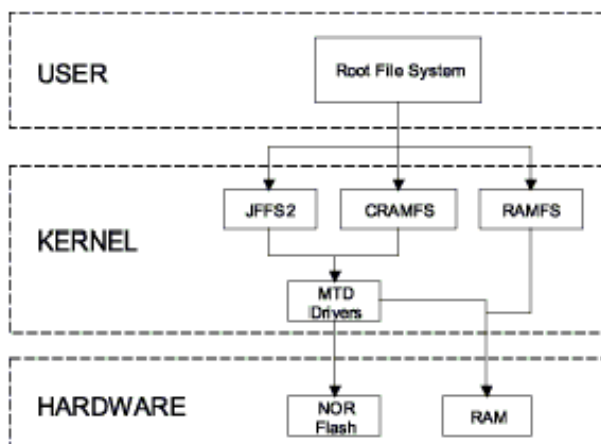
On some boards, the Quad SPI NOR - N25Q256A is equipped, while on some other boards S25FL128S is equipped. Check the Quad SPI NOR type on the boards and then configure it properly.

The N25Q256A is a high-performance multiple input/output serial Flash memory device. The innovative, high-performance, dual and quad input/output instructions enable double or quadruple the transfer bandwidth for READ and PROGRAM operations. The memory is organized as 512 (64 KB) main sectors and can be erased 64 KB sectors at a time. The device features 3-byte or 4-byte address modes to access memory beyond 128 MB. When 4-byte address mode is enabled, all commands requiring an address must be entered and exited with a 4-byte address mode command: ENTER 4-BYTE ADDRESS MODE command and EXIT 4-BYTE ADDRESS MODE command. The 4-byte address mode can also be enabled through the nonvolatile configuration register. The memory can be operated with three different protocols:Extended SPI (standard SPI protocol upgraded with dual and quad operations), Dual I/O SPI and Quad I/O SPI. Each protocol contains unique commands to perform READ operations in DTR mode. This enables high data throughput while running at lower clock frequencies.

The S25FL128S device is flash non-volatile memory product. It connects to a host system via a Serial Peripheral Interface (SPI). Traditional SPI single bit serial input and output (SIngle I/O or SIO) is supported as well as optional two bit (Dual I/O or DIO) and four bit (Quad I/O or QIO) serial commands. It also adds support for Double Data Rate (DDR) read commands for SIO, DIO, and QIO that transfer address and read data on both edges of the clock.

3.5.3 Software Operation

In a Flash-based embedded Linux system, a number of Linux technologies work together to implement a file system. The following figure illustrates the relationships between some of the standard components.

Figure 3-1. Components of a Flash-Based File System

The MTD subsystem for Linux OS is a generic interface to memory devices, such as Flash and RAM, providing simple read, write, and erase access to physical memory devices. Devices called mtdblock devices can be mounted by JFFS, JFFS2, and CRAMFS file systems. The Quad SPI NOR MTD driver is based on the MTD data Flash driver in the kernel by adding SPI access. In the initialization phase, the Quad SPI NOR MTD driver detects a data Flash by reading the JEDEC ID. Then the driver adds the MTD device. The SPI NOR MTD driver also provides the interfaces to read, write, and erase NOR Flash.

3.5.4 Driver Features

This Quad NOR driver implementation supports the following feature:

- Provides necessary information for the upper-layer MTD driver.

3.5.5 Source Code Structure

Table 3-6. QuadSPI Driver File

File	Description
drivers/mtd/spi-nor/spi-nor.c	SPI-NOR framework
drivers/mtd/spi-nor/fsl-quadspi.c	Quad SPI Driver

3.5.6 Menu Configuration Options

To enable the Quad SPI driver, the following options must be set:

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

- Device Drivers >Memory Technology Device (MTD) support > Freescale Quad SPI controller

For the configuration have these set to enable Quad Spi.

- CONFIG_MTD_SPI_NOR: Framework for the SPI NOR which can be used by the SPI device drivers and the SPI-NOR device driver.
- CONFIG_SPI_FSL_QUADSPI: Enables support for the Quad SPI controller in master mode.

3.6 SATA

3.6.1 Introduction

The SATA AHCI driver is based on the LIBATA layer of the block device infrastructure of the Linux kernel. The detailed hardware operation of SATA is detailed in the Synopsys DesignWare Cores SATA AHCI documentation, named SATA_Data_Book.pdf.

3.6.2 Board Configuration Options

With the power off, install the SATA cable and hard drive.

3.6.3 Software Operation

The details about the libata APIs, see the libATA Developer's Guide.

The SATA AHCI driver is based on the LIBATA layer of the block device infrastructure of the Linux kernel. i.MX integrated AHCI linux driver combined the standard AHCI drivers handle the details of the integrated i.MX SATA AHCI controller, while the LIBATA layer understands and executes the SATA protocols. The SATA device, such as a hard disk, is exposed to the application in user space by the /dev/sda* interface. Filesystems are built upon the block device. The AHCI specified integrated DMA engine, which assists the SATA controller hardware in the DMA transfer modes.

3.6.4   Source Code Structure

The source code of the i.MX AHCI SATA driver is located in drivers/ata The standard AHCI and AHCI platform drivers are used to do the actual SATA operations.

Table 3-7.   SATA Driver Files

File	Description
drivers/ata/ahci_imx.c	i.MX AHCI SATA Driver
drivers/ata/ahci.c	Standard AHCI drivers
drivers/ata/ahci-platform.c	Standard AHCI platform drivers

3.6.5   Menu Configuration Options

The following Linux kernel configurations are provided for SATA driver:

```
    Symbol: AHCI_IMX
[=y]
    Type  :
tristate

    Prompt: Freescale i.MX AHCI SATA
support


Location:

        -> Device
Drivers

        -> Serial ATA and Parallel ATA drivers (ATA
[=y])
          -> Platform AHCI SATA support (SATA_AHCI_PLATFORM
[=y])
```

In busybox, enable "fdisk" under "Linux System Utilities".

3.6.6   Programming Interface

The application interface to the SATA driver is the standard POSIX device interface (for example: open, close, read, write, and ioctl) on /dev/sda*.

3.6.7 Usage Example

NOTE
There may be a known error message when few kinds of SATA disks are initialized, such as:

ata1.00: serial number mismatch '090311PB0300QKG3TB1A' ! = "

ata1.00: revalidation failed (errno=-19)

This should be ignored.

1. After building the kernel and the SATA AHCI driver and deploying, boot the target, and log in as root.
2. Make sure that the AHCI and AHCI platform drivers are built in the kernel or loaded into the kernel.

You should see messages similar to the following:

```
ahci: SSS flag set, parallel bus scan disabled
ahci ahci: AHCI 0001.0300 32 slots 1 ports 3 Gbps 0x1 impl platform mode
ahci ahci: flags: ncq sntf stag pm led clo only pmp pio slum part ccc apst
scsi0 : ahci_platform
ata1: SATA max UDMA/133 mmio [mem 0x02200000-0x02203fff] port 0x100 irq 71
ata1: SATA link up 3.0 Gbps (SStatus 123 SControl 300)
ata1.00: ATA-8: SAMSUNG HM100UI, 2AM10001, max UDMA/133
ata1.00: 1953525168 sectors, multi 0: LBA48 NCQ (depth 31/32)
ata1.00: configured for UDMA/133
scsi 0:0:0:0: Direct-Access     ATA      SAMSUNG HM100UI  2AM1 PQ: 0 ANSI: 5
sd 0:0:0:0: [sda] 1953525168 512-byte logical blocks: (1.00 TB/931 GiB)
sd 0:0:0:0: [sda] 4096-byte physical blocks
sd 0:0:0:0: [sda] Write Protect is off
sd 0:0:0:0: [sda] Write cache: enabled, read cache: enabled, doesn't support DPO or FUA
sda: sda1
sd 0:0:0:0: [sda] Attached SCSI disk
```

You may use standard Linux utilities to partition and create a file system on the drive (for example: fdisk and mke2fs) to be mounted and used by applications.

The device nodes for the drive and its partitions appears under /dev/sda*. For example, to check basic kernel settings for the drive, execute hdparm /dev/sda.

3.6.8 Usage Example

Create Partitons

The following command can be used to find out the capacities of the hard disk. If the hard disk is pre-formatted, this command shows the size of the hard disk, partitions, and filesystem type:

```
$fdisk -l /dev/sda
```

If the hard disk is not formatted, create the partitions on the hard disk using the following command:

```
$fdisk /dev/sda
```

After the partition, the created files resemble /dev/sda[1-4].

Block Read/Write Test: The command, dd, is used for for reading/writing blocks. Note this command can corrupt the partitions and filesystem on Hard disk.

To clear the first 5 KB of the card, do the following:

```
$dd if=/dev/zero of=/dev/sda1 bs=1024 count=5
```

The response should be as follows:

5+0 records in

5+0 records out

To write a file content to the card enter the following text, substituting the name of the file to be written for file_name, do the following:

```
$dd if=file_name of=/dev/sda1
```

To read 1KB of data from the card enter the following text, substituting the name of the file to be written for output_file, do the following:

```
$dd if=/dev/sda1 of=output_file bs=1024 count=1
```

Files System Tests

Format the hard disk partitons using mkfs.vfat or mkfs.ext2, depending on the filesystem:

```
$mkfs.ext2 /dev/sda1
$mkfs.vfat /dev/sda1
```

Mount the file system as follows:

```
$mkdir /mnt/sda1
$mount -t ext2 /dev/sda1 /mnt/sda1
```

After mounting, file/directory, operations can be performed in /mnt/sda1.

Unmount the filesystem as follows:

```
$umount /mnt/sda1
```

3.7 Smart Direct Memory Access (SDMA) API

3.7.1 Overview

The Smart Direct Memory Access (SDMA) API driver controls the SDMA hardware.

It provides an API to other drivers for transferring data between MCU memory space and the peripherals. It supports the following features:

- Loading channel scripts from the MCU memory space into SDMA internal RAM
- Loading context parameters of the scripts
- Loading buffer descriptor parameters of the scripts
- Controlling execution of the scripts
- Callback mechanism at the end of script execution

3.7.2 Hardware Operation

The SDMA controller is responsible for transferring data between the MCU memory space and peripherals and includes the following features:

- Multichannel DMA supporting up to 32 time-division multiplexed DMA channels.
- Powered by a 16-bit Instruction-Set micro-RISC engine.
- Each channel executes specific script.
- Very fast context-switching with two-level priority based preemptive multitasking.
- 4 Kbytes ROM containing startup scripts (that is, boot code) and other common utilities that can be referenced by RAM-located scripts.
- 8 Kbyte RAM area is divided into a processor context area and a code space area used to store channel scripts that are downloaded from the system memory.

3.7.3 Software Operation

The driver provides an API for other drivers to control SDMA channels. SDMA channels run dedicated scripts according to peripheral and transfer types. The SDMA API driver is responsible for loading the scripts into SDMA memory, initializing the channel descriptors, and controlling the buffer descriptors and SDMA registers.

The table below provides a list of drivers that use SDMA and the number of SDMA physical channels used by each driver. A driver can specify the SDMA channel number that it wishes to use, static channel allocation, or can have the SDMA driver provide a free SDMA channel for the driver to use, dynamic channel allocation. For dynamic channel allocation, the list of SDMA channels is scanned from channel 32 to channel 1. When a free channel is found, that channel is allocated for the requested DMA transfers.

Table 3-8.   SDMA Channel Usage

Driver Name	Number of SDMA Channels	SDMA Channel Used
SDMA CMD	1	Static Channel allocation-uses SDMA channels 0
SSI	2 per device	Dynamic channel allocation
UART	2 per device	Dynamic channel allocation
SPDIF	2 per device	Dynamic channel allocation
ESAI	2 per device	Dynamic channel allocation

3.7.4   Source Code Structure

The dmaengine.h (header file for SDMA API) is available in the directory linux/include/ linux

The following table shows the source files available in the directory drivers/dma.

Table 3-9.   SDMA API Source Files

File	Description
drivers/dma/dmaengine.c	SDMA management routine
drivers/dma/imx-sdma.c	SDMA implement driver
drivers/dma/imx-dma.c	i.MX DMA driver

The following table shows the image files available in the directory firmware/imx/sdma for 4.1 and 4.9 kernels. For 4.14 kernel, the sdma firmware is provided with the firmware-imx package and not in the kernel source tree.

Table 3-10.   SDMA Script Files

File	Description
sdma-imx6q.bin	SDMA RAM scripts for i.MX 6
sdma-imx7d.bin	SDMA RAM scripts for i.MX 7 and i.MX 8M

3.8   SPI NOR Flash Memory Technology Device (MTD)

3.8.1 Introduction

The SPI NOR Flash Memory Technology Device (MTD) driver provides the support to the data Flash though the SPI interface.

By default, the SPI NOR Flash MTD driver creates static MTD partitions to support data Flash.

3.8.2 Hardware Operation

On some boards, the SPI NOR - AT45DB321D is equipped, while on some boards M25P32 is equipped. Check the SPI NOR type on the boards and then configure it properly.

The AT45DB321D is a 2.7 V, serial-interface sequential access Flash memory. The AT45DB321D serial interface is SPI compatible for frequencies up to 66 MHz. The memory is organized as 8,192 pages of 512 bytes or 528 bytes. The AT45DB321D also contains two SRAM buffers of 512/528 bytes each which allow receiving of data while a page in the main memory is being reprogrammed, as well as writing a continuous data stream.

The M25P32 is a 32 Mbit (4M x 8) Serial Flash memory, with advanced write protection mechanisms, accessed by a high-speed SPI-compatible bus up to 75 MHz. The memory is organized as 64 sectors, each containing 256 pages. Each page is 256 bytes wide. Therefore, the whole memory can be viewed as consisting of 16384 pages, or 4,194,304 bytes. The memory can be programmed 1 to 256 bytes at a time using the Page Program instruction. The whole memory can be erased using the Bulk Erase instruction, or a sector at a time, using the Sector Erase instruction.

Unlike conventional Flash memories that are accessed randomly, these two SPI NOR access data sequentially. They operate from a single 2.7-3.6 V power supply for program and read operations. They are enabled through a chip select pin and accessed through a three-wire interface: Serial Input, Serial Output, and Serial Clock.

3.8.3 Software Operation

In a Flash-based embedded Linux system, a number of Linux technologies work together to implement a file system. The figure below illustrates the relationships between some of the standard components.

Figure 3-2. Components of a Flash-Based File System

The MTD subsystem for Linux OS is a generic interface to memory devices, such as Flash and RAM, providing simple read, write, and erase access to physical memory devices. Devices called mtdblock devices can be mounted by JFFS, JFFS2 and CRAMFS file systems. The SPI NOR MTD driver is based on the MTD data Flash driver in the kernel by adding SPI access. In the initialization phase, the SPI NOR MTD driver detects a data Flash by reading the JEDEC ID. Then the driver adds the MTD device. The SPI NOR MTD driver also provides the interfaces to read, write, and erase NOR Flash.

3.8.4 Source Code Structure

The following table shows the driver files:

Table 3-11.   SPI NOR MTD Driver Files

File	Description
drivers/mtd/devices/m25p80.c	Source file
drivers/mtd/spi-nor/spi-nor.c	Source file

3.8.5 Menu Configuration Options

In menu configuration enable the following module:

- CONFIG_MTD_M25P80: This config enables access to most modern SPI flash chips, used for program and data storage.
- Device Drivers > Memory Technology Device (MTD) support >Self-contained MTD device drivers > Support most SPI Flash chips (AT26DF, M25P, W25X, and so on)

Chapter 4
Connectivity

4.1 ADC

4.1.1 ADC Introduction

The features of the ADC-Digital are as follows:

- Two 12-bit ADCs
- Linear successive approximation algorithm with up to 12-bit resolution with 10/11 bit accuracy
- Up to 1 MS/s sampling rate
- Up to 8 single-ended external analog inputs
- Single or continuous conversion (automatic return to idle after single conversion)
- Output Modes: (in right-justified unsigned format)
    - 12-bit
    - 10-bit
    - 8-bit
- Configurable sample time and conversion speed/power
- Conversion complete and hardware average complete flag and interrupt
- Input clock selectable from up to four sources
- Asynchronous clock source for lower noise operation with option to output the clock
- Selectable asynchronous hardware conversion trigger with hardware channel select
- Selectable voltage reference, Internal, External, or Alternate
- Operation in low power modes for lower noise operation
- Hardware average function
- Self-calibration mode

4.1.2 ADC External Signals

- ADC_VREFH: Voltage reference high

- ADC_VREHL: Voltage reference low
- ADC1_IN0: Analog channel 1 input 0
- ADC1_IN1: Analog channel 1 input 1
- ADC1_IN2: Analog channel 1 input 2
- ADC1_IN3: Analog channel 1 input 3
- ADC2_IN0: Analog channel 2 input 0
- ADC2_IN1: Analog channel 2 input 1
- ADC2_IN2: Analog channel 2 input 2
- ADC2_IN3: Analog channel 2 input 3

The ADC pin settings should be done in the ADCx_PCTL register. No other extra IOMUX settings are required.

4.1.3 ADC Driver Overview

The ADC driver is developed under the Linux IIO (Industrial I/O) driver frame. The ADC driver only provides the basic functions. The following features are supported:

- Four external inputs for each ADC controller channel
- 12 bit ADC
- Single conversion
- Hardware average
- Low power mode of ADC
- Sample rate changes in the available sample rate group

4.1.4 Source Code Structure

Table 4-1. ADC Driver Files

File	Description
drivers/iio/adc/vf610_adc.c	i.MX 6UltraLite and i.MX 6SoloX ADC functions.
drivers/iio/adc/imx7d_adc.c	i.MX 7Dual ADC functions.
drivers/iio/adc/imx8qxp_adc.c	i.MX 8QXP ADC functions.

4.1.5 Menu Configuration Options

Configure the kernel option to enable the module by menuconfig:

Device Drivers > Industrial I/O support > Analog to digital converters > Freescale vf610 ADC driver

Device Drivers > Industrial I/O support > Analog to digital converters > i.MX 7Dual ADC driver

Device Drivers > Industrial I/O support > Analog to digital converters > i.MX 8QXP ADC driver

4.1.6 Programming Interface

Linux IIO provides some system interface to get the raw ADC data from the related input. Users can also set the sample rate in the available sample rate group. The ADC controllers system interface is located:

/sys/devices/soc0/soc.1/2200000.aips-bus/2280000.adc/iio:device0:

/sys/devices/soc0/soc.1/2200000.aips-bus/2284000.adc/iio:device1:

The following table lists the software interfaces.

Table 4-2. Software Interfaces

Software interface	Description
in_voltage0_raw~ in_voltage3_raw	**cat in_voltage0_raw** to get raw ADC data
sampling_frequency_available	**cat sampling_frequency_available** to get available sample rate group
in_voltage_sampling_frequency	**cat in_voltage_sampling_frequency** to show current sample rate **echo value** > **in_voltage_sampling_frequency** to set the sample rate

4.2 ENET IEEE-1588

4.2.1 Introduction

ENET IEEE-1588 driver performs a set of functions that enabling precise synchronization of clocks in network communication.

The driver requires a protocol stack to complete IEEE-1588 full protocol. It complies with the LinuxPTP stack.

To allow for IEEE 1588 or similar time synchronization protocol implementations, the ENET MAC is combined with a time-stamping module to support precise time stamping of incoming and outgoing frames.
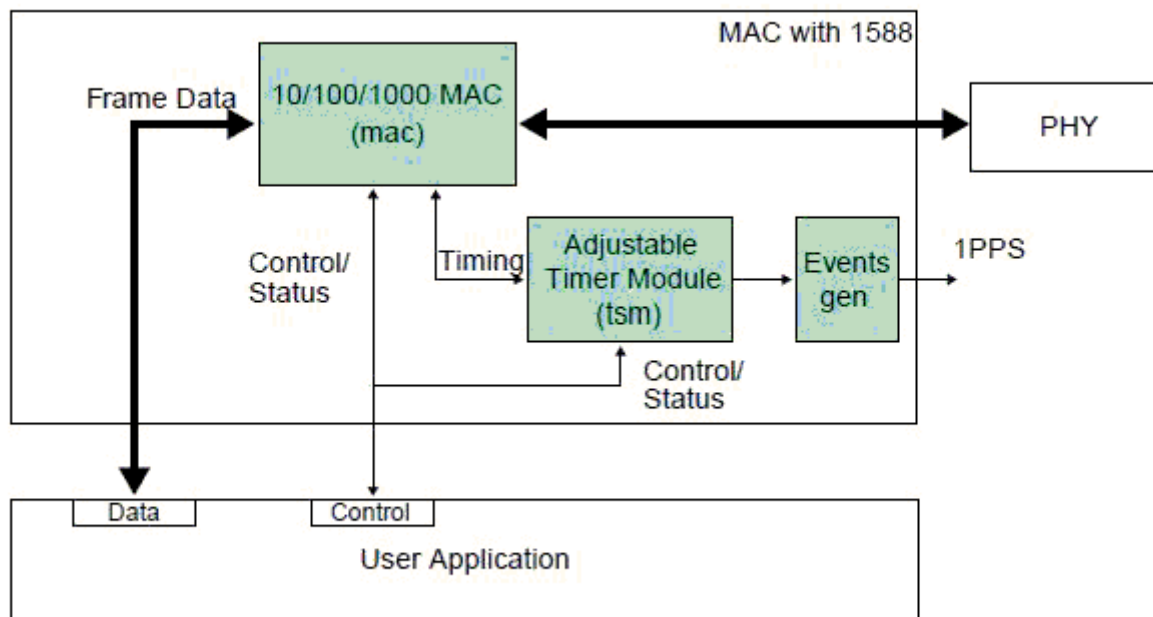
Figure 4-1. IEEE 1588 Functions Overview

4.2.1.1 Transmit Timestamping

On transmit, only 1588 event frames need to be time-stamped. The Client application (for example, the MAC driver) should detect 1588 event frames and set the signal ff_tx_ts_frm together with the frame.

For every transmitted frame, the MAC returns the captured timestamp on tx_ts (31:0) with the frame sequence number (tx_ts_id(3:0)) and the transmit status. The transmit status bit tx_ts_stat (5) indicates that the application had the ff_tx_ts_frm signal asserted for the frame.

If ff_tx_ts_frm is set to '1', the MAC additionally memorizes the timestamp for the frame in the register TS_TIMESTAMP. The interrupt bit EIR (TS_AVAIL) is set to indicate that a new timestamp is available.

Software would implement a handshaking procedure by setting the ff_tx_ts_frm signal when it transmits the frame it needs a timestamp for and then waits on the EIR (TS_AVAIL) interrupt bit to know when the timestamp is available. It then can read the timestamp from the TS_TIMESTAMP register. This is done for all event frames; other frames do not use the ff_tx_ts_frm indicator and hence do not interfere with the timestamp capture.

4.2.1.2 Receive Timestamping

When a frame is received, the MAC latches the value of the timer when the frame SFD field is detected and provides the captured timestamp on ff_rx_ts(31:0). This is done for all received frames.

The DMA controller has to ensure that it transfers the timestamp provided for the frame into the corresponding field within the receive descriptor for software access.

4.2.2 Software Operation

The 1588 Driver has the functions listed below:

- Module initialization-Initializes the module with the device-specific structure, and registers a character driver.
- Interrupt servicing routine-Supports events, such as TS_AVAIL, TS_TIMER. The driver shares interrupt servicing routine with FEC driver.

4.2.2.1 Source Code Structure

Table below lists the source files in drivers/net/ethernet/freescale directory.

Table 4-3. ENET 1588 File List

File	Description
drivers/net/ethernet/frescale/fec.h	Header file defining registers
drivers/net/ethernet/freescale/fec_ptp.c	ENET 1588 timer

4.2.2.2 Menu Configuration Options

By default, ENET 1588 is enabled.

4.2.2.3 Programming Interface

The 1588 driver complies with the Linuxptp protocol stack interface.

Stack-specific defines are added to the header file (fec.h).

4.2.3 1588 Stack Introduction

This release supports the following type of the 1588 Stack:

- Linuxptp stack

  This software is an implementation of the Precision Time Protocol (PTP) according to IEEE standard 1588 for Linux OS. The dual design goals are to provide a robust implementation of the standard and to use the most relevant and modern Application Programming Interfaces (API) offered by the Linux OS kernel. Supporting legacy APIs and other platforms is not a goal. The software is copyrighted by the authors and is licensed under the GNU General Public License.

The software development is hosted at Source Forge: sourceforge.net/projects/linuxptp/

4.2.3.1 Linuxptp Stack Features

Linuxptp support the following features:

- Ordinary/Boundary Clock
- Best master clock algorithm
- Transport over UDP/IPv4, UDP/IPv6, and IEEE 802.3
- Transparent clock (E2E/P2P)
- Slave only
- Supporting IEEE 802.1AS-2011 in the role of end station

4.2.3.2 Using Linuxptp

Run ptp4 1588 stack binary with the following commands.

Linuxptp:

```
Transport on UDP IPV4 with E2E delay mechanism: ptp4l -A -4 -H -m -i eth0
Transport on UDP IPV4 with P2P delay mechanism: ptp4l -P -A -4 -H -m -i eth0
Transport on UDP IPV6 with E2E delay mechanism: ptp4l -A -6 -H -m -i eth0
Transport on UDP IPV6 with P2P delay mechanism: ptp4l -P -A -6 -H -m -i eth0
Transport on IEEE 802.3 with E2E delay mechanism: ptp4l -A -2 -H -m -i eth0
Transport on IEEE 802.3 with P2P delay mechanism: ptp4l -P -A -2 -H -m -i eth0
```

4.3 Enhanced Configurable Serial Peripheral Interface (ECSPI)

4.3.1 Introduction

The ECSPI driver implements a standard Linux driver interface to the ECSPI controllers.

It supports the following features:

- Interrupt-driven transmit/receive of bytes
- Multiple master controller interface
- Multiple slaves select
- Multiclient requests

ECSPI is used for fast data communication with fewer software interrupts than conventional serial communications. >Each ECSPI is equipped with a data FIFO and is a master/slave configurable serial peripheral interface module, allowing the processor to interface with external SPI master or slave devices.

The primary features of the ECSPI includes:

- Master/slave-configurable
- Four chip select signals to support multiple peripherals
- Up to 32-bit programmable data transfer
- 64 x 32-bit FIFO for both transmit and receive data
- Configurable polarity and phase of the Chip Select (SS) and SPI Clock (SCLK)

The ECSPI module supports the following features:

- Implements each of the functions required by a ECSPI module to interface to Linux OS
- Multiple SPI master controllers
- Multiclient synchronous requests

4.3.2 Software Operation

The following sections describe the ECSPI software operation.

4.3.3   SPI Sub-System in Linux OS

The ECSPI driver layer is located between the client layer (SPI-NOR Flash are examples of clients) and the hardware access layer. The figure below shows the block diagram for SPI subsystem in Linux OS.

The SPI requests go into I/O queues. Requests for a given SPI device are executed in FIFO order and they complete asynchronously through completion callbacks. There are also some simple synchronous wrappers for those calls including the ones for common transaction types such as writing a command and then reading its response.
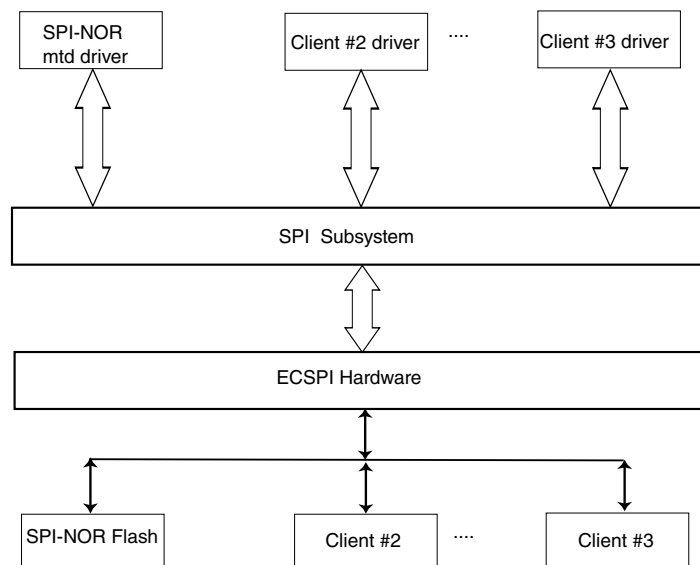
Figure 4-2. SPI Subsystem

All SPI clients must have a protocol driver associated with them and they all must be sharing the same controller driver. Only the controller driver can interact with the underlying SPI hardware module. The figure below shows how the different SPI drivers are layered in the SPI subsystem.
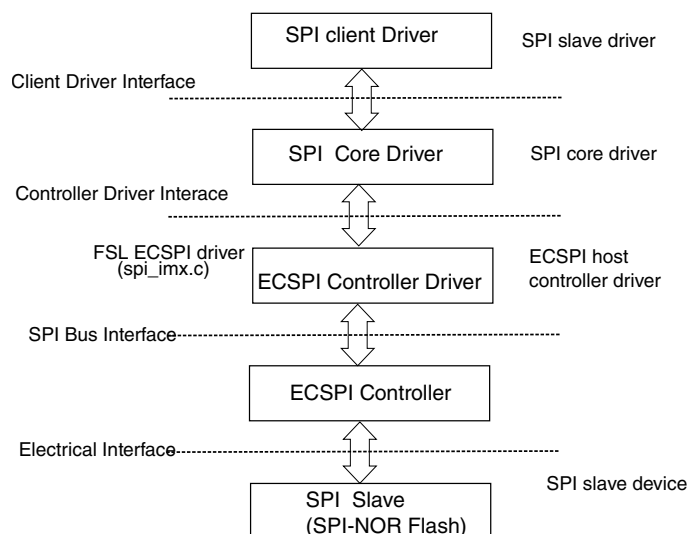
Figure 4-3. Layering of SPI Drivers in SPI Subsystem

4.3.4 Software Limitations

The ECSPI driver limitations are as follows:

- Does not currently have SPI slave logic implementation
- Does not support a single client connected to multiple masters
- Does not currently implement the user space interface with the help of the device node entry but supports sysfs interface

4.3.5 Standard Operations

The ECSPI driver is responsible for implementing standard entry points for init, exit, chip select, and transfer. The driver implements the following functions:

- Init function spi_imx_init() registers the device_driver structure.
- Probe function spi_imx_probe() performs initialization and registration of the SPI device-specific structure with SPI core driver. The driver probes for memory and IRQ resources. Configures the IOMUX to enable ECSPI I/O pins, requests for IRQ and resets the hardware.
- Chip select function spi_imx_chipselect() configures the hardware ECSPI for the current SPI device. Sets the word size, transfer mode, data rate for this device.
- SPI transfer function spi_imx_transfer() handles data transfers operations.
- SPI setup function spi_imx_setup() initializes the current SPI device.
- SPI driver ISR spi_imx_isr() is called when the data transfer operation is completed and an interrupt is generated.

4.3.6 ECSPI Synchronous Operation

The figure below shows how the ECSPI provides synchronous read/write operations.
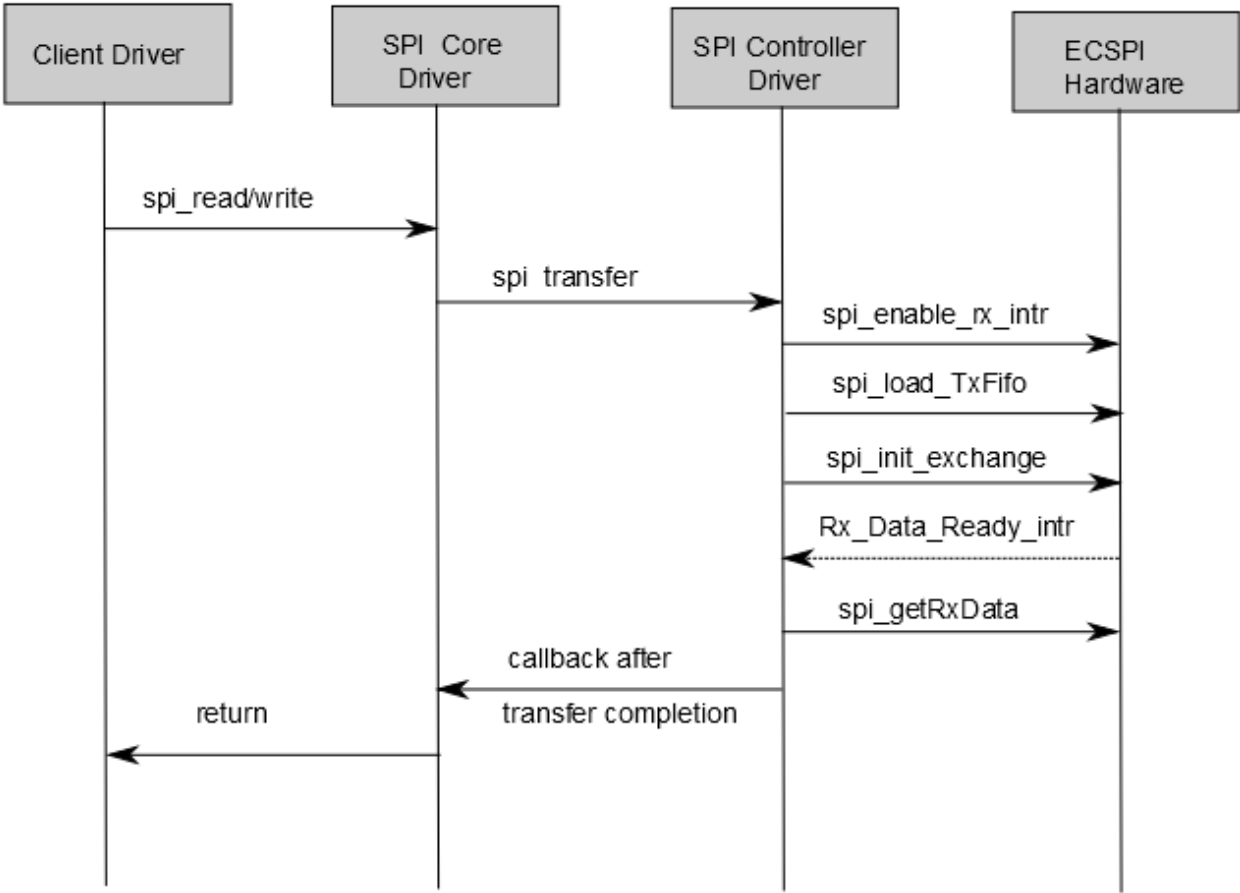
Figure 4-4. ECSPI Synchronous Operation

4.3.7 Source Code Structure

Table below shows the source files available in the drivers/spi directory:

Table 4-4. ECSPI Driver Files

File	Description
driveers/spi/spi-imx.c	SPI Master Controller driver

4.3.8 Menu Configuration Options

In menu configuration enable the following module:

- CONFIG_SPI build support for the SPI core. In menuconfig, this option is available under:
    - Device Drivers > SPI Support.
- CONFIG_BITBANG is the Library code that is automatically selected by drivers that need it. SPI_IMX selects it. In menuconfig, this option is available under:
    - Device Drivers > SPI Support > Utilities for Bitbanging SPI masters.
- CONFIG_SPI_IMX implements the SPI master mode for ECSPI. In menuconfig, this option is available under:
    - Device Drivers > SPI Support > Freescale i.MX SPI controllers.

4.3.9 Programming Interface

This driver implements all the functions that are required by the SPI core to interface with the ECSPI hardware.

For more information, see the Linux document generated from build: make htmldocs.

4.3.10 Interrupt Requirements

The SPI interface generates interrupts.

ECSPI interrupt requirements are listed in table below.

Table 4-5. ECSPI Interrupt Requirements

Parameter	Equation	Typical	Worst Case
BaudRate/ Transfer Length	(BaudRate/(TransferLength)) * (1/Rxtl)	31250	1500000

The typical values are based on a baud rate of 1 Mbps with a receiver trigger level (Rxtl) of 1 and a 32-bit transfer length. The worst-case is based on a baud rate of 12 Mbps (max supported by the SPI interface) with a 8-bits transfer length.

4.4 Fast Ethernet Controller (FEC)

4.4.1 Introduction

The Fast Ethernet Controller (FEC) driver performs the full set of IEEE 802.3/Ethernet CSMA/CD media access control and channel interface functions.

The FEC requires an external interface adapter and transceiver function to complete the interface to the Ethernet media. It supports half or full-duplex operation on 10 Mbps, 100 Mbps, and 1000 Mbps-related Ethernet networks.

The FEC driver supports the following features:

- Full/Half duplex operation
- Link status change detect
- Auto-negotiation (determines the network speed and full or half-duplex operation)
- Transmits features such as automatic retransmission on collision and CRC generation
- Obtaining statistics from the device such as transmit collisions

The network adapter can be accessed through the ifconfig command with interface name ethx. The driver auto-probes the external adaptor (PHY device).

4.4.2 Hardware Operation

The FEC is an Ethernet controller that interfaces the system to the LAN network.

The FEC supports different standard MAC-PHY (physical) interfaces for connection to an external Ethernet transceiver. The FEC supports the 10/100 Mbps MII, 10/100 Mbps RMII, and 10/100/1000 Mbps RGMII. In addition, the FEC supports 1000 Mbps RGMII, which uses 4-bit reduced GMII operating at 125 MHz.

A brief overview of the device functionality is provided here. For details, see the FEC chapter of the Applications Processor Reference Manual

In MII mode, there are 18 signals defined by the IEEE 802.3 standard and supported by the EMAC. MII, RMII and RGMII modes uses a subset of the 18 signals. These signals are listed in table below.

Table 4-6. Pin Usage in MII, RMII and RGMII Modes

Direction	EMAC Pin Name	RMII Usage	RGMII Usage (not supported by i.MX 6SoloLite or i.MX 6UltraLite)
In/Out	FEC_MDIO	Management Data Input/output	Management Data Input/Output
Out	FEC_MDC	General output	Management Data Clock

Table continues on the next page...

Table 4-6. Pin Usage in MII, RMII and RGMII Modes (continued)

Direction	EMAC Pin Name	RMII Usage	RGMII Usage (not supported by i.MX 6SoloLite or i.MX 6UltraLite)
Out	FEC_TXD[0]	Data out, bit 0	Data out, bit 0
Out	FEC_TXD[1]	Data out, bit 1	Data out, bit 1
Out	FEC_TXD[2]	Not Used	Data out, bit 2
Out	FEC_TXD[3]	Not Used	Data out, bit 3
Out	FEC_TX_EN	Transmit Enable	Transmit Enable
Out	FEC_TX_ER	Not Used	Not Used
In	FEC_CRS	Not Used	Not Used
In	FEC_COL	Not Used	Not Used
In	FEC_TX_CLK	Not Used	Synchronous clock reference (REF_CLK, can connect from PHY)
In	FEC_RX_ER	Receive Error	Not Used
In	FEC_RX_CLK	Not Used	Synchronous clock reference (REF_CLK, can connect from PHY)
In	FEC_RX_DV	Receive Data Valid and generate CRS	RXDV XOR RXERR on the falling edge of FEC_RX_CLK.
In	FEC_RXD[0]	Data in, bit 0	Data in, bit 0
In	FEC_RXD[1]	Data in, bit 1	Data in, bit 1
In	FEC_RXD[2]	Not Used	Data in, bit 2
In	FEC_RXD[3]	Not Used	Data in, bit 3

The MII management interface consists of two pins, FEC_MDIO, and FEC_MDC. The FEC hardware operation can be divided in the parts listed below. For details, see the Applications Processor Reference Manuals.

- Transmission-The Ethernet transmitter is designed to work with almost no intervention from software. Once ECR[ETHER_EN] is asserted and data appears in the transmit FIFO, the Ethernet MAC is able to transmit onto the network. When the transmit FIFO fills to the watermark (defined by the TFWR), the MAC transmit logic asserts FEC_TX_EN and starts transmitting the preamble (PA) sequence, the start frame delimiter (SFD), and then the frame information from the FIFO. However, the controller defers the transmission if the network is busy (FEC_CRS asserts).
- Before transmitting, the controller waits for carrier sense to become inactive, then determines if carrier sense stays inactive for 60 bit times. If the transmission begins after waiting an additional 36 bit times (96 bit times after carrier sense originally became inactive), both buffer (TXB) and frame (TXF) interrupts may be generated as determined by the settings in the EIMR.

- Reception-The FEC receiver is designed to work with almost no intervention from the host and can perform address recognition, CRC checking, short frame checking, and maximum frame length checking. When the driver enables the FEC receiver by

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

asserting ECR[ETHER_EN], it immediately starts processing receive frames. When FEC_RX_DV asserts, the receiver checks for a valid PA/SFD header. If the PA/SFD is valid, it is stripped and the frame is processed by the receiver. If a valid PA/SFD is not found, the frame is ignored. In MII mode, the receiver checks for at least one byte matching the SFD. Zero or more PA bytes may occur, but if a 00 bit sequence is detected prior to the SFD byte, the frame is ignored.

- After the first six bytes of the frame have been received, the FEC performs address recognition on the frame. During reception, the Ethernet controller checks for various error conditions and once the entire frame is written into the FIFO, a 32-bit frame status word is written into the FIFO. This status word contains the M, BC, MC, LG, NO, CR, OV, and TR status bits, and the frame length. Receive Buffer (RXB) and Frame Interrupts (RXF) may be generated if enabled by the EIMR register. When the receive frame is complete, the FEC sets the L bit in the RxBD, writes the other frame status bits into the RxBD, and clears the E bit. The Ethernet controller next generates a maskable interrupt (RXF bit in EIR, maskable by RXF bit in EIMR), indicating that a frame has been received and is in memory. The Ethernet controller then waits for a new frame.

- Interrupt management-When an event occurs that sets a bit in the EIR, an interrupt is generated if the corresponding bit in the interrupt mask register (EIMR) is also set. The bit in the EIR is cleared if a one is written to that bit position; writing zero has no effect. This register is cleared upon hardware reset. These interrupts can be divided into operational interrupts, transceiver/network error interrupts, and internal error interrupts. Interrupts which may occur in normal operation are GRA, TXF, TXB, RXF, RXB. Interrupts resulting from errors/problems detected in the network or transceiver are HBERR, BABR, BABT, LC, and RL. Interrupts resulting from internal errors are HBERR and UN. Some of the error interrupts are independently counted in the MIB block counters. Software may choose to mask off these interrupts as these errors are visible to network management through the MIB counters.

- PHY management-phylib was used to manage all the FEC PHY-related operation such as PHY discovery, link status, and state machine.MDIO bus will be created in FEC driver and registered into the system. See Documentation/networking/phy.txt under the Linux OS source directory for more information.

4.4.3 Software Operation

The FEC driver supports the following functions:

- Module initialization-Initializes the module with the device-specific structure
- Rx/Tx transmition
- Interrupt servicing routine

- PHY management
- FEC management such init/start/stop
- i.MX 6 FEC module use little-endian format

4.4.4 Source Code Structure

The table below shows the source files.

They are available at the drivers/net/ethernet/freescale directory.

Table 4-7. FEC Driver Files

File	Description
drivers/net/ethernet/freescale/fec.h	Header file defining registers
drivers/net/ethernet/freescale/fec_main.c	Linux driver for Ethernet LAN controller
drivers/net/ethernet/freescale/fec_fixup.c	Linux driver for SoC and PHY special implement

4.4.5 Menu Configuration Options

Configure the kernel to provide for this module:

- CONFIG_FEC is provided for this module. This option is available under:
    - Device Drivers > Network device support > Ethernet (10, 100 or 1000 Mbit) > FEC Ethernet controller.
    - To mount NFS-rootfs through FEC, disable the other Network config in the menuconfig if need.

4.4.6 Programming Interface

Device-specific defines are added to the header file (fec.h) and they provide common board configuration options.

fec.h defines the struct for the register access and the struct for the buffer descriptor. For example,

```
/*
 *      Define the buffer descriptor structure.
 */
struct bufdesc {
        unsigned short          cbd_datlen;     /* Data length */
        unsigned short          cbd_sc;         /* Control and status info */
        unsigned long           cbd_bufaddr;    /* Buffer address */

};
```

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

```
struct bufdesc_ex {
        struct bufdesc desc;
        unsigned long cbd_esc;
        unsigned long cbd_prot;
        unsigned long cbd_bdu;
        unsigned long ts;
        unsigned short res0[4];
};

/*
 *      Define the register access structure.
 */
#define FEC_IEVENT              0x004 /* Interrupt event reg */
#define FEC_IMASK               0x008 /* Interrupt mask reg */
#define FEC_R_DES_ACTIVE        0x010 /* Receive descriptor reg */
#define FEC_X_DES_ACTIVE        0x014 /* Transmit descriptor reg */
#define FEC_ECNTRL              0x024 /* Ethernet control reg */
#define FEC_MII_DATA            0x040 /* MII manage frame reg */
#define FEC_MII_SPEED           0x044 /* MII speed control reg */
#define FEC_MIB_CTRLSTAT        0x064 /* MIB control/status reg */
#define FEC_R_CNTRL             0x084 /* Receive control reg */
#define FEC_X_CNTRL             0x0c4 /* Transmit Control reg */
#define FEC_ADDR_LOW            0x0e4 /* Low 32bits MAC address */
#define FEC_ADDR_HIGH           0x0e8 /* High 16bits MAC address */
#define FEC_OPD                 0x0ec /* Opcode + Pause duration */
#define FEC_HASH_TABLE_HIGH     0x118 /* High 32bits hash table */
#define FEC_HASH_TABLE_LOW      0x11c /* Low 32bits hash table */
#define FEC_GRP_HASH_TABLE_HIGH 0x120 /* High 32bits hash table */
#define FEC_GRP_HASH_TABLE_LOW  0x124 /* Low 32bits hash table */
#define FEC_X_WMRK              0x144 /* FIFO transmit water mark */
#define FEC_R_BOUND             0x14c /* FIFO receive bound reg */
#define FEC_R_FSTART            0x150 /* FIFO receive start reg */
#define FEC_R_DES_START         0x180 /* Receive descriptor ring */
#define FEC_X_DES_START         0x184 /* Transmit descriptor ring */
#define FEC_R_BUFF_SIZE         0x188 /* Maximum receive buff size */
#define FEC_MIIGSK_CFGR         0x300 /* MIIGSK config register */
#define FEC_MIIGSK_ENR          0x308 /* MIIGSK enable register */
```

4.4.6.1 Getting a MAC Address

The MAC address can be set through the kernel command line, kernel device tree DTS
file, OCOTP, or MAC registers set by bootloader, such as U-Boot. The FEC driver uses it
to configure the MAC address for the network device. In general, use kernel command
line in a form of fec.macaddr=0x00,0x04,0x9f,0x01,0x30,0xe0 to set the MAC address.
Due to certain pin conflicts (FEC RMII mode needs to use GPIO_16 or RGMII_TX_CTL
pin as reference clock input/output channel), the one of the both pins cannot connect to
branch lines for other modules use because the branch lines have serious influence on
clock.

4.5 FlexCAN

4.5.1   Introduction

FlexCAN is a communication controller implementing the CAN protocol according to the CAN 2.0B protocol specification.

The CAN protocol was primarily designed to be used as a vehicle serial data bus meeting the specific requirements of this field such as real-time processing, reliable operation in the EMI environment of a vehicle, cost-effectiveness, and required bandwidth. The standard and extended message frames are supported. The maximum message buffer is 64. The driver is a network device driver of PF_CAN protocol family.

For detailed information, see lwn.net/Articles/253425 or Documentation/networking/can.txt in Linux source directory.

The FlexCAN on the i.MX 8QuadMax/8QuadXPlus supports CAN FD protocol.

4.5.1.1   Software Operation

The CAN driver is a network device driver. For the common information on software operation, refer to the documents in the kernel source directory Documentation/networking/can.txt.

The CAN network device driver interface provides a generic interface to setup, configure and monitor CAN network devices. The user can then configure the CAN device, like setting the bit-timing parameters, via the netlink interface using the program "ip" from the "IPROUTE2" utility suite.

Starting and stopping the CAN network device.

A CAN network device is started or stopped as usual with the command "ifconfig canX up/down" or "ip link set canX up/down". Be aware that you *must* define proper bit-timing parameters for real CAN devices before you can start it to avoid error-prone default settings:
  • ip link set canX up type can bitrate 125000

The iproute2 tool also provides some other configuration capbilities for can bus such as bit-timing setting. For details, see kernel doc: Documentation/networking/can.txt

4.5.1.2   Source Code Structure

Table below shows the driver source file available in drivers/net/can

Table 4-8. FlexCAN Driver Files

File	Description
drivers/net/can/flexcan.c	FlexCAN driver

4.5.1.3 Menu Configuration Options

The following kernel configuration options are provided for this module.

- CONFIG_CAN - Build support for PF_CAN protocol family. In menuconfig, this option is available under

  Networking > CAN bus subsystem support.

- CONFIG_CAN_RAW - Build support for Raw CAN protocol. In menuconfig, this option is available under

  Networking > CAN bus subsystem support > Raw CAN Protocol (raw access with CAN-ID filtering).

- CONFIG_CAN_BCM - Build support for Broadcast Manager CAN protocol. In menuconfig, this option is available under

  Networking > CAN bus subsystem support > Broadcast Manager CAN Protocol (with content filtering).

- CONFIG_CAN_VCAN - Build support for Virtual Local CAN interface (also in Ethernet interface). In menuconfig, this option is available under

  Networking > CAN bus subsystem support > CAN Device Driver > Virtual Local CAN Interface (vcan).

- CONFIG_CAN_DEBUG_DEVICES - Build support to produce debug messages to the system log to the driver. In menuconfig, this option is available under

  Networking > CAN bus subsystem support > CAN Device Driver > CAN devices debugging messages.

- CONFIG_CAN_FLEXCAN - Build support for FlexCAN device driver. In menuconfig, this option is available under

  Networking > CAN bus subsystem support > CAN Device Driver > Freescale FlexCAN.

4.6 Inter-IC (I2C)

4.6.1 Introduction

LPI2C is a bidirectional serial bus that provides a simple, efficient method of data exchange, minimizing the interconnection between devices.

The LPI2C driver for Linux OS has two parts:

- Bus driver-low level interface that is used to communicate with the LPI2C bus
- Chip driver-interface between other device drivers and the LPI2C bus driver

The I2C bus driver is a low-level interface that is used to interface with the I2C bus. This driver is invoked by the I2C chip driver and it is not exposed to the user space. The standard Linux kernel contains a core I2C module that is used by the chip driver to access the bus driver to transfer data over the I2C bus. This bus driver supports:

- Compatibility with the I2C bus standard
- Bit rates up to 400 Kbps
- Standard I2C master mode
- Power management features by suspending and resuming I2C.

4.6.2 LPI2C Bus Driver Overview

The LPI2C bus driver is invoked only by the chip driver and is not exposed to the user space. The standard Linux kernel contains a core I2C module that is used by the chip driver to access the LPI2C bus driver to transfer data over the LPI2C bus. The chip driver uses a standard kernel space API that is provided in the Linux kernel to access the core I2C module. The standard I2C kernel functions are documented in the files available under Documentation/i2c in the kernel source tree. This bus driver supports the following features:

- Compatible with the I2C bus standard
- Interrupt-driven, byte-by-byte data transfer
- Standard I2C master mode

4.6.3 I2C Device Driver Overview

The I2C device driver implements all the Linux I2C data structures that are required to communicate with the LPI2C bus driver. It exposes a custom kernel space API to the other device drivers to transfer data to the device that is connected to the LPI2C bus.

Internally, these API functions use the standard I2C kernel space API to call the I2C core module. The I2C core module looks up the LPI2C bus driver and calls the appropriate function in the LPI2C bus driver to transfer data. This driver provides the following functions to other device drivers:

- Read function to read the device registers
- Write function to write to the device registers

4.6.4 Software Operation

The I2C driver for Linux OS has two parts: an I2C bus driver and an I2C chip driver.

4.6.5 I2C Bus Driver Software Operation

The I2C bus driver is described by a structure called i2c_adapter. The most important field in this structure is struct i2c_algorithm *algo. This field is a pointer to the i2c_algorithm structure that describes how data is transferred over the I2C bus. The algorithm structure contains a pointer to a function that is called whenever the I2C chip driver wants to communicate with an I2C device.

During startup, the I2C bus adapter is registered with the I2C core when the driver is loaded. Certain architectures have more than one I2C module. If so, the driver registers separate i2c_adapter structures for each I2C module with the I2C core. These adapters are unregistered (removed) when the driver is unloaded.

During normal communication, it times out and returns an error when the transfer has some error condition, such as NACK is detected. When error condition occurs, I2C driver should stop current transfer.

4.6.6 I2C Device Driver Software Operation

The I2C driver controls an individual I2C device on the I2C bus. A structure, i2c_driver, describes the I2C chip driver. The fields of interest in this structure are flags and attach_adapter. The flags field is set to a value I2C_DF_NOTIFY so that the chip driver can be notified of any new I2C devices, after the driver is loaded. When the I2C bus driver is loaded, this driver stores the i2c_adapter structure associated with this bus driver so that it can use the appropriate methods to transfer data.

4.6.7   Driver Features

The LPI2C driver supports the following features:

- I2C communication protocol
- I2C master mode of operation

NOTE

The LPI2C driver does not support the slave mode.

4.6.8   Source Code Structure

Table below shows the driver source file available in drivers/i2c/busses

Table 4-9.   I2C Driver Files

File	Description
drivers/i2c/busses/imx-lpi2c.c	LPI2C Bus Driver for i.MX 8 and i.MX 8X
drivers/i2c/busses/imx-i2c.c	I2C Bus Driver for i.MX 6, i.MX 7 and i.MX 8M

4.6.9   Menu Configuration Options

Configure the kernel option to enable the module by menuconfig:

For i.MX 6, i.MX 7 and i.MX 8M select Device Drivers > I2C support > I2C Hardware Bus support > IMX I2C interface.

For i.MX 8 and i.MX 8X select Device Drivers > I2C support > I2C Hardware Bus support > IMX Low Power I2C interface.

4.6.10   Programming Interface

The LPI2C device driver can use the standard SMBus interface to read and write the registers of the device connected to the LPI2C bus. For more information, see include/linux/i2c.h.

4.7   Media Local Bus

4.7.1   Introduction

MediaLB is an on-PCB or inter-chip communication bus specifically designed to standardize a common hardware interface and software API library.

This standardization allows an application or multiple applications to access the MOST Network data or to communicate with other applications with minimum effort. MediaLB supports all the *MOST Network data transport* methods: synchronous stream data, asynchronous packet data, and control message data. MediaLB also supports an isochronous data transport method. For detailed information about the MediaLB, see the Media Local Bus Specification.

The MediaLB module implements the Physical Layer and Link Layer of the MediaLB specification, interfacing the i.MX to the MediaLB controller.
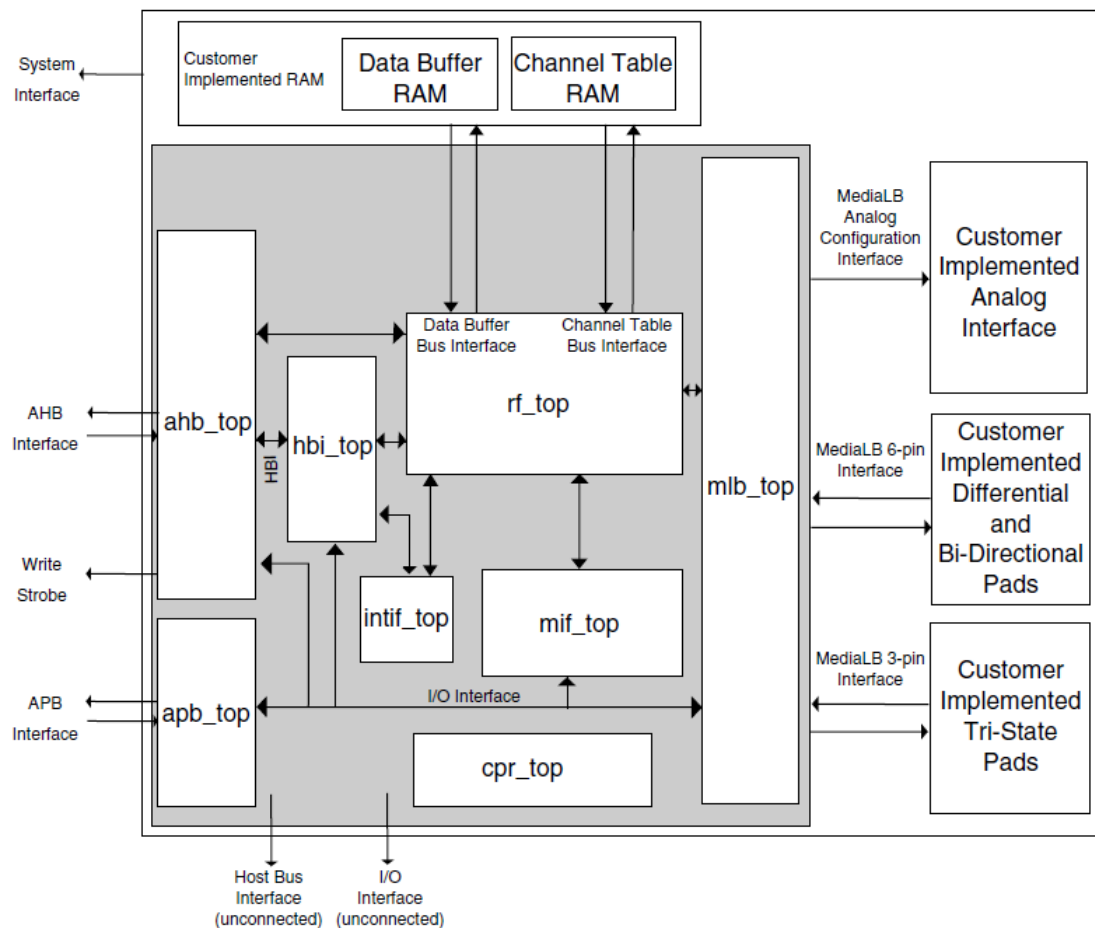
Figure 4-5. MLB Device Top-Level Block Diagram

The MLB implements the 3-pin MediaLB mode and can run at speeds up to 1024Fs. It does not implement MediaLB controller functionality. All MediaLB devices support a set of physical channels for sending data over the MediaLB. Each physical channel is 4 bytes in length (quadlet) and grouped into logical channels with one or more physical channels allocated to each logical channel. These logical channels can be any combination of channel type (synchronous, asynchronous, control, or isochronous) and direction (transmit or receive).

The MLB provides support for up to 64 logical channels and up to 64 physical channels. Each logical channel is referenced using an unique channel address and represents a unidirectional data path between a MediaLB device transmitting the data and the MediaLB device(s) receiving the data.

The supported features are the following.

- Synchronous, asynchronous, control, and isochronous channel.
- Up to 64 logical channels and 64 physical channels running at a maximum speed of 1024Fs.
- Transmission of commands and data and reception of receive status when functioning as the transmitting device associated with a logical channel address.
- Reception of commands and data and transmission as receive status responses when functioning as the receiving device associated with a logical channel address.
- MediaLB lock detection.
- System channel command handling.
- 256Fs, 512Fs and 1024Fs frame rates.
- Asynchronous, control, synchronous, and isochronous channel types.
- The following configurations to MLB device module:
    - Frame rate
    - Device address
    - Channel address
- MLB channel exception get interface. All the channel exceptions are sent and handled by the application.

4.7.2 MLB Driver Overview

The MLB driver is designed as a common Linux OS character driver. It implements one asynchronous and one control channel device with Ping-Pong buffering operation mode. The supported frame rates are 256, 512, and 1024Fs. The MLB driver uses common read/write interfaces to receive/send packets and uses the ioctl interface to configure the MLB device module.

The MLB driver architecture is shown in the figure below.
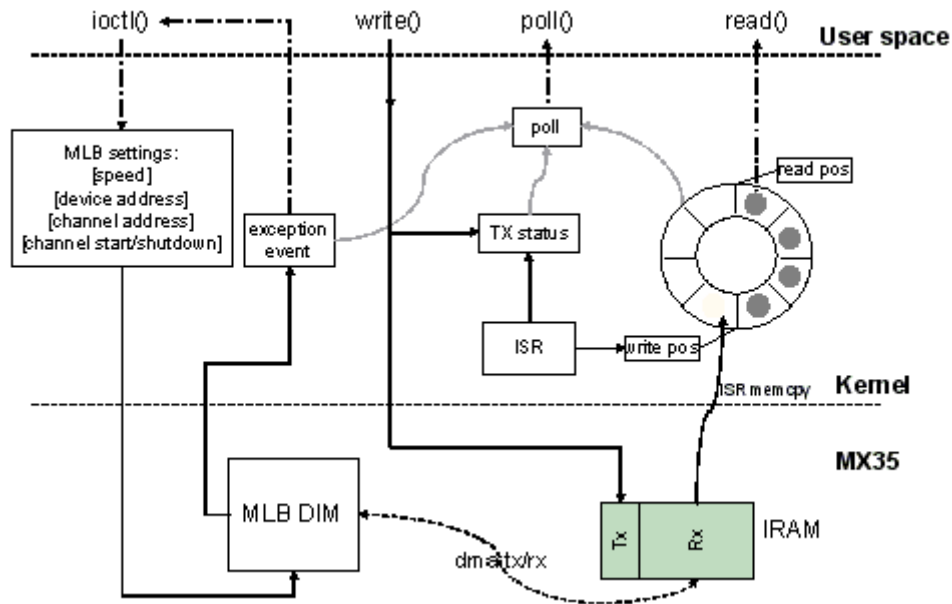
Figure 4-6. MLB Driver Architecture Diagram

The MLB driver creates four minor devices. These four devices support control Tx/Rx channel, asynchronous Tx/Rx channel, synchronous Tx/Rx channel, and isochronous Tx/Rx channel. Their device files are /dev/ctrl, /dev/async, /dev/sync, and /dev/isoc. Each minor device has the same interfaces, and handle both Tx and Rx operation. The following description is for both control and asynchronous device.

The driver uses IRAM as MLB device module Tx/Rx buffer. All the data transmission and reception between module and IRAM is handled by the MLB module DMA. The driver is responsible for configuring the buffer start and end pointer for the MLB module.

For reception, the driver uses a ring buffer to buffer the received packet for read. When a packet arrives, the MLB module puts the received packet into the IRAM Rx buffer, and notifies the driver by interrupt. The driver then copy the packet from the IRAM to one ring buffer node indicated by the write position, and updates the write position with the next empty node. Finally the packet reader application is notified, and it gets one packet from the node indicated by the read position of ring buffer. After the read is completed, it updates the read position with the next available buffer node. There is no received packet in the ring buffer when the read and write position is the same.

For transmission, the driver writes the packet given by the writer application into the IRAM Tx buffer, updates the Tx status and sets MLB device module Tx buffer pointer to start transmission. After transmission completes, the driver is notified by interrupt and updates the Tx status to accept the next packet from the application.

The driver supports NON BLOCK I/O. User applications can poll to check if there are packets or exception events to read, and also they can check if a packet can be sent or not. If there are exception events, the application can call ioctl to get the event. The ioctl also provides the interface to configure the frame rate, device address, and channel address.

4.7.3 Software Operation

The MLB driver provides a common interface to application.

- Packet read/write-BLOCK and NONBLOCK Packet I/O modes are supported. Only one packet can be read or written at once. The minimum read length must be greater or equal to the received packet length, meanwhile the write length must be shorter than 1024 Bytes.
- Polling-The MLB driver provide polling interface which polls for three status, application can use select to get current I/O status:
  - Packet available for read (ready to read)
  - Driver is ready to send next packet (ready to write)
  - Exception event comes (ready to read)
- ioctl-MLB driver provides the following ioctl:

`MLB_SET_FPS`

Argument type: unsigned int

Set frame rate, the argument must be 256, 512 or 1024.

`MLB_GET_VER`

Argument type: unsigned long

Get MLB device module version, which is 0x02000202 by default on the i.MX35.

`MLB_SET_DEVADDR`

Argument type: unsigned char

Set MLB device address, which is used by the system channel MlbScan command.

`MLB_CHAN_SETADDR`

Argument type: unsigned int

Set the corresponding channel address [8:1] bits. This ioctl combines both tx and rx channel address, the argument format is: tx_ca[8:1] << 16 | rx_ca[8:1].

`MLB_CHAN_STARTUP`

Startup the corresponding type of channel for transmit and reception.

`MLB_CHAN_SHUTDOWN`

Shutdown the corresponding type of channel.

```
MLB_CHAN_GETEVENT
```

Argument type: unsigned long

Get exception event from MLB device module, the event is defined as a set of enumeration:

```
MLB_EVT_TX_PROTO_ERR_CUR
MLB_EVT_TX_BRK_DETECT_CUR
MLB_EVT_RX_PROTO_ERR_CUR
MLB_EVT_RX_BRK_DETECT_CUR
```

4.7.4 Source Code Structure

Table below lists the MLB Driver source files.

Table 4-10.   MLB Driver Source Files

File	Description
drivers/mxc/mlb/mxc_mlb.c	Source file for MLB driver
include/linux/mxc_mlb.h	Include file for MLB driver

4.7.5 Menu Configuration Options

In menu configuration enable the following module:

Device Drivers > MXC support drivers > MXC Media Local Bus Driver > MLB support.

4.8 PCI Express Root Complex

4.8.1 Terminology and Conventions

The following terminologies and conventions are used in this document:

- Bridge

  A Function that virtually or actually connects a PCI/PCI-X segment or PCI Express Port with an internal component interconnect or with another PCI/PCI-X bus segment or PCI Express Port.

- Downstream

- The relative position of an interconnect/System Element (Port/component) that is farther from the Root Complex. The Ports on a Switch that are not the Upstream Port are Downstream Ports. All Ports on a Root Complex are Downstream Ports. The Downstream component on a Link is the component farther from the Root Complex.
- A direction of information flow where the information is flowing away from the Root Complex.

- Endpoint

  One of several defined System Elements. A Function that has a Type 00h Configuration Space header.

- Host

  The entity comprising of one (or more) Central Processing Unit(s) (CPU) and resources, such as Memory (RAM) that can be shared across multiple PCIe nodes connected through a Root Complex.

- Lane

  A set of differential signal pairs, one pair for transmission and one pair for reception.

- Link

  The collection of two Ports and their interconnecting Lanes. A Link is a dual simplex communications path between two components.

- PCIe Fabric

  A topology comprised of various PCI Express nodes, also referred as devices. A device in the fabric can be Root Complex, Endpoint, PCIe-PCI/PCI-X Bridge or a Switch.

- Port
    - Logically, an interface between a component and a PCI Express Link.
    - Physically, a group of Transmitters and Receivers located on the same chip that define a Link.
- Root Complex

  RC A defined System Element that includes a Host Bridge, zero or more Root Complex Integrated Endpoints, zero or more Root Complex Event Collectors, and one or more Root Ports.

- Root Port

A PCI Express Port on a Root Complex that maps a portion of the Hierarchy through an associated virtual PCI-PCI Bridge.

- Upstream
    - The relative position of an interconnect/System Element (Port/component) that is closer to the Root Complex. The Port on a Switch that is closest topologically to the Root Complex is the Upstream Port. The Port on a component that contains only Endpoint or Bridge Functions is an Upstream Port. The Upstream component on a Link is the component closer to the Root Complex.

Any element of the fabric which is relatively closer towards RC is treated as 'Upstream'. All PCIe Endpoint ports (including termination points for bridges) and Switch ports, which are closer to RC are called Upstream Ports on that device. An Upstream Flow is the communication moving towards RC.

4.8.2 PCIe Topology on i.MX

There is one PCIe port on the i.MX. Currently, only the RC mode is enabled in the Linux BSP.

The following figure describes the diagram of the PCIe RC port on i.MX.
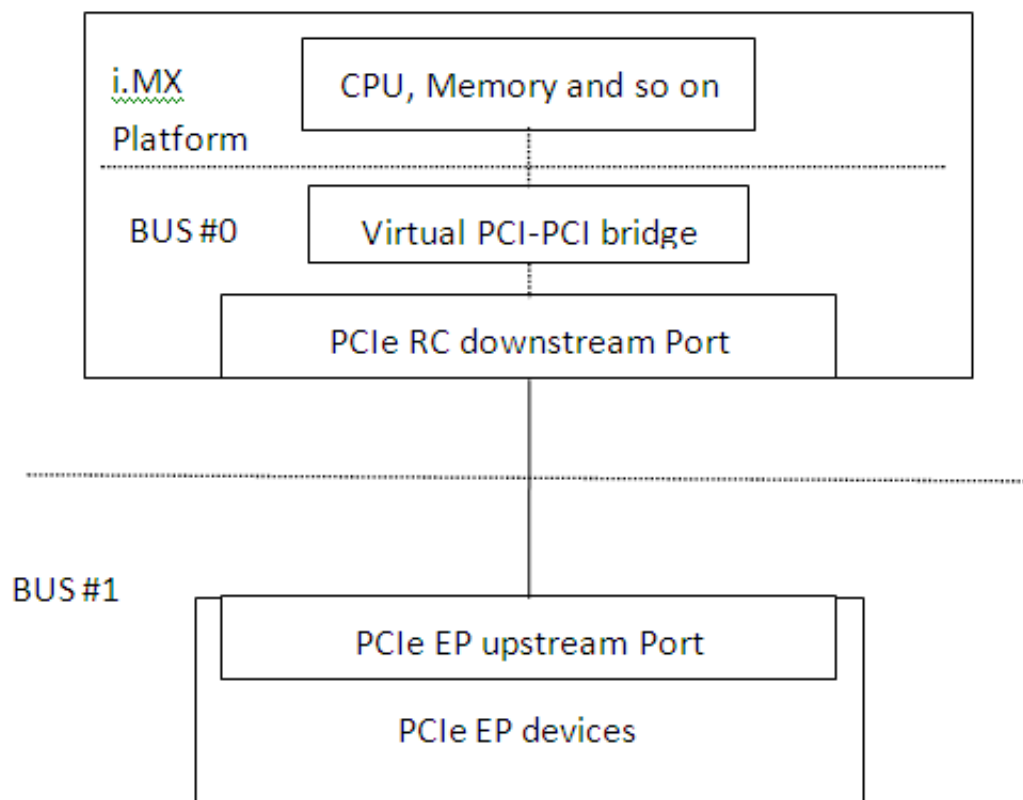
Figure 4-7. diagram of the PCIe RC port on i.MX

PCI Enumeration Mapping

As PCI Express is point to point topology, to maintain compatibility with legacy PCI Bus - Device notion used for Software Enumeration, we introduce the following concepts which allow various nodes and their internals to be identified (e.g., PCIe Switches) in terms of PCI devices/functions:

- Host Bridge: A bridge, integrated into RC to have PCI compatible connection to Host. The PCI side of this bridge is Bus #0 always. This means, the device on this bus will be the host itself.
- Virtual PCI-PCI Bridge: Each PCI Express port which is part of RC or a Switch is treated as a virtual PCI-PCI bridge. This means each port has a primary and secondary PCI bus and the downstream is mapped into the remote configuration space.
- Root port associated virtual bridge has Bus #0 on the primary side with secondary bus on the downstream.
- Each PCIe Switch is viewed as collection of as many virtual PCI-PCI bridges as number of downstream ports, connected to a virtual PCI bus which is actually secondary bus of another PCI-PCI bridge forming the upstream port of the switch.
- The upstream port of each EP can either be part of the secondary bus segment of virtual PCI-PCI Bridge representing downstream port of a switch or of the root port.

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

4.8.3 Features

The following are the various features supported by i.MX as a PCI Express Root Complex driver.

- Express Base Specification Revision 2.0-compliant.
- Gen2 operation with x1 link supporting 5 GT/s raw transfer rate in single direction.
- Support Legacy Interrupts (INTx) and MSI.
- Max_Payload_Size size (128 bytes).
- It fits into Linux PCI Bus framework to provide PCI compatible software enumeration support.
- In addition, it provides interface to Endpoint Drivers to access the respective devices detected downstream.
- The same interface can be used by the PCI Express Port Bus Driver framework in Linux OS to handle AER, ASP, and so on.
- Interrupt handling facility for EP drivers either as Legacy Interrupts (INTx).
- Access to EP I/O BARs through generic I/O accessories in Linux PCI subsystem.
- Seamless handling of PCIe errors.
- Supports the L0, L0s, L1, and L1.1 ASPM power management.

4.8.4 Linux OS PCI Subsystem and RC driver

In Linux OS, the PCI implementation can roughly be divided into the following main components: PCI BIOS architecture-specific Linux OS implementation, Host Controller (RC) Module, and Core.

- PCI BIOS Architecture-specific Linux OS implementation to kick off PCI bus initialization. It interfaces with PCI Host Controller code as well as the PCI Core to perform bus enumeration and allocation of resources such as memory and interrupts. The successful completion of BIOS execution assures that all the PCI devices in the system are assigned parts of available PCI resources and their respective drivers (referred as Slave Drivers). PCI can take control of them using the facilities provided by PCI Core. It is possible to skip resource allocation (if they were assigned before Linux OS was booted, for example PC scenario).
- Host Controller (RC) Module handles hardware (SoC + Board) specific initialization and configuration and it invokes PCI BIOS. It should provide callback functions for BIOS as well as PCI Core, which will be called during PCI system initialization and accessing PCI bus for configuration cycles. It provides resources information for available memory/IO space, INTx interrupt lines, MSI. It should also facilitate IO

space access (as supported) through in _x_ () out _x_ () You may need to provide indirect memory access (if supported by h/w) through read _x_ () write _x_ ().

- Core creates and initializes the data structure tree for bus devices as well as bridges in the system, handles bus/device numberings, creates device entries and proc/sysfs information, provides services for BIOS and slave drivers and provides hot plug support (optional/as supported by h/w). It targets (EP) driver interface query and initializes corresponding devices found during enumeration. It also provides MSI interrupt handling framework and PCI express port bus support. It provides Hot-Plug support (if supported), advanced error reporting support, power management event support, and virtual Channel support to run on PCI express ports (if supported).

4.8.5 PCIe Driver Source Files

Table 4-11. Source Files

File	Description
drivers/pci/host/pci-imx6.c	MX6 PCIe source

4.8.6 System Resource: Memory Layout
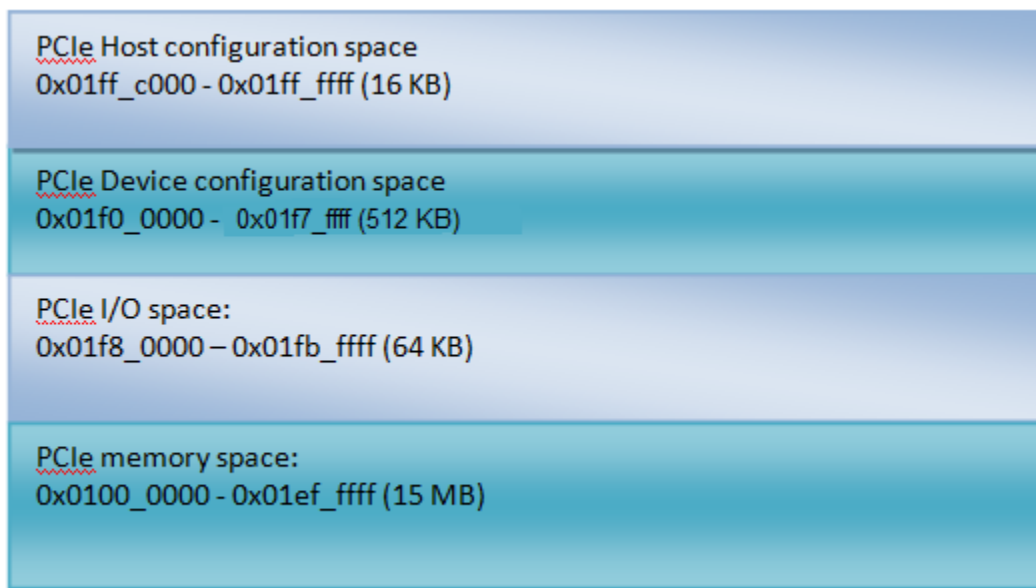
Figure 4-8. Memory Layout (i.MX 6Quad/6DualLite/6Solo)
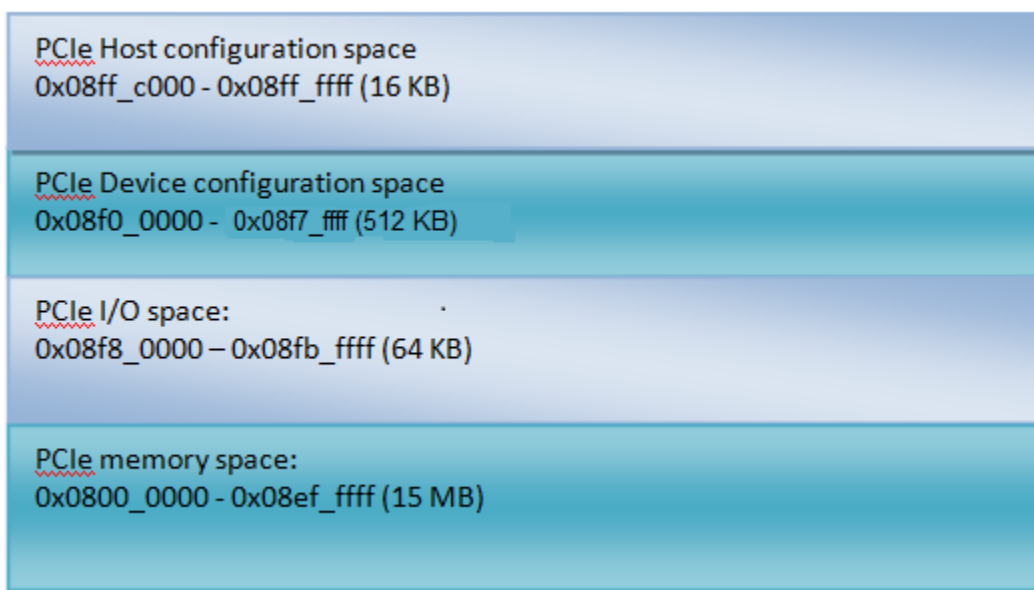
Figure 4-9. Memory Layout (i.MX 6SoloX)
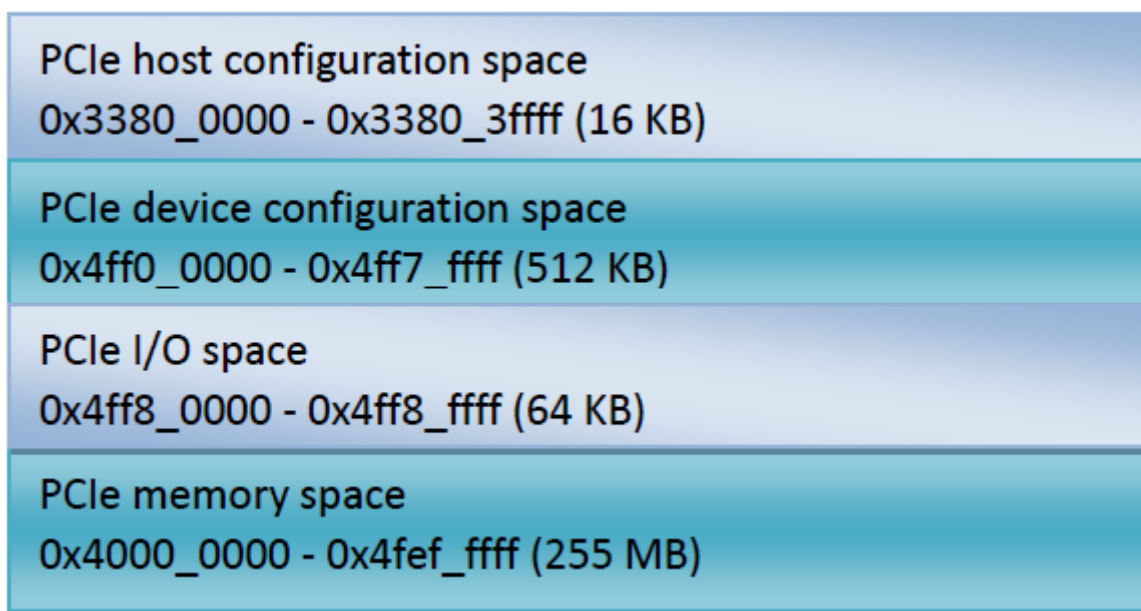
Figure 4-10. Memory Layout (i.MX 7Dual)

- IO and memory spaces are two address spaces used by the devices to communicate with their device driver running in the Linux kernel on CPU.
- The upper 16 KB PCIe host configuration space.
  - This memory segment is used to map the configuration space of PCIe RC. SW can access PCIe RC core configuration space through the DBI interface.
- PCIe device configuration space.
  - Used to map the configuration spaces of PCIe EP devices that are inserted to the RC downstream port.

i.MX 8QuadMax/8QuadXPlus:

i.MX 8QuadMax has both PCIeA and PCIeB, while i.MX 8QuadXPlus has only PCIeB.

- PCIeA

    - PCIe host configuration space: 0x5f00_0000 – 0x5f00_ffff (64K bytes)
    - PCIe device configuration space: 0x6ff0_0000 – 0x6ff7_ffff (512K bytes)
    - PCIe IO space: 0x6ff8_0000 – 0x6ff8_ffff (64K bytes)
    - PCIe memory space: 0x6000_0000 – 0x6fef_ffff (255M bytes)

- PCIeB

    - PCIe host configuration space: 0x5f01_0000 – 0x5f01_ffff (64K bytes)
    - PCIe device configuration space: 0x7ff0_0000 – 0x7ff7_ffff (512K bytes)
    - PCIe IO space: 0x7ff8_0000 – 0x7ff8_ffff (64K bytes)
    - PCIe memory space: 0x7000_0000 – 0x7fef_ffff (255M bytes)

i.MX 8M Quad:

- PCIe0

    - PCIe host configuration space: 0x3380_0000 – 0x33bf_ffff (4Mbytes)
    - PCIe device configuration space: 0x1ff0_0000 – 0x1ff7_ffff (512K bytes)
    - PCIe IO space: 0x1ff8_0000 – 0x1ff8_ffff (64K bytes)
    - PCIe memory space: 0x1800_0000 – 0x1fef_ffff (127M bytes)

- PCIE1

    - PCIe host configuration space: 0x33c0_0000 – 0x33ff_ffff (4Mbytes)
    - PCIe device configuration space: 0x27f0_0000 – 0x27f7_ffff (512K bytes)
    - PCIe IO space: 0x27f8_0000 – 0x27f8_ffff (64K bytes)
    - PCIe memory space: 0x2000_0000 – 0x27ef_ffff (127M bytes)

4.8.7 System Resource: Interrupt lines

i.MX Root Complex driver uses interrupt line 152 for MSI INT on i.MX 6 platforms, and 154 for MSI INT on i.MX 7Dual platforms.

4.9 USB

4.9.1   Introduction

The universal serial bus (USB) driver implements a standard Linux driver interface to the CHIPIDEA USB-HS OTG controller.

The USB provides a universal link that can be used across a wide range of PC-to-peripheral interconnects. It supports plug-and-play, port expansion, and any new USB peripheral that uses the same type of port.

The CHIPIDEA USB controller is Enhanced Host Controller Interface (EHCI)-compliant. This USB driver has the following features:

- High-speed OTG core supported
- High-speed Host Only core (Host1), high-speed, full speed, and low devices are supported
- High-speed Inter-Chip core (Host2 & Host3)
- High-speed Host Only core (OTG2), high-speed, full speed, and low devices are supported. A USB2Pci bridge is connected to OTG2 by default. Therefore, users may not be able to connect other USB devices on this port.
- High-speed Inter-Chip core (Host2)
- Host mode-Supports HID (Human Interface Devices), MSC (Mass Storage Class)
- Peripheral mode-Supports MSC, and CDC (Communication Devices Class) drivers, which include Ethernet and serial support
- Embedded DMA controller

4.9.2   Architectural Overview

The USB host system is composed of a number of hardware and software layers.

The figure below shows a conceptual block diagram of the building block layers in a host system that support USB 2.0.
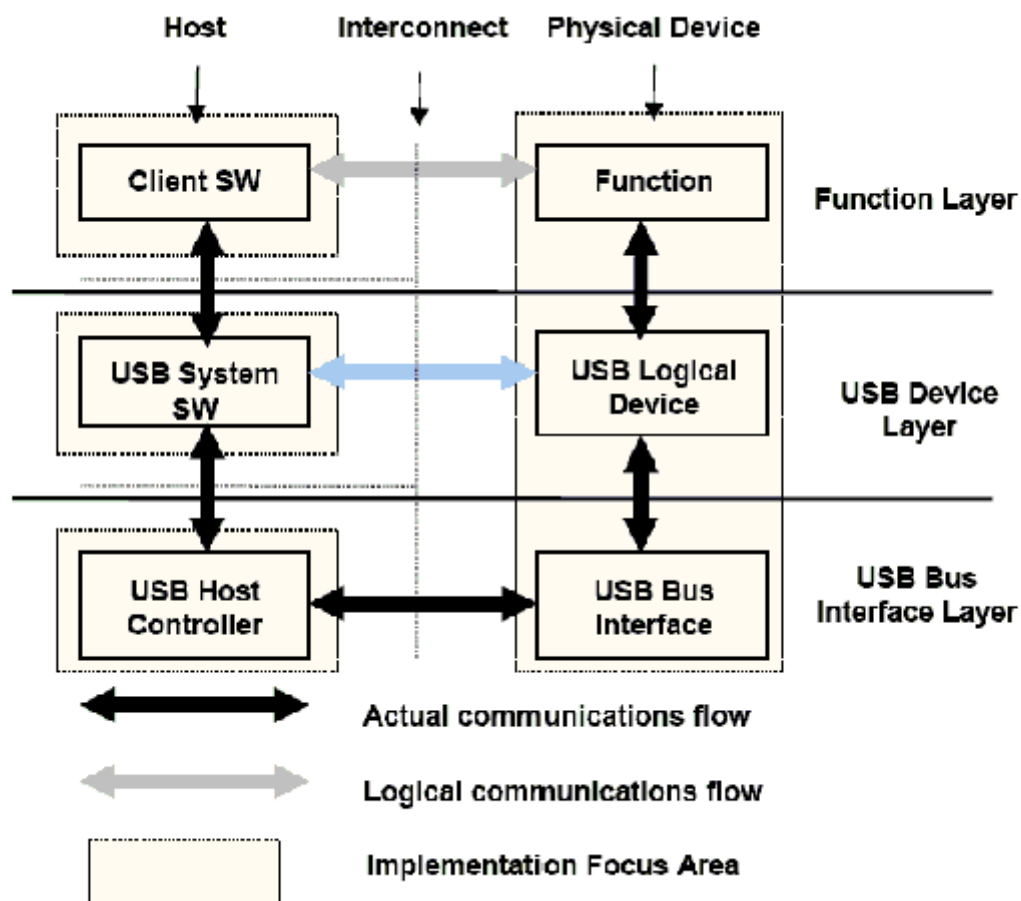
Figure 4-11. USB Block Diagram

4.9.3 Hardware Operation

For information on hardware operations, refer to the EHCI spec.ehci-r10.pdf.

The spec is available at Enhanced Host Controller Interface for USB 2.0: Specification

4.9.4 Software Operation

The Linux OS contains a USB driver, which implements the USB protocols. For the USB host, it only implements the hardware specified initialization functions. For the USB peripheral, it implements the gadget framework. For OTG, ID dynamic switch host/device modes are supported. Currently, the runtime suspend for USB is supported, that is to say when the USB is not in use (both for host and peripheral mode), the USB will enter low power mode.

4.9.5 Source Code Structure

The table below describes the USB source in drivers/usb.

Table 4-12.   Chipidea USB Driver Files

File	Description
drivers/usb/chipidea/core.c	Chipidea IP core driver
drivers/usb/chipidea/udc.c	Chipidea peripheral driver
drivers/usb/chipidea/host.c	Chipidea host driver
drivers/usb/chipidea/otg.c	Chipidea OTG driver
drivers/usb/chipidea/otg_fsm.c	Chipidea OTG HNP and SRP driver
drivers/usb/chipidea/ci_hdrc_imx.c	i.MX glue layer
drivers/usb/chipidea/usbmisc_imx.c	i.MX SoC abstract layer
drivers/usb/phy/phy-mxs-usb.c	i.MX 6 USB physical driver

4.9.6 Menu Configuration Options

In menu configuration enable the following modules.

Device Drivers > [*] USB support > EHCI HCD (USB 2.0) support and ChipIdea Highspeed Dual Role Controller [*] USB Physical Layer drivers --->

Device Drivers > USB Physical Layer drivers > Freescale MXS USB PHY support

Device Drivers > USB Gadget Support

1. CONFIG_USB-Build Support for Host-side USB
2. CONFIG_USB_EHCI_HCD EHCI HCD (USB 2.0) support

   Default y

3. CONFIG_USB_CHIPIDEA- ChipIdea high-speed Dual Role Controller

   Default y

4. CONFIG_USB_CHIPIDEA_UDC - ChipIdea device controller

   Default y

5. CONFIG_USB_CHIPIDEA_HOST - ChipIdea host controller

   Default y

6. CONFIG_USB_GADGET - USB Gadget Support

Default y

7. CONFIG_USB_MXS_PHY - Freescale MXS USB PHY support

Default y

4.9.7   USB Wakeup Usage

The following example is for the OTG port and the first EHCI device.

Controller wakeup setting, after the following settings, the VBUS and ID will be wakeup source.

```
echo enabled > /sys/bus/platform/devices/20c9000.usbphy/power/wakeup
echo enabled > /sys/bus/platform/devices/2184000.usb/power/wakeup
echo enabled > /sys/bus/platform/devices/ci_hdrc.0/power/wakeup
```

EHCI wakeup setting, after the following settings, the host will have wakeup ability, such as remote wakeup and connect/disconnect wakeup

```
echo enabled > /sys/bus/usb/devices/usb1/power/wakeup
echo enabled > /sys/bus/usb/devices/1-1/power/wakeup
```

NOTE
When the OTG mode switches from the host to the device, it will delete the EHCI wakeup, and the user needs to set it again before the system suspending.

4.9.8   How to Close the USB Child Device Power

The following code string outlines how to close the USB child device power:

```
echo auto > /sys/bus/usb/devices/1-1/power/control
echo auto > /sys/bus/usb/devices/1-1.1/power/control (If there is a hub at USB device)
```

4.9.9   Changing the Controller Operation Mode

To change the default settings, the use can modify the DTS file as follows:

```
            dr_mode =  "host"  /* Set controller as gadget-only mode */
            dr_mode =  "peripheral" /* Set controller as host-only mode */
            dr_mode =  "otg" /* Set controller as otg mode */
```

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

4.9.10 Loadable Module Support

The modprobe utility will automatically load the modules which have dependency among all modules.

The loading command is as follows:

```
modprobe phy_mxs_usb
modprobe ci_hdrc_imx
```

The unloading command is as follows:

```
modprobe -r ci_hdrc_imx
modprobe -r phy_mxs_usb
```

4.9.11 USB Charger Detection

i.MX SoC has USB charger detection ability, but it has no charging ability. The user can use the `/sys` entry to know the USB charger type, charging current, and whether the charger exists, as shown in the following three lines:

```
cat /sys/class/power_supply/imx6_usb_charger/type
cat /sys/class/power_supply/imx6_usb_charger/current_max
cat /sys/class/power_supply/imx6_usb_charger/present
```

Currently, the i.MX 6 Sabre-SD board does not support the USB charger detection function. i.MX 6 Sabre-Auto and i.MX 6SoloLite EVK support the function.

4.9.12 Embeded Host Certification

4.9.12.1 Adding TPL-Support Property

To pass embeded host USB certification, "tpl-support" should be added in DTS to enable Targeted Peripheral List (TPL). For example, to enable TPL on the Host port of i.MX 6UltraLite EVK board (imx6ul-14x14-evk.dts):

```
&usbotg2 {
        dr_mode = "host";
        disable-over-current;
        tpl-support;
        status = "okay";
};
```

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

4.9.12.2 VBUS Control

The VBUS should be kept off until the Linux USB host function is ready. For example, on the i.MX 6UltraLite EVK board, because the pin is multiplexed with the touch function, you need to rework the board to make the GPIO (GPIO1_IO02) selected for VBUScontrol.

Disable the touch function in its DTS file (imx6ul-14x14-evk.dts) as follows:

```
&tsc {
        pinctrl-names = "default";
        pinctrl-0 = <&pinctrl_tsc>;
        xnur-gpio = <&gpio1 3 0>;
        measure_delay_time = <0xffff>;
        pre_charge_time = <0xfff>;
        status = "disabled";
};
```

Add VBUS GPIO pinctrl and its regulator node:

```
pinctrl_usb_otg2: usbotg2grp {
            fsl,pins = <
                MX6UL_PAD_GPIO1_IO02__GPIO1_IO02          0xb0
            >;
   };

reg_usb_otg2_vbus: regulator@2 {
            compatible = "regulator-fixed";
            reg = <2>;
            pinctrl-names = "default";
            pinctrl-0 = <&pinctrl_usb_otg2>;
            regulator-name = "usb_otg2_vbus";
            regulator-min-microvolt = <5000000>;
            regulator-max-microvolt = <5000000>;
            gpio = <&gpio1 2 GPIO_ACTIVE_HIGH>;
            enable-active-high;
};

&usbotg2 {
    vbus-supply = <&reg_usb_otg2_vbus>;
    dr_mode = "host";
    disable-over-current;
    tpl-support;
    status = "okay";
};
```

4.10 USB3

4.10.1 Introduction

For i.MX 8 and i.MX 8X families, a super-speed USB IP from Cadence is provided supporting USB 3.0 which includes a new transfer rate referred to as Super Speed (SS) USB with higher transfer rates and significantly faster than the USB 2.0 standard.

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

The supported features the following.

- Host mode is implemented with a Linux OS standard XHCI driver with super-speed supported and tested.
- For Device Mode only single queue is supported. Mass storage, ether, and serial are supported.

4.10.2   Source Code Structure

Table 4-13.   USB3 Driver Source Files

File	Description
drivers/usb3/cdns3/cdns3-nxp-reg-def.h	Register definitions
drivers/usb3/cdns3/core.c	USB3 core driver
drivers/usb3/cdns3/core.h	USB3 Core header
drivers/usb3/cdns3/dev-regs-macro.h	USB3 Macros
drivers/usb3/cdns3/dev-regs-map.h	USB3 Register mapping
drivers/usb3/cdns3/gadget.c	USB3 Gadget
drivers/usb3/cdns3/gadget.h	USB3 Gadget header
drivers/usb3/cdns3/gadget-export.h	USB3 Gadget Export header
drivers/usb3/cdns3/host.c	USB3 Host
drivers/usb3/cdns3/host-export.h	USB3 Host Export header
drivers/usb3/cdns3/io.h	USB3 IO

4.11   Low Power Universal Asynchronous Receiver/Transmitter (LPUART)

4.11.1   Introduction

The low-level UART driver interfaces the Linux serial driver API to all the UART ports.

It has the following features:

- Interrupt-driven and eDMA-driven transmit/receive of characters
- Standard Linux baud rates up to 4 Mbps
- Transmit and receive characters with 7-bit, 8-bit, 9-bit, or 10-bit character length
- Transmits one or two stop bits
- Supports TIOCMGET IOCTL to read the modem control lines. Only supports the constants TIOCM_CTS and TIOCM_CAR, plus TIOCM_RI in DTE mode only

- Supports TIOCMSET IOCTL to set the modem control lines. Supports the constants TIOCM_RTS and TIOCM_DTR only
- Odd and even parity
- XON/XOFF software flow control. Serial communication using software flow control is reliable when communication speeds are not too high and the probability of buffer overruns is minimal
- CTS/RTS hardware flow control-both interrupt-driven software-controlled hardware flow and hardware-driven hardware-controlled flow
- Send and receive break characters through the standard Linux serial API
- Recognizes frame and parity errors
- Ability to ignore characters with break, parity and frame errors
- Get and set UART port information through the TIOCGSSERIAL and TIOCSSERIAL TTY IOCTL. Some programs like setserial and dip use this feature to make sure that the baud rate was set properly and to get general information on the device. The UART type should be set to 52 as defined in the serial_core.h header file.
- Power management feature by suspending and resuming the UART ports
- Standard TTY layer IOCTL calls

All the UART ports can be accessed from the device files /dev/ttyLP0 to /dev/ttyLP1.

4.11.2   Hardware Operation

To determine the number of UART modules available on the device see the Applications Processor Reference Manual associated with SoC.

Each UART hardware port is capable of standard RS-232 serial communication.

Each UART contains a 64-byte transmitter FIFO and a 32-half-word deep receiver FIFO. Each UART also supports a variety of maskable interrupts when the data level in each FIFO reaches a programmed threshold level and when there is a change in state in the modem signals.

4.11.3   Software Operation

The Linux OS contains a core UART driver that manages many of the serial operations that are common across UART drivers for various platforms.

The low-level UART driver is responsible for supplying information such as the UART port information and a set of control functions to the core UART driver. These functions are implemented as a low-level interface between the Linux OS and the UART hardware.

They cannot be called from other drivers or from a user application. The control functions used to control the hardware are passed to the core driver through a structure called uart_ops, and the port information is passed through a structure called uart_port. The low level driver is also responsible for handling the various interrupts for the UART ports, and providing console support if necessary.

Each UART can be configured to use DMA for the data transfer by enabling the DMA channel in the DTS file.

The driver requests two DMA channels for the UARTs that need DMA transfer. On a receive transaction, the driver copies the data from the DMA receive buffer to the TTY Flip Buffer.

While using DMA to transmit, the driver copies the data from the UART transmit buffer to the DMA transmit buffer and sends this buffer to the DMA system. For more information, see the Linux documentation on the serial driver in the kernel source tree.

4.11.4 Driver Features

The UART driver supports the following features:

- Baud rates up to 4 Mbps
- Recognizes frame and parity errors only in interrupt-driven mode; does not recognize these errors in DMA-driven mode
- Sends, receives, and appropriately handles break characters
- Recognizes the modem control signals
- Ignores characters with frame, parity, and break errors if requested to do so
- Implements support for hardware flow control
- Get and set the UART port information; certain flow control count information is not available in hardware-driven hardware flow control mode
- Power management
- Interrupt-driven and DMA-driven data transfer

4.11.5 Source Code Structure

Table below shows the UART driver source files.

Table 4-14. UART Driver Files

File	Description
drivers/tty/serial/fsl_lpuart.c	LP UART driver

For the i.MX 8, i.MX 8X and i.MX 8M configuration options are specified in the device trees located in arch/arm64/boot/dts directory.

4.11.6 Menu Configuration Options

The UART driver is enabled by default.

The menu configuration option is located at:

Device Drivers > Character devices > Serial drivers > Freescale LPUART serial port support [*] Console on Freescale LPUART serial port

4.11.7 Programming Interface

The UART driver implements all the methods required by the Linux serial API to interface with the UART port and provides a set of control methods to the Linux core UART driver. For more information about the methods implemented in the driver, see the API document.

4.11.8 Interrupt Requirements

The UART driver interface generates only one interrupt.

The status is used to determine which kinds of interrupt occurs, such as RX or TX.

4.12 Bluetooth

4.12.1 Bluetooth Wireless Technology Introduction

Bluetooth technology is low-cost, low-power, short-range wireless technology. It was designed as a replacement for cables and other short-range technologies like IrDA. Bluetooth wireless technology operates in personal area range that typically extends up to 10 meters. For more information about Bluetooth wireless technology, see www.bluetooth.com/.

For i.MX, bluetooth is supported with multiple vendors listed below:

- i.MX 7Dual supports WiFi chip on the board.
- i.MX 6 supports WiFi chip Murata module based on BCM 9339.

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

- i.MX 7ULP supports Wi-Fi chip with Murata 1PJ module based on Qualcomm QCA9377-3. The QCA9377-3 is a single-die wireless local area network (WLAN) and Bluetooth (BT) combination solution to support 1 × 1 IEEE 802.11a/b/g/n/ac WLAN standards and BT 4.1 + HS, enabling seamless integration of LAN/BT and lowenergy technology.
- i.MX 8M Quad supports Wi-Fi chip with Murata 1CQ module based on Qualcomm QCA6174. The QCA6174 is a single-die wireless local area network (WLAN) and Bluetooth combo solution to support 2 x 2 multi-user multiple input, multiple output (MU-MIMO) with two spatial streams IEEE802.11 a/b/g/n/ac WLAN standards and Bluetooth 4.2 + HS, designed to deliver superior integration of WLAN/Bluetooth and low energy technology.
- i.MX 8M Mini supports Wi-Fi chip with Murata 1PJ module based on Qualcomm QCA9377.

QCA supports WiFi through firmware provided with firmware-imx software and BCM supports firmware provided with imx-firmware at imx-firmware github

4.12.2   Bluetooth Driver Overview

i.MX uses the open source Bluetooth driver. The Bluetooth software is divided into four parts as follows:
- 4-wire UART and TTY driver: It is the communication interface with the Bluetooth module.
- Bluetooth HCI device driver: UART (H4) is a serial protocol for communication between the Bluetooth device and host. This protocol is required for most Bluetooth devices with the UART interface.
- Bluetooth kernel stack: Bluetooth framework and protocols implementation.
- Bluetooth user stack: Supplies several user-space utilities and integrate many profiles for use cases.

4.12.3   Bluetooth Driver Files

The Bluetooth driver source files are available in the kernel source directory.
- Bluetooth HCI device driver:
    - drivers/bluetooth/hci_h4.c
    - drivers/bluetooth/hci_ldisc.c
- Bluetooth kernel stack:
    - net/bluetooth/*

4.12.4 Bluetooth Stack

BlueZ is the official Linux standard Bluetooth protocol stack, it is the latest version of 5.x and it is a Bluetooth stack for Linux kernel-based family of operating systems. Its goal is to program an implementation of the Bluetooth wireless standards specifications for Linux. To use Linux Bluetooth subsystem, you need several user-space utilities like hciconfig and bluetoothd. These utilities and updates to Bluetooth kernel modules are provided in the BlueZ packages. For more information, see www.bluez.org/.

BlueZ source code are available in the git: git://git.kernel.org/pub/scm/bluetooth/bluez.git. The current BSP package tests pass with BlueZ 5.49.

4.12.5 Menu Configuration Options

The following Linux kernel configuration option is provided for this module:
- UART interface:
  - CONFIG_SERIAL_IMX
  - CONFIG_TTY
- HCI interface:
  - CONFIG_BT_HCIUART
  - CONFIG_BT_HCIUART_H4
- Bluetooth Stack:
  - CONFIG_BT
  - CONFIG_BT_RFCOMM
  - CONFIG_BT_RFCOMM_TTY
  - CONFIG_BT_BNEP
  - CONFIG_BT_BNEP_MC_FILTER
  - CONFIG_BT_BNEP_PROTO_FILTER
  - CONFIG_BT_HIDP

4.13 Wi-Fi

4.13.1 Introduction

WiFi on i.MX is supported by Qualcomm QCA and Broadcom BCRM

- For i.MX 7Dual the Broadcom BCRM is supported on the reference board
- For i.MX 6 i.MX 8 the Broadcom BCRM is supported with a Murata plugin module.

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

- For i.MX 8 and i.MX 8X use the M.2 interface to connect with the QCA 1CQ Wi-Fi module.
- For i.MX 7ULP had the QCA6174 chip on the reference board.
- For i.MX 8M the QCA6174 chip is supported on the reference board.
- For i.MX 8M Mini, the QCA9377 chip is supported on the reference board.

The QCA6174 is a single-die wireless local area network (WLAN) and Bluetooth $^®$ combo solution to support 2 x 2 multi-user multiple input, multiple output (MU-MIMO) with two spatial streams IEEE802.11 a/b/g/n/ac WLAN standards and Bluetooth 4.2 + HS, designed to deliver superior integration of WLAN/Bluetooth and low energy technology.

4.13.2   Software Operation

The BSP uses the QCA CLD LEA1.0 Wi-Fi driver, which is released by Qualcomm, and the firmware and binaries are limited to LA_OPT_BASE_LICENSE and Qualcomm Atheros License.

4.13.3   Driver features

The QCA CLD is a CFG80211 driver, which supports both the station and AP mode of operation.

The driver requires firmware that runs on the chip's network processor. The following directory designates the firmware location in rootfs: /lib/firmware/.

4.13.4   Source Code Structure

The QCA CLD driver source files are available at codeaurora.org.

4.13.5   Menu Configuration Options

The following Linux kernel configuration option is provided for this module:
- CONFIG_MAC80211=y
- COCONFIG_NL80211_TESTMODE=y
- CONFIG_CFG80211_WEXT=y
- CONFIG_HOSTAP=y
- CONFIG_CFG80211_INTERNAL_REGDB=y

4.13.6 Device Tree Binding

For device tree, the ATH10K driver requires the following nodes to be defined in the device tree. For example,

```
&pcie0{
        pinctrl-names = "default";
        pinctrl-0 = <&pinctrl_pcie0>;
        clkreq-gpio = <&gpio5 20 GPIO_ACTIVE_LOW>;
        disable-gpio = <&gpio5 29 GPIO_ACTIVE_LOW>;
        reset-gpio = <&gpio5 28 GPIO_ACTIVE_LOW>;
        ext_osc = <1>;
        hard-wired = <1>;
        status = "okay";
};

regulators {
    compatible = "simple-bus";
    #address-cells = <1>;
    #size-cells = <0>;

    epdev_on: fixedregulator@100 {
        compatible = "regulator-fixed";
        regulator-min-microvolt = <3300000>;
        regulator-max-microvolt = <3300000>;
        regulator-name = "epdev_on";
        gpio = <&gpio4 9 0>;
        enable-active-high;
    };
}
```

4.13.7 Configuring WLAN from User Space

4.13.7.1 Connecting AP in Station Mode

The following command group is used to connect WLAN to a given SSID.

```
head -n 4 /etc/wpa_supplicant.conf  > /etc/wpa_supplicant.conf.tmp
wpa_passphrase ssid password >> /etc/wpa_supplicant.conf.tmp
mv /etc/wpa_supplicant.conf  /etc/wpa_supplicant.conf.bak
mv /etc/wpa_supplicant.conf.tmp  /etc/wpa_supplicant.conf
wpa_supplicant -B -i wlp1s0 -c /etc/wpa_supplicant.conf -D nl80211
```

Here is an example of wpa_supplicant.conf:

```
ctrl_interface=/var/run/wpa_supplicant
ctrl_interface_group=0
update_config=1
network={
    ssid="NETGEAR73"
    #psk="freshbutter"
    psk=eb0376fc14ee5d1e6ce129ad54da038adab……
}
```

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

4.13.7.2   Obtaining an IP address

The following command is used to get an IP address for wlan0:

```
udhcpc -i wlan0
```

Chapter 5
Graphics

5.1 Graphics Processing Unit (GPU)

5.1.1 Introduction

The Graphics Processing Unit (GPU) is a graphics accelerator targeting embedded 2D/3D graphics applications.

The 3D graphics processing unit (GPU3D) is an embedded engine that accelerates user level graphics Application Programming Interface (APIs) such as OpenGL ES 1.1, OpenGL ES 2.0, and OpenGL ES 3.0 and OpenCL 1.1EP. The 2D graphics processing unit (GPU2D) is an embedded 2D graphics accelerator targeting graphical user interfaces (GUI) rendering boost. The VG graphics processing unit (GPUVG) is an embedded vector graphic accelerator for supporting the OpenVG 1.1 graphics API and feature set. The GPU driver kernel module source is in the kernel source tree, but the libraries are delivered as binary only.

Graphics Processing Unit	Hardware	Applicable Platform
3D	Vivante dual-GC7000XSVX	8QuadMax
3D	Vivante GC7000Lite	8QuadXPlus/8M Quad
3D	Vivante GC7000 Nano Ultra	7ULP and 8M Mini
3D	Vivante GC2000	6Quad/6Dual
3D	Vivante GC2000+	6QuadPlus/6DualPlus
3D	Vivante GC880	6DualLite/6Solo
3D/2D	Vivante GC400T	6SoloX
2D	Vivante GC320	6Quad/6Dual/6DualLite/6Solo/6SoloLite
Vector	Vivante GC355	6Quad/6Dual/6SoloLite
2D	Vivante GC328	8M Quad and 8M Mini

NOTE
- GC400T does not support OpenGL ES 3.0.
- GC880/GC400T does not support OpenCL 1.1EP. GC2000 and GC2000+ support OpenCL 1.1 EP.
- GC7000XSVX supports OpenCL 1.2 FP, OpenVX 1.0.1, and Vulkan 1.0.

5.1.2 Driver Features

The GPU driver enables this board to provide the following software and hardware support:

- EGL (EGL is an interface between Khronos rendering APIs such as OpenGL ES or OpenVG and the underlying native platform window system) 1.5 API defined by Khronos Group.
- OpenGL ES (OpenGL® ES is a royalty-free, cross-platform API for full-function 2D and 3D graphics on embedded systems) 1.1 API defined by Khronos Group.
- OpenGL ES 2.0 API defined by Khronos Group.
- OpenGL ES 3.0/3.1/3.2 API defined by Khronos Group.
- OpenVG (OpenVG is a royalty-free, cross-platform API that provides a low-level hardware acceleration interface for vector graphics libraries such as Flash and SVG) 1.1 API defined by Khronos Group.
- OpenCL (OpenCL is the first open, royalty-free standard for cross-platform, parallel programming of modern processors.) 1.1 EP API defined by Khronos Group.
- OpenGL 2.1 API defined by Khronos Group.
- Automatic 3D core slowing down, when hot notification from thermal driver is active, 3D core will run at 1/64 clock.
- OpenCL1.1/1.2FP API defined by Khronos Group.
- OpenVX 1.0.1 API defined by Khronos Group.
- Vulkan 1.0 API defined by Khronos Group.

5.1.3 Hardware Operation

For detailed hardware operations, see the GPU chapters in the Applications Processor Reference Manual specific to SoC.

5.1.4   Software Operation

The GPU driver is divided into two layers. The first layer is running in kernel mode and acts as the base driver for the whole stack. This layer provides the essential hardware access, device management, memory management, command queue management, context management and power management. The second layer is running in user mode, implementing the stack logic and providing the following APIs to the upper layer applications:

- OpenGL ES 1.1, 2.0, and 3.0 API
- EGL 1.5 API
- OpenGL ES11/20/30/31/32
- OpenCL 1.1/1.2 FP
- OpenVX 1.0.1
- Vulkan 1.0
- OpenGL 4.0
- WebGL 1.0.2
- OpenVG 1.1 API
- OpenCL 1.1 EP API

5.1.5   Source Code Structure

Table below lists GPU driver kernel module source structure:

`drivers/mxc/gpu-viv`

Table 5-1.   GPU Driver Files

File	Description
Kconfig Kbuild config	Kernel configure file and makefile
hal/kernel/arch	Hardware-specific driver code for GC2000, GC880, GC400T, and GC320
hal/kernel/archvg	Hardware-specific driver code for GC355
hal/kernel	Kernel mode HAL driver
hal/os/linux/kernel	OS layer HAL driver

NOTE

If you replace the whole content in this directory, the GPU kernel driver can be upgraded.

5.1.6   Library Structure

Table below lists GPU driver user mode library structure:

<ROOTFS>/usr/lib

Table 5-2.   GPU Library Files

File	Description
libCLC.so	OpenCL frontend compiler library
libEGL.so**	EGL1.4 library
libGAL.so	GAL user mode driver
libGLES_CL.so	OpenGL ES 1.1 common lite library (without EGL API, no float point support API)
libGL.so**	OpenGL 2.1 common library
libGLES_CM.so	OpenGL ES 1.1 common library (without EGL API, include float point support API)
libGLESv1_CL.so**	OpenGL ES 1.1 common lite library (with EGL API, no float point support API)
libGLESv1_CM.so**	OpenGL ES 1.1 common library (with EGL API, include float point support API)
libGLESv2.so**	OpenGL ES 2.0/3.0/3.1/3.2 library
libGLSLC.so	OpenGL ES shader language compiler library
libVSC.so	OpenGL front-end compiler library
libVivanteOpenCL.so	Vivante
libOpenCL.so	OpenCL ICD wrapper library
libOpenVG.so*	OpenVG 1.1 library
libVDK.so	VDK wrapper library.
libVIVANTE.so	Vivante user mode driver.
xorg/modules/drivers/vivante_drv.so	EXA library for X11 acceleration.
libwayland-viv.so	Wayland server-side library for Vivante's EGL driver
libgc_wayland_protocol.so	Vivante Wayland Protocol Extension Library
libOpenVX.so*	OpenVX 1.0 library
libvulkan..so*	Vulkan 1.0 library

**SONAME is used for libEGL.so, libGLESv2.so, libGLESv1_CM.so, libGLESv1_CL.so, libGL.so.

*For libOpenVG.so, there are two libraries for the OpenVG feature. libOpenVG.3d.so is the GC7000XSVX/GC2000+/GC2000/GC880/GC400T-based OpenVG library. libOpenVG.2d.so is the gc355 based OpenVG library.
  • For i.MX 6DualPlus/QuadPlus and i.MX 6Dual/Quad, both libOpenVG.3d.so and libOpenVG.2d.so can be used.
  • For i.MX 6DualLite, and i.MX 6SoloX, only libOpenVG.3d.so can be used.

- For i.MX 6SoloLite, only libOpenVG.2d.so can be used.
- If no SOC limitation, for the x11 backend, libOpenVG.3d.so is linked by default.
- If no SOC limitation, for framebuffer, directFB, and Wayland backends, the default openVG library is linked to libOpenVG.2d.so.

This can be done by using the following sequence of commands:

```
cd <ROOTFS>/usr/lib
sudo ln -s libOpenVG_355.so libOpenVG.so
```

5.1.7 API References

See the following web sites for detailed specifications:

- OpenGL ES 1.1, 2.0, and 3.0 API: www.khronos.org/opengles/
- OpenCL 1.1 EP www.khronos.org/opencl/
- EGL 1.4 API: www.khronos.org/egl/
- OpenVG 1.1 API: www.khronos.org/openvg/
- OpenGL ES API: www.khronos.org/gles/
- OpenCL API: www.khronos.org/opencl/
- OpenVX API: www.khronos.org/openvx/
- Vulkan API: www.khronos.org/vulkan/
- OpenGL API: www.khronos.org/opengl/
- WebGL API: www.khronos.org/webgl/

5.1.8 Menu Configuration Options

In menu configuration enable the following module for the GPU driver:

CONFIG_MXC_GPU_VIV is a configuration option for GPU driver. In the menuconfig this option is available under Device Drivers > MXC support drivers > MXC Vivante GPU support > MXC Vivante GPU support.

On the screen displayed, select **Configure the kernel**, select Device Drivers > MXC support drivers > MXC Vivante GPU support > MXC Vivante GPU support, and then exit. When the next screen appears, select the following options to enable the GPU driver:

- Package list > imx-gpu-viv
- This package provides proprietary binary libraries, and test code built from the GPU for framebuffer

5.2 Wayland

5.2.1 Introduction

Wayland is a protocol for a compositor to talk to its clients as well as a C library implementation of that protocol. The compositor can be a standalone display server running on Linux kernel modesetting and evdev input devices, an X application, or a Wayland client itself. The clients can be traditional applications, X servers or other display servers.

Part of the Wayland project is also the Weston reference implementation of a Wayland compositor. The Weston compositor is a minimal and fast compositor and is suitable for many embedded and mobile use cases.

This chapter describes how to enable Wayland/Weston support on an i.MX series device.

5.2.2 Software Operation

This release is based on the Wayland 1.15 version and Weston 4.0 version.

5.2.3 Yocto Build Instructions

The instructions for Yocto Project build are as follows:

1. Prepare a Yocto build directory and follow the setup instructions in the *i.MX Yocto Project User's Guide* (IMXLXYOCTOUG) for DISTRO Wayland.
2. Set up Yocto for Wayland in the build directory:

```
$ MACHINE = <your-machine> DISTRO=fsl-imx-xwayland source fsl-setup-release.sh -b build-wayland
```

3. Build an image.

```
$ bitbake fsl-image-gui
```

5.2.4 Customizing Weston

The i.MX Weston includes two compositors. One is the EGL3D compositor, which is accelerated by the 3D core. The other is G2D compositor accelerated by the 2D BLT engines.

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

Weston options can be updated in the file "/etc/init.d/weston".

Table 5-3.   Common options for Weston

Weston option	Description
tty	default to current tty.
device	"/dev/fb0", default frame buffer , Multi display supported in G2D compositor.
use-gl	EGL accelerated, defaults to be "1".
use-g2d	G2D accelerated, defaults to be "0".
idle-time	Idle time in seconds.

5.2.4.1   Multi display supported in Weston

Multi display was supported in G2D compositor only. Add these options to start Weston:

```
weston --tty=1 --device=/dev/fb0,/dev/fb2 --use-g2d=1 &
```

5.2.4.2   Multi buffer supported in Weston

The Weston server supports both single buffering and multi buffering. In single buffering, the damage area is rendered to the offscreen surface and blits to front buffer.The offscreen surface is used to avoid flickering. By default, the Weston server starts with single buffering.

In multi buffering, instead of rendering to offscreen, the damage area is rendered to back buffer and does the flip, but the frame rate will be restricted to the display rate. A maximum of three buffers are supported.

Before starting the Weston server, export FB_MULTI_BUFFER to control the number of buffers to be used.

Environment variables for single buffering:

```
export FB_MULTI_BUFFER=1
```

Environment variables for double buffering:

```
export FB_MULTI_BUFFER=2
```

5.2.5   Running Weston

Perform the following operations to run Weston:

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

1. Boot the i.MX device.
2. To run clients, the second button in the top bar will run weston-terminal, from which you can run clients. There are a few demo clients available in the Weston build directory, but they are all pretty simple and mostly for testing specific features in the Wayland protocol:
   - 'weston-terminal' is a simple terminal emulator, not very compliant, but works well enough for bash.
   - 'weston-flower' draws a flower on the screen, testing the frame protocol.
   - 'weston-smoke' tests SHM buffer sharing.
   - 'weston-image' loads the image files passed on the command line and shows them.

5.3 X Windows Acceleration

5.3.1 Introduction

X-Windows System (aka X11 or X) is a portable, client-server based, graphics display system. X11 is only supported for i.MX 6.

X-Windows system can run with a default frame buffer driver which handles all drawing operations to the main display. As there is a 2D GPU (graphics processing unit) available, then some drawing operations can be accelerated. High-level X operations may get decomposed into low level drawing operations which are accelerated for X-Windows System.

5.3.2 Hardware Operation

X-Windows System acceleration on i.MX with GPU utilizes the Vivante GC320 2D GPU.

Acceleration is also dependent on the frame buffer memory.

5.3.3 Software Operation

X-Windows acceleration is supported for X.org X Server version 1.11.x and later versions supporting the EXA interface version 2.5.

The following list summarizes the types of operations that are accelerated for X11. All operations involve frame buffer memory which may be on screen or off screen:

- Solid fill of a rectangle.
- Upload image in system memory into video memory.
- Copy of a rectangle with same pixel format with possible source-target rectangle overlap.
- Copy of a rectangle supporting most XRender compositing operations with these options:
    - Pixel format conversion.
    - Repeating pattern source.
    - Porter-Duff blending of source with target.
    - Source alpha masking.

The following list includes additional features supported as part of the X-Windows acceleration:

- Allocation of X pixmaps directly in frame buffer memory.
- EGL swap buffers where the EGL window surface is an X-window.
- X-window can be composited into an X pixmap which can be used directly as any EGL surface.

5.3.4 X-Windows Acceleration Architecture

The following block diagram shows the components that are involved in the acceleration of X-Windows System:
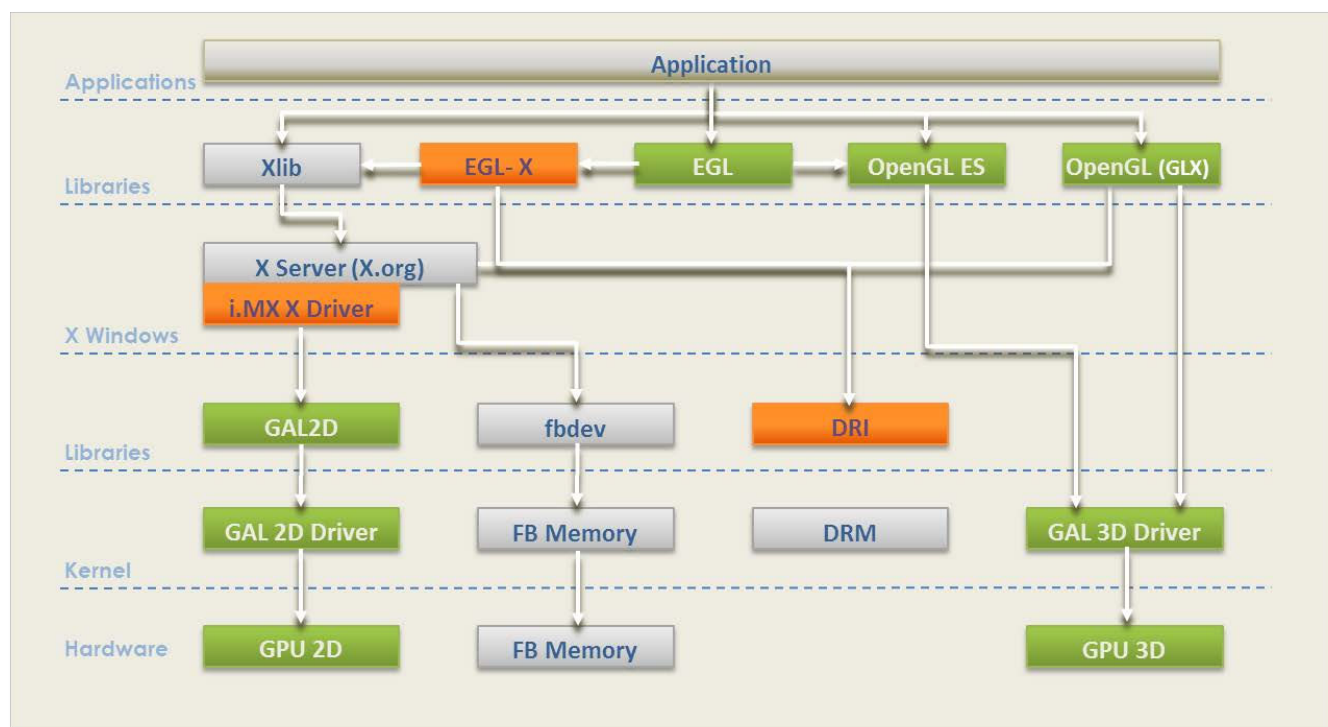
Figure 5-1. X Driver Architecture

The components shown in green are those provided as part of the Vivante 2D/3D GPU driver support which includes OpenGL/ES and EGL, though some i.MX processors, such as i.MX 6SoloLite do not contain 3D HW module. The components shown in light gray are the standard components in the X-Windows System without acceleration. The components shown in orange are those added to support X-Windows System acceleration and briefly described here.

The **i.MX X Driver** library module (`vivante-drv.so`) is loaded by the X server and contains the high-level implementation of the X-Windows acceleration interface for i.MX platforms containing the GC320 2D GPU core. The entire linearly contiguous frame buffer memory in `/dev/fb0` is used for allocating pixmaps for X both on screen and off screen. The driver supports a custom X extension which allows X clients to query the GPU address of any X pixmap stored in frame buffer memory.

The **libGAL.so** library module (`libGAL.so`) contains the register level programming interface to the GC320 GPU module. This includes the storing of register programming commands into packets which can be streamed to the device. The functions in the libGAL.so library are called by the i.MX X Driver code.

The **EGL-X** library module (`libEGL.so`) contains the X-Windows implementation of the low level EGL platform-specific support functions. This allows X-window and X pixmap objects to be used as EGL window and pixmap surfaces. The EGL-X library uses Xlib function calls in its implementation along with the i.MX X Driver module's X extension for querying the GPU address of X pixmaps stored in frame buffer memory.

5.3.5   i.MX Driver for X-Windows System

The i.MX X Driver, referred to as vivante-drv.so, implements the EXA interface of the X server in providing acceleration.

The Vivante X Driver, referred to as vivante-drv.so, implements the EXA interface of the X server to provide acceleration.

The following list describes details particular to this implementation:

- The implementation builds upon the source from the fbdev frame buffer driver for X so that it can be the fallback when the acceleration is disabled.
- The implementation is based on X server EXA version 2.5.0.
- The EXA solid fill operation is accelerated, except for source/target drawables containing less than 300x300 pixels in which case fallback is to software rendering.
- The EXA copy operation is accelerated, except for source/target drawables containing less than 400x120 pixels in which case fallback is to software rendering.
- EXA putimage (upload into video memory) is accelerated, except for source drawables containing less than 400x400 pixels in which case fallback is to software rendering. For EXA solid fill and copy operations, only solid plane masks and only GXcopy raster-op operations are accelerated.
- For EXA copy operation, the raster-op operations (GXandInverted, GXnor, GXorReverse, GXorInverted, and GXnand) are not accelerated.
- EXA composite allows for many options and combinations of source/mask/target for rendering.
- Most of the (commonly used) EXA composite operations are accelerated.

The following types of EXA composite operations are accelerated:

- Composite operations for source/target drawables containing at least 640 pixels. If less than 640 pixels, the composite path falls to software.
- Simple source composite operations are used when source/target drawables contain more than 200x200 pixels (operations with mask not supported).
- Constant source (with or without alpha mask) composite with target.
- Repeating pattern source (with or without alpha mask) composite with target.

- Only these blending functions: SOURCE, OVER, IN, IN-REVERSE, OUT-REVERSE, and ADD (some of these are needed to support component-alpha blending which is accelerate).
- In general, the following types of (less commonly used) EXA composite operations are not accelerated:
  - Transformed (that is, scaled, rotated) sources and masks
  - Gradient sources
  - Alpha masks with repeating patterns

The implementation handles all pixmap allocation for X through the EXA callback interface. A first attempt is made to allocate the memory where it can be accessed by a physical GPU address. This attempt can fail if there is insufficient GPU accessible memory remaining, but it can also fail when the bits per pixel being requested for the pixmap is less than eight (8). If the attempt to allocate from the GPU accessible memory fails, then the memory is allocated from the system. If the pixmap memory is allocated from the system, then this pixmap cannot be involved in a GPU accelerated option. The number of pitch bytes used to access the pixmap memory may be different depending on whether it was allocated from GPU accessible memory or from the system. Once the memory for an X pixmap has been allocated, whether it is from GPU accessible memory or from the system, the pixmap is locked and can never migrate to the other type of memory. Pixmap migration from GPU accessible memory to system memory is not necessary since a system virtual address is always available for GPU accessible memory. Pixmap migration from system memory to GPU accessible memory is not currently implemented, but would only help in situations where there was insufficient GPU accessible memory at initial allocation but more memory becomes available (through de-allocation) at a later time. The GPU accessible memory pitch (horizontal) alignment for Vivante 2D GPUs is 8 pixels. Because the memory can be allocated from GPU accessible memory, these pixels could be used in EGL for OpenGL/ES drawing operations. All of the memory allocated for /dev/fb0 is made available to an internal linear offscreen memory manager based on the one used in EXA. The portion of this memory beyond the screen memory is available for allocation of X pixmap, where this memory area is GPU accessible. The amount of memory allocated to /dev/fb0 needs to be several MB more than the amount needed for the screen. The actual amount needed depends on the number of X-Windows and pixmaps used, the possible usage of X pixmaps as textures, and whether X-Windows are using the XComposite extension. An X extension, i.e., VIVEXT shown in Fig. 1, is provided so that X clients can query the physical GPU address associated with an X pixmap, if that X pixmap was allocated in the GPU accessible memory.

5.3.6 i.MX Direct Rendering Infrastructure (DRI) for X-Windows System

The Direct Rendering Infrastructure, also known as the DRI, is a framework for allowing direct access to graphics hardware under the X Window System in a safe and efficient manner. It includes changes to the X server, to several client libraries, and to the kernel (DRM, Direct Rendering Manager). The most important activity for the DRI is to create fast OpenGL and OpenGL ES implementations that render to framebuffer memory directly. Without DRI, the OpenGL driver has to depend on X server for final rendering (indirect rendering), which degrades the overall performance significantly.

The components of Vivante's DRI OpenGL implementation include:

- The Direct Rendering Manager (DRM) is a kernel module that provides APIs to userland to synchronize access to hardware and to manage different classes of video memory buffers. Vivante's DRI implementation uses selected DRM APIs for opening/closing DRI device, and locking/unlocking FB. Most other buffer management and DMA management functions are handled by Vivante's specific kernel module: galcore.ko.
- The EXA driver is a DRI-enabled DDX 2D driver which initializes the DRM when X server starts. As all X Window pixmap buffers are allocated by the EXA driver from GPU memory, the GPU can render directly into these buffers if the buffer information is passed from the X server process to the X client processes (GL or GLES applications) properly.
- The Vivante-specific X extension "vivext" passes buffer information from X server to X clients. This Vivante X extension includes the following three interfaces:
  - DrawableFlush, which enables X clients to notify X server to flush the GPU cache for a drawable surface.
  - DrawableInfo, which enables X clients to query the drawable information (position, size, physical address, stride, cliplist, etc.) from the X server.
  - PixmapPhysAddr, which enables X clients to query the physical address and stride of a pixmap buffer from X server.

The integration of GL/GLES application windows with Ubuntu Unity2D desktop is achieved by following steps:

- GL/GLES applications render a frame into the pixmap buffers that are allocated in the EXA driver.
- In the SwapBuffers implementation, the driver notifies X server that the pixmap buffer region is damaged through Xdamage and Xfixes APIs.
- Then the X server will present the latest pixmap buffer to the Unity2D desktop while maintaining the proper window overlap characteristics relative to the other windows on the desktop.

On a compositing X desktop, such as Ubuntu Unity 2D, GLES/GL applications always render into the full rectangular back buffer of a window. There is no window clipping required. So the Vivante DRI implementation can take advantage of the GPU's resolve function and render into the window back buffer directly.

On a legacy X window desktop, such as Gnome, Xwin, etc., GLES/GL applications have to render onto the frame buffer surface directly. Thus, the DRI driver uses the DrawableInfo interface in the VIVEXT extension to obtain the cliplist of the window, then copies the sub-regions of the render target to the frame buffer according to the cliplist. This will ensure that the GLES/GL windows overlap with other windows on the desktop properly. However, the copying of the render target sub-regions to the frame buffer has to be done by the CPU as the sub-regions' starting address and alignment may not meet GPU copy requirements.

The Vivante DRI implementation can detect the type of X window manager (compositing desktop manager or legacy desktop manager) at run-time, and use appropriate DRI rendering paths for GLES/GL applications.

5.3.7 EGL- X Library

The EGL-X library implements the low level EGL interface when used in X Window System. The following list describes details particular to this implementation:

- The eglDisplay native display type is "Display*" in X.
- The eglWindowSurfacenative window surface type is "Window" in X.
- The eglPixmapSurface native pixmap surface type is "Pixmap" in X.

When an eglWindowSurface is created, the back buffers used for double-buffering can have different representations from the window surface (based on the selected eglConfig). An attempt is made to create each back buffer using the representation which provides the most efficient blit of the back buffer contents to the window surface when eglSwapBuffers is called.

The back buffer is allocated by creating an X pixmap of the necessary size. Use the X extension for the Vivante X Driver module to query the physical frame buffer address for this X pixmap if it was allocated in the offscreen frame buffer memory.

5.3.8 xorg.conf for i.MX

The `/etc/X11/xorg.conf` file must be properly configured to use the i.MX 6 X Driver.

The /etc/X11/xorg.conf file must be properly configured to use the Vivante X Driver. This configuration appears in a "Device" section of the file which contains some required entries and some entries that are optional. The following example shows a preferred configuration for using the Vivante X Driver:

```
Section "ServerLayout"
        Identifier      "Default Layout"
        Screen          "Default Screen"
EndSection

Section "Module"
        Load            "dbe"
        Load            "extmod"
        Load            "freetype"
        Load            "glx"
        Load            "dri"
EndSection

Section "InputDevice"
        Identifier      "Generic Keyboard"
        Driver          "kbd"
        Option          "XkbLayout" "us"
        Option          "XkbModel" "pc105"
        Option          "XkbRules" "xorg"
EndSection

Section "InputDevice"
        Identifier      "Configured Mouse"
        Driver          "mouse"
        Option          "CorePointer"
EndSection

Section "Device"
        Identifier      "Your Accelerated Framebuffer Device"
        Driver          "vivante"
        Option          "fbdev"         "/dev/fb0"
        Option          "vivante_fbdev" "/dev/fb0"
        Option          "HWcursor"       "false"
EndSection

Section "Monitor"
        Identifier      "Configured Monitor"
EndSection

Section "Screen"
        Identifier      "Default Screen"
        Monitor         "Configured Monitor"
        Device          "Your Accelerated Framebuffer Device"
        DefaultDepth    24
EndSection

Section "DRI"
        Mode 0666
EndSection
```

Mandatory Strings

Some important entries recognized by the Vivante X Driver are described as follows.

Device Identifier and Screen Device String

The mandatory Identifier entry in the Device section specifies the unique name to associate with this graphics device.

```
Section "Device"
        Identifier      "Your Accelerated Framebuffer Device"
```

The following entry ties a specific graphics device to a screen. The Device Identifier string must match the Device string in a Screensection of the xorg.conf file. For example:

```
Section "Screen"
        Identifier      "Default Screen"
        <other entries>
        Device          "Your Accelerated Framebuffer Device"
        <other entries>
EndSection
```

Device Driver String

The mandatory Driver entry specifies the name of the loadable Vivante X driver.

Driver "vivante"

Device fbdevPath Strings

The mandatory entries fbdev and vivante_dev specify the path for the frame buffer device to use.

```
Section "Device"
        Identifier      "Your Accelerated Framebuffer Device"
        Driver          "vivante"
        Option          "fbdev"         "/dev/fb0"
        Option          "vivante_fbdev" "/dev/fb0"
        <other entries>
EndSection
```

5.3.9 Setting Up X-Windows System Acceleration on Yocto

Prerequisites:

- xserver-xorg-video-imx-viv-(ver).tar.gz, which is Vivante EXA plugin source code based on GPU driver
- drm-update-arm.patch, which is a patch with adding the Arm lock implementation for libdrm xf86drm.h. Note that the original xh86drm.h header file from libdrm does not have lock for supporting Arm architecture. This patch is located in the community Yocto Project layers Yocto_build/sources/meta-freescale/recipes-graphics/drm/libdrm/mx6, and shown below: drm-update-arm.patch:

```
+#elif defined(__arm__)
+       #undef DRM_DEV_MODE
+       #define DRM_DEV_MODE     (S_IRUSR|S_IWUSR|S_IRGRP|S_IWGRP|S_IROTH|S_IWOTH)
```

```
+
+        #define DRM_CAS(lock,old,new,__ret)              \
+        do {                                             \
+              __asm__ __volatile__ (                     \
+                   "1: ldrex %0, [%1]\n"                 \
+                   "   teq %0, %2\n"                     \
+                   "   strexeq %0, %3, [%1]\n"           \
+              : "r" (__ret)                              \
+              : "r" (lock), "r" (old), "r" (new)         \
+              : "cc","memory");                          \
+        } while (0)
+
 #endif /* architecture */
 #endif /* __GNUC__ >= 2 */
```

Build and install instructions:

- Install the prerequisites modules or patches in the appropriate locations and with right recipes in Yocto environment.
- Build XServer with correct drm header file (xf86drm.h). The purpose is to create correct dri module
- Build GPU EXA module with the command 'bitbake xf86-video-imxfb-vivante'. vivante_drv.so will be generated with successful build, and then install it together with xorg and libdri library in target board rootfs in /usr/lib/xorg/modules/
- Install the pre-Yocto-built imx-gpu-viv binary in target board rootfs. For accelerating X11, the X11 backend is required
- Now ready to run the X11 applications in target board.

NOTE
x11 applications are down if the Arm core version xf86drm.h is not used.

5.3.10 Setting Up X Window System Acceleration

- Install any packages appropriate for your platform.
- Verify that the device file /dev/galcore is present.
- Verify that the file /etc/X11/xorg.conf contains the correct entries as described in the previous section.
- Assuming the above steps have been performed, do the following to verify that X Window System acceleration is indeed operating.
- Examine the log file /var/log/Xorg.0.log and confirm that the following lines are present.

```
[   41.752] (II) Loading /usr/lib/xorg/modules/drivers/vivante_drv.so
    [   41.752] (II) VIVANTE(0): using default device
    [   41.752] (II) VIVANTE(0): Creating default Display subsection in Screen
section "Default Screen" for depth/fbbpp 24/32
    [   41.752] (**) VIVANTE(0): Depth 24, (--) framebufferbpp 32
    [   41.752] (==) VIVANTE(0): RGB weight 888
    [   41.752] (==) VIVANTE(0): Default visual is TrueColor
    [   41.753] (==) VIVANTE(0): Using gamma correction (1.0, 1.0, 1.0)
```

```
        [    41.753] (II) VIVANTE(0): hardware: DISP3 BG (video memory: 8100kB)
        [    41.753] (II) VIVANTE(0): checking modes against framebuffer device...
        [    41.753] (II) VIVANTE(0): checking modes against monitor...
        [    41.753] (--) VIVANTE(0): Virtual size is 1920x1080 (pitch 1920)
        [    41.753] (**) VIVANTE(0):  Built-in mode "current": 148.5 MHz, 67.5 kHz,
60.0 Hz
        [    41.753] (II) VIVANTE(0): Modeline "current"x0.0  148.50  1920 2008 2052
2200  1080 1084 1089 1125 +hsync +
        vsync -csync (67.5 kHz)
        [    41.753] (==) VIVANTE(0): DPI set to (96, 96)
        [    41.753] (II) Loading sub module "fb"
        [    41.753] (II) LoadModule: "fb"
        [    41.754] (II) Loading /usr/lib/xorg/modules/libfb.so
        [    41.755] (II) Module fb: vendor="X.Org Foundation"
        [    41.755] compiled for 1.10.4, module version = 1.0.0
        [    41.755] ABI class: X.Org ANSI C Emulation, version 0.4
        [    41.755] (II) Loading sub module "exa"
        [    41.755] (II) LoadModule: "exa"
        [    41.756] (II) Loading /usr/lib/xorg/modules/libexa.so
        [    41.756] (II) Module exa: vendor="X.Org Foundation"
        [    41.756] compiled for 1.10.4, module version = 2.5.0
        [    41.756] ABI class: X.Org Video Driver, version 10.0
        [    41.756] (--) Depth 24 pixmap format is 32 bpp
        [    41.797] (II) VIVANTE(0): FB Start = 0x33142000  FB Base = 0x33142000  FB
Offset = (nil)
        [    41.797] (II) VIVANTE(0): test Initializing EXA
        [    41.798] (II) EXA(0): Driver allocated offscreenpixmaps
        [    41.798] (II) EXA(0): Driver registered support for the following
operations:
        [    41.798] (II)         Solid
        [    41.798] (II)         Copy
        [    41.798] (II)         Composite (RENDER acceleration)
        [    41.798] (II)         UploadToScreen
        [    42.075] (==) VIVANTE(0): Backing store disabled
        [    42.084] (==) VIVANTE(0): DPMS enabled
```

5.3.11 Troubleshooting

1. Framebuffer devices can be specified by environment variable. This is especially
   useful when there are multiple framebuffer devices.

   ```
   export FB_FRAMEBUFFER_0=/dev/fb2
   ```

2. If the above does not resolve the issue:
   - If DRM booted up properly, check the /var/log/X11.n log file (n will represent
     instance number) for more information.
   - If DRM did not boot properly, check your kernel mode driver installation. (See
     sections 6.4.2 and 6.4.3 above).

3. Window is created, but nothing is drawn
   - If you run an OpenGL application and find a window was created, but nothing
     was drawn, try to export the ${__GL_DEV_FB} environment variable:

   ```
   export __GL_DEV_FB=$FB_FRAMEBUFFER_0.
   ```

4. Cannot open Display message
   - If you have a message similar to "Cannot open Display," use the following
     command to check whether X is running at :0 or at :1 instance, use:

   ```
   $ ps -ef|grep X
   ```

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

- Then depending on the returned instance number, add the following environment variable

```
export DISPLAY=:n
```

- Then run it again.

5. UART terminal cannot run GPU application with lightdm
   - Use ssh terminal instead.
6. EXA build script failure
   - Check the log file and make sure your system time is set correctly.
7. Invalid MIT-MAGIC-COOKIE-1 Key error message
   - Some GPU applications are not permitted to run using root. Use an alternate account instead.
8. Segment fault occurs while running GPU application
   - Check the attribute for dev/galcore should be updated to 666.
   - To update this attribute automatically on system boot,
   - Locate and edit file /etc/udev/rules.d/<bsp-specific.rules>.
   - Add: "KERNEL=="galcore",MODE="0666""
   - Lastly, make sure your kernel and GPU drivers are matched.
9. Check whether Compiz is running
   - If your host or target has issues after installing the OpenGL Development Packages in Table 6, check whether compiz is running with the following command:

```
$ ps -ef|grep compiz
```

   - If compiz is running, then Ubuntu is using Unity3D by default. To set the default window manager to Unity2D:
   - Locate and edit file /var/lib/AccountsService/users/<username>.
   - Change ubuntu to ubunto-2d.

Chapter 6
Video

6.1 Capture Overview

6.1.1 Introduction

The i.MX capture driver support is through the V4L2 interface with camera sensor controllers and interfaces. Applications cannot use the camera driver directly. Instead, the applications use the V4L2 capture driver to open and close the camera for preview and image capture, controlling the camera, getting images from camera, and starting the camera preview.

The list of capture controllers are the following:

- Camera Serial Interface - CSI
- IPU-CSI
- Video Interface Unit - VIU
- Image Sensor Interface - ISI

The list of capture interfaces for transfering image data are the following

- Parallel-CSI
- MIPI-CSI2
- HDMI RX
- TV Deocder

This chapter will explain the differences between the various controllers and interfaces.

NOTE
The i.MX 6 with IPU uses internaldev for V4L2 interface while all others use subdev for V4L2 interface.

The following table describes the different controllers and interfaces combinations.

Table 6-1. Camera Controllers and Interfaces

SoC	Controller	Interface
6SoloLite/6SLL	CSI	Parallel CSI
6SoloX	VIU	Parallel CSI and TV Decoder
6UltraLite/6ULL	CSI	Parallel CSI
6DualLite/Solo	IPU-CSI	Parallel CSI interaldev IPU
6QuadPlus/Quad/Dual	IPU-CSI	Parallel CSI internaldev IPU
7Dual/Solo	CSI	MIPI-CSI2 using Samsung and Parallel CSI
7ULP	VIUI	MIPI-CSI2 using Mixel
8QuadMax	ISI	MIPI-CSI2 using Mixel and HDMI Receiver using Cadence
8QuadXPlus	ISI	MIPI-CSI2 using Mixel and Parallel CSI using IMX8
8M Quad	CSI	MIPI-CSI2 using Mixel and Parallel CSI
8M Mini	CSI	MIPI-CSI2 using Samsung

Some additional details are listed below:

- The ISI controller is a new controller used for i.MX 8 and i.MX 8X families but not with the 8M family.
- The i.MX 6 SoC without IPU, i.MX 7Dual and i.MX 8M use the same CSI controller driver.
- The i.MX 8 and i.MX 8X families use a newer IMX8 CSI driver.
- The i.MX 6 with IPU use a customized CSI that interfaces with IPU hardware.
- Each SoC can support one or more interfaces as described in the previous table. The interfaces align with Video for Linux V4L2 APIs.
- In some cases the capture controller is not interfacing to a camera but a video input unit. Some also interface to HDMI Receivers

6.1.2 Omnivision Camera

The Omnivision Camera supports multiple interfaces and types of cameras. The Omnivision camera is a small camera sensor and lens module with low power consumption.

The Omnivision camera uses the serial camera control bus (SCCB) interface to control the sensor operation working as an I^2C client for control operations. This camera supports transfer modes of CSI, MIPI-CSI2 and Parallel-CSI interfaces When using MIPI mode, OV5640 connects to i.MX chip through the MIPI CSI-2 interface. MIPI receives the sensor data and transfers them to CSI.

Below the table lists the differnt Omnivision cameras and which interface is supported.

Table 6-2. Camera Controllers and Interfaces

Camera	Controller Interface
OV5640	CSI controller/MIPI-CSI2/Parallel CSI
OV5642	Parallel CSI
OV10635	MIPI-CSI2/Parallel CSI

The Omnivision menu configuration has multiple options based on the support

Top level selection is the following:

Device Drivers > Multimedia support (MEDIA_SUPPORT [=y]) > V4L platform devices (V4L_PLATFORM_DRIVERS)

The next level selections are based on the different interfaces for each SoC

- For MX6 wiith IPU select both > MXC Camera/V4L2 PRP Features support and > OmniVision ov5640 camera support (MXC_CAMERA_OV5640)
- For MX6 without IPU select > OmniVision ov5640 camera support (MXC_CAMERA_OV5640_V2)
- For MX7 select > OmniVision ov5640 camera support using mipi (MXC_CAMERA_OV5640_MIPI_V2)
- For MX8 select > IMX8 Camera ISI/MIPI Features support (VIDEO_MX8_CAPTURE) > IMX8 Camera Controller (IMX8_CAPTURE_DRIVER) and Maxim OV5640_V3 driver support (MXC_CAMERA_OV5640_V3)
- For MX8M select > OmniVision ov5640 camera support (MXC_CAMERA_OV5640_V2) and OmniVision ov5640 camera support using mipi (MXC_CAMERA_OV5640_MIPI_V2)

The following table describes the supported camera features for each interface.

Table 6-3. Capture Interface Features

Interface	Features
IPU-CSI	- Parallel CSI with 2 Ports, 20 bits+ 8 bits - Playback 1080i/p + D1@30fps @ 30fps - Record 1080p @ 30fps 2-way 720@30fps - De-interlacing high quality motion adaptive algorithm. - Resizing - fully flexible - Rotation/inversion support - Color conversion-fully flexibl - Memory interface - AXI split transaction 64-bit 266Mhz - Memory Bus - selective read for combining

Table continues on the next page...

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

Table 6-3.   Capture Interface Features (continued)

Interface	Features
	• Control capabilities - display and dma controller, internal synchronization • Synchronization - double/triple buffering, frame-by-frame or tight sub-frame with internal memory
CSI	• MIPI-CSI-2 with 2 ports, 4 lanes x 1.5 Gbps • Playback - two 1080p30 • Record - 1080p30x2 • 2-way - 1080p30x2 • Deinterlacing-simple bob and weave • Memory interface throughput - 64-bit • Controller capabilities - DMA • Synchronizaton - double buffer
ISI-8QuadXPlus	• Parallel CSI - one port 24 bits • MIPI-CSI-2 with 2 ports, 4 lanes x 1.5 Gbps • Playback - 1080p30x2 • Record - 1080p30x2 • 2-way - 1080p30x2 • Deinterlacing-simple bob and weave • Resizing • JPEG Encode/Decode • Memory interface throughput - 124-bit, 400MHz • Controller capabilities - DMA • Synchronizaton - double buffer
ISI-8QuadMax	• MIPI-CSI-2 with 2 ports, 4 lanes x 1.5 Gbps • HDMI Receiver - 1 port HDMI 2.0 4K30 • Playback - 4K30x1 or 1080px30x4 • Record - 4K30x1 or 1080px30x4 • 2-way - 4K30x1 or 1080px30x4 • Deinterlacing-simple bob and weave • Resizing • JPEG Encode/Decode • Memory interface throughput - 124-bit, 400MHz • Controller capabilities - DMA • Synchronizaton - double buffer

6.1.3 Parallel CSI

The Parallel CSI driver enables a direct connection to external CMOS sensors and CCIR656 video sources. The CSI and sensor drivers are implemented in the Video for Linux Two (V4L2) driver framework consisting of the image capture driver and the video output driver.

The driver initializes the CSI interface and configures and operates with the hardware registers for the CSI module. The following features are supported:

- Configurable interface logic to support most commonly available CMOS sensors.
- Full control of 8-bit/pixel, 10-bit/pixel or 16-bit/pixel data format to 32-bit receive FIFO packing.

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

- 128x32 FIFO to store received image pixel data.
- Receive FIFO overrun protection mechanism.
- Embedded DMA controllers to transfer data from receive FIFO or statistic FIFO through AHB bus.
- Support for double buffering two frames in the external memory.
- Single interrupt source to interrupt controller from maskable interrupt sources: Start of Frame, End of Frame and so on.
- Configurable master clock frequency output to sensor.

The V4L2 CSI capture device includes two interfaces: the capture and overlay interfaces. The capture and overlay interface use the CSI embedded DMA controller to implement the function using V4L2 APIs. The following is the data flow of capture and overlay.

1. The camera sends the data to the CSI receive FIFO, through the 8-bit/10-bit data port.
2. The embedded DMA controllers transfer data from the receive FIFO to external memory through the AHB bus.
3. The data is save to user space memory or output to the frame buffer directly.

i.MX 6 with IPU use a IPU-CSI driver that interfaces with the IPU directly. i.MX Quad Plus/Quad/Dual have support for two IPU-CSI senaros. i.MX 6 without IPU and i.MX 7Dual/Solo use a separate CSI sensor driver that interfaces directly to the sensor.

6.1.4 MIPI Camera Serial Interface (MIPI CSI)

There are four blocks in the MIPI CSI-2 D-PHY: PHY adaptation layer, packet analyzer, image date interface, and register bank.

MIPI CSI-2 is a MIPI-Camera Serial Interface Host Controller with a high performance serial interconnect bus for mobile application which connects camera sensors to the host system. The CSI-2 Host Controller is a digital core that implements all protocol functions defined in the MIPI CSI-2 Specification. In doing so, it provides an interface between the system and the MIPI D-PHY and allows communication with MIPI CSI-2-compliant Camera Sensor.

The MIPI CSI2 driver is used to manage the MIPI D-PHY and lets it work with both MIPI sensor and IPU CSI. MIPI CSI2 driver implements functions as follows:

- MIPI CSI-2 low-level interface for managing the mipi D-PHY register and clock
- MIPI CSI-2 common API for communication between MIPI sensor and MIPI D-PHY

By calling MIPI common APIs, MIPI sensor can set certain information about sensor (such as datatype, lanes number, etc.) to MIPI CSI2 driver to configure D-PHY. In order for the IPU CSI module driver to have the correct configuration, receive appropriate data, and process it correctly, it is necessary for it to receive information about sensor (such as datatype, virtual channel, IPU ID, CSI ID, etc.) from the MIPI CSI2 driver.

Functions and operations are listed as follows:

- PHY Adaptation Layer is responsible for managing the D-PHY interface including PHY error handling;
- Packet Analyzer is responsible for data lane merging if required, together with header decoding, error detection and correction, frame size error detection and CRC error detection;
- Image Date Interface separates CSI-2 packet header information and reorders data according to memory storage format. It also generates timing accurate video synchronization signals. Several error detections are also performed at frame-level and line-level;
- Register Bank is accessible through a standard AMBA-APB slave interface and provides access to the CSI-2 Host Controller register for configuration and control. There is also a fully programmable interrupt generator to inform the system upon certain events;

MIPI CSI2 driver for Linux OS has two parts: MIPI CSI2 driver initialize operation which initializes mipi_csi2_info struct, and MIPI CSI2 common APIs which exports APIs for CSI module driver and MIPI sensor driver.

6.1.5 HDMI

HDMI video interfaces with the Image Sensor Inerface (ISI).

On i.MX 8QuadMax, the HDMI receiver video interface support onne port HDMI 2.0 4K30.

6.1.6 Software Operation

The V4L2 opteratons for capture support modes, picture formats and picture sizes varying for each capture interface.

The imx-test repo has unit tests for these interfaces in the mxc_v4l2_test. See README for details on how to run tests.

6.1.7   V4L2 Capture

Video for Linux Two (V4L2) is a Linux standard. The API specification is available at http://v4l2spec.bytesex.org/spec/.

The V4L2 capture device includes two interfaces: the capture and overlay interfaces using the V4L2 API for capture and overlay devices.

The following are some sample use cases for the V4L2 capture APIs:

1. Sets the capture pixel format and size using IOCTL VIDIOC_S_FMT.
2. Sets the control information using IOCTL VIDIOC_S_CTRL, for rotation.
3. Requests a buffer using IOCTL VIDIOC_REQBUFS.
4. Memory maps the buffer to its user space.
5. Executes the IOCTL VIDIOC_DQBUF.
6. Passes the data that requires post-processing to the buffer.
7. Queues the buffer using the IOCTL command VIDIOC_QBUF.
8. Starts the stream by executing IOCTL VIDIOC_STREAMON.
   - VIDIOC_STREAMON and VIDIOC_OVERLAY cannot be enabled simultaneously.

The following tables lists the V4L2 capture ioctls used in the i.MX Capture Drivers. For more information see the V4l2 Chaper.

Table 6-4.   V4L2 Capture API IOCTLs

IOCTL	Description
VIDIOC_QUERYCAP	Query Device Capabiities
VIDIOC_G_FMT VIDIOC_S_FMT	Get or Set Data format
VIDIOC_S_DEST_CROP	Set cropping rectange
VIDIOC_REQBUFS	Initiate Memory Mapping
VIDIOC_QueryBUF	Query status of buffer
VIDIOC_QBUF,VIDIOC_DQBUF	Exchange buffer with driver
VIDIOC_STREAMON,VIDIOC_STREAMOFF	Start or stop streaming
VIDIOC_G_CTRL,VIDIOC_S_CTRL	
VIDIOC_CROPCAP	Query cropping capabilities
VIDIOC_G_CROP,VIDIOC_S_CROP	Get or set Cropping
VIDIOC_OVERLAY	Start or stop video overlay
VIDIOC_G_FBUF, VIDIOC_S_FBUF	Get or set frame buffer ovrelay parameters
VIDIOC_G_PARM,VIDIOC_S_PARM	Get or set streaming parameters
VIDIOC_G_STD,VIDIOC_S_STD	Get or Set the video standard
VIDIOC_G_OUTPUT,VIDIOC_S_OUTPUT	Get or Set the video output

Table continues on the next page...

Table 6-4. V4L2 Capture API IOCTLs (continued)

IOCTL	Description
VIDIOC_G_INPUT,VIDIOC_S_INPUT	Get or set the video input
VIDIOC_ENUMSTD	Enumerate video standards
VIDIOC_ENUMOUTPUT,VIDIOC_ENUMINPUT	Enumerate output and inputs
VIDIOC_ENUM_FMT	Enumerate image formats
VIDIOC_ENUM_FRAMESIZE,VIDIOC_ENUM_FRAMEINTERVALSS	Enumerate frame sizes and intervals
VIDIOC_DBG_G_CHIP_IDENT	Chip Identification

6.1.8 Source Code Structure

The table below shows the capture driver source files. For i.MX 6 and i.MX 7 the source files are in drivers/media/platform/mxc/capture. For i.MX 8 series the source files are in drivers/media/platform/imx8. For MIPI-CSI the source files are in drivers/mxc/mipi.

Table 6-5. Omnivision V4L2 Camera Driver Files

File	Description
• drivers/media/platform/imx8/mipi-csi2.c • drivers/media/platform/imx8/mipi-csi2.h • drivers/media/platform/imx8/mipi-csi2-yav.c	i.MX 8 MIPI-CSI2 Capture Interface driver
• drivers/media/platform/imx8/parallel-csi.c • drivers/media/platform/imx8/parallel-csi.h	i.MX 8 MIPI-CSI2 Parallel-CSI Interface driver
• drivers/media/platform/imx8/mxc-isi-core.c • drivers/media/platform/imx8/mxc-isi-cap.c • drivers/media/platform/imx8/mxc-isi-hw.c	i.MX 8 ISI Capture Controller driver
• drivers/media/platform/imx8/ov5640_v3.c • drivers/media/platform/imx8/max9286.c	i.MX 8 Omnivision Camera V3 Camera interface
• drivers/media/platform/imx8/mxc-jpeg-hw.c • drivers/media/platform/imx8/mxc-jpeg-hw.c	i.MX 8 JEPG hardware interface
• drivers/mxc/mipi/mipi-csi2.c • drivers/mxc/mipi/mipi-csi2.h	MX6 and MX7 MIPI-CSI2 interface core driver
• drivers/media/platform/mxc/capture/ ipu_bg_overlay_sdc.c • drivers/media/platform/mxc/capture/ipu_csi_enc.c • drivers/media/platform/mxc/capture/ ipu_fg_overlay_sdc.c • drivers/media/platform/mxc/capture/ipu_prp_enc.c • drivers/media/platform/mxc/capture/ ipu_prp_vf_sdc_bg.c • drivers/media/platform/mxc/capture/ipu_prp_vf_sdc.c • drivers/media/platform/mxc/capture/ipu_still.c • drivers/media/platform/mxc/capture/v4l2-int-device	i.MX 6 IPU V4L2 plugin
• drivers/media/platform/mxc/capture/mx6s_capture.c	CSI Omnivision Camera V4L2 plugin

Table continues on the next page...

Table 6-5. Omnivision V4L2 Camera Driver Files (continued)

File	Description
• drivers/media/platform/mxc/capture/ov5640.c • drivers/media/platform/mxc/capture/ ov5640_camera_int.c	
• drivers/media/platform/mxc/capture/ov5640_v2.c • drivers/media/platform/mxc/capture/ ov5640_camera_v2.c	Paralllel CSI Omnivision Camera V4L2 plugin
• drivers/media/platform/mxc/capture/ov5640_mipi.c • drivers/media/platform/mxc/capture/ ov5640_camera_mipi_int.c	MIPI-CSI Omnivision Camera V4L2 plugin
• drivers/media/platform/mxc/capture/ov5640_mipi_v2.c • drivers/media/platform/mxc/capture/ ov5640_camera_mipi_v2.c	MIPI-CSI2 Omnivision Camera V4L2 plugin
• drivers/media/platform/mxc/capture/adv7180.c • drivers/media/platform/mxc/capture/adv7180_tvin.c	TV Decoder ADV7180 V4L2

6.2 Display Overview

6.2.1 Introduction

The i.MX Display systems uses display controllers to optimize video data movement to display interfaces and graphics processing. Each display controller is implemented through a Linux driver and into a display framework either framebuffer or DRM. In some cases a display controller includes authentication ensuring a secure video pipeline. In others the display controller will include additional features for scaling, de-interlacing, tiling and color conversion during tranfers. For i.MX 8 supporting multiple displays is done with use of two controllers working together. This chapter provides a high level overview of i.MX display controllers and interfaces and the difference between framebuffer and DRM display drivers. The following display controllers are used.

- IPU
- PXP
- eLCDIF
- DPU
- DCSS - on i.MX 8M only

A display interface will interface to the display controller, display panel and in some cases encoders display bridges. The following display interfaces are supported.

- EPDC - supporting EInk displays
- Parallel - supporting LCD displays

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

- LVDS - supporting LVDS displays
- HDMI - supporting both on chip and external HDMI
- Display Port - supporting eDP panels
- MIPI-DSI - supporting MIPI displays

NOTE

Analog display is no longer supported. Analog interface was used i.MX 37 and i.MX 5 families.

The following HDMI display bridges/encoder are supported.

- Parallel to HDMI - using Silicon Image si902x
- LVDS to HDMI - using ITE it6263
- MIPI-DSI to HDMI - using Analog Devices adv7535

Each SOC supports different display features. Some of these are configured in the device trees located at arch/arm/boot/dts and arch/arm64/boot/dts. Go to the hardware reference manual for more details on the following.

- Throughput
  - Number of outputs
  - Pixel clock rate
  - Max number of displays and corresponding resolution
  - Resolution at 60 Hz.
- Interface
  - Parallel - number of ports and bit size
  - LVDS - number of lanes and channels.
  - MIPI-DSI - number of ports, lanes channels and speed
- Processing
  - On the fly combining including high resolution displays
  - Off-line combining speed

6.2.2 Frame Buffer

Frame buffer drivers are supported for i.MX 6 and i.MX 7 but not for i.MX 8. The frame buffer drivers are supported using the the imxfb driver in drivers/video/fbdev. The frame buffer kernel fbdev structure is defined here or here on kernel.org. For more information on i.MX V4LS go the V4L2 chapter.

The panels are supported with the framebuffer driver for the TRULY and EInk panels in the video/fbdev/mxc folder. See the panels supported by seaching for PANEL in the imx_v7_defconfig. The Trully panels are only supported with the MIPI DSI interface. The Eink panels are only supported with the EPDC interface.

6.2.3 Direct Render Model (DRM)

Direct Render Model (DRM) is the new display driver use for i.MX. The i.MX DRM driver is in drivers/gpu/drm/imx. Other components have DRM interfaces such as GPU and DCSS. The DRM framework is documented here on kernel.org.

The i.MX DRM drivers are implemented with the following drivers.

- Hardware library support files
- Core DRM drivers
- Hardware dependent DRM drivers
- HDMI DRM drivers supporting hdp HDMI/Display Port

The DRM driver uses the DPU on the i.MX 8QuadMax and i.MX 8QuadXPlus and uses LCDIF for the i.MX 8M Quad and i.MX 8M Mini.

The i.MX DRM framework also includes panel drivers which exist in driver/gpu/panel. The supported drm panels are Simple and Raydium RM67191.

6.2.4 Display Resolution

The display resoluton calculation uses the following factors.

- Frame Width
- Frame Height
- Frame rate (fps)
- Blanking Interval - provided in the display's DS up to 35% (1.35) - use minn values

The pixel clock [MHz] is calculated according to Frame Width x Frame Height x Frames Rate x Blanking Interval

Things to consider are the following

- Data format (pixel per clock)
- Display's source clock (DI#_CLK_EXT bit
- The load on the display controller (DC)

6.2.5 Authentication

Display authentication allows hardware processing to ensure display content is not compromised. This is done through a display authentication CRC using the authentication hardware This hardware is the DCIC integrated through the frame buffer display framework on i.MX 6 and the DPU implemented in the DRM display framework for i.MX 8.

Display authentication CRC is supported on the following SoC.

- i.MX 6 Solox supports authentication using DCIC for 1 display.
- i.MX 6 QuadPlus/Quad/Dual support authentication using DCIC with 2 displays.
- i.MX 8 8QuadXPlus can authenticate 2 display using the DPU.
- i.MX 8QuadPlus can authenticate 4 displays for i.MX 8QuadMax using DPU.

6.2.6 Tiling

Tiling through hardware provides optimized video data display. This is implemented through different hardware blocks. The newest feature is the Display Prefetch Resolve (DPR) which increases performance on the i.MX 6 QuadPlus, i.MX 8QuadMax and i.MX 8QuadXPlus.

Tile support is enabled on the following

- i.MX 6Quad/Dual supports tiling using Video Data Order Adapter (VDOA).
- i.MX 6QuadPlus supports both tiling VDOA and Display Prefetch Resolve (DPR) version 1
- i.MX 8QuadXPlus and i.MX 8QuadMax supports tiling using Display Prefetch Resolve (DPR) version2

6.3 Display Controllers

6.3.1 Display Processing Unit (DPU)

6.3.1.1 Introduction

The display processing unit (DPU) is designed to support video and graphics processing functions and to interface with video and still display sensors and displays. The DPU driver provides internel kernel-level APIs to manipulate logical channels. A logical

channel represents a complete DPU processing flow. For example, a complete DPU processing flow (logical channel) might consist of reading a YUV buffer from memory and displaying it to an external interface. The DPU API consists of a set of common functions for all channels. Its functions are to initialize channels, set up buffers, enable and disable channels and set up interrupts.

Typical logical channels include:

- CSI direct to memory
- Memory to synchronous frame buffer background
- Memory to synchronous frame buffer foreground

The higher level drivers are responsible for memory allocation and providing user-level API. DPU interfaces are available for capture in the V4L2 framework and for display using the DRM display framework. DPU interfaces with LVDS, MIPI-DSI, HDMI and Parallel display interfaces.

The DPU display controller supports a 32bit display composition engine that includes the following:

- 2 Display output streams on independent panesl.
- Two layer composition
- Automatic safety stream pani plus detection using CRC matching using a Region CRC checker

The DPU display controller supports a 2D composition enginer which provides efficiency, performance and safety. The DPU 2D graphics engine support reduces the burden on the GPU so it only does 3D GPU. Video efficiency with overlay native video and graphics uses minimal access to system memory. Power efficiencies allow the 3D engine to be off for windowing GUI's like the Android Hardware Composer.

The DPU also supports the following for authentication.

- CRC checker with 8 stackable regions maskable, exclusive top-to-bottom priority
- CRC check can be inserted after any stage in the post-processing pipe
- CRC failure can generate SW interrupt, or switch the Frame Gen to either Safety Stream or Constant Plane

The DPU display interfacce cache supports the following.

- Each display plane has a multi line cache
- This contains 8 lines of pixels for each plane
- RGB, YUV etc formats supported
- Supports Video and GPU tile formats
- Contents are fetched from memory to fill cache ahead of time

- Horizontal and vertical fetches supported
- Warp fetches not supported, require bypass

6.3.1.2 DRM

The display processing unit (DPU) interfaces with the DRM driver supporting video display.

6.3.1.3 Source Code Structure

The DPU drivers are separating into DRM, blitting and main processing. Common functions are provided in the drivers/gpu/drm/imx/dpu and drivers/gpu/imx/dpu-blit while the main driver exists in drivers/gpu/imx/dpu. The following table lists the source files.

Table 6-6. DPU Driver source

File	Description
DRM Source	
drivers/gpu/drm/imx/dpu/dpu-plane.	DRM DPU Plane
drivers/gpu/drm/imx/dpu/dpu-crtc	DRM DPU CRTC
drivers/gpu/drm/imx/dpu/dpu-blit	DRM DPU blitter
drivers/gpu/drm/imx/dpu/dpu-kms	DRM DPU KMS
DPU Blitter Source	
drivers/gpu/imx/dpu-blit/dpu-blit	DPU Bliter
drivers/gpu/imx/dpu-blit/dpu-blit-registers.h	DPU Blit registers
DRM Core Source	
drivers/gpu/imx/dpu/dpu-vscaler.c	DPU VScaler
drivers/gpu/imx/dpu/dpu-fetchwarp.c	DPU Fetchwarp
drivers/gpu/imx/dpu/constframe.c	DPU Const Frame
drivers/gpu/imx/dpu/dpu-prv.h	DPU Private headers
drivers/gpu/imx/dpu/dpu-disengcfg.c	DPU Display Configurations
drivers/gpu/imx/dpu/dpu-fetchunit.c	DPU Fetch Unit
drivers/gpu/imx/dpu/dpu-framegen.c	DPU Frame Generator
drivers/gpu/imx/dpu/dpu-hscaler.c	DPU HScaler
drivers/gpu/imx/dpu/dpu-extdst.c	DPU External Destination
drivers/gpu/imx/dpu/dpu-common.c	DPU Common
drivers/gpu/imx/dpu/dpu-fetchlayer.c	DPU Fetch Layer
drivers/gpu/imx/dpu/dpu-tcon.c	DPU TCon
drivers/gpu/imx/dpu/dpu-layerblend.c	DPU Layer Blend
drivers/gpu/imx/dpu/dpu-fetcheco.c	DPU Fetch Encode

Table continues on the next page...

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

Table 6-6.   DPU Driver source (continued)

File	Description
drivers/gpu/imx/dpu/dpu-fetchdecode.c	DPU Decode

6.3.1.4   Menu Configuration Options

The following Linux kernel configuration options are provided for the DPU module.

Device Drivers -> i.MX DPU core support

6.3.2   Image Processing Unit (IPU)

6.3.2.1   Introduction

The image processing unit (IPU) is designed to support video and graphics processing functions and to interface with video and still image sensors and displays. The IPU driver provides a kernel-level API to manipulate logical channels. A logical channel represents a complete IPU processing flow. For example, a complete IPU processing flow (logical channel) might consist of reading a YUV buffer from memory, performing post-processing, and writing an RGB buffer to memory. A logical channel maps one to three IDMA channels and maps to either zero or one IC tasks. A logical channel can have one input, one output, and one secondary input IDMA channel. The IPU API consists of a set of common functions for all channels. Its functions are to initialize channels, set up buffers, enable and disable channels, link channels for auto frame synchronization, and set up interrupts.

The IPU is a display controller and supports the following display interfaces which are supported through the framebuffer display framework. The access is only exposed through the framebuffer fbdev application framework.

- Parallel
- LVDS
- HDMI
- MIPI-DSI

Typical logical channels include:

- CSI direct to memory
- CSI to viewfinder pre-processing to memory

- Memory to viewfinder pre-processing to memory
- Memory to viewfinder rotation to memory
- Previous field channel of memory to video deinterlacing and viewfinder pre-processing to memory
- Current field channel of memory to video deinterlacing and viewfinder pre-processing to memory
- Next field channel of memory to video deinterlacing and viewfinder pre-processing to memory
- CSI to encoder pre-processing to memory
- Memory to encoder pre-processing to memory
- Memory to encoder rotation to memory
- Memory to post-processing rotation to memory
- Memory to synchronous frame buffer background
- Memory to synchronous frame buffer foreground
- Memory to synchronous frame buffer DC
- Memory to synchronous frame buffer mask

The IPU API has some additional functions that are not common across all channels, and are specific to an IPU sub-module. The types of functions for the IPU sub-modules are as follows:

- Synchronous frame buffer functions
- Panel interface initialization
- Set foreground positions
- Set local/global alpha and color key
- Set gamma
- CSI functions
- Sensor interface initialization
- Set sensor clock
- Set capture size
- Enable or disable prefetching linear frames by using PRE/PRG
- Enable or disable resolving tiled frames by using PRE/PRG

The higher level drivers are responsible for memory allocation, chaining of channels, and providing user-level API.

6.3.2.2   Hardware Operation

The detailed hardware operation of the IPU is discussed in the Applications Processor Reference Manual. The following figure shows the IPU hardware modules.
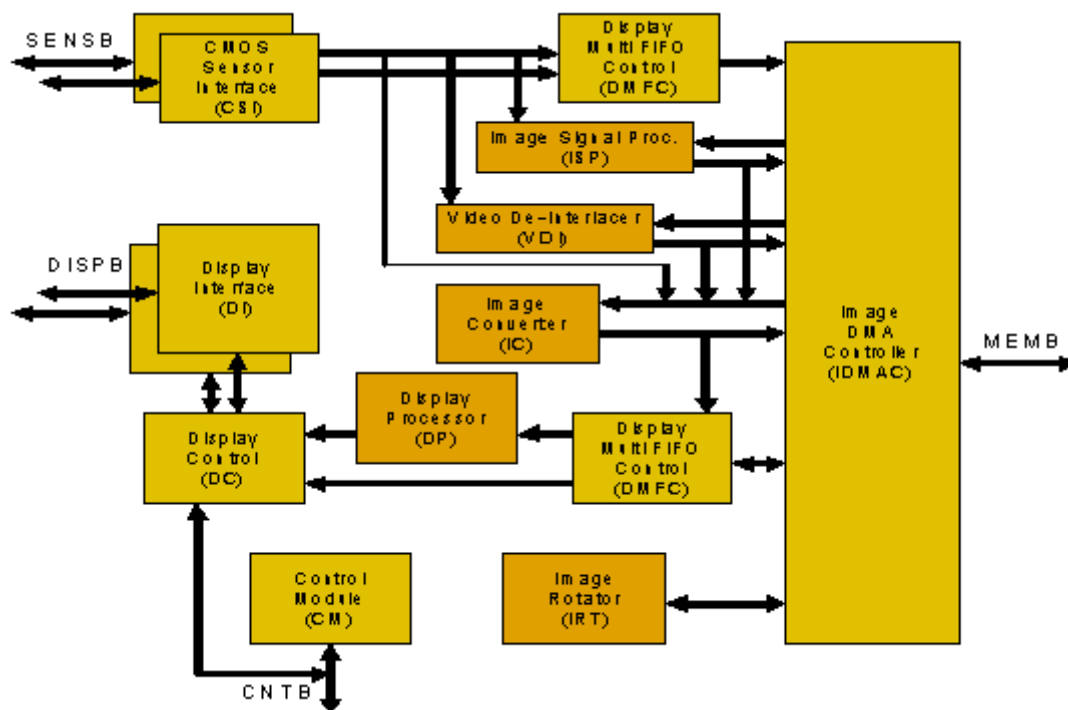
Figure 6-1. IPUv3EX/IPUv3H IPU Module Overview

6.3.2.3   Software Operation

The IPU driver is a self-contained driver module in the Linux kernel.

It consists of a custom kernel-level API for the following blocks:

• Synchronous frame buffer driver
• Display Interface (DI)
• Display Processor (DP)
• Image DMA Controller (IDMAC)
• CMOS Sensor Interface (CSI)
• Image Converter (IC)
• Prefetch/Resolve Engine/Gasket (PRE/PRG)

Figure below shows the interaction between the different graphics/video drivers and the IPU.
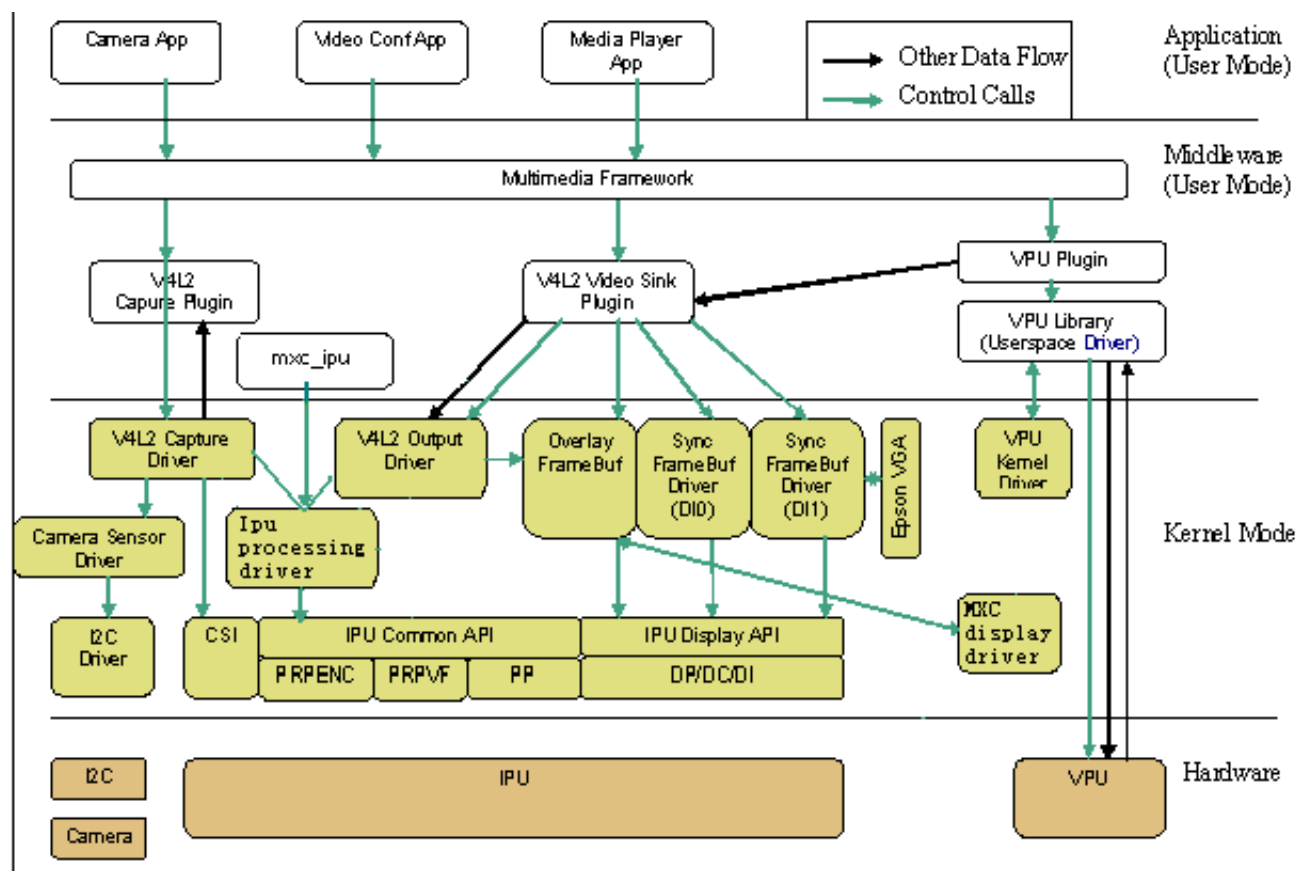
Figure 6-2. Graphics/Video Drivers Software Interaction for IPUv3

The drivers for IPUv1 are named simply ipu. Drivers for IPUv3 contain 3 or v3 in the name. The IPU drivers are sub-divided as follows:

- Device drivers-include the frame buffer driver for the synchronous frame buffer, the frame buffer driver for the displays, V4L2 capture drivers for IPU pre-processing, the V4L2 output driver for IPU post-processing, and the ipu processing driver which provide system interface to user space or V4L2 drivers. The frame buffer device drivers are available in `drivers/video/mxc`. The V4L2 device drivers are available in `drivers/media/platform/mxc`.
- The MXC display driver is introduced as a simple framework to manage interaction between IPU and display device drivers (e.g., LCD, LVDS, HDMI, MIPI, etc.)
- Low-level library routines-interface to the IPU hardware registers. They take input from the high-level device drivers and communicate with the IPU hardware. The low-level libraries are available in the directory of the Linux kernel.

6.3.2.4   IPU Frame Buffer Drivers Overview

The frame buffer device provides an abstraction for the graphics hardware. It represents the frame buffer video hardware, and allows application software to access the graphics hardware through a well-defined interface, so that the software is not required to know anything about the low-level hardware registers.

The driver is enabled by selecting the frame buffer option under the graphics parameters in the kernel configuration. To supplement the frame buffer driver, the kernel builder may also include support for fonts and a startup logo. This device depends on the virtual terminal (VT) console to switch from serial to graphics mode. The device is accessed through special device nodes, located in the /dev directory, as /dev/fb*. fb0 is generally the primary frame buffer.

Other than the physical memory allocation and LCD panel configuration, the common kernel video API is utilized for setting colors, palette registration, image blitting, and memory mapping. The IPU reads the raw pixel data from the frame buffer memory and sends it to the panel for display.

6.3.2.5   IPU Frame Buffer Hardware Operation

The frame buffer interacts with the IPU hardware driver module.

6.3.2.6   IPU Frame Buffer Software Operation

A frame buffer device is a memory device, such as /dev/mem, and it has features similar to a memory device. Users can read it, write to it, seek to some location in it, and mmap() it (the main use). The difference is that the memory that appears in the special file is not the whole memory, but the frame buffer of some video hardware.

/dev/fb* also interacts with several IOCTLs, which allows users to query and set information about the hardware. The color map is also handled through IOCTLs. For more information on what IOCTLs exist and which data structures they use, see `include/uapi/linux/fb.h`. The following are a few of the IOCTLs functions:

- Request general information about the hardware, such as name, organization of the screen memory (planes, packed pixels, and so on), and address and length of the screen memory.
- Request and change variable information about the hardware, such as visible and virtual geometry, depth, color map format, timing, and so on. The driver suggests

values to meet the hardware capabilities (the hardware returns EINVAL if that is not possible) if this information is changed.

- Get and set parts of the color map. Communication is 16 bits-per-pixel (values for red, green, blue, transparency) to support all existing hardware. The driver does all the calculations required to apply the options to the hardware (round to fewer bits, possibly discard transparency value).

The hardware abstraction makes the implementation of application programs easier and more portable. The only thing that must be built into the application programs is the screen organization (bitplanes or chunky pixels, and so on), because it works on the frame buffer image data directly.

The MXC frame buffer driver (drivers/video/mxc/mxc_ipuv3_fb.c) interacts closely with the generic Linux frame buffer driver (`drivers/video/fbdev/core/fbmem.c`).

6.3.2.7 Synchronous Frame Buffer Driver

The synchronous frame buffer screen driver implements a Linux standard frame buffer driver API for synchronous LCD panels or those without memory. The synchronous frame buffer screen driver is the top level kernel video driver that interacts with kernel and user level applications. This is enabled by selecting the Synchronous Panel Frame buffer option under the graphics support device drivers in the kernel configuration. To supplement the frame buffer driver, the kernel builder may also include support for fonts and a startup logo. This depends on the VT console for switching from serial to graphics mode.

Except for physical memory allocation and LCD panel configuration, the common kernel video API is utilized for setting colors, palette registration, image blitting, and memory mapping. The IPU reads the raw pixel data from the frame buffer memory and sends it to the panel for display.

The frame buffer driver supports different panels as a kernel configuration option. Support for new panels can be added by defining new values for a structure of panel settings.

The frame buffer interacts with the IPU driver using custom APIs that allow:

- Initialization of panel interface settings
- Initialization of IPU channel settings for LCD refresh
- Changing the frame buffer address for double buffering support

The following features are supported:

- Configurable screen resolution

- Configurable RGB 16, 24, or 32 bits per pixel frame buffer
- Configurable panel interface signal timings and polarities
- Palette/color conversion management
- Power management
- LCD power off/on
- Enable/disable PRE/PRG features

User applications utilize the generic video API (the standard Linux frame buffer driver API) to perform functions with the frame buffer. These include the following:

- Obtaining screen information, such as the resolution or scan length
- Allocating user space memory using mmap for performing direct blitting operations

A second frame buffer driver supports a second video/graphics plane.

6.3.2.8 IPU Backlight Driver

IPU drivers also control the backlight. The IPU backlight driver implements IPU PWM backlight control for panels. It exports a sys control file under /sys/class/backlight/pwm-backlight.0/brightness to user space. The default backlight intensity value is 128.

6.3.2.9 IPU Device Driver

IPU (processing) device driver provide image processing features: resizing/rotation/CSC/ combination/deinterlacing based on IC/IRT modules in IPUv3.

The IPU device driver is task based, user just need prepare task setting, queue task, then block wait task finish. The driver now supports only blocking method, and the non-block method will be added in the future. The task structures are as follows:

```
struct ipu_task {
        struct ipu_input input;
        struct ipu_output output;

        bool overlay_en;
        struct ipu_overlay overlay;

#define IPU_TASK_PRIORITY_NORMAL 0
#define IPU_TASK_PRIORITY_HIGH   1
        u8      priority;

#define IPU_TASK_ID_ANY 0
#define IPU_TASK_ID_VF  1
#define IPU_TASK_ID_PP  2
#define IPU_TASK_ID_MAX 3
        u8      task_id;

        int     timeout;
};
```

```
struct ipu_input {
        u32 width;
        u32 height;
        u32 format;
        struct ipu_crop crop;
        dma_addr_t paddr;

        struct ipu_deinterlace deinterlace;
        dma_addr_t paddr_n; /*valid when deinterlace enable*/
};

struct ipu_overlay {
        u32 width;
        u32 height;
        u32 format;
        struct ipu_crop crop;
        struct ipu_alpha alpha;
        struct ipu_colorkey colorkey;
        dma_addr_t
paddr;

};




struct ipu_output
{


        u32 width;
        u32 height;
        u32 format;
        u8 rotate;
        struct ipu_crop crop;
        dma_addr_t paddr;
};
```

To prepare the task, the user just needs to fill task.input, task.overlay (if need combine) and task.output parameters, and then queue task either by `int ipu_queue_task(struct ipu_task *task);` if from the kernel level (V4L2 driver for example), or by IPU_QUEUE_TASK ioctl under /dev/mxc_ipu if from the application level.

6.3.2.10 Source Code Structure

The source files associated with the IPU, Sensor, V4L2, and Panel drivers are available in the following folders.

- drivers/mxc/ipu3
- drivers/video/mxc
- drivers/video/fbdev/mxc
- drivers/video/backlight

See the V4L2 chapter for more information on the IPU V4L2 driver files

Table 6-7. IPU Driver Files

File	Description
driveers/mxc/ipu3/ipu_common.c	IPU common library functions
driveers/mxc/ipu3/ipu_common.c	IPU common library functions
drivers/mxc/ipu3/ipu_ic.c	IPU IC base driver
drivers/mxc/ipu3/ipu_device.c	IPU driver device interface and fops functions.
drivers/mxc/ipu3/ipu_capture.c	IPU CSI capture base driver
drivers/mxc/ipu3/ipu_disp.c	IPU display functions
drivers/mxc/ipu3/ipu_calc_stripes_sizes.c	Multistripes method functions for ipu_device.c
drivers/mxc/ipu3/pre.c	i.MX 6 QuadPlus Prefetch/Resolve the engine driver
drivers/mxc/ipu3/prg.c	i.MX 6 QuadPlus Prefetch/Resolve the Gasket driver
drivers/mxc/ipu3/mxc_ipuv3_fb.c	Driver for synchronous frame buffer
drivers/mxc/ipu3/vdoa.c	VDOA post-processing driver, used by ipu_device.c
drivers/video/fbdev/mxc/mxc_lcdif.c	Display Driver for CLAA-WVGA and SEIKO-WVGA LCD support
drivers/video/fbdev/mxc/mxc_hdmi.c	Display Driver for HDMI interface
drivers/video/fbdev/mxc/ldb.c	Driver for synchronous frame buffer for on chip LVDS
drivers/video/fbdev/mxc/mxc_dispdrv.c	Display Driver framework for synchronous frame buffer
drivers/video/fbdev/mxc/mxc_edid.c	Driver for EDID

Table 6-8 lists the header files associated with the IPU and Panel drivers.

Table 6-8. IPU Global Header Files

File	Description
drivers/mxc/ipu3/ipu_param_mem.h	Hellper functions for IPU parameter memory access
drivers/mxc/ipu3/ipu_prv.h	Header file for Pre-processing drivers
drivers/mxc/ipu3/ipu_regs.h	IPU register definitions
drivers/mxc/ipu3/pre-regs.h	Prefetch/Resolve Engine register definitions
drivers/mxc/ipu3/prg-regs.h	Prefetch/Resolve Gasket register definitions
drivers/mxc/ipu3/vdoa.h	Header file for VDOA drivers
drivers/video/fbdev/mxc/mxc_dispdrv.h	Header file for display driver
include/linux/uapi/mxcfb.h	Header file for the synchronous framebuffer driver
include/linux/uapi/ipu.h	Header file for IPU APIs

6.3.2.11   Menu Configuration Options

The following Linux kernel configuration options are provided for the IPU module.

In menu configuration enable the following module:

- CONFIG_MXC_IPU_V3 - Includes support for the Image Processing Unit. In menuconfig, this option is available under:

  Device Drivers > MXC support drivers > Image Processing Unit Driver

  By default, this option is Y for all architectures.

  If ARCH_MXC is true, CONFIG_MXC_IPU_V3 will be set.

- CONFIG_MXC_IPU_V3_PRG - This enables support for the IPUv3 prefetch gasket engine to support double buffer handshake control between IPUv3 and prefetch engine (PRE), snoop the AXI interface for display refresh requests to memory, and modify the request address to fetch the double buffered row of blocks in OCRAM.

  Device Drivers > MXC support drivers > i.MX IPUv3 prefetch gasket engine

  This option depends on CONFIG_MXC_IPU_V3 and CONFIG_MXC_IPU_V3_PRE.

- CONFIG_MXC_IPU_V3_PRE - This enables support for the IPUv3 prefetch engine to improve the system memory performance. The engine has the capability to resolve framebuffers in tile pixel format to linear.

  Device Drivers > MXC support drivers > i.MX IPUv3 prefetch engine

  This option depends on CONFIG_MXC_IPU_V3. Enabling this option selects CONFIG_MXC_IPU_V3_PRG.

- CONFIG_MXC_CAMERA_OV5640_MIPI - Option for both the OV 5640 mipi sensor driver and the use case driver. This option is dependent on the VIDEO_MXC_CAPTURE option. In menuconfig, this option is available under:

  Device Drivers > Multimedia support > V4L platform devices > MXC Video For Linux Video Capture > MXC Camera/V4L2 PRP Features support > OmniVision 5640 Camera support using mipi

- CONFIG_MXC_CAMERA_OV5640 - Option for both the OV5640 sensor driver and the use case driver. This option is dependent on the VIDEO_MXC_CAPTURE option. In menuconfig, this option is available under:

  Device Drivers > Multimedia platform > V4L platform devices > MXC Video For Linux Video Capture > MXC Camera/V4L2 PRP Features support > OmniVision ov5640 camera support

  Only one sensor should be installed at a time.

- CONFIG_MXC_IPU_PRP_VF_SDC - Option for the IPU (here the > symbols illustrates data flow direction between HW blocks):

CSI > IC > MEM MEM > IC (PRP VF) > MEM

Use case driver for dumb sensor or

CSI > IC(PRP VF) > MEM

for smart sensors. In menuconfig, this option is available under:

Multimedia devices > Video capture adapters > MXC Video For Linux Camera > MXC Camera/V4L2 PRP Features support > Pre-Processor VF SDC library

By default, this option is M for all.

- CONFIG_MXC_IPU_PRP_ENC - Option for the IPU:

  Use case driver for dumb sensors

  CSI > IC > MEM MEM > IC (PRP ENC) > MEM

  or for smart sensors

  CSI > IC(PRP ENC) > MEM.

  In menuconfig, this option is available under:

  Device Drivers > Multimedia Devices > Video capture adapters > MXC Video For Linux Camera > MXC Camera/V4L2 PRP Features support > Pre-processor Encoder library

  By default, this option is set to M for all.

- CONFIG_VIDEO_MXC_CAMERA - This is configuration option for V4L2 capture Driver. This option is dependent on the following expression:

  VIDEO_DEV && MXC_IPU && MXC_IPU_PRP_VF_SDC && MXC_IPU_PRP_ENC

  In menuconfig, this option is available under:

  Device Drivers > Multimedia devices > Video capture adapters > MXC Video For Linux Camera

  By default, this option is M for all.

- CONFIG_VIDEO_MXC_OUTPUT - This is configuration option for V4L2 output Driver. This option is dependent on VIDEO_DEV && MXC_IPU option. In menuconfig, this option is available under:

  Device Drivers > Multimedia devices > Video capture adapters > MXC Video for Linux Video Output

By default, this option is Y for all.

- CONFIG_FB - This is the configuration option to include frame buffer support in the Linux kernel. In menuconfig, this option is available under:

  Device Drivers > Graphics support > Support for frame buffer devices

  By default, this option is Y for all architectures.

- CONFIG_FB_MXC - This is the configuration option for the MXC Frame buffer driver. This option is dependent on the CONFIG_FB option. In menuconfig, this option is available under:

  Device Drivers > Graphics support > MXC Framebuffer support

  By default, this option is Y for all architectures.

- CONFIG_FB_MXC_SYNC_PANEL - This is the configuration option that chooses the synchronous panel framebuffer. This option is dependent on the CONFIG_FB_MXC option. In menuconfig, this option is available under:

  Device Drivers > Graphics support > MXC Framebuffer support > Synchronous Panel Framebuffer

  By default this option is Y for all architectures.

- CONFIG_FB_MXC_LDB - This configuration option selects the LVDS module on i.MX 6 chip. This option is dependent on CONFIG_FB_MXC_SYNC_PANEL and CONFIG_MXC_IPUV3 || FB_MXS options. In menuconfig, this option is available under:

  Device Drivers > Graphics support > MXC Framebuffer support > Synchronous Panel Framebuffer > MXC LDB

- CONFIG_FB_MXC_SII9022 - This configuration option selects the SII9022 HDMI chip. This option is dependent on CONFIG_FB_MXC_SYNC_PANEL option. In menuconfig, this option is available under:

  Device Drivers > Graphics support > MXC Framebuffer support > Synchronous Panel Framebuffer > Si Image SII9022 DVI/HDMI Interface Chip

6.3.3 Pixel Pipeline (PxP)

6.3.3.1 Introduction

The PxP is a display controller that works wtih the EPDC display interface. The Pixel Pipeline (PxP) DMA engine driver provides a unique API, which are implemented as a dmaengine client that smooths over the details of different hardware offload engine implementations. Typically, the users of PxP DMA-ENGINE driver include EPDC driver, V4L2 Output driver, and the PxP user-space library.

The PxP driver uses PxP registers to interact with the hardware. For detailed hardware operations, see the Applications Processor Reference Manual document associated with SoC.

6.3.3.2 Software Operation

There are different versions of PxP IP. To ease the maintenance for the new version of PxP used on i.MX 7Dual, which has new features mainly for EPDC like hardware collision detection, E Ink Gen-II waveform algorithm (REAGL/-D) processing in hardware, and hardware dithering support, there are different drivers (drivers/dma/pxp/pxp_dma_v3.c). However, each version uses the DMA Engine framework.

6.3.3.3 Key Data Structs

The PxP DMA Engine driver implementation depends on the DMA Engine Framework. There are three important structs in the DMA Engine Framework which are extended by the PxP driver: struct dma_device, struct dma_chan, struct dma_async_tx_descriptor. The PxP driver implements several callback functions which are called by the DMA Engine Framework (or DMA slave) when a DMA slave (client) interacts with the DMA Engine.

The PxP driver implements the following callback functions in struct dma_device:

device_alloc_chan_resources / allocate resources and descriptors */*

device_free_chan_resources / release DMA channel's resources */*

device_tx_status / poll for transaction completion */*

device_issue_pending / push pending transactions to hardware */*

and,

device_prep_slave_sg / prepares a slave DMA operation */*

device_terminate_all/ manipulate all pending operations on a channel, returns zero or error code */*

The first four functions are used by the DMA Engine Framework, the last two are used by the DMA slave (DMA client). Notably, *device_issue_pending* is used to trigger the start of a PxP operation.

The PxP DMA driver also implements the interface *tx_submit in struct dma_async_tx_descriptor*, which is used to prepare the descriptor(s) which will be executed by the engine. When tasks are received in pxp_tx_submit, they are not configured and executed immediately. Rather, they are added to a task queue and the function call is allowed to return immediately.

6.3.3.4   Channel Management

Although ePxP does not have multiple channels in hardware, the virtual channels are supported in the driver. This provides flexibility in the multiple instance/client design. At any time, a user can call *dma_request_channel*() to get a free channel, and then configure this channel with several descriptors. A descriptor is required for each input plane and for the output plane. When the PxP is no longer being used, the channel should be released by calling *dma_release_channel*(). Detailed elements of channel management are handled by the driver and are transparent to the client.

6.3.3.5   Descriptor Management

The DMA Engine processes the task based on the descriptor. One DMA channel is usually associated with several descriptors. Descriptors are recycled resources, under control of the offload engine driver, to be reused as operations complete. The extended TX descriptor packet (pxp_tx_desc), allows the user to pass PxP configuration information to the driver. This includes everything that the PxP needs to execute a processing task.

6.3.3.6   Completion Notification

There are two ways for an application to receive notification that a PxP operation has completed.

- Call dma_wait_for_async_tx(). This call causes the CPU to spin while it polls for the completion of the operation.
- Specify a completion callback.

The latter method is recommended. After the PxP operation completes, the PxP output buffer data can be retrieved.

For general information for DMA Engine Framework, see*Documentation/dmaengine.txt* in the Linux kernel source tree.

6.3.3.7   Limitations

- The driver currently does not support scatterlist objects in the way they are traditionally used. Instead of using the scatterlist parameter object to provide a chain of memory sources and destinations, the driver currently uses it to provide the input and output buffers (and overlay buffers, if needed) for one transfer.
- The PxP driver may not properly execute a series of transfers that is queued in rapid sequence. It is recommended to wait for each transfer to complete before submitting a new one.

6.3.3.8   Menu Configuration Options

The following Linux kernel configuration option is provided for this module:

For i.MX 7Dual select Device Drivers > DMA Engine support > [*] MXC PxP V3 support > [*] MXC PxP Client Device

For i.MX 6 select Device Drivers > DMA Engine support > [*] MXC PxP V2 support > [*] MXC PxP Client Device

6.3.3.9   Source Code Structure

The PxP driver source code is located in drivers/dma/pxp.

Table 6-9.   Pxp source

File	Description
drivers/dma/pxp/pxp_device.c	PxP Device
drivers/dma/pxp/pxp_dma_v2.c	PxP DMA for i.MX 6
drivers/dma/pxp/pxp_dma_v3.c	PxP DMA for i.MX 7
drivers/dma/pxp/regs-pxp_v2.h	PxP registers for i.MX 6
drivers/dma/pxp/regs-pxp_v3.h	Pxp registers for i.MX 7
drivers/dma/pxp/regs-pxp_v3.h	Pxp registers for i.MX 7
include/linux/drivers/dma/pxp/pxp_dma_v3.c	PxP Device for i.MX 7
include/linux/pxp_dma.h	PxP DMA kernel header
include/linux/pxp_device.h	PxP Device kernel header
include/linux/uapi/pxp_dma.h	PxP DMA user space header
include/linux/uapi/pxp_device.h	PxP Device user space header

Table continues on the next page...

Table 6-9. Pxp source (continued)

File	Description
drivers/media/platform/mxc/output/mxc_pxp_v4l2.c	PxP V4L2
drivers/media/platform/mxc/output/mxc_pxp_v4l2.c	PxP V4L2 header
drivers/media/platform/mxc/output/mxc_vout.c	i.MX V4L2 Output driver

6.3.4   ELCDIF Frame Buffer

6.3.4.1   Introduction

The eLCDIF is a display controller that works with the Parallel LCD interface. The driver is implemented as a display subsystem driver either frame buffer or DRM which controls generic LCD low-level operations allowing low level hardware control. Only DOTCLK mode of the ELCDIF is tested, so theoretically the ELCDIF frame buffer driver can work with a sync LCD panel driver to support a frame buffer device. The sync LCD driver is organized in a flexible and extensible manner and is abstracted from any specific sync LCD panel support. To support another sync LCD panel, the user can write a sync LCD driver by referring to the existing ones.

6.3.4.2   Software Operation

For the eLCDIF implemented as a framebuffer driver the frame buffer device is a memory device similar to /dev/mem. It can be read from, written to, or some location in it can be sought and mapped using mmap(). The difference is that the memory available to the user is not the entire allocated memory, but only the frame buffer of the video hardware. The device is accessed through special device nodes, usually located in the /dev directory, /dev/fb*. /dev/fb* also has several IOCTLs which act on it and through which information about the hardware can be queried and set. The color map handling operates through IOCTLs as well. See linux/fb.h for more information on which IOCTLs there are and which data structures are used.

The i.MX ELCDIF frame buffer driver implementation is abstracted from the actual hardware. The default panel driver is picked up by video mode defined in platform data or passed in with 'video=mxc_elcdif_fb:resolution, bpp=bits_per_pixel' kernel bootup command during probing. The resolution should be in the common frame buffer video mode pattern and bits_per_pixel should be the frame buffer's color depth.

6.3.4.3 Menu Configuration Options

The following menu options will configure the MXC ELCDIF frame buffer driver. This option depends on FB and (ARCH_MXS || ARCH_MXC).

Frame buffer Devices > MXS LCD framebuffer support (CONFIG_FB_MXS)

6.3.4.4 Source Code Structure

The source for frame buffer is in drivres/video/fbdev/mxc and the drm driver is in drivers/gpu/drm/imx/lcdif.

Table 6-10. ELCIF source

File	Description
drivers/video/fbdev/mxsfb.c	ELCDIF frame buffer driver
drivers/video/fbdev/mxc/mxc_lcdif.c	ELCDIF frame buffer driver
drivers/gpu/drm/imx/lcdif/lcdif-crtc.c	ELCDIF DRM Authentication
drivers/gpu/drm/imx/lcdif/lcdif-kms.c	ELCDIF DRM KMS
drivers/gpu/drm/imx/lcdif/lcdif-kms.h	ELCDIF DRM KMS Header
drivers/gpu/drm/imx/lcdif/lcdif-plane.c	ELCDIF DRM Plane
drivers/gpu/drm/imx/lcdif/lcdif-plane.h	ELCDIF DRM Plane header

6.3.5 Display Control Subsystem (DCSS)

6.3.5.1 Introduction

The Display control subsystem (DCSS) is a display control for i.MX 8M Quad that integrates through the DRM display framework. The DCSS provides a mechanmism to display frame buffers in memory out to UltraHD or HDTVs with the capabillity to combine up to 3 layers of graphics or video overlay to the HDMI output. The key featuers of the DCSS controller include:

- Supports up to 3 layers of graphics or video
    - Arbitrary offset
    - One plane can be graphics with 8 bit alpha support
    - Upscale 1920x1080p60 video or graphics to 3840x2160p60
    - Downscale 3840x2160p30 video to 1920x1080p30 or 1280x720p30
- HDR support:

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

- HDR10 with 2084 and 2020 color spaces
- Dolby Vision single and dual layer formats
- HLG
- HDMI 2.0a supporting one display:
  - Resolutions of: 640x480p60, 720x480p60, 1280x720p60, 1920x1080p60,3840x2160p60, 4096x2160p60
  - HDCP 2.2 and HDCP 1.4
- Pixel clock up to 596 MHz
- Output can also go to MIPI DSI output
- Frame Buffer Compression – Lossless compression of buffers

6.3.5.2   Source Code Structure

The DCSS drm driver is located in drivers/gpu/drm/imx/dcss and the DCSS core driver are in drivers/gpu/imx/dcss

Table 6-11.   DCSS Driver source

File	Description
drivers/gpu/drm/dcss/dcss-plane	DRM DCSS Plane
drivers/gpu/drm/dcss/dcss-kms	DRM DCSS KMS
drivers/gpu/drm/dcss/dcss-crtc	DRM DCSS CRTC header
drivers/gpu/imc/dcss/dcss-dec400d.c	DCSS dec400d
drivers/gpu/imx/dcss/dcss-scaler	DCSS Scaler
drivers/gpu/imx/dcss/dcss-ss.c	DCSS ss
drivers/gpu/imx/dcss/dcss-hdr10.c	DCSS hdr10
drivers/gpu/imx/dcss/dcss-wtsc1.c	DCSS wtsc1
drivers/gpu/imx/dcss/dcss-dtg.c	DCSS dtg
drivers/gpu/imx/dcss/dcss-common.c	DCSS common
drivers/gpu/imx/dcss/dcss-ctx1d.c	DCSS ctx1d
drivers/gpu/imx/dcss/dcss-dtrc.c	DCSS DTRC
drivers/gpu/imx/dcss/dcss-dpr.c	DCSS ctx1d
drivers/gpu/imx/dcss/dcss-blkctr.c	DCSS ctx1d

6.4   Display Interfaces

6.4.1   Parallel LCD Interface

6.4.1.1 Introduction

The Parallel interface supports display to LCDs. The Parallel Display interface is supported through the display controllers and implemented using the display framework which is fbdev framework on i.MX 6 and i.MX 7 and drm framework for i.MX8. .

The following controllers support the parallel interface

- IPU on i.MX with IPU
- DPU on all i.MX8
- ElCDIF on i.MX with PxP

The Parallel interface supports at least one port on i.MX SoC that enable the parallel interface and supports two ports for i.MX with IPU. The enabled SoC have varying bitrates from 18bit to 24 bits per port. On i.MX 6 with IPU the Parallel interface also supports a synchronous mode for display refresh and asynchronous mode to memory and is very flexible with a glue-less connection to RAM-less displays, display controllers and TV encoders.

6.4.2 MIPI DSI Interface

6.4.2.1 Software Operation

The MIPI DSI driver has two parts: MIPI DSI IP driver and MIPI DSI display panel driver.

The MIPI DSI IP driver has a private structure called mipi_dsi_info. The instance to which the MIPI DSI IP is attached is described in field int dev_id while the DI instance inside IPU is described in the field int disp_id.

During startup, the MIPI DSI IP driver is registered with the framebuffer driver through the field struct mxc_dispdrv_handle when the driver is loaded. It also registers a framebuffer event notifier with framebuffer core to perform the display panel blank/unblank operation. The field struct fb_videomode *mode and struct mipi_lcd_config *lcd_config are received from the display panel callback. The MIPI DSI IP needs this infomation to configure the MIPI DSI hardware registers.

After initializing the MIPI DSI IP controller and the display module, the MIPI DSI IP gets the pixel streams from IPU through DPI-2 interface and serializes pixel data and video event through high-speed data links for display. When there is an framebuffer blank/unblank event, the registered notifier will be called to enter/leave low power mode.

The MIPI DSI IP driver provides 3 APIs for MIPI DSI display panel driver to configure display module.

The drivers uses the APIs provided by the MIPI DSI IP driver to read/write the display module registers. Usually, there is a MIPI DSI slave controller integrated on the display panel. After power on reset, the MIPI DSI display panel needs to be configured through standard MIPI DCS command or MIPI DSI Generic command according to the manufacturer's specification.

6.4.2.2   Source Code Structure

Table below shows the MIPI DSI driver source files available in drivers/video/fbdev/mxc.

Table 6-12.   MIPI DSI Driver Files

File	Description
drveirs/video/fbdev/mxc/mipi_dsi.c	MIPI DSI IP Frame buffer driver source file
drivers/video/fbdev/mxc/mipi_dsi.h	MIPI DSI IP Frame bufferdriver header file
drivers/video/fbdev/mxc/mxcfb_hx8369_wvga.c	MIPI DSI Frame bufferDisplay Panel driver source file
drivers/video/fbdev/mxc/mipi_dsi_samsung.c	MIPI DSI Frame buffer Samsung source file
drivers/video/fbdev/mxc/mipi_dsi_northwest.c	MIPI DSI Frame buffer Northwest source file
drivers/video/fbdev/mxc/mxcfb_hx8363_wvga.c	i.MX 7 Frame buffer Truly WVGA Panel TFT3P5581E
drivers/video/fbdev/mxc/mxcfb_hx8369_wvga.c	i.MX 6 Frame buffer Truly WVGA sync panel
drivers/video/fbdev/mxc/mxcfb_otm808b_wvga.c	Truly Frame buffer WVGA Panel TFT3P5079E
drivers/gpu/drm/imx/sec_mipi_dsmi-imx.c	Samsung DRM driver
drivers/gpu/drm/imx/nwl_dsi-imx.c	Northwest DRM driver

6.4.2.3   Menu Configuration Options

In menu configuration enable the following module:

Device Drivers > Graphics support > MXC Framebuffer support > Synchronous Panel Framebuffer > MXC MIPI_DSI

Device Drivers > Graphics support > MXC Framebuffer support > Synchronous Panel Framebuffer > MXC MIPI_DSI_SAMSUNG

Device Drivers > Graphics support > DRM Support for Freescale i.MX > Support for Northwest Logic MIPI DSI displays

Device Drivers > Graphics support > DRM Support for Freescale i.MX > Support for Samsung MIPI DSIM displays

6.4.3 LVDS Interface

6.4.3.1 Introduction

Low Voltage Differential Signalling (LVDS) supports high bandwidth and high definitiion graphics and fast frame rate with lower power consumption. The implentation uses paris of wires where each wire in the pair carries inverse signal of the other. This creates less interference and noise. The LVDS interferace uses four, six or eight paris of wirse with additional ones carrying clock and ground wires.

The purpose of the LVDS interface is to support the flow of synchronous RGB data from the display controller to external display devices through the LVDS interface.

This support covers all aspects of these activities:

1. Connectivity to relevant devices - Displays with LVDS receivers.
2. Data arrangement required by the external display receiver and by LVDS display standards.
3. Synchronization and control capabilities.

The LVDS interface supports multiple controllers listed below.

- LDB - double on i.MX 6 with IPU
- Mixel on i.MX 8QuadMax
- Mixel Combo on i.MX 8QuadXPlus

The LVDS drivers works with the supported display framework which is framebuffer for i.MX 6 and i.MX 7 and drm for i.MX 8.

The LVDS interface has the following structure of support

- Channels - usually 2 channels
- Each channel supports a number of data pairs
- Data pixel rate which can vary on each data pair
- Control signals for HSYNC,VSYNC, DE

The LVDS interface supports the following displays.

- IT6263 LVDS to HDMI bridge - implemented with our LDB driver
- LVDS dual channel panel

The relevant standards for LVDS are the following.

- PHY Standard: ANSI EIA-644A
- Display Protocol Standards:
    - SPWG   Standard Panel Working Group Specification 3.8 (May 2007)
    - VESA PSWG – Panel Standardization Working Group – set of standards for panels using LVDS.
    - JEIDA/JEITA DISM Standard JEIDA-59-1999
    - OpenLDI (National) – Revision 0.95 13/May/1999. *Only* Unbalanced operating mode supported (aligned with vast majority of LCD vendors).

The LVDS interface is supported through the framebuffer framework on i.MX 6 and i.MX 7 and the drm framework on i.MX 8.

6.4.3.2   Software Operation

The LVDS driver is functional if the driver is built-in and the device tree status is set to "okay".

When the LVDS device driver is probed properly, the driver configures the clocks for the LVDS. The LVDS driver probe function sets the default mode to 1080p60. The LVDS channel mapping mode and bit mapping mode are set to use 30-bit JEIDA mode.

The driver takes the following steps to enable an LVDS channel:

1. Enable the power to the LVDS.
2. Set ldb_di_clk's parent clk and the parent clk's rate.
3. Set ldb_di_clk's rate.
4. Enable both ldb_di_clk and its parent clk.
5. Set the LVDS in a proper mode including display signals' polarities, channel mapping mode, and bit mapping mode.
6. Enable related i.MX LVDS channels.

6.4.3.3   Source Code Structure

Table 6-13.   LVDS Source

File	Description
drivers/video/fbdev/mxc/imx_lvds.c	LVDS driver
drivers/gpu/drm/imx-ldb.c	LDB driver with Mixel information

6.4.3.4 Menu Configuration Options

In menu configuration enable the following module:

Device Drivers > Graphics support > DRM Support for Freescale i.MX > Support for LVDS displays

6.4.4 LVDS Display Bridge (LDB)

6.4.4.1 Introduction

This section describes the LVDS Display Bridge (LDB) driver which controls the LDB module to connect with the external display devices with the LVDS interface. The purpose of the LDB is to support flow of synchronous RGB data from IPU or LCDIF to external display devices through LVDS interface.

This support covers the following:

- Connectivity to relevant devices - Displays with LVDS receivers.
- Arranging data as required by the external display receiver and by LVDS display standards.
- Synchronization and control capabilities.

6.4.4.2 Software Operation

The LDB driver is functional if the driver is built-in.

When the LDB device is probed properly, the driver configures the LDB reference resistor mode and the LDB regulator by using platform data information. The LDB driver probe function tries to match video modes for external display devices to LVDS interface. The display signal polarities control bits of the LDB are set according to the matched video modes. LVDS channel mapping mode and bit mapping mode of the LDB are set according to the LDB device tree node set by the user. The LDB is fully enabled in probe function if the driver identifies a display device with LVDS interface as the primary display device.

The steps the driver takes to enable an LVDS channel are:

1. Set ldb_di_clk's parent clk and the parent clk's rate.
2. Set ldb_di_clk's rate.
3. Enable both ldb_di_clk and its parent clk.

4. Set the LDB in a proper mode including display signals' polarities, LVDS channel mapping mode, bit mapping mode, and reference resistor mode.
5. Enable related LVDS channels.

6.4.4.3   Source Code Structure

Table 6-14.   LDB Source

File	Description
drivers/video/fbdev/mxc/ldb.c	LDB Framebuffer driver

6.4.4.4   Menu Configuration Options

The following Linux kernel configuration options are provided for this module.

In menu configuration enable the following module:

Device Drivers -> Graphics support -> MXC Framebufer support ->Synchronous Panel Framebuffer -> MXC LDB

6.4.5   Electrophoretic Display Controller (EPDC) Interface

6.4.5.1   Introduction

The Electrophoretic Display Controller (EPDC) is a direct-drive active matrix EPD controller designed to drive E Ink EPD panels supporting a wide variety of TFT backplanes. The EPDC framebuffer driver acts as a standard Linux frame buffer device. This driver supports a set of custom API extensions, accessible from user space (via IOCTL) or another kernel module (via direct function call) in order to provide the user with access to EPD-specific functionality. The EPDC driver is abstracted from any specific E Ink® panel type, providing flexibility to work with a range of E Ink panel types and specifications.

The EPDC driver supports the following features:

- EPDC driver as a loadable or built-in module.
- RGB565, RGB24, RGB32 and Y8 frame buffer formats.
- Full and partial EPD screen updates.
- Up to 256 panel-specific waveform modes.
- Automatic optimal waveform selection for a given update.

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

- Synchronization by waiting for a specific update request to complete.
- Screen updates from an alternate (overlay) buffer.
- Automated collision handling.
- 64 simultaneous update regions.
- Ppixel inversion in a Y8 frame buffer format.
- 90, 180, and 270 degree HW-accelerated frame buffer rotation.
- Panning (y-direction only).
- Automated full and partial screen updates through the Linux fb_deferred_io mechanism.
- Three EPDC driver display update schemes: Snapshot, Queue, and Queue and Merge.
- Setting the ambient temperature through either a one-time designated API call or on a per-update basis.
- User control of the delay between completing all updates and powering down the EPDC.

6.4.5.2 EPDC Frame Buffer Driver Overview

The frame buffer device provides an abstraction for the graphics hardware. It represents the frame buffer video hardware and allows application software to access the graphics hardware through a well-defined interface, abstracting from software how to manage the low-level hardware registers. The EPDC driver supports this model with one key caveat: the contents of the frame buffer are not automatically updated to the E Ink display. Instead, a custom API function call is required to trigger an update to the E Ink display. The details of this process are explained in the EPDC Frame Buffer Driver Extensions.

The frame buffer driver is enabled by selecting the frame buffer option under the graphics parameters in the kernel configuration. To supplement the frame buffer driver, the kernel builder may also includes support for fonts and a startup logo. The frame buffer device depends on the virtual terminal (VT) console to switch from serial to graphics mode. The device is accessed through special device nodes, located in the /dev directory, as /dev/fb*. fb0 is generally the primary frame buffer.

A frame buffer device is a memory device, such as /dev/mem, and has features similar to a memory device. Users can read it, write to it, seek to some location in it, and mmap() it (the main use). The difference is that the memory that appears in the special file is not the whole memory, but the frame buffer of some video hardware.

The EPDC frame buffer driver (drivers/video/fbdev/mxc/mxc_epdc_fb.c on i.MX 6SoloLite or i.MX 6DualLite or drivers/video/fbdev/mxc/mxc_epdc_v2_fb.c for generation-II EPDC on i.MX 7Dual) interacts closely with the generic Linux frame buffer driver (drivers/video/fbmem.c).

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

For additional details on the frame buffer device, see documentation in the Linux kernel found in Documentation/fb/framebuffer.txt.

6.4.5.3 EPDC Frame Buffer Driver Extensions

E Ink display technology, in conjunction with the EPDC, has several features that distinguish it from standard LCD-based frame buffer devices. These differences introduce the need for API extensions to the frame buffer interface. The EPDC refreshes the E Ink display asynchronously and supports partial screen updates. Therefore, the EPDC requires notification from the user when the frame buffer contents have been modified and which region needs updating. Another unique characteristic of EPDC updates to the E Ink display is the long screen update latencies (between 300-980 ms), which introduces the need for a mechanism to allow the user to wait for a given screen update to complete.

The custom API extensions to the frame buffer device are accessible both from user space applications and from within kernel space. The standard device IOCTL interface provides access to the custom API for user space applications. The IOCTL extensions, along with relevant data structures and definitions, can be found in include/linux/mxcfb_epdc.h. A full description of these IOCTLs can be found in the Programming Interface section Software Operation.

For kernel mode access to the custom API extensions, the IOCTL interface should be bypassed in favor of direct access to the underlying functions.

6.4.5.4 EPDC Panel Configuration

The EPDC driver is designed to flexibly support E Ink panels with a variety of panel resolutions, timing parameters, and waveform modes. The EPDC driver is kept panel-agnostic through the use of an EPDC panel mode structure, imx_epdc_fb_mode, which can be found in include/linux/mxcfb_epdc.h.

```
struct imx_epdc_fb_mode {
                struct fb_videomode *vmode;
                int vscan_holdoff;
                int sdoed_width;
                int sdoed_delay;
                int sdoez_width;
                int sdoez_delay;
                int gdclk_hp_offs;
                int gdsp_offs;
                int gdoe_offs;
                int gdclk_offs;
                int num_ce;
};
```

The imx_epdc_fb_mode structure consists of an fb_videomode structure reference and a set of EPD timing parameters. The fb_videomode structure defines the panel resolution and the basic timing parameters (pixel clock frequency, hsync and vsync margins) and the additional timing parameters in imx_epdc_fb_mode define EPD-specific timing parameters, such as the source and gate driver timings. For details on how to configure E Ink panel timing parameters, see the EPDC programming model section in the *i.MX 6SoloLite Applications Processor Reference Manual* (IMX6SLRM), *i.MX 6DualLite Applications Processor Reference Manual* (IMX6DLRM), or *i.MX 7Dual Applications Processor Reference Manual* (IMX7DRM).

In addition to the EPDC panel mode data, functions may be passed to the EPDC driver to define how to handle the EPDC pins when the EPDC driver is enabled or disabled. These functions should disable the EPDC pins for purposes of power savings.

6.4.5.5 Boot Command Line Parameters

Additional configuration for the EPDC driver is provided through boot command line parameters. The format of the command line option is

```
epdc video=mxcepdcfb:[panel_name],bpp=16
```

.

The EPDC driver parses these options and tries to match panel_name to the name of video mode specified in the imx_epdc_fb_mode panel mode structure. If no match is found, then the first panel mode provided in the platform data is used by the EPDC driver. The bpp setting from this command line sets the initial bits per pixel setting for the frame buffer. A setting of 32 or 24 selects the RGB888 pixel format, one of 16 selects RGB565 pixel format, while a setting of 8 selects 8-bit grayscale (Y8) format.

6.4.5.6 EPDC Waveform Loading

The EPDC driver requires a waveform file for proper operation. This waveform file contains the waveform information needed to generate the waveforms that drive updates to the E Ink panel. A pointer to the waveform file data is programmed into the EPDC before the first update is performed.

There are two options for selecting a waveform file:

1. Select one of the default waveform files included in this BSP release.
2. Use a new waveform file that is specific to the E Ink panel being used.

The waveform file is loaded by the EPDC driver using the Linux firmware APIs.

6.4.5.7 Using a Default Waveform File

The quickest and easiest way to get started using an E Ink panel and the EPDC driver is to use one of the default waveform files provided in the Linux BSP. This should enable updates to several different types of E Ink panel without a panel-specific waveform file. The drawback is that optimal quality should not be expected. Typically, using a non-panel-specific waveform file for an E Ink panel results in more ghosting artifacts and overall poorer color quality.

The following default waveform files included in the BSP reside in /lib/firmware/imx/epdc:

- epdc_E60_V110.fw - Default waveform for the 6.0 inch V110 E Ink panel.
- epdc_E60_V220.fw - Default waveform for the 6.0 inch V220 E Ink panel (supports animation mode updates).
- epdc_E97_V110.fw - Default waveform for the 9.7 inch V110 E Ink panel.
- epdc_E060SCM.fw - Default waveform for the 6.0 inch Pearl E Ink panel (supports animation mode updates).
- epdc_ED060XH2C1.fw - Default waveform for the 6.0 inch E Ink panel (No Reagl/-D Support by default. For Reagl/-D support, contact NXP support.)

The EPDC driver attempts to load a waveform file with the name "epdc_[panel_name].fw" under the directory /lib/firmware/imx/epdc in rootfs, where panel_name refers to the string specified in the fb_videomode `name` field. This panel_name information should be provided to the EPDC driver through the kernel command line parameters described in the preceding chapter. For example, to load the epdc_E060SCM.fw default firmware file for a Pearl panel, set the EPDC kernel command line paratmeter to the following:

```
video=mxcepdcfb:E060SCM,bpp=16
```

6.4.5.8 Using a Custom Waveform File

To ensure the optimal E Ink display quality, use a waveform file specific to E Ink panel being used. The raw waveform file type (.wbf) requires conversion to a format that can be understood and read by the EPDC. This conversion script is not included as part of the BSP. Therefore, contact NXP to acquire this conversion script.

Once the waveform conversion script has been run on the raw waveform file, the converted waveform file should be renamed so that the EPDC driver can find it and load it. The driver is going to search for a waveform file with the name "epdc_[panel_name].fw" under the directory /lib/firmware/imx/epdc in rootfs, where

panel_name refers to the string specified in the fb_videomode `name` field. For example, if the panel is named "E60_ABCD", then the converted waveform file should be named epdc_E60_ABCD.fw.

NOTE

If the EPDC driver searches for a firmware waveform file that matches the names of one of the default waveform files (see preceding chapter), it will choose the default firmware files that are built into the BSP over any firmware file that has been added in the firmware search path. Therefore, if you leave the BSP so that it uses the default firmware files, make sure to use a panel name other than those associated with the default firmware files, as those default waveform files will be preferred and selected over a new waveform file placed in the firmware search path.

6.4.5.9   EPDC Panel Initialization

The framebuffer driver will not typically (see note below for exceptions) go through any hardware initialization steps when the framebuffer driver module is loaded. Instead, a subsequent user mode call must be made to request that the driver initialize itself for a specific EPD panel. To initialize the EPDC hardware and E Ink panel, an FBIOPUT_VSCREENINFO ioctl call must be made, with the xres and yres fields of the fb_var_screeninfo parameter set to match the X and Y resolution of a supported E Ink panel type. To ensure that the EPDC driver receives the initialization request, the activate field of the fb_var_screeninfo parameter should be set to FB_ACTIVATE_FORCE.

NOTE

The exception is when the FB Console driver is included in the kernel. When the EPDC driver registers the framebuffer device, the FB Console driver will subsequently make an FBIOPUT_VSCREENINFO ioctl call. This will in turn initialize the EPDC panel.

6.4.5.10   Grayscale Framebuffer Selection

The EPDC framebuffer driver supports the use of 8-bit grayscale (Y8) and 8-bit inverted grayscale (Y8 inverted) pixel formats for the framebuffer (in addition to the more common RGB565 pixel format). In order to configure the framebuffer format as 8-bit grayscale, the application would call the FBIOPUT_VSCREENINFO framebuffer ioctl. This ioctl takes an fb_var_screeninfo pointer as a parameter. This parameter specifies the

attributes of the framebuffer and allows the application to request changes to the framebuffer format. There are two key members of the fb_var_screeninfo parameter that must be set in order to request a change to 8-bit grayscale format: bits_per_pixel and grayscale. bits_per_pixel must be set to 8 and grayscale must be set to one of the 2 valid grayscale format values: GRAYSCALE_8BIT or GRAYSCALE_8BIT_INVERTED.

The following code snippet demonstrates a request to change the framebuffer to use the Y8 pixel format:

```
fb_screen_info screen_info;
screen_info.bits_per_pixel = 8;
screen_info.grayscale = GRAYSCALE_8BIT;
retval = ioctl(fd_fb0, FBIOPUT_VSCREENINFO, &screen_info);
```

6.4.5.11 Software Operation

The EPDC Frame Buffer is accessible from user space and from kernel space. A single set of functions describes the EPDC Frame Buffer driver extension. There are two modes for accessing these functions with user space using the IOCTL interface and kernel space using funcions directly. Each IOCTL and function combination is described next.

MXCFB_SET_WAVEFORM_MODES / mxc_epdc_fb_set_waveform_modes()

Description:

Defines a mapping for common waveform modes.

Parameters:

mxcfb_waveform_modes *modes*

Pointer to a structure containing the waveform mode values for common waveform modes. These values must be configured in order for automatic waveform mode selection to function properly.

MXCFB_SET_TEMPERATURE / mxc_epdc_fb_set_temperature

Description:

Set the temperature to be used by the EPDC driver in subsequent panel updates.

Parameters:

int*32_t temperature*

Temperature value, in degrees Celsius. Note that this temperature setting may be overridden by setting the temperature value parameter to anything other than TEMP_USE_AMBIENT when using the MXCFB_SEND_UPDATE ioctl.

MXCFB_SET_AUTO_UPDATE_MODE / mxc_epdc_fb_set_auto_update

Description:

Select between automatic and region update mode.

Parameters:

__u32 mode

In region update mode, updates must be submitted via the MXCFB_SEND_UPDATE IOCTL.

In automatic mode, updates are generated automatically by the driver by detecting pages in frame buffer memory region that have been modified.

MXCFB_SET_UPDATE_SCHEME / mxc_epdc_fb_set_upd_scheme

Description:

Select a scheme that dictates how the flow of updates within the driver.

Parameters:

__u32 scheme

Select of the following updates schemes:

UPDATE_SCHEME_SNAPSHOT - In the Snapshot update scheme, the contents of the framebuffer are immediately processed and stored in a driver-internal memory buffer. By the time the call to MXCFB_SEND_UPDATE has completed, the framebuffer region is free and can be modified without affecting the integrity of the last update. If the update frame submission is delayed due to other pending updates, the original buffer contents will be displayed when the update is finally submitted to the EPDC hardware. If the update results in a collision, the original update contents will be resubmitted when the collision has cleared.

UPDATE_SCHEME_QUEUE - The Queue update scheme uses a work queue to asynchronously handle the processing and submission of all updates. When an update is submitted via MXCFB_SEND_UPDATE, the update is added to the queue and then processed in order as EPDC hardware resources become available. As a result, the framebuffer contents processed and updated are not guaranteed to reflect what was present in the framebuffer when the update was sent to the driver.

UPDATE_SCHEME_QUEUE_AND_MERGE - The Queue and Merge scheme uses the queueing concept from the Queue scheme, but adds a merging step. This means that, before an update is processed in the work queue, it is first compared with other pending updates. If any update matches the mode and flags of the current update and also overlaps

the update region of the current update, then that update will be merged with the current update. After attempting to merge all pending updates, the final merged update will be processed and submitted.

MXCFB_SEND_UPDATE / mxc_epdc_fb_send_update

Description:

Request a region of the frame buffer be updated to the display.

Parameters:

mxcfb_update_data *upd_data*

Pointer to a structure defining the region of the frame buffer, waveform mode, and collision mode for the current update. This structure also includes a flags field to select from one of the following update options:

EPDC_FLAG_ENABLE_INVERSION - Enables inversion of all pixels in the update region.

EPDC_FLAG_FORCE_MONOCHROME - Enables full black/white posterization of all pixels in the update region.

EPDC_FLAG_USE_ALT_BUFFER - Enables updating from an alternate (non-framebuffer) memory buffer.

If enabled, the final *upd_data* parameter includes detailed configuration information for the alternate memory buffer.

MXCFB_WAIT_FOR_UPDATE_COMPLETE / mxc_epdc_fb_wait_update_complete

Description:

Block and wait for a previous update request to complete.

Parameters:

mxfb_update_marker_data *marker_data*

The update_marker value used to identify a particular update (passed as a parameter in MXCFB_SEND_UPDATE IOCTL call) should be re-used here to wait for the update to complete. If the update was a collision test update, the collision_test variable will return the result indicating whether a collision occurred.

MXCFB_SET_PWRDOWN_DELAY / mxc_epdc_fb_set_pwrdown_delay

Description:

Set the delay between the completion of all updates in the driver and when the driver should power down the EPDC and the E Ink display power supplies.

Parameters:

int32_t *delay*

Input delay value in milliseconds. To disable EPDC power down altogether, use FB_POWERDOWN_DISABLE (defined below).

MXCFB_GET_PWRDOWN_DELAY / mxc_epdc_fb_get_pwrdown_delay

Description:

Retrieve the driver's current power down delay value.

Parameters:

int32_t *delay*

Output delay value in milliseconds.

6.4.5.12   Structures and Defines

```
#define GRAYSCALE_8BIT                                              0x1
#define GRAYSCALE_8BIT_INVERTED                                     0x2


#define AUTO_UPDATE_MODE_REGION_MODE                                0

#define AUTO_UPDATE_MODE_AUTOMATIC_MODE                             1

#define UPDATE_SCHEME_SNAPSHOT                                      0
#define UPDATE_SCHEME_QUEUE                                         1
#define UPDATE_SCHEME_QUEUE_AND_MERGE                               2

#define UPDATE_MODE_PARTIAL                                         0x0
#define UPDATE_MODE_FULL                                            0x1

#define WAVEFORM_MODE_AUTO                                          257

#define TEMP_USE_AMBIENT                                            0x1000

#define EPDC_FLAG_ENABLE_INVERSION                                  0x01
#define EPDC_FLAG_FORCE_MONOCHROME                                  0x02
#define EPDC_FLAG_USE_ALT_BUFFER                                    0x100
#define EPDC_FLAG_TEST_COLLISION                                    0x200

#define FB_POWERDOWN_DISABLE                                        -1

struct mxcfb_rect {
    __u32 left; /* Starting X coordinate for update region */
    __u32 top; /* Starting Y coordinate for update region */
    __u32 width; /* Width of update region */
    __u32 height; /* Height of update region */
};

struct mxcfb_waveform_modes {
    int mode_init; /* INIT waveform mode */
```

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

```
    int mode_du; /* DU waveform mode */
    int mode_gc4; /* GC4 waveform mode */
    int mode_gc8; /* GC8 waveform mode */
    int mode_gc16; /* GC16 waveform mode */
    int mode_gc32; /* GC32 waveform mode */
};

struct mxcfb_alt_buffer_data {
    __u32 phys_addr; /* physical address of alternate image buffer */
    __u32 width; /* width of entire buffer */
    __u32 height; /* height of entire buffer */
    struct mxcfb_rect alt_update_region; /* region within buffer to update */
};

struct mxcfb_update_data {
    struct mxcfb_rect update_region; /* Rectangular update region bounds */
    __u32 waveform_mode; /* Waveform mode for update */
    __u32 update_mode; /* Update mode selection (partial/full) */
    __u32 update_marker; /* Marker used when waiting for completion */
    int temp; /* Temperature in Celsius */
    uint flags; /* Select options for the current update */
    struct mxcfb_alt_buffer_data alt_buffer_data; /* Alternate buffer data */
};
struct mxcfb_update_marker_data { __u32 update_marker; __u32 collision_test; };
```

6.4.5.13 Source Code Structure

Table below lists the source files associated with the EPDC driver and headers for programming access.

Table 6-15. EPDC Source

File	Description
drivers/video/fbev/mxc/mxc_epdc_v2_fb.c	EPDC Generation-II V2 frame buffer driver for i.MX 7Dual
drivers/video/fbdev/mxc/epdc_v2_regs.h	EPDC Generation-II Register definition
drivers/video/fbdev/mxc/mxc_epdc_fb.c	Generation-I EPDC frame buffer driver for i.MX 6Sololite, 6SLL, and 6DualLite
drivers/video/fbdev/mxc/epdc_regs.h	EPDC Generation-IRegister definitions
drivers/video/fbdev/mxc/epdc_v2_regs.h	Generation-II EPDC v2 register definitions
include/linux/uapi/mxcfb.h	Header file for the EPDC IOCTLs and frame buffer driver
include/linux/mxcfb_epdc.h	Header file for direct kernel access to the EPDC API extension

6.4.5.14 Menu Configuration Options

The following Linux kernel configuration options are provided for the EPDC module:

- CONFIG_FB_MXC_EINK_PANEL - support for the Electrophoretic Display Controller. In menuconfig, select Device Drivers > Graphics Support > E-Ink Panel Framebuffer

- CONFIG_FB_MXC_EINK_V2_PANEL - support for v2 Electrophoretic Display Controller. In menuconfig, this option is available with Device Drivers > Graphics support > E-Ink Panel Framebuffer based on EPDC V2
- CONFIG_FB - includes frame buffer support and is enabled by default. In menuconfig select Device Drivers > Graphics support > Support for frame buffer devices
- CONFIG_MXC_PXP_V2 - support for the PxP and required by the EPDC driver for processing (color space conversion, rotation, auto-waveform selection) framebuffer update regions. In menuconfig select Device Drivers > DMA Engine support > MXC PxP support
- CONFIG_MXC_PXP_V3 - support for next level PxP and required by Generation-II EPDC driver for processing framebuffer update regions. In menuconfig select Device Drivers > DMA Engine support > MXC PxP V3 support

6.4.6 High-Definition Multimedia Interface (HDMI) and Display Port (DP) Overview

6.4.6.1 Introduction

High-Definition Multimedia Interface (HDMI) and Display Port (DP) present high defintion video. The HDMI module is supported on some i.MX chips either with on chip solutions or external solutions. The Display Port DP provides an embedded Display Port (eDP) Transmitter including HDMI Tranmit (TX) Controller and PHY.

The following are compliance versions.

- HDMI 1.4 and 2.0
- DVI 1.0
- DP 1.3
- eDP 1.4
- HDCP 1.4/2.2

Each SoC HDMI solution is presented in separate chapters. Display Port on i.MX uses the same IP block but has a different specification. The following table lists which SOC support HDMI and Display Port and its supported version.

Table 6-16. HDMI Support

SoC	Features
i.MX 6QuadPlus/Quad/Dual	HDMI 1.4 on chip
i.MX 6SoloLite	HDMI 1.4 extenral chip

Table continues on the next page...

Table 6-16. HDMI Support (continued)

SoC	Features
i.MX 7ULP	HDMI 1.4 external chip
i.MX 8M Quad	HDMI 2.0/Display Port 1.3 on chip
i.MX 8QuadMax	HDMI 2.0/Display Port 1.3 on chip

HDMI Audio data source comes from S/PDIF TX.

6.4.6.2 Software Operation

The HDMI driver is divided into sub-components based on its two primary purposes: providing video and audio to an HDMI sink device.

The video display driver component and audio driver component require an additional core driver component to manage common HDMI resources, including the HDMI registers, clocks, and IRQ.

6.4.6.3 Core

The onchip HDMI i.MX solutions support a core driver that manages resources that must be shared between the HDMI audio and video drivers. The HDMI audio and video drivers depend on the HDMI core driver, and the HDMI core driver should always be loaded and initialized before audio and video. The core driver serves the following functions:

- Map the HDMI register region and provide APIs for reading and writing to HDMI registers.
- Perform one-time initialization of key HDMI registers.
- Initialize the HDMI IRQ and provide shared APIs for enabling and disabling the IRQ.
- Provide a means for sharing information between the audio and video drivers (e.g., the HDMI pixel clock).
- Provide a means for synchronization between HDMI video and HDMI audio while blank/unblank, plug in/plug out events happen. HDMI audio cannot start work while HDMI cable is in the state of plug out or HDMI is in state of blank. Every time HDMI audio starts a playback, HDMI audio driver should register its PCM into core driver and unregister PCM when the playback is finished. Once HDMI video blank or cable plug out event happens, core driver would pause HDMI audio DMA controller if its PCM is registered. When HDMI is unblanked or cable plug in event happens, core driver would firstly check if the cable is in the state of plug in, the

video state is unblank and the PCM is registered. If items listed above are all yes, core driver would restart HDMI audio DMA.

6.4.6.4 Display Device Registration and Initialization

The following sequence of software activities occurs in the OS boot flow to connect the HDMI display device to the i.MX Frame Buffer driver through the MXC Display Driver system:

1. During the HDMI video driver initialization, `mxc_dispdrv_register()` is called to register the HDMI module as a display device and to set the `mxc_hdmi_disp_init()` function as the display device init callback.
2. When the i.MX Frame Buffer driver is initialized, `mxc_dispdrv_init()` is called. This results in an init call to all registered display devices.
3. The `mxc_hdmi_disp_init()` callback is executed. The HDMI driver receives a structure from the i.MX Frame Buffer driver containing frame buffer information (fbi). The HDMI driver registers itself to receive notifications for FB driver events. Finally, the HDMI driver completes initialization by configuring the HDMI to receive a hotplug interrupt.

NOTE
All display device drivers must be initialized before the i.MX Frame Buffer driver in order for all display devices to be registered as MXC Display Driver devices.

6.4.6.5 Hotplug Handling and Video Mode Changes

Once the connection between the i.MX frame buffer driver and the HDMI has been established through the MXC Display Driver interface, the HDMI video driver waits for a hotplug interrupt indicating that a valid HDMI sink device is connected and ready to receive HDMI video data. Subsequent communications between the HDMI and i.MX Frame Buffer Driver are conducted through the Linux Frame Buffer APIs. The following list demonstrates the software flow to recognize a HDMI sink device and configure the ELCDIF FB driver to drive video output:

1. The HDMI video driver receives a hotplug interrupt and reads the EDID from the HDMI sink device constructing a list of video modes from the retrieved EDID information. Using either the video mode string from the Linux kernel command line (for the initial connection) or the most recent video mode (for a later HDMI cable connection), the HDMI driver selects a video mode from the mode list that is the closest match.

2. The HDMI video driver calls `fb_set_var()` to change the video mode in the i.MX Frame Buffer driver. The i.MX Frame Buffer driver completes its reconfiguration for the new mode.
3. As a result of calling fb_set_var(), a Frame Buffer notification is sent back to the HDMI driver indicating that an FB_EVENT_MODE_CHANGE has occurred. The HDMI driver configures the HDMI hardware for the new video mode.
4. Finally, the HDMI module is enabled to generate output to the HDMI sink device.

The i.MX Frame Buffer Driver will align to the display interface specific to each SoC as noted for each SoC HDMI chapter.

6.4.6.6 Audio

Since the HDMI Tx audio driver uses the ALSA SoC framework, it is broken into several files as listed in the source code structure sections of each hdmi chapter. Most of the code is in the platform DMA driver (sound/soc/imx/imx-hdmi-dma.c) and the CODEC driver (sound/soc/codecs/mxc_hdmi.c). The machine driver (sound/soc/imx/imx-hdmi.c) allocates the SoC audio device and links all the SoC components together. The DAI driver (sound/soc/imx/imx-hdmi-dai.c) is a SoC requirements. It is primarily used to get the platform data.

The HDMI CODEC driver does most of the initialization of the HDMI audio sampler. Note that the HDMI Tx block only implements the AHB DMA audio and not the other audio interfaces (SSI, S/PDIF, etc). The other main function of the HDMI CODEC driver is to set up a struct of the IEC header information which needs to go into the audio stream. Since the struct is hooked into the ALSA layer, IEC settings will be accessible in userspace using the 'iecset' utility.

The platform DMA driver handles the HDMI Tx block DMA engine. Note that HDMI audio uses the HDMI block DMA as well as SDMA. SDMA is used to implement the multi-buffer mechanism. Since the HDMI Tx block does not automatically merge the IEC audio header information into the audio stream, the platform DMA driver does the merging by using hdmi_dma_copy() (for no memory map use) or hdmi_dma_mmap_copy() (for memory map mode use) function before sending the buffers out. Note that, due to IEC audio header adding operation, it is possible that the user space application may not be able to get enough CPU periods to feed the data into HDMI audio driver in time, especially when system loading is high. In this case, some spark noise will be heard. In a different audio framework (ALSA LIB, or PULSE AUDIO), a different log about this noise may be printed. For example, in ALSA LIB, logs like "underrung!!! at least * ms is lost" are printed.

HDMI audio playback depends on HDMI pixel clock. Therefore, while in the state of HDMI blank and cable plug out, HDMI audio is either stopped or can't be played. See detailed information in software_operation_core.

Note that, because HDMI audio driver needs to add the IEC header, the driver needs to know the amount of data already written into the HDMI audio driver. If application is not able to decipher the amount of data written, for example DMIX plugin in ALSA LIB, the HDMI audio driver is not able to work properly. There will be no sound heard.

The HDMI audio supports the features below:

- Playback sample rate
    - 32k, 44.1k, 48k, 88.2k, 96k, 176.4k, 192k
    - HDMI sink capability
- Playback Channels:
    - 2, 4, 6, 8
    - HDMI sink capability
- Playback audio formats:
    - SNDRV_PCM_FMTBIT_S16_LE

6.4.6.7   i.MX 8 Display Port

6.4.6.7.1   Software Operation

The HDMI driver is divided into sub-components based on its two primary purposes: providing HDP DRM driver and Core API driver.

The HDP DRM driver require a Core API driver component to the configured HDMI FW.

6.4.6.7.2   Source Code Structure

The HDMI driver has three software components: HDP core API driver, HDP display driver, and HDMI audio driver.

The Core API source code is available in the drivers/mxc/hdp directory.

Table 6-17.   HDP Core API Driver File List

Files	Description
• drivers/mxc/hdp/API_HDCP.c • drivers/mxc/hdp/API_HDMITX.c • drivers/mxc/hdp/API_HDMIRX.c • drivers/mxc/hdp/API_Infoframe.c • drivers/mxc/hdp/API_AVI.c	HDMI Display Port Driver

Table continues on the next page...

Table 6-17.   HDP Core API Driver File List (continued)

Files	Description
• drivers/mxc/hdp/API_General.c • drivers/mxc/hdp/API_DPTX.c	
• drivers/mxc/hdp/API_HDMI_Audio.c • drivers/mxc/hdp/API_HDMI_RX_Audio.c	HDMI Display Port Audio Driver
• drivers/gpu/drm/imx/hdp/API_AFE_t28hpc_hdmitx.c • drivers/gpu/drm/imx/hdp/ API_AFE_ss28fdsoi_kiran_hdmitx.c • drivers/gpu/drm/imx/hdp/imx-hdp-audio.c • drivers/gpu/drm/imx/hdp/imx-hdp.c • drivers/gpu/drm/imx/hdp/imx-hdmi.c • drivers/gpu/drm/imx/hdp/t28hpc_hdmitx_table.c • drivers/gpu/drm/imx/hdp/ss28fdsoi_hdmitx_table.c • drivers/gpu/drm/imx/hdp/imx-arc.c	HDMI DRM driver
• drivers/mxc/hdp-cec/imx-hdp-cec.c	HDMI Display port CEC authentiation
• sound/soc/fsl/fsl_hdmi.c • sound/soc/fsl/imx-hdmi.c • sound/soc/fsl/hdmi_pcm.S • sound/soc/fsl/imx-hdmi-dma.c • sound/soc/fsl/imx-cdnhdmi.c	HDMI Sound Driver

6.4.6.7.3   Menu Configuration Options

There are three main Linux kernel configuration options used to select and include HDMI driver functionality in the Linux OS image.

There are four main Linux kernel configuration options used to select and include HDMI driver functionality in the Linux OS image.

The CONFIG_MX8_HDP option provides support for the HDP Core API driver, and can be selected in menuconfig at the following menu location:

Device Drivers > MXC support drivers > IMX8 HDP API

The CONFIG_DRM_IMX_HDP option provides support for the HDP DRM video driver, and can be selected in menuconfig at the following menu location:

Device Drivers > Graphics support > IMX8 HD Display Controller

The CONFIG_IMX_HDP_CEC option provides support for the HDMI CEC driver, and can be selected in menuconfig at the following menu location:

Device Drivers > Graphics support > IMX8 HD Display Controller > Enable IMX HDP CEC support

The CONFIG_SND_SOC_IMX_CDNHDMI option provides support for HDMI audio through the ALSA/SoC subsystem, and can be found in menuconfig at the following location:

Device Drivers > Sound card support > Advanced Linux Sound Architecture > ALSA for SoC audio support > SoC Audio support for CDN - HDMI

6.4.6.8   i.MX 6 On Chip High-Definition Multimedia Interface (HDMI)

6.4.6.8.1   Introduction

The High-Definition Multimedia Interface (HDMI) driver supports the on-chip DesignWare HDMI hardware module on the i.MX 6QuadPlus, 6Quad and 6Dual SoC, This driver provides the capability to transfer uncompressed video, audio, and data using a single cable.

The HDMI driver is divided into four sub-components:

- Video display device driver that integrates with the Linux Frame Buffer API
- Audio driver that integrates with the ALSA/SoC sub-system
- CEC driver
- Multifunction device (MFD) driver which manages the shared software and hardware resources of the HDMI driver.

The HDMI driver supports the following features:

- Integration with the MXC Display Device framework (for managing display device connections with the IPU(s))
- HDMI video output up to 1080p60 resolution
- Support for reading EDID information from an HDMI sink device
- Hotplug detection
- Support CEC
- Automated clock management to minimize power consumption
- Support for system suspend/resume
- HDMI audio playback (2, 4, 6, or 8 channels, 16-bit, for sample rates 32-KHz to 192-KHz)
- IEC audio header information exposed through ALSA using 'iecset' utility

The HDMI module receives video data from the Image Processing Unit (IPU), audio data from the external memory interface, and control data from the CPU, as shown in the figure below. Output data is transmitted via three Transition-Minimized Differential Signaling (TMDS) channels to an HDMI sink device external to the SoC. The HDMI also carries a VESA Data Display Channel (DDC). The DDC is an I2C interface which allows

the HDMI source to query the HDMI sink for Extended Displa-y Identification Data (EDID). A CEC channel provides optional high-level control functions between the source and sink device.
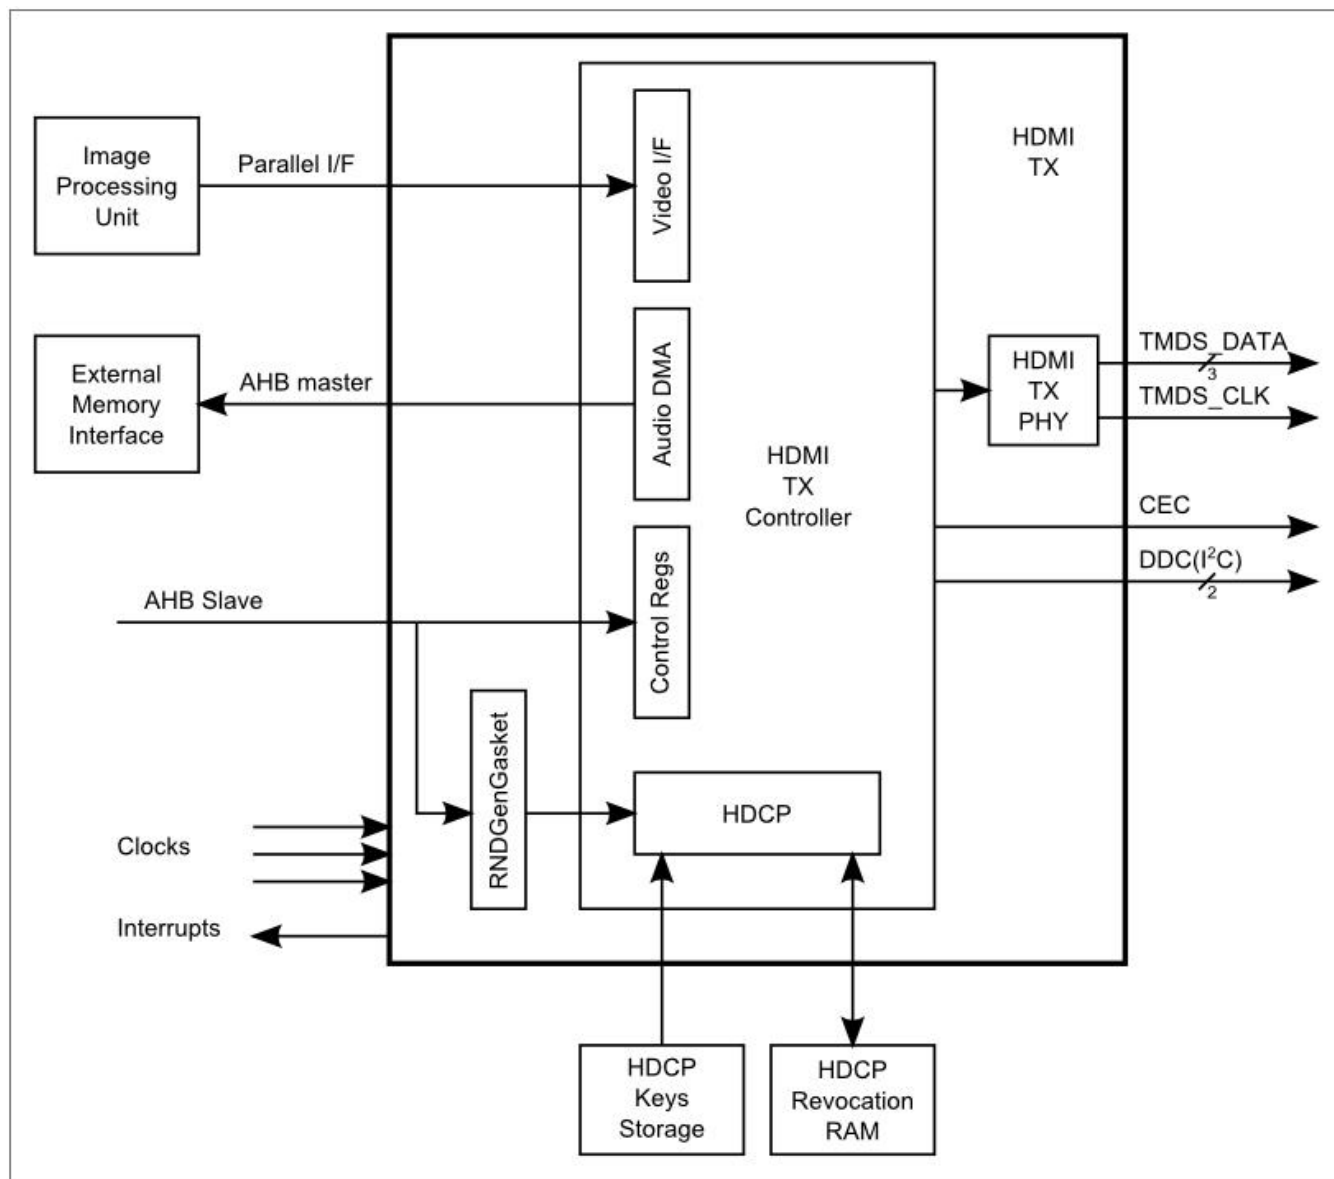
Figure 6-3. HDMI HW Integration

The video input to the HDMI is configurable and may come from either of the two IPU modules in the i.MX 6 serials and from either of the two Display Interface (DI) ports of the IPU, DI0 or DI1. This configuration is controlled through the IOMUX module using the HDMI_MUX_CTRL register field. See the figure below for an illustration of this interconnection.
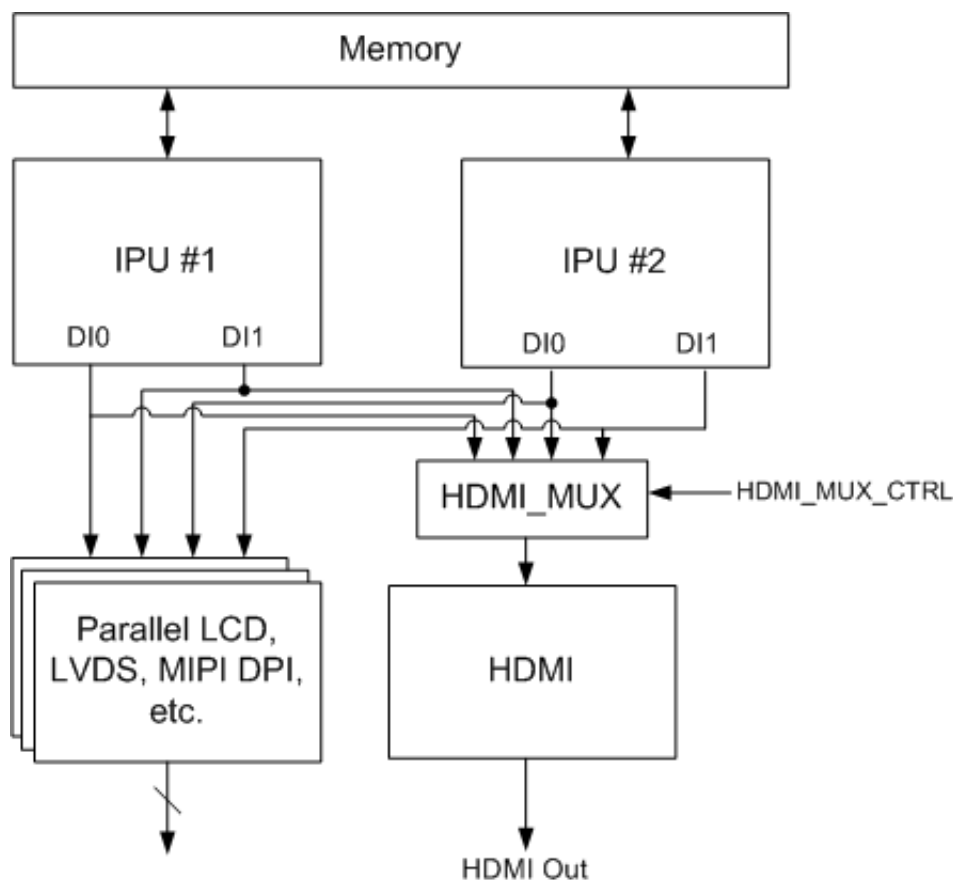
Figure 6-4. IPU-HDMI Hardware Interconnection

6.4.6.8.2 Software Operation

The HDMI driver is divided into sub-components based on its two primary purposes: providing video and audio to an HDMI sink device.

The video display driver component and audio driver component require an additional core driver component to manage common HDMI resources, including the HDMI registers, clocks, and IRQ. The following diagram illustrates both the interconnection between the various HDMI sub-drivers and the interconnection between the HDMI video driver and the Linux Frame Buffer subsystem.
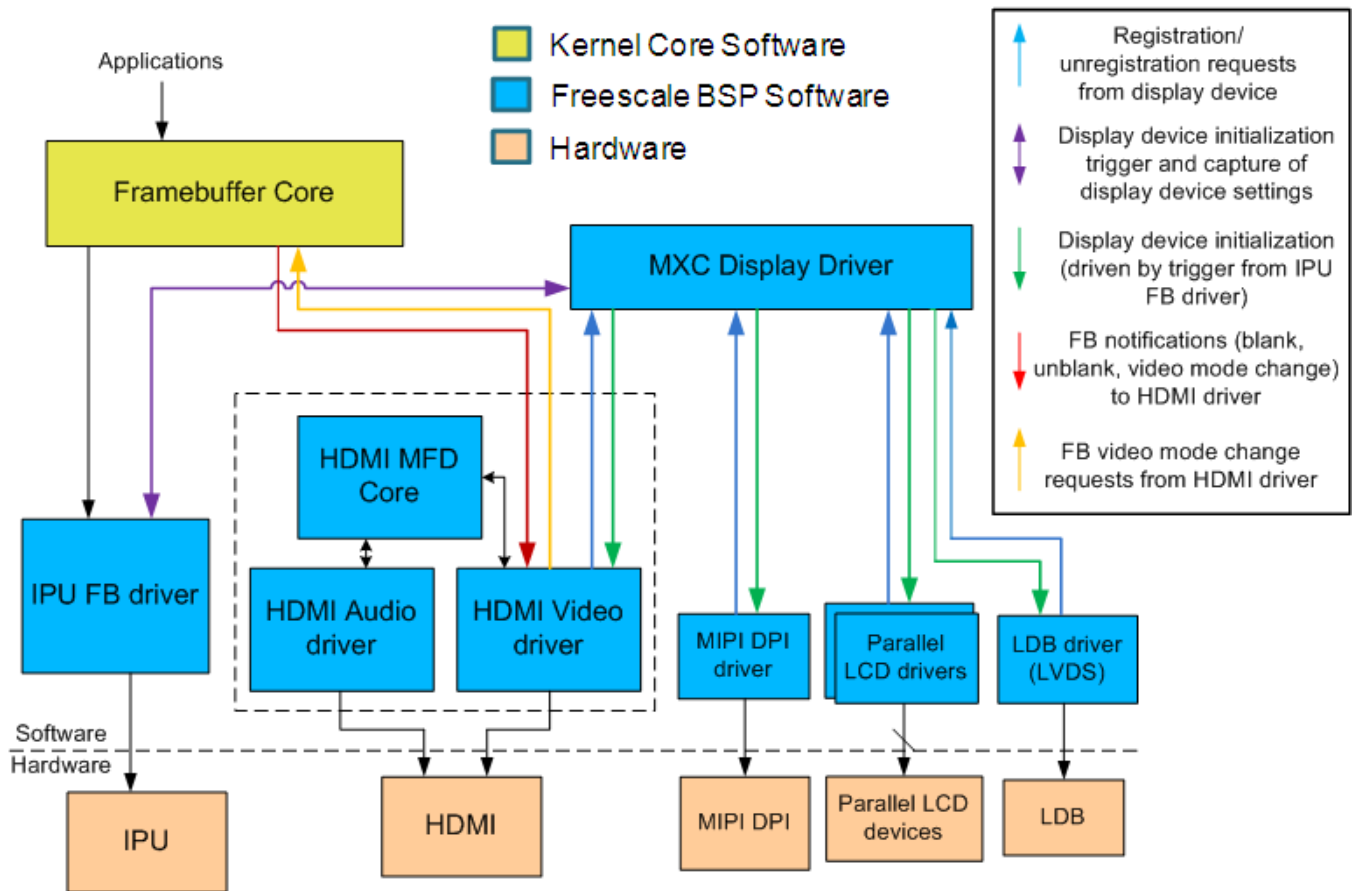
Figure 6-5. HDMI Video SW Architecture

The i.MX 6Dual/6Quad/6QuadPlus/6Solo/6DualLite supports many different types of display output devices (e.g., LVDS, LCD, HDMI and MIPI displays) connected to and driven by the IPU modules. The MXC Display Driver API provides a system for registering display devices and configuring how they should be connected to each of the IPU DIs. The HDMI driver registers itself as a display device using this API in order to receive the correct video input from the IPU.

6.4.6.8.3   CEC

HDMI CEC is a protocol that provides high-level control functions between all of the various audiovisual products is a user's environment. The HDMI CEC driver implements software part of HDMI CEC low Level protocol. It includes getting Logical address, CEC message sending and receiving, error handle, message re-transmitting, etc.
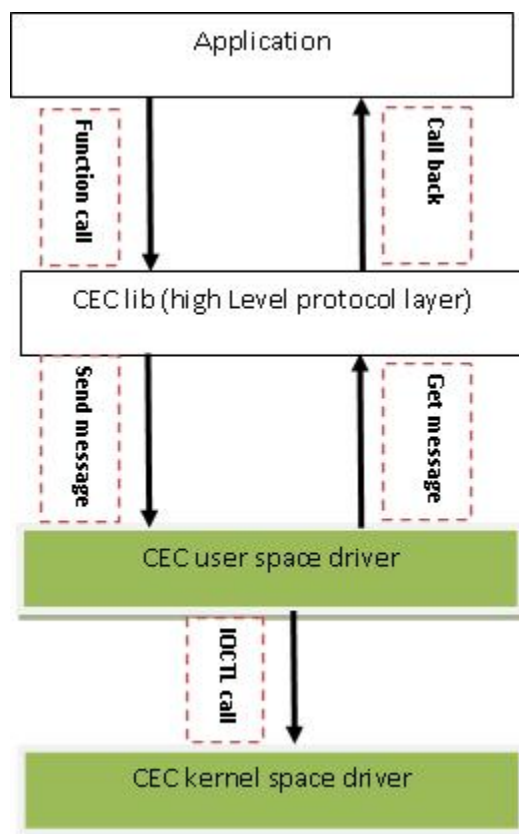
Figure 6-6. HDMI CEC SW Architecture

6.4.6.8.4 Source Code Structure

The HDMI source code is provided in the HDMI core driver, the HDMI display driver, and the HDMI audio driver.

Table 6-18. HDMI Source

File	Description
drivers/mfd/mxc-hdmi-core.c	HDMI core driver implemention
include/linux/mfd/mxc-hdmi-core.h	HDMI core driver header file
drivers/video/fbdev/mxc/mxc_hdmi.c	HDMI display driver implemention
sound/soc/fsl/fsl_hdmi.c	HDMI Audio SoC DAI driver implementation
sound/soc/fsl/imx-hdmi-dma.c	HDMI Audio SoC platform DMA driver implementation
drivers/mxc/hdmi-cec/hdmi-cec.c	HDMI CEC driver implemention. The HDMI CEC library files are provided in imx-lib repo on codeaurraforum.
drivers/mxc/hdmi-cec/hdmi-cec.c	HDMI CEC driver implemention. The HDMI CEC library files are provided in imx-lib on code aurora forum.

6.4.6.8.5   Menu Configuration Options

There are three main Linux kernel configuration options used to select and include HDMI driver functionality in the Linux OS image.

HDMI video support is dependent on support for the Synchronous Panel Framebuffer and also on the inclusion of IPUv3 support.

CONFIG_FB_MXC_HDMI provides support for the HDMI video driver and can be selected with Device Drivers > Graphics support > Support for frame buffer devices > MXC HDMI driver support

CONFIG_SND_SOC_IMX_HDMI provides support for HDMI audio through the ALSA/SoC subsystem, and can be selected with Device Drivers > Sound card support > Advanced Linux Sound Architecture > ALSA for SoC audio support > SoC Audio support for IMX - HDMI

Selecting either of the previous two configuration options will cause the MXC HDMI Core configuration option, CONFIG_MFD_MXC_HDMI, to be selected. This option can be selected wtih Device Drivers > Multifunction device drivers > MXC HDMI Core

CONFIG_MXC_HDMI_CEC option provides support for the HDMI CEC driver, and can be selected with Device Drivers > MXC support drivers > MXC HDMI CEC (Consumer Electronic Control) support

6.4.6.9   External HDMI

6.4.6.9.1   Introduction

The High Definition Multimedia Interface (HDMI) driver supports the external SiI9022 HDMI hardware module providing capability to transfer uncompressed video, audio, and data using a single cable.

The HDMI driver is divided into two sub-components: a video display device driver that integrates with the Linux Frame Buffer API and an S/PDIF audio driver that transfers S/PDIF audio data to SiI9022 HDMI hardware module.

The HDMI driver is only for demo application and supports the following features:

- HDMI video output supports 1080p60 and 720p60 resolutions.
- Support for reading EDID information from an HDMI sink device for video.
- Hotplug detection
- HDMI audio playback (2 channels, 16/24 bit, 44.1 KHz sample rate)

External HDMI is supported on i.MX 6 SoloLite and 7ULP SoC.

Output data is transmitted via three Transition-Minimized Differential Signaling (TMDS) channels to an HDMI sink device external to the SoC. Additionally, the HDMI carries a VESA Data Display Channel (DDC). DDC is an I2C interface which allows the HDMI source to query the HDMI sink for Extended Display Identification Data (EDID). A CEC channel provides optional high-level control functions between the source and sink devices.

6.4.6.9.2   Software Operation

The HDMI driver is divided into sub-components based on its two primary purposes: providing video and audio to an HDMI sink device.

The audio output depends on video display.

6.4.6.9.3   Source Code Structure

The source code for the HDMI driver is divided into the HDMI display driver and HDMI audio driver.

The HDMI display driver source is available in drivers/video/fbdev/mxc. The HDMI Audio driver source is in sound/soc/fsl

Table 6-19.   HDMI Source

File	Description
drivers/video/fbdev/mxc/mxsfb_sii902x.c	HDMI display driver implementation.
sound/soc/fsl/imx-spdif.c	S/PDIF Audio SoC Machine driver implementation.
sound/soc/fsl/fsl_spdif.c	S/PDIF Audio SoC DAI driver implementation.

6.4.6.9.4   Menu Configuration Options

There are two main Linux kernel configuration options used to select and include HDMI driver functionality in the Linux OS image.

The following configuration options are required to enable HDMI support.

The CONFIG_FB_MXS_SII902X option provides support for the Sii902x HDMI video driver and can be selected with Device Drivers > Support for frame buffer devices > Si Image SII9022 DVI/HDMI Interface Chip.

HDMI video on i.MX 6Sololite is dependent on MXC ELCDIF Framebuffer.

The CONFIG_SND_SOC_IMX_SII902X option provides support for the HDMI Audio driver and can be selected with Device Drivers > Sound card support > ALSA for SoC audio support > ommon SoC Audio options for Freescale CPUs: > SoC Audio support for i.MX boards with sii902x

6.5 Video for Linux 2 (V4L2)

6.5.1 Introduction

The Video for Linux Two (V4L2) driver is plug-in for the V4L2 framework that enables support for camera capture and display.

Some i.MX SoC support V4L2 based on the associated images processing units and capture hardware.

For more information on V4L2 go to the API specification for Linux Video for Linux 2 available at Linux Media Subsystem Documentation.

The V4L2 APIs enable camera and display controls but i.MX 8 only supports V4L2 capture and not display using the DPU instead for display control. i.MX 6 and i.MX 7 use both capture and display V4L2.

6.5.1.1 i.MX 8 DPU V4L2

The Video for Linux Two (V4L2) driver on i.MX 8 is plug-in for the V4L2 framework that enables support for camera capture only with the Display Processing Unit (DPU).

The V4L2 DPU camera driver supports only basic capture. The V4l2 capture device takes incoming video images, either from a camera or a TV decoder, and captures the images to memory. The features supported by the V4L2 driver are as follows:

- RGB 24-bit and YUV 4:2:2 interleaved formats for capture interface
- Plug-in of different sensor drivers
- Streaming (queued) input buffer
- Programmable input and output pixel format and size
- RGB 16, 24, and 32-bit, YUV 4:2:0, and 4:2:2 interleaved input formats

The command modprobe mxc_v4l2_capture must be run before using V4L2 camera functions.

6.5.1.2   PxP V4L2

The Video for Linux Two (V4L2) drivers for PxP are used for dsiplay output only.

6.5.1.3   i.MX 6 with IPU V4L2

The Video for Linux Two (V4L2) drivers are plug-ins to the V4L2 framework that enable support for camera and preprocessing functions, as well as video and post-processing functions. The V4L2 camera driver implements support for all camera-related functions. The V4l2 capture device takes incoming video images, either from a camera or a stream, and manipulates them. The output device takes video and manipulates it, then sends it to a display or similar device.

The features supported by the IPU V4L2 driver are the follows:

- Direct preview and output to SDC foreground overlay plane (with synchronized to LCD refresh)
- Direct preview to graphics frame buffer (without synchronized to LCD refresh)
- Color keying or alpha blending of frame buffer and overlay planes
- Streaming (queued) capture from IPU encoding channel
- Direct (raw Bayer) still capture (sensor dependent)
- Programmable pixel format, size, frame rate for preview and capture
- Programmable rotation and flipping using custom API
- RGB 16-bit, 24-bit, and 32-bit preview formats
- Raw Bayer (still only, sensor dependent), RGB 16, 24, and 32-bit, YUV 4:2:0 and 4:2:2 planar, YUV 4:2:2 interleaved, and JPEG formats for capture
- Control of sensor properties including exposure, white-balance, brightness, contrast, and so on
- Plug-in of different sensor drivers
- Link post-processing resize and CSC, rotation, and display IPU channels
- Streaming (queued) input buffer
- Double buffering of overlay and intermediate (rotation) buffers
- Configurable 3+ buffering of input buffers
- Programmable input and output pixel format and size
- Programmable scaling and frame rate
- RGB 16, 24, and 32-bit, YUV 4:2:0 and 4:2:2 planar, and YUV 4:2:2 interleaved input formats
- TV output

The command modprobe mxc_v4l2_capture must be run before V4L2 functions.

6.5.1.4 IPU V4L2 Capture Device

The V4L2 capture device includes two interfaces:

- Capture interface-uses IPU pre-processing ENC channels to record the YCrCb video stream
- Overlay interface-uses the IPU device driver to display the preview video to the SDC foreground and background panel.

V4L2 capture support can be selected during kernel configuration. The driver includes two layers. The top layer is the common Video for Linux driver, which contains chain buffer management, stream API and other ioctl interfaces. The files for this device are located in

```
drivers/media/platform/mxc/capture/
```

.

The V4L2 capture device driver is in the mxc_v4l2_capture.c file. The low level overlay driver is in the ipu_fg_overlay_sdc.c, ipu_bg_overlay_sdc.c

This code (ipu_prp_enc.c) interfaces with the IPU ENC hardware, and ipu_still.c interfaces with the IPU CSI hardware. Sensor frame rate control is handled by VIDIOC_S_PARM ioctl. Before the frame rate is set, the sensor turns on the AE and AWB turn on. The frame rate may change depending on light sensor samples.

Drivers for specific cameras can be found in

```
drivers/media/platform/mxc/capture/
```

.

6.5.2 V4L2 Capture Device

The V4L2 capture device includes two interfaces:

- Capture interface-uses i.MX processing engine to record the YCrCb video stream
- Overlay interface-uses i.MX processing engine to display the preview video to the SDC foreground and background panel.

The driver includes two layers. The top layer is the common Video for Linux driver, which contains chain buffer management, stream API and other ioctl interfaces. The low level layer is the i.MX SoC implementation for the display engine associated with the SoC detailed in each V4l2 SoC chapter.

6.5.2.1   V4L2 Capture IOCTLs

Currently, the memory map stream API is supported. Supported V4L2 IOCTLs include the following:

- VIDIOC_QUERYCAP
- VIDIOC_G_FMT
- VIDIOC_S_FMT
- VIDIOC_REQBUFS
- VIDIOC_QUERYBUF
- VIDIOC_QBUF
- VIDIOC_DQBUF
- VIDIOC_STREAMON
- VIDIOC_STREAMOFF
- VIDIOC_OVERLAY
- VIDIOC_G_FBUF
- VIDIOC_S_FBUF
- VIDIOC_G_CTRL
- VIDIOC_S_CTRL
- VIDIOC_CROPCAP
- VIDIOC_G_CROP
- VIDIOC_S_CROP
- VIDIOC_S_PARM
- VIDIOC_G_PARM
- VIDIOC_ENUMSTD
- VIDIOC_G_STD
- VIDIOC_S_STD
- VIDIOC_ENUMOUTPUT
- VIDIOC_G_OUTPUT
- VIDIOC_S_OUTPUT

V4L2 control code has been extended to provide support for rotation. The ID is V4L2_CID_PRIVATE_BASE. Supported values include:

- 0-Normal operation
- 1-Vertical flip
- 2-Horizontal flip
- 3-180° rotation
- 4-90° rotation clockwise
- 5-90° rotation clockwise and vertical flip
- 6-90° rotation clockwise and horizontal flip
- 7-90° rotation counter-clockwise

Figure below shows a block diagram of V4L2 Capture API interaction.
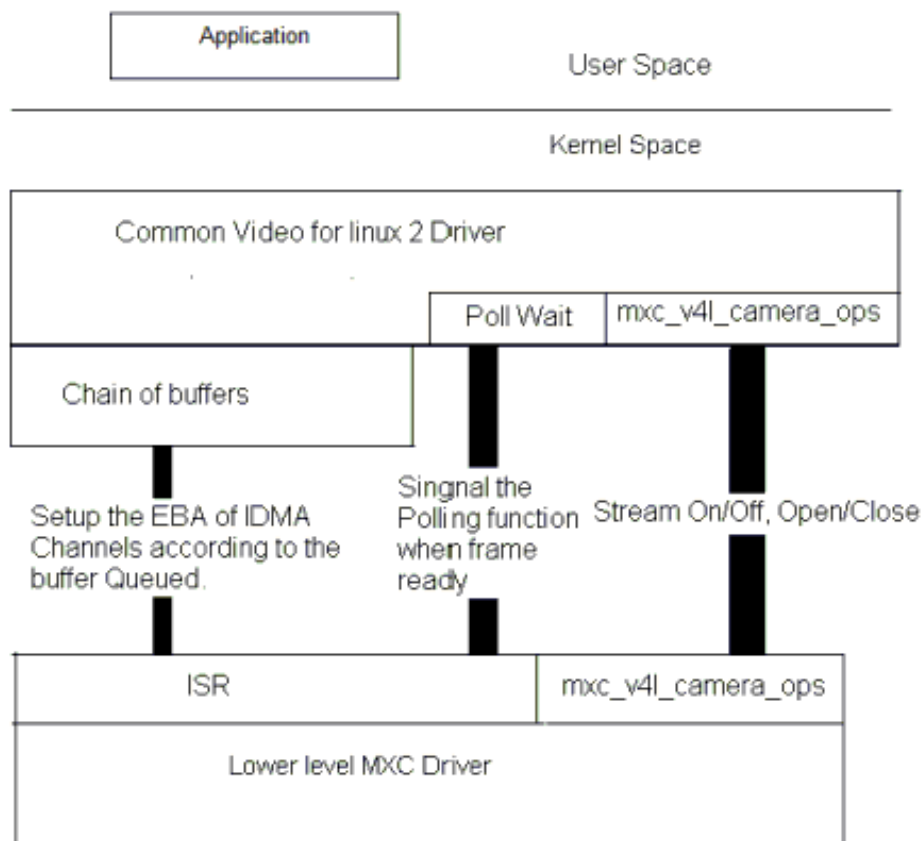
Figure 6-7. Video4Linux2 Capture API Interaction

6.5.2.2 Use of the V4L2 Capture APIs

This section describes a sample V4L2 capture process. The application completes the following steps:

1. Sets the capture pixel format and size by IOCTL VIDIOC_S_FMT.
2. Sets the control information by IOCTL VIDIOC_S_CTRL for rotation usage.
3. Requests a buffer using IOCTL VIDIOC_REQBUFS. The common V4L2 driver creates a chain of buffers (currently the maximum number of frames is 3).
4. Memory maps the buffer to its user space.
5. Queues buffers using the IOCTL command VIDIOC_QBUF.
6. Starts the stream using the IOCTL VIDIOC_STREAMON. This IOCTL enables the i.MX Processing Enginee tasks and the IDMA channels. When the processing is completed for a frame, the driver switches to the buffer that is queued for the next frame. The driver also signals the semaphore to indicate that a buffer is ready.
7. Takes the buffer from the queue using the IOCTL VIDIOC_DQBUF. This IOCTL blocks until it has been signaled by the ISR driver.

8. Stores the buffer to a YCrCb file.
9. Replaces the buffer in the queue of the V4L2 driver by executing VIDIOC_QBUF again.

For the V4L2 still image capture process, the application completes the following steps:

1. Sets the capture pixel format and size by executing the IOCTL VIDIOC_S_FMT.
2. Reads one frame still image with YUV422.

FOr the V4L2 overlay support use case, the application completes the following steps:

1. Sets the overlay window by IOCTL VIDIOC_S_FMT.
2. Turns on overlay task by IOCTL VIDIOC_OVERLAY.
3. Turns off overlay task by IOCTL VIDIOC_OVERLAY.

6.5.3 V4L2 Output Device

The driver implements the standard V4L2 API for output devices. V4L2 output device support can be selected during kernel configuration. The driver is available at

```
drivers/media/platform/mxc/output/mxc_vout.c
```

.

6.5.3.1 V4L2 Output IOCTLs

Currently, the memory map stream API is supported. Supported V4L2 IOCTLs include the following:

- VIDIOC_QUERYCAP
- VIDIOC_REQBUFS
- VIDIOC_G_FMT
- VIDIOC_S_FMT
- VIDIOC_QUERYBUF
- VIDIOC_QBUF
- VIDIOC_DQBUF
- VIDIOC_STREAMON
- VIDIOC_STREAMOFF
- VIDIOC_G_CTRL
- VIDIOC_S_CTRL
- VIDIOC_CROPCAP
- VIDIOC_G_CROP

- VIDIOC_S_CROP
- VIDIOC_ENUM_FMT

The V4L2 control code has been extended to provide support for de-interlace motion. For this use, the ID is V4L2_CID_MXC_MOTION. Supported values include the following:

- 0-Medium motion
- 1-Low motion
- 2-High motion

6.5.3.2   Use of the V4L2 Output APIs

This section describes a sample V4L2 output process that uses the V4L2 output APIs. The application completes the following steps:

1. Sets the input pixel format and size using IOCTL VIDIOC_S_FMT.
2. Sets the control information using IOCTL VIDIOC_S_CTRL, for rotation, de-interlace motion(if needed).
3. Sets the output information using IOCTL VIDIOC_S_CROP.
4. Requests a buffer using IOCTL VIDIOC_REQBUFS. The common V4L2 driver creates a chain of buffers (not allocated yet).
5. Memory maps the buffer to its user space.
6. Executes the IOCTL VIDIOC_QUERYBUF to query buffers.
7. Passes the data that requires post-processing to the buffer.
8. Queues the buffer using the IOCTL command VIDIOC_QBUF.
9. Executes the IOCTL VIDIOC_DQBUF to dequeue buffers.
10. Starts the stream by executing IOCTL VIDIOC_STREAMON.
11. Stop the stream by excuting IOCTL VIDIOC_STREAMOFF.

6.5.4   Software Operatoins

6.5.4.1   Source Code Structure

The following table lists the source and header files associated with the V4L2 drivers.

These files are available in drivers/media/platform/mxc

Table 6-20. V4L2 Driver Files

File	Description
drivers/media/platform/mxc/output/mxc_vout.c	MX6 and MX7 V4L2 output device driver
drivers/media/platform/mxc/output/mxc_pxp_v4l2.c	V4L2 PxP output device driver
drivers/media/platform/mxc/output/mxc_pxp_v4l2.h	V4L2 PxP output device driver header
drivers/media/platform/mxc/capture/mxc_v4l2_capture.c	V4L2 capture device driver
drivers/media/platform/mxc/capture/mxc_v4l2_capture.h	Header file for V4L2 capture device driver
drivers/media/platform/mxc/capture/ipu_bg_overlay_sdc.c	IPU synchronous background driver
drivers/media/platform/mxc/capture/ipu_fg_overlay_sdc.c	IPU synchronous forground driver
drivers/media/platform/mxc/capture/ipu_prp_sw.h	IPU Pre-processing header
drivers/media/platform/mxc/capture/ipu_still.c	IPU Pre-processing still image capture driver
drivers/media/platform/mxc/capture/ipu_prp_vf_sdc_bg.c	IPU Pre-processing view finder (synchronous background)
drivers/media/platform/mxc/capture/ipu_prp_enc.c	IPU Pre-processing Encoder driver
drivers/media/platform/mxc/capture/ipu_csi_enc.c	IPU CSI interface driver

Drivers for V4L2 cameras can be found in divers/media/platform/mxc/capture.

Drivers for V4L2 output can be found in drivers/media/platform/mxc/output

6.5.4.2 Menu Configuration Options

The kernel configuration options are provided below.

Device Drivers > V4L platform devices > MXC Video For Linux Video Capture

Device Drivers > V4L platform devices > MXC Video For Linux Video Output

6.6 Video Analog-to-Digital Converter (VADC)

6.6.1 Introduction

The video analog-to-digital converter (VADC) consists of an analog video front end (AFE), and a digital video decoder. The AFE accepts NTSC or PAL input from a device, such as an analog camera.

The two parts are configured in the VADC driver. The video decoder outputs the YUV444-formatted data.

The Video ADC has the following features:

- Internal voltage and current reference generator

- 10-bit resolution (9.5 bit ENOB at 66.5 Msps)
- 4 analog inputs, with all inputs available for CVBS
- Programmable anti-aliasing filter, gain, and clamp

The video decoder has the following features:

- NTSC/PAL decoder
- Direct data path (no complex resampling)
- Automatic standards detection
- 2D adaptive comb filter
- Datapath/clocking architecture encompasses a time base corrector for VCR signals
- Luma passband is flat to > 6 MHz

6.6.2 Software Operation

The VADC driver is located under the Linux V4L2 architecture and it implements the V4L2 capture interfaces. Applications cannot use the camera driver directly. Instead, the applications use the V4L2 capture driver to open and close the camera for image capture.

The V4L2 capture supports the following operation:
- Capture stream mode

The following picture format is supported:
- YUV444

The following picture sizes are supported:
- PAL
- NTSC

6.6.3 Source Code Structure

Table below shows the VADC driver source files available in the drivers/media/platform/mxc/capture.

Table 6-21.   VADC Driver Files

File	Description
drivers/media/platform/mxc/capture/mxc_vadc.c	VADC driver source code
drivers/media/platform/mxc/capture/mxc_vadc.h	VADC driver Header

6.6.4   Menu Configuration Options

In menu configuration enable the following module:

Device Drivers > Multimedia devices > Video capture adapters > MXC Video For Linux Camera > MXC VADC support

6.6.5   DTS Configuration

VADC analog inputs can choose [0-3]. CSI1 or CSI2 can be used to capture the VADC data. They can be configured in the DTS file.

For example:

```
vadc_in = <0>;  /* VADC input select */
csi_id = <1>; /* CSI select */
```

The VADC input selected to vin1 and CSI2 is used to capture the VADC data.

6.7   Video Processing Unit (VPU)

6.7.1   Introduction

The VPU hardware performs all of the codec computation and most of the bitstream parsing/packeting. Therefore, the software takes advantage of less control and effort to implement a complex and efficient multimedia codec system.

Different VPUs are supported on i.MX 6 and i.MX 8 SoC. The following table lists the different VPUs.

Table 6-22.   VPU

SoC	VPU	Library
i.MX 6	Chips and Media	imx-vpu.so
i.MX 8M Quad and 8M Mini	Hantro	imx-hantro.so
i.MX 8QuadMax and i.MX 8QuadXPlus	Malone	no library

6.7.2 Software Operation

The VPU software can be divided into two parts: the kernel driver and the user-space library as well as the application in user space. The kernel driver takes responsibility for system control and reserving resources (memory/IRQ). It provides an IOCTL interface for the application layer in user-space as a path to access system resources. The application in user-space calls related IOCTLs and codec library functions to implement a complex codec system.

The VPU kernel driver includes the following functions:

- Module initialization which initializes the module with the device-specific structure
- Device initialization which initializes the VPU clock and hardware and request the IRQ
- Interrupt servicing routine which supports events that one frame has been finished
- File operation routine which provides the following interfaces to user space:
- File open
- File release
- File IOCTL to provide interface for memory allocating and releasing
- Memory map for register and memory accessing in user space

The VPU user space driver has the following functions:

- Codec lib
- Initializes codec system
- Sets codec system configuration
- Controls codec system by command
- Reports codec status and result
- System I/O operation
- Requests and frees memory
- Maps and unmaps memory/register to user space
- Device management

User space application for simple verification:

- Read video raw data
- YUV file dump
- General options to configure the codec behavior

The following figure shows a simple workflow shown in the H.264 example.
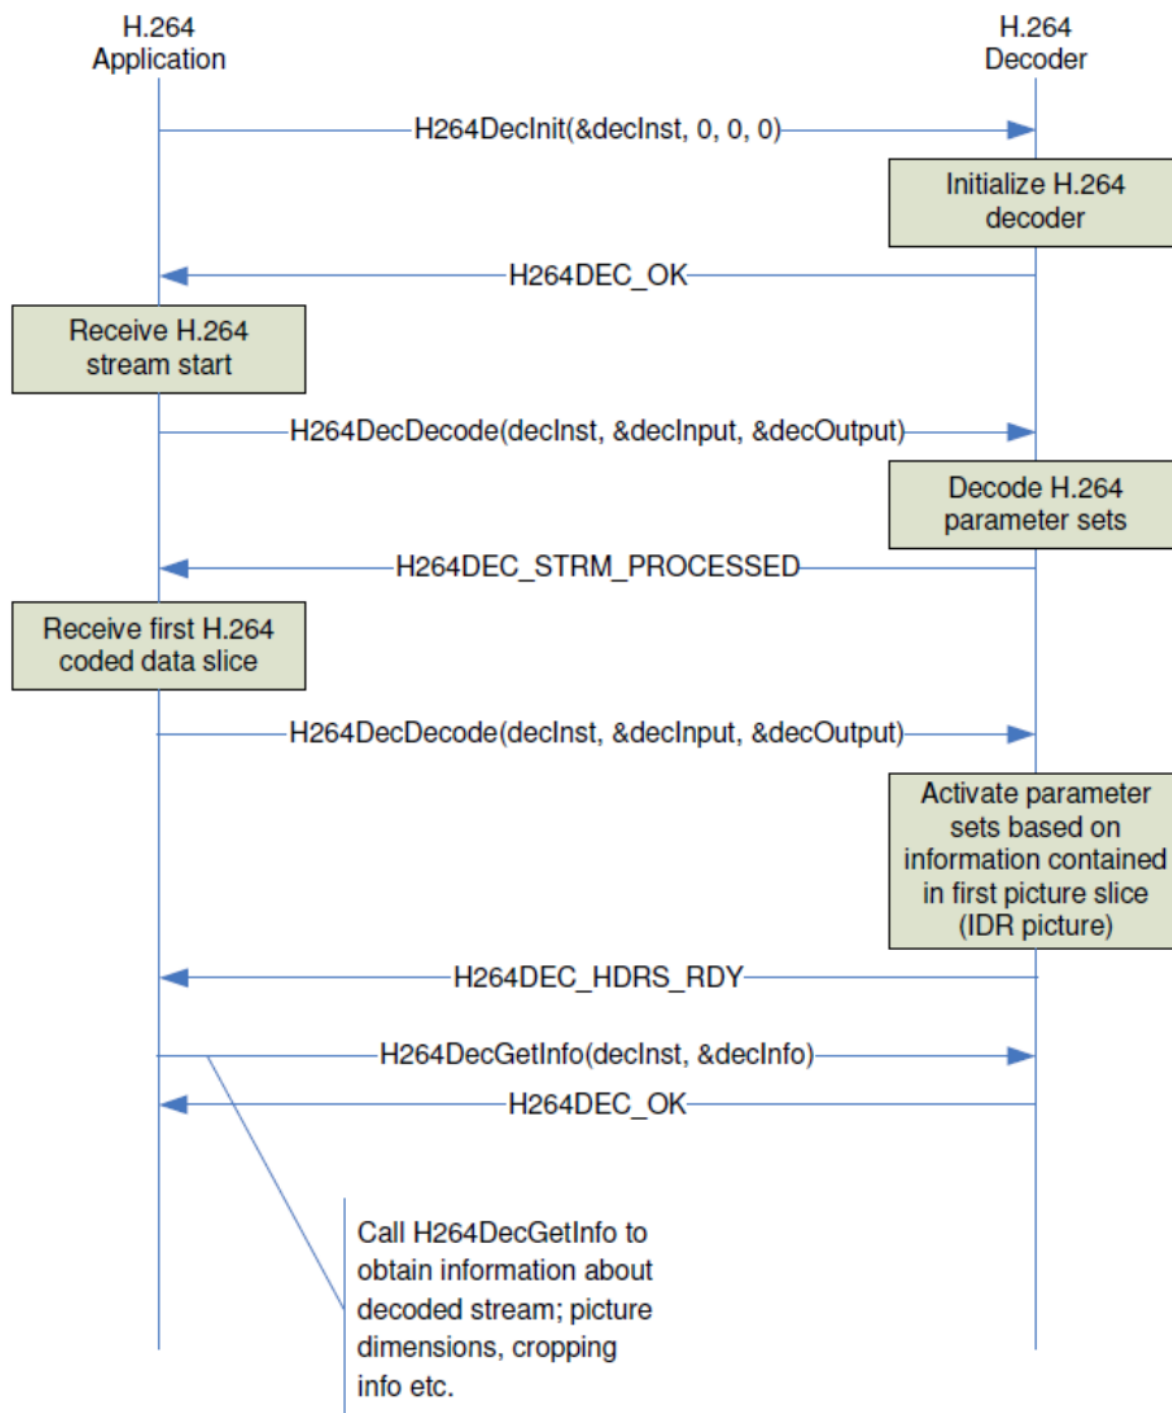
Figure 6-8. Simple Workflow Shown in the H.264 Example
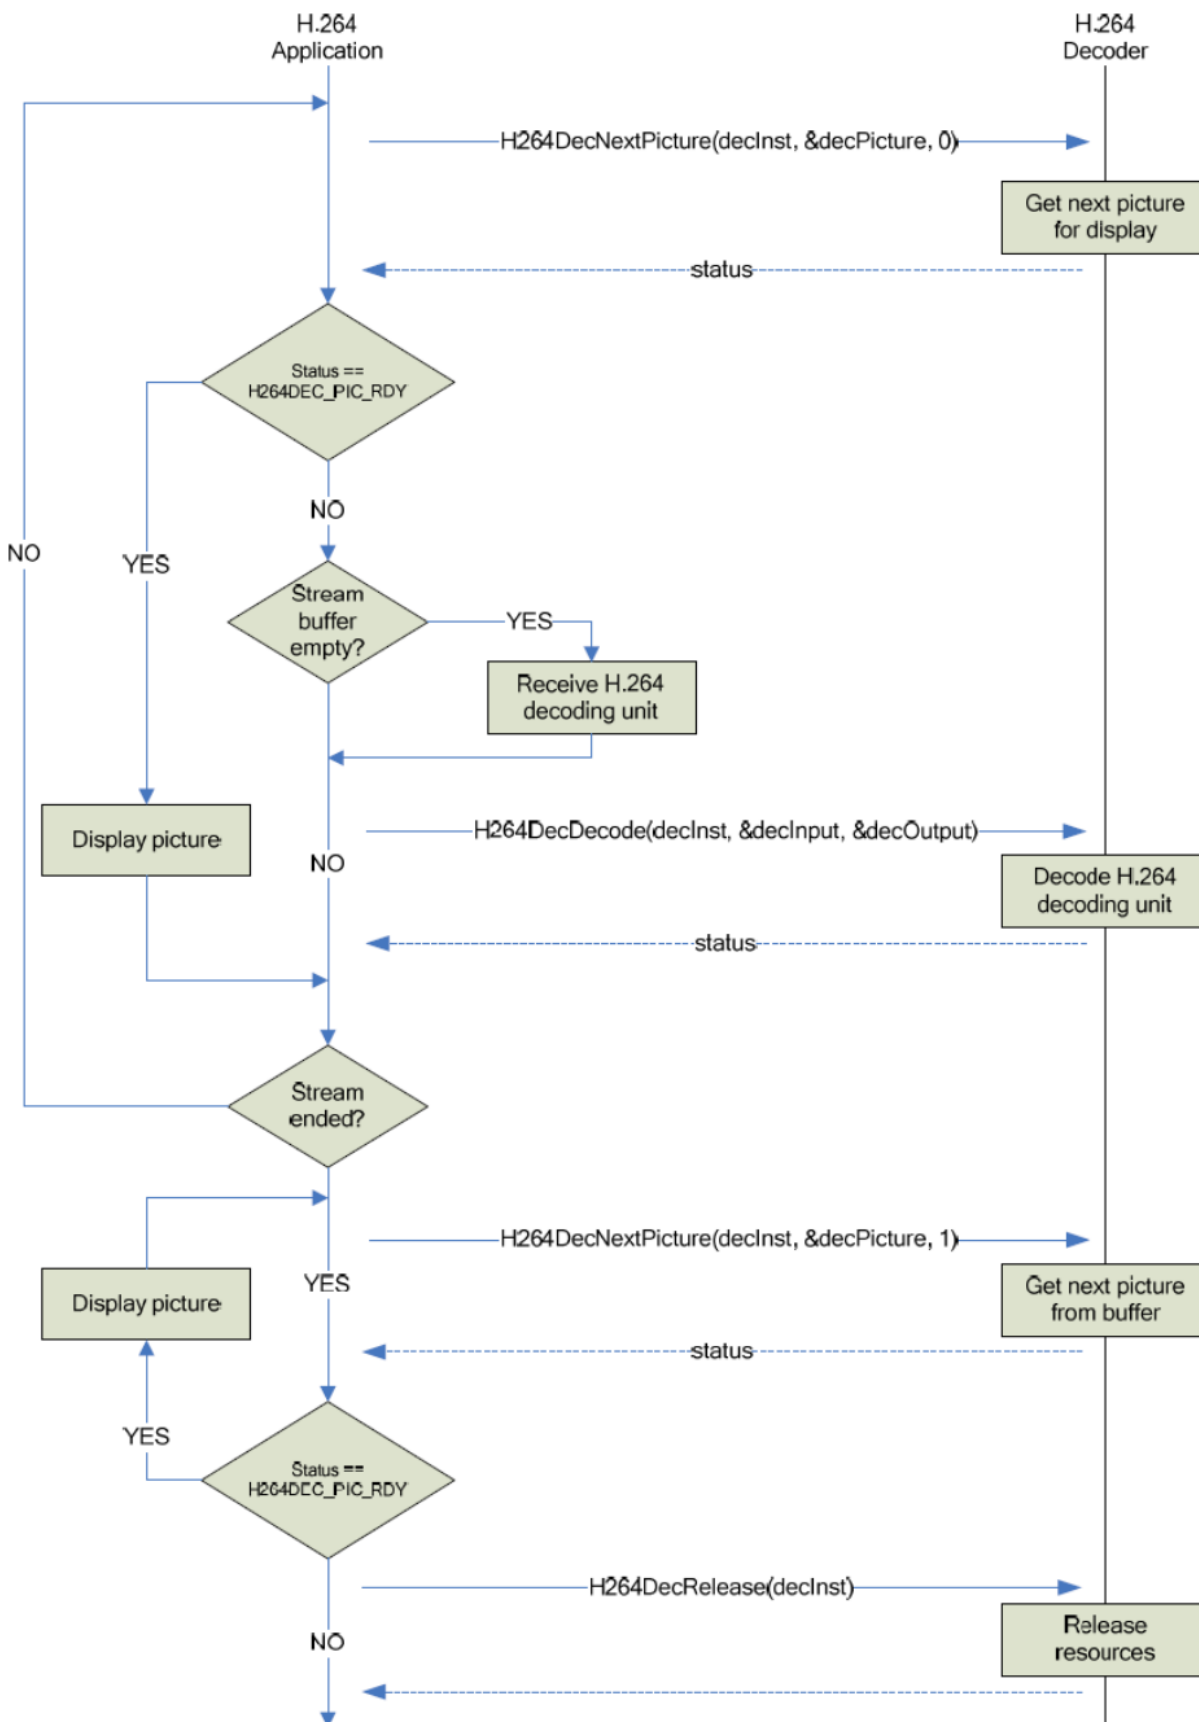
There is only a user-space programming interface for the VPU module. A user in the application layer cannot access the kernel driver interface directly. The VPU library accesses the kernel driver interface for users.

There is one unified interface to wrap all different video formats. The following are the related APIs:

```
CODEC_STATE decoder_decode_xxx(CODEC_PROTOTYPE * arg,STREAM_BUFFER * buf, OMX_U32 *
consumed,FRAME * frame);
CODEC_STATE decoder_getinfo_xxx(CODEC_PROTOTYPE * arg,STREAM_INFO * pkg);
CODEC_STATE decoder_setppargs_xxx(CODEC_PROTOTYPE * codec,PP_ARGS * args);
CODEC_STATE decoder_setframebuffer_xxx(CODEC_PROTOTYPE * arg, BUFFER *buff,OMX_U32
available_buffers);
CODEC_STATE decoder_pictureconsumed_xxx(CODEC_PROTOTYPE * arg, BUFFER *buff);
CODEC_STATE decoder_getframe_mpeg4(CODEC_PROTOTYPE * arg, FRAME * frame,OMX_BOOL eos);
FRAME_BUFFER_INFO decoder_getframebufferinfo_xxx(CODEC_PROTOTYPE * arg);
CODEC_STATE decoder_endofstream_xxx(CODEC_PROTOTYPE * arg)
OMX_S32 decoder_scanframe_xxx(CODEC_PROTOTYPE * arg, STREAM_BUFFER * buf,OMX_U32 * first,
OMX_U32 * last);
CODEC_STATE decoder_abort_xxx(CODEC_PROTOTYPE * arg);
CODEC_STATE decoder_abortafter_xxx(CODEC_PROTOTYPE * arg);
CODEC_STATE decoder_setnoreorder_xxx(CODEC_PROTOTYPE * arg, OMX_BOOL no_reorder);
static void decoder_destroy_xxx(CODEC_PROTOTYPE * arg)
```

6.7.3 Menu Configuration Options

In menu configuration enable the following module for the VPU driver:

For i.MX 6 with VPU select Device Drivers > MXC support drivers > Support for MXC VPU(Video Processing Unit)

For i.MX 8M select Device Drivers > MXC support drivers > MXC HANTRO (Video Processing Unit) support

For i.MX 8QuadMax and i.MX 8QuadXPlus select Device Drivers > MXC support drivers > Support for MXC VPU(Video Processing Unit) DECODER

Chapter 7
Audio

7.1 Advanced Linux Sound Architecture (ALSA) System on a Chip (ASoC) Sound

7.1.1 ALSA Sound Driver Introduction

The Advanced Linux Sound Architecture (ALSA), now the most popular architecture in Linux system, provides audio and MIDI functionality to the Linux operating system.

ALSA has the following significant features:

- Efficient support for all types of audio interfaces, from consumer sound cards to professional multichannel audio interfaces.
- Fully modularized sound drivers.
- SMP and thread-safe design.
- User space library (alsa-lib) to simplify application programming and provide higher level functionality.
- Support for the older Open Sound System (OSS) API, providing binary compatibility for most OSS programs.

ALSA System on Chip (ASoC) layer is designed for SoC audio. The overall project goal of the ASoC layer provides better ALSA support for embedded system on chip processors and portable audio CODECs.

The ASoC layer also provides the following features:
- CODEC independence. Allows reuse of CODEC drivers on other platforms and machines.
- Easy I2S/PCM audio interface setup between CODEC and SoC. Each SoC interface and CODEC registers its audio interface capabilities with the core.

- Dynamic Audio Power Management (DAPM). DAPM is an ASoC technology designed to minimize audio subsystem power consumption no matter what audio-use case is active. DAPM guarantees the lowest audio power state at all times and is completely transparent to user space audio components. DAPM is ideal for mobile devices or devices with complex audio requirements.
- Pop and click reduction. Pops and clicks can be reduced by powering the CODEC up/down in the correct sequence (including using digital mute). ASoC signals the CODEC when to change power states.
- Machine-specific controls. Allow machines to add controls to the sound card, for example, volume control for speaker amp.
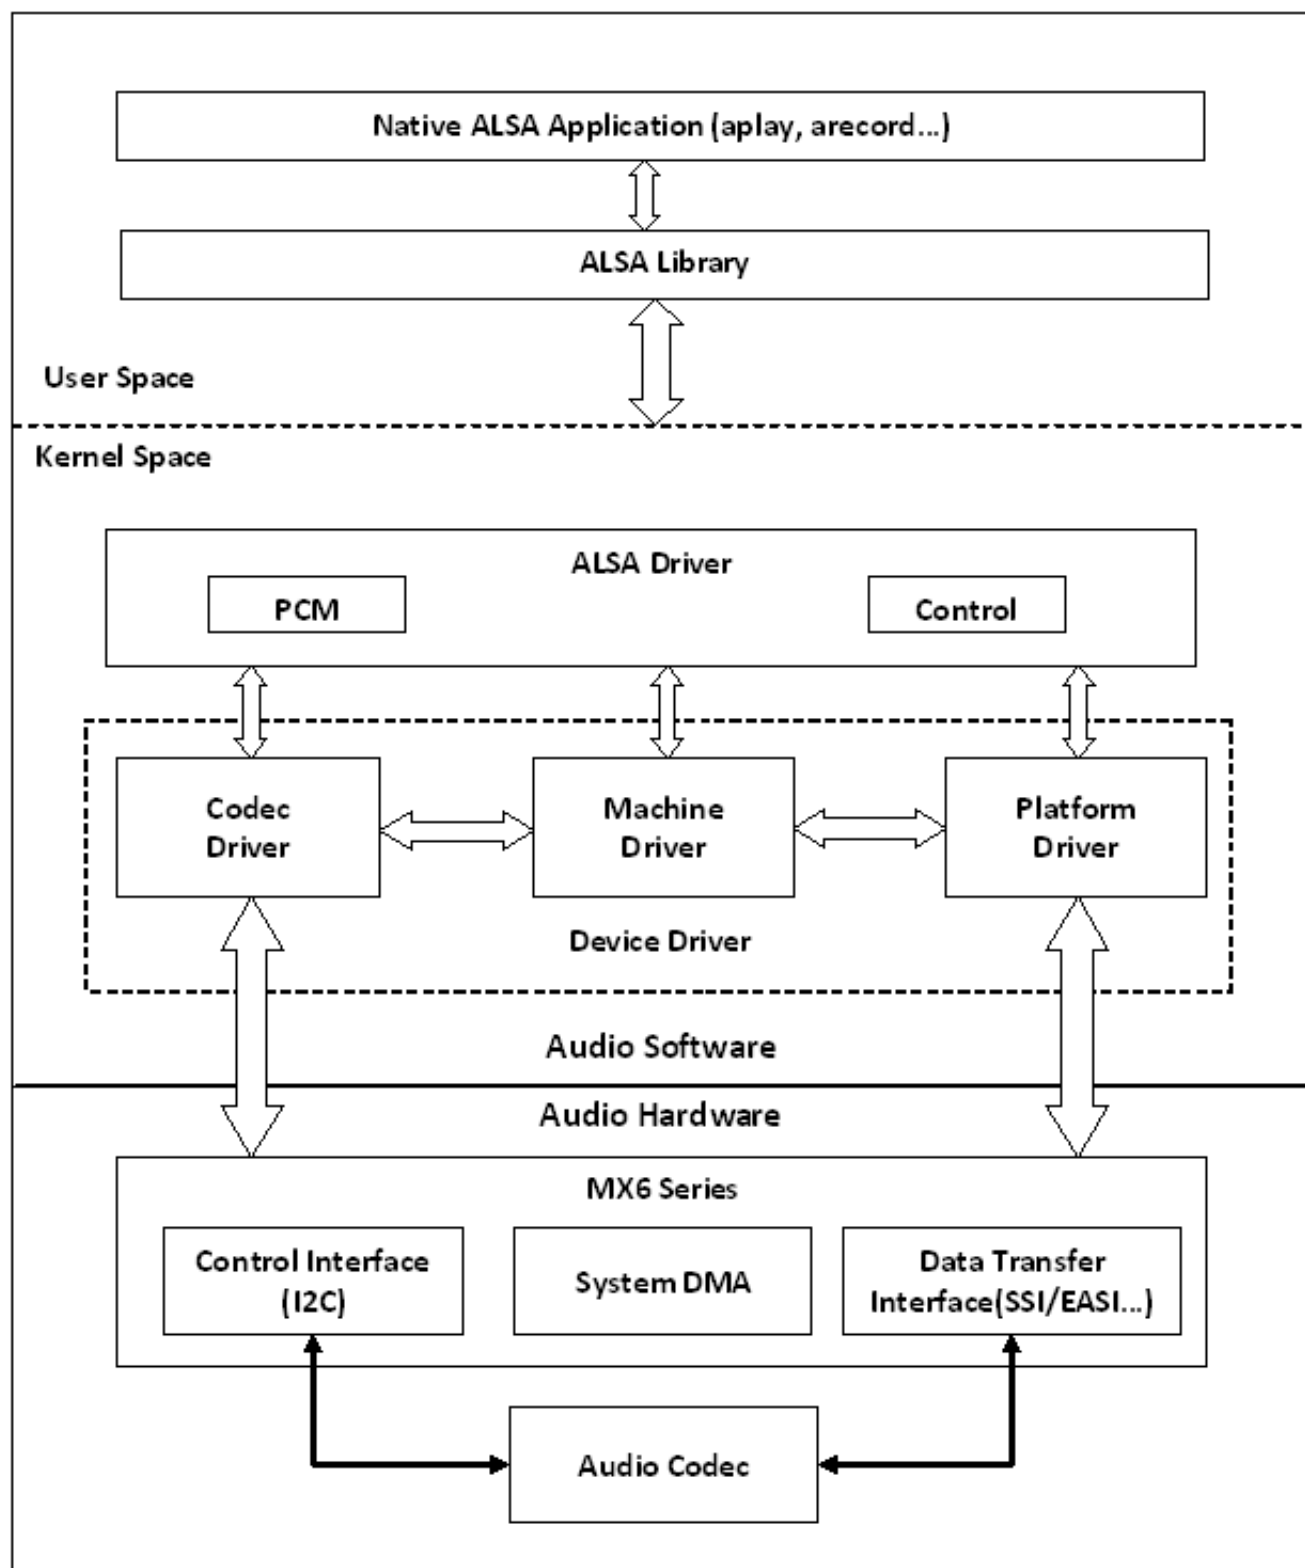
Figure 7-1. ALSA SoC Software Architecture

ASoC basically splits an embedded audio system into 3 components:

- Machine driver-handles any machine-specific controls and audio events, such as turning on an external amp at the beginning of playback.
- Platform driver-contains the audio DMA engine and audio interface drivers (for example, I^2S, AC97, PCM) for that platform.
- CODEC driver-platform independent and contains audio controls, audio interface capabilities, the CODEC DAPM definition, and CODEC I/O functions.

More detailed information about ASoC can be found in the Linux kernel documentation in the Linux OS source tree at linux/Documentation/sound/alsa/soc and at www.alsa-project.org/main/index.php/ASoC.

7.1.2 SoC Sound Card

Currently, the stereo CODEC (WM8958, WM8960, WM8962, and WM8524), 7.1 CODEC (cs42888), and AM/FM CODEC (si4763) drivers are implemented using ASoC architecture.

These sound card drivers are built in independently. The stereo sound card supports stereo playback and capture. The 7.1 sound card supports up to eight channels of audio playback. While enabling ASRC, 7.1 sound card only supports 2 or 6 channels audio playback. The AM/FM sound card supports radio PCM capture.

NOTE

The 7.1 CODEC is only supported on the i.MX 6Quad and i.MX 6Solo SABRE Auto platform.

The AM/FM CODEC is only supported on the i.MX 6Quad and i.MX 6Solo SABRE Auto platform.

7.1.2.1 Stereo CODEC Features

The stereo CODEC supports the following features:

- Sample rates for playback and capture are 8 KHz, 32 KHz, 44.1 KHz, 48 KHz, and 96 KHz
- Channels:
  - Playback: supports two channels.
  - Capture: supports two channels.
- Audio formats:
  - Playback:
    - SNDRV_PCM_FMTBIT_S16_LE

- SNDRV_PCM_FMTBIT_S20_3LE
- SNDRV_PCM_FMTBIT_S24_LE
  - Capture:
    - SNDRV_PCM_FMTBIT_S16_LE
    - SNDRV_PCM_FMTBIT_S20_3LE
    - SNDRV_PCM_FMTBIT_S24_LE

7.1.2.2   7.1 Audio Codec Features

- Sample rates for playback and record:
  - 48 KHz, 96 KHz, 192 KHz
  - Playback: 5.512 k, 8 k, 11.025 k, 16 k, 22 k, 32 k, 44.1 k, 48 k, 64 k, 88.2 k, 96 k, 176.4 k, 192 k (ASRC enabled)
- Channels:
  - Playback: 2, 4, 6, 8 channels
  - Playback(ASRC enabled): 2, 6 channels
  - Capture: 2, 4 channels
- Audio formats:
  - Playback:
    - SNDRV_PCM_FMTBIT_S16_LE
    - SNDRV_PCM_FMTBIT_S20_3LE
    - SNDRV_PCM_FMTBIT_S24_LE
  - Playback(ASRC enabled):
    - SNDRV_PCM_FMTBIT_S16_LE
    - SNDRV_PCM_FMTBIT_S24_LE
  - Capture:
    - SNDRV_PCM_FMTBIT_S16_LE
    - SNDRV_PCM_FMTBIT_S20_3LE
    - SNDRV_PCM_FMTBIT_S24_LE

7.1.2.3   AM/FM Codec Features

- Supported sample rate for Capture: 48 KHz
- Supported channels:
  - Capture: supports two channels.
- Supported audio formats:
  - Capture: SNDRV_PCM_FMTBIT_S16_LE

7.1.2.4 Sound Card Information

The registered sound card information can be listed as follows using the commands aplay -l and arecord -l. For example, the stereo sound card is registered as card 0.

```
root@freescale /$ aplay -l
**** List of PLAYBACK Hardware Devices ****
card 0: wm8962audio [wm8962-audio], device 0: HiFi wm8962-0 []
 Subdevices: 1/1
 Subdevice #0: subdevice #0
```

7.1.3 Hardware Operation

The following sections describe the hardware operation of the ASoC driver.

7.1.3.1 Stereo Audio CODEC

The stereo audio CODEC is controlled by the I^{2}C interface. The audio data is transferred from the user data buffer to/from the SSI FIFO through the DMA channel. The DMA channel is selected according to the audio sample bits. AUDMUX is used to set up the path between the SSI port and the output port which connects with the CODEC. The CODEC works in master mode and provides the BCLK and LRCLK. The BCLK and LRCLK can be configured according to the audio sample rate.

The WM8958, WM8960, and WM8962 ASoC CODEC driver exports the audio record/ playback/mixer APIs according to the ASoC architecture.

The CODEC driver is generic and hardware independent code that configures the CODEC to provide audio capture and playback. It does not contain code that is specific to the target platform or machine. The CODEC driver handles:

- CODEC DAI and PCM configuration
- CODEC control I/O-using I^{2}C
- Mixers and audio controls
- CODEC audio operations
- DAC Digital mute control

The WM8958, WM8960, and WM8962 CODEC are registered as an I^{2}C client when the module initializes. The APIs are exported to the upper layer by the structure snd_soc_dai_ops .

Headphone insertion/removal can be detected through a GPIO interrupt signal.

SSI dual FIFO features are enabled by default.

7.1.3.2   7.1 Audio Codec

The 7.1 audio codec includes 8-channel DAC and 4-channel ADC, which are controlled by the I2C interface. The audio data is transferred from the user data buffer to the ESAI fifo, through a DMA channel. The DMA channel is selected according to audio sample bits. The codec works in slave mode as the ESAI provides the BCLK and LRCLK. The BCLK and LRCLK can be configured according to the audio sample rate. The ESAI supports up to eight audio output ports. While enabling ASRC, 7.1 audio codec supports 2 or 6 channel playback through ASRC. On the i.MX 6 Sabre ARD board, a CS42888 codec with 4 audio in port is used, each port receive two channels of data in the I2S format(network mode), providing 8-channel of playback functionality. This codec also has 2 audio output port connected with ESAI, providing 4-channel of recording functionality.

The codec driver is generic and hardware independent code that configures the codec to provide audio capture and playback. It does not contain code that is specific to the target platform or machine. The codec driver handles:

- Codec DAI and PCM configuration
- Codec control I/O-using I2C
- Mixers and audio controls
- Codec audio operations
- DAI Digital mute control

The CS42888 codec is registered as an I2C client when the module initializes. The APIs are exported to the upper layer by the structure snd_soc_dai_ops.

7.1.3.3   AM/FM Codec

The AM/FM codec is a virtual codec, it only has a PCM interface connected to the Tuner device. The audio data is transferred from the user data buffer to or from the SSI FIFO through the DMA channel. The DMA channel is selected according to the audio sample bits. AUDMUX is used to set up the path between the SSI port and the output port which connects with the codec. The codec works in master mode as it provides the BCLK and LRCLK. The BCLK and LRCLK can be configured according to the audio sample rate.

7.1.4   Software Operation

The following sections describe the software operation of the ASoC driver.

7.1.4.1   ASoC Driver Source Architecture

File imx-pcm-dma.c is shared by the stereo ALSA SoC driver, the 7.1 ALSA SoC driver and other CODEC driver. This file is responsible for preallocating DMA buffers and managing DMA channels.

The stereo CODEC is connected to the CPU through the SSI interface. fsl_ssi.c registers the CPU DAI driver for the stereo ALSA SoC and configures the on-chip SSI interface. wm8962.c registers the stereo CODEC and hifi DAI drivers. The direct hardware operations on the stereo codec are in wm8994.c, wm8960.c, and wm8962.c. imx-wm8958.c, imx-wm8960.c and imx-wm8962.c are the machine layer codes, which create the driver device and register the stereo sound card.

The multichannel codec is connected to the CPU through the ESAI interface. fsl_esai.c registers the CPU DAI driver for the stereo ALSA SoC and configures the on-chip ESAI interface. cs42888.c registers the multichannel CODEC and hifi DAI drivers. The direct hardware operations on the multichannel CODEC are in cs42888.c. imx-cs42888.c is the machine layer code which creates the driver device and registers the stereo sound card.

The AM/FM CODEC is connected to the CPU through the SSI interface. fsl_ssi.c registers the CPU DAI driver for the stereo ALSA SoC and configures the on-chip SSI interface. si476x.c registers the Tuner CODEC and Tuner DAI drivers. The direct hardware operations on the CODEC are in si476x.c. imx-si476x.c is the machine layer code which creates the driver device and registers the sound card.

7.1.4.2   Sound Card Registration

The codecs have the same registration sequence:

1. The codec driver registers the codec driver, DAI driver, and their operation functions.
2. The platform driver registers the PCM driver, CPU DAI driver and their operation functions, pre allocates buffers for PCM components and sets playback and capture operations as applicable.
3. The machine layer creates the DAI link between codec and CPU registers the sound card and PCM devices.

7.1.4.3   Device Open

The ALSA driver performs the following functions:

- Allocates a free substream for the operation to be performed.
- Opens the low level hardware device.
- Assigns the hardware capabilities to ALSA runtime information (the runtime structure contains all the hardware, DMA, and software capabilities of an opened substream).
- Configures DMA read or write channel for operation.
- Configures CPU DAI and CODEC DAI interface.
- Configures CODEC hardware.
- Triggers the transfer.

After triggering for the first time, the subsequent DMA read/write operations are configured by the DMA callback.

7.1.4.4   Device Tree Binding

See the following documents:

- Documentation/devicetree/bindings/sound/fsl,ssi.txt
- Documentation/devicetree/bindings/sound/fsl-sai.txt
- Documentation/devicetree/bindings/sound/fsl,esai.txt
- Documentation/devicetree/bindings/sound/fsl,asrc.txt
- Documentation/devicetree/bindings/sound/wm8962.txt
- Documentation/devicetree/bindings/sound/wm8960.txt
- Documentation/devicetree/bindings/sound/wm8994.txt
- Documentation/devicetree/bindings/sound/cs42xx8.txt
- Documentation/devicetree/bindings/sound/imx-audmux.txt
- Documentation/devicetree/bindings/sound/imx-audio-wm8962.txt
- Documentation/devicetree/bindings/sound/imx-audio-cs42888.txt
- Documentation/devicetree/bindings/sound/imx-audio-si476x.txt

7.1.4.5   Source Code Structure

The following table shows the stereo codec SoC driver source in sound/soc/fsl.

Table 7-1.   Stereo Codec SoC Driver Files

File	Description
sound/soc/fsl/imx-wm8958.c sound/soc/fsl/imx-wm8960.c sound/soc/fsl/imx-wm8962.c	Machine layer for stereo CODEC ALSA SoC (CODEC as I2S Master)
sound/soc/fsl/imx-pcm-dma.c	Platform layer for stereo CODEC ALSA SoC

Table continues on the next page...

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

Table 7-1.   Stereo Codec SoC Driver Files (continued)

File	Description
sound/soc/fsl/imx-pcm.h	Header file for PCM driver and AUDMUX register definitions
sound/soc/fsl/fsl_ssi.c	SSI CPU DAI driver for stereo CODEC ALSA SoC
sound/soc/fsl/fsl_ssi.h	Header file for SSI CPU DAI driver and SSI register definitions
sound/soc/fsl/fsl_sai.c	SAI CPU DAI driver for stereo CODEC ALSA SoC
sound/soc/fsl/fsll_sai.h	Header file for SAI CPU DAI driver and SAI register definitions
codecs/wm8994.c codecs/wm8960.c codecs/wm8962.c	CODEC layer for stereo CODEC ALSA SoC
codecs/wm8994.h codecs/wm8960.h codecs/wm8962.h	Header file for stereo CODEC driver

Table below lists the AM/FM codec SoC driver source files. These files are under `sound/soc`.

Table 7-2.   AM/FM Codec SoC Driver Source Files

File	Description
sound/soc/fsl/imx-si476x.c	Machine layer for stereo CODEC ALSA SoC (CODEC as I2S Slave)
sound/soc/fsl/imx-pcm-dma.c	Platform layer for stereo CODEC ALSA SoC
sound/soc/fsl/imx-pcm.h	Header file for pcm driver and AUDMUX register definitions
sound/soc/fsl/fsl_ssi.c	SSI CPU DAI driver for stereo CODEC ALSA SoC
sound/soc/fsl/fsl_ssi.h	Header file for SSI CPU DAI driver and SSI register definitions
sound/soc/codecs/si476x.c	Codec layer for stereo CODEC ALSA SoC

Table below shows the multiple-channel ADC SoC driver source files.

Table 7-3.   CS42888 ASoC Driver Source File

File	Description
sound/soc/fsl/imx-cs42888.c	Machine layer for multiple-channel CODEC ALSA SoC
sound/soc/fsl/imx-pcm-dma.c	Platform layer for multiple-channel CODEC ALSA SoC
sound/soc/fsl/imx-pcm.h	Header file for pcm driver
sound/soc/fsl/fsl_esai.c	ESAI CPU DAI driver for multiple-channel CODEC ALSA SoC
sound/soc/fsl/fsl_esai.h	Header file for ESAI CPU DAI driver
sound/soc/codecs/cs42xx8.c	CODEC layer for multiple-channel codec ALSA SoC
sound/soc/>codecs/cs42xx8.h	Header file for multiple-channel CODEC driver

Table continues on the next page...

Table 7-3.   CS42888 ASoC Driver Source File (continued)

File	Description
sound/soc/fsl/fsl_asrc.c	CPU DAI driver of ASRC P2P
sound/soc/fsl/fsl_asrc.h	Header file for CPU DAI driver of ASRC P2P
sound/soc/fsl/fsl_asrc_pcm.c	Platform layer for ASRC P2P

7.1.4.6   Menu Configuration Options

The following Linux kernel configuration options are provided for this module.

- SoC Audio supports for WM8958, WM8960, and WM8962 CODEC. In menuconfig, this option is available:

```
-> Device Drivers
    -> Sound card support
        -> Advanced Linux Sound Architecture
            -> ALSA for SoC audio support
                -> SoC Audio for Freescale CPUs
                    -> SoC Audio support for i.MX boards with wm8962 (or
wm8958, wm8960)
```

- SoC Audio supports for i.MX cs42888. In menuconfig, this option is available:

```
-> Device Drivers
    -> Sound card support
        -> Advanced Linux Sound Architecture
            -> ALSA for SoC audio support
                -> SoC Audio for Freescale CPUs
                    -> SoC Audio support for i.MX boards with cs42888
```

- SoC Audio supports for AM/FM. In menuconfig, this option is available:

```
-> Device Drivers
    -> Sound card support
        -> Advanced Linux Sound Architecture
            -> ALSA for SoC audio support
                -> SoC Audio for Freescale CPUs
                    -> SoC Audio support for i.MX boards with si476x
```

7.2   Asynchronous Sample Rate Converter (ASRC)

7.2.1 Introduction

The Asynchronous Sample Rate Converter (ASRC) converts the sampling rate of a signal to a signal of different sampling rate. The ASRC supports concurrent sample rate conversion of up to 10 channels. The sample rate conversion of each channel is associated to a pair of incoming and outgoing sampling rates. The ASRC supports up to three sampling rate pairs simultaneously.

7.2.1.1 Hardware Operation

ASRC includes the following features:

- Supports ratio (Fsin/Fsout) range between 1/24 to 8.
- Designed for rate conversion between 44.1 KHz, 32 KHz, 48 KHz, and 96 KHz.
- Other input sampling rates in the range of 8 KHz to 100 KHz are also supported, but with less performance (see IC spec for more details).
- Other output sampling rates in the range of 30 KHz to 100 KHz are also supported, but with less performance.
- Automatic accommodation to slow variations in the incoming and outgoing sampling rates.
- Tolerant to sample clock jitter.
- Designed mainly for real-time streaming audio usage. Can be used for non-realtime streaming audio usage when the input sampling clocks are not available.
- In any usage case, the output sampling clocks must be activated.
- In case of real-time streaming audio, both input and output clocks need to be available and activated.
- In case of non-realtime streaming audio, the input sampling rate clocks can be avoided by setting ideal-ratio values into ASRC interface registers.

The ASRC supports polling, interrupt and DMA modes, but only DMA mode is used in the platform for better performance. The ASRC supports following DMA channels:

- Peripheral to peripheral, for example: ASRC to ESAI
- Memory to peripheral, for example: memory to ASRC
- Peripheral to memory, for example: ASRC to memory

For more information, see the ASRC chapter in the Applications Processor documentation associated with the SoC.

7.2.2   Software Operation

As an assistant component in the audio system, the ASRC driver implementation depends on the use cases in the platform.

Currently, ASRC is used in two scenarios.

- Memory > ASRC > Memory, ASRC is controlled by the user application or ALSA plug-in.
- Memory > ASRC > peripheral, ASRC is controlled directly by other ALSA drivers.
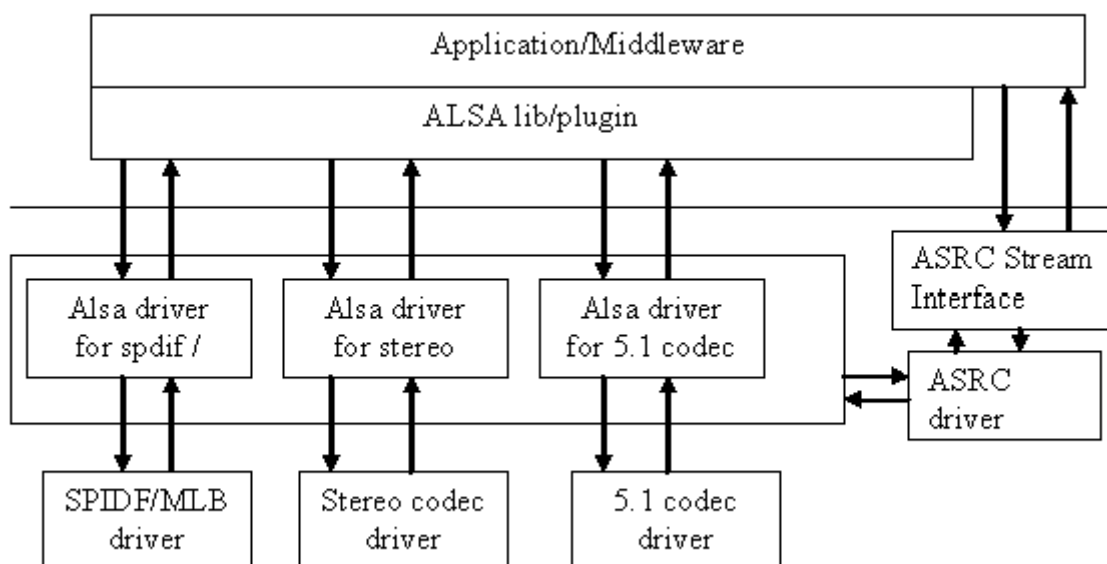
Figure 7-2. Audio Driver Interactions

As illustrated in the figure above, the ASRC stream interface provides the interface for the user space. The ASRC registers itself under /dev/mxc_asrc and creates proc file /proc/driver/asrc when the module is inserted. proc is used to track the channel number for each pair. If all the pairs are not used, users can adjust the channel number through the proc file. The number of the total channels should be ten, or else the adjusted value cannot be saved properly.

7.2.2.1   Sequence for Memory to ASRC to Memory
- Open /dev/mxc_asrc device
- Request ASRC pair - ASRC_REQ_PAIR
- Configure ASRC pair - ASRC_CONIFG_PAIR
- Start ASRC - ASRC_START_CONV
- Write the raw audio data (to be converted) into the user maintained input buffer. Fill asrc_convert_buffer struct with input/output buffer length and address. Driver would

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

copy output data to user maintained output buffer address according to the output buffer size. Repeat this step until all data is converted. -ASRC_CONVERT
* Stop ASRC conversion - ASRC_STOP_CONV
* Release ASRC pair - ASRC_RELEASE_PAIR
* Close /dev/mxc_asrc device

7.2.2.2 Sequence for Memory to ASRC to Peripheral

Memory to ASRC to peripheral audio path is involved in 7.1 audio codec driver. In 7.1 audio sound card, a new device with the name "cs42888audio [cs42888-audio], device 1: HiFi-ASRC-FE (*)" is specified for playback and capture with ASRC. The steps below show the flow of calling ASRC to memory to peripheral:

* The sound device(PCM) has been registered and start to enable the DMA channel in ALSA driver
* Request ASRC pair - fsl_asrc_request_pair
* Configure ASRC pair - fsl_asrc_config_pair
* Enable the DMA channel from Memory to ASRC and from ASRC to Memory
* Start DMA channel and start ASRC conversion - fsl_asrc_start_pair
* When audio data playback complete, stop DMA channel and ASRC - fsl_asrc_stop_pair
* Release ASRC pair - fsl_asrc_release_pair

7.2.2.3 Source Code Structure

The table below lists the source files available in sound/soc/fsl.

Table 7-4.   ASRC Source File List

File	Description
sound/soc/fsl/fsl_asrc_m2m.c	ASRC M2M driver implementation codes
sound/soc/fsl/imx-cs42888.c	Memory to ASRC to ESAI TX implementation in 7.1 audio codec machine driver
sound/soc/fsl/imx-pcm-dma.c	Memory to ASRC to ESAI TX implementation in 7.1 audio codec platform driver
sound/soc/fsl/fsl_esai.c	Memory to ASRC to ESAI TX implementation in 7.1 audio codec CPU driver
sound/soc/fsl/cs42xx8	Memory to ASRC to ESAI TX implementation in 7.1 audio codec codec driver
sound/soc/fsl/fsl_asrc.c	ALSA CPU DAI driver of ASRC P2P
sound/soc/fsl/fsl_asrc.h	Header file for ALSA CPU DAI driver of ASRC P2P
sound/soc/fsl/fsl_asrc_dma.c	ALSA platform layer for ASRC P2P

Table 7-4. ASRC Source File List
(continued)

File	Description
sound/soc/fsl/sound/soc/fsl/fsl_asrc_dma.c	ALSA platform layer for ASRC M2M

7.2.2.4 Menu Configuration Options

The menu configuration options are as follows:

```
-> Device Drivers
   -> Sound card support
      -> Advanced Linux Sound Architecture
         -> ALSA for SoC audio support
            -> SoC Audio for Freescale i.MX CPUs
               -> Asynchronous Sample Rate Converter (ASRC) module support
```

Then the ASRC driver can only be configured with the build-in module.

7.2.2.5 Device Tree Binding

The functions of device tree bindings for ASRC M2M are as follows:

- compatible: Compatible list, must contain "fsl,imx6q-asrc".
- reg: Offset and length of the register set for the device.
- interrupts: Contains the asrc interrupt.
- clocks: Contains an entry for each entry in clock-names.
- clock-names: Must contain "mem", "ipg", "asrck", and "dma". (Generally, "dma" is used for SPBA clock.)
- dmas: Generic dma devicetree binding as described in Documentation/devicetree/bindings/dma/dma.txt.
- dma-names: Six dmas have to be defined, "txa", "rxa", "txb", "rxb", "txc", "rxc".
- fsl,clk-map-version: the mapping relationship in different SOC is different. This version number can be used to indicate clock map information.
- fsl,clk-channel-bits: indicates the channel bit information.

The functions of device tree bindings for ASRC P2P are as follows:

- compatible: Compatible list, must contain "fsl,imx6q-asrc-p2p".
- fsl,p2p-rate: A valid sample rate for Back-End (I2S) playback and record.
- fsl,p2p-width: A valid sample width for Back-End (I2S) playback and record.
- fsl,asrc-dma-rx-events: Contains three SDMA event numbers for ASRC Rx.
- fsl,asrc-dma-tx-events: Contains three SDMA event numbers for ASRC Tx.

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

7.2.2.6　Programming Interface (Exported API and IOCTLs)

The ASRC Exported API allows the ALSA driver to use ASRC services.

The ASRC IOCTLs below are used for user space applications:

ASRC_REQ_PAIR:

Apply a pair from ASRC driver. Once a pair is allocated, ASRC core clock is enabled.

ASRC_CONFIG_PAIR:

Configure ASRC pair allocated. User is responsible for providing parameters defined in struct asrc_config. Items in asrc_config is listed below:
- pair: ASRC pair allocated by the IOCTL(ASRC_REQ_PAIR).
- channel_num: channel number.
- buffer_num: buffer number need for input and output buffer use.The input/output buffers are allocated inside ASRC driver. User is responsible for remap it into user space.
- dma_buffer_size: buffer size for input and output buffers. The buffer size should be in the unit of page size. Usually, 4k bytes is used.
- input_sample_rate: input sampling rate. Input sample rate should be in 5.512k, 8k, 11.025k, 16k, 22k, 32k, 44.1k, 48k, 64k, 88.2k 96k, 176.4k, 192k.
- output_sample_rate: output sampling rate. Output sampling rate should be in 32k, 44.1k, 48k, 64k, 88.2k, 96k, 176.4k 192k.
- input_word_width: word width of input audio data. The input data word width can be 16 bit or 24 bit.
- output_word_width: word width of output audio data. The output data word width can be 16 bit or 24 bit.
- inclk: the input clock source can be ESAI RX clock, SSI1 RX clock, SSI2 RX clock, SPDIF RX clock, MLB_clock, ESAI TX clock, SSI1 TX clock, SSI2 TX clock, SPDIF TX clock, ASRCLK1 clock, NONE. If using clock except NONE, user should make sure that the clock is available.
- outclk: the output clock source is the same as the input clock source.

ASRC_CONVERT:

Convert the input data into output data according to the parameters set by ASRC_CONFIG_PAIR. Driver would copy input_buffer_length bytes data from the input_buffer_vaddr for convert. After convert, driver fill the output_buffer_length according to data number generated by ASRC and copy output_buffer_length to output_buffer_vaddr. However, before calling ASRC_CONVERT, User is responsible for filling the output_buffer_length according to the ratio of input sample rate and output

sample rate. If the generated buffer size is larger than user filled output_buffer_size, driver would only copy user filled output_buffer_size to output_buffer_vaddr. If the generated buffer size is smaller than user filled output_buffer_size(the difference should be less than 64 bytes.), calling ASRC_CONVERT would fail.

- input_buffer_vaddr: virtual address of input buffer.
- output_buffer_vaddr: virtual address of output buffer.
- input_buffer_length: length of input buffer(bytes).
- output_buffer_length: length of output buffer(bytes).

ASRC_START_CONV:

Start ASRC pair convert.

ASRC_STOP_CONV:

Stop ASRC pair convert.

ASRC_STATUS:

Query ASRC pair status.

7.3 HDMI Audio

7.3.1 Introduction

HDMI Audio is covered in the HDMI overview chapter in video. See HDMI Audio for more details.

7.4 The Sony/Philips Digital Interface (S/PDIF)

7.4.1 Introduction

The Sony/Philips Digital Interface (S/PDIF) audio module is a stereo transceiver that allows the processor to receive and transmit digital audio. The S/PDIF transceiver allows the handling of both S/PDIF channel status (CS) and User (U) data. The frequency measurement block allows the S/PDIF RX section to derive the receive clock from the incoming S/PDIF stream.

7.4.1.1 S/PDIF Overview

Figure below shows the block diagram of the S/PDIF interface.
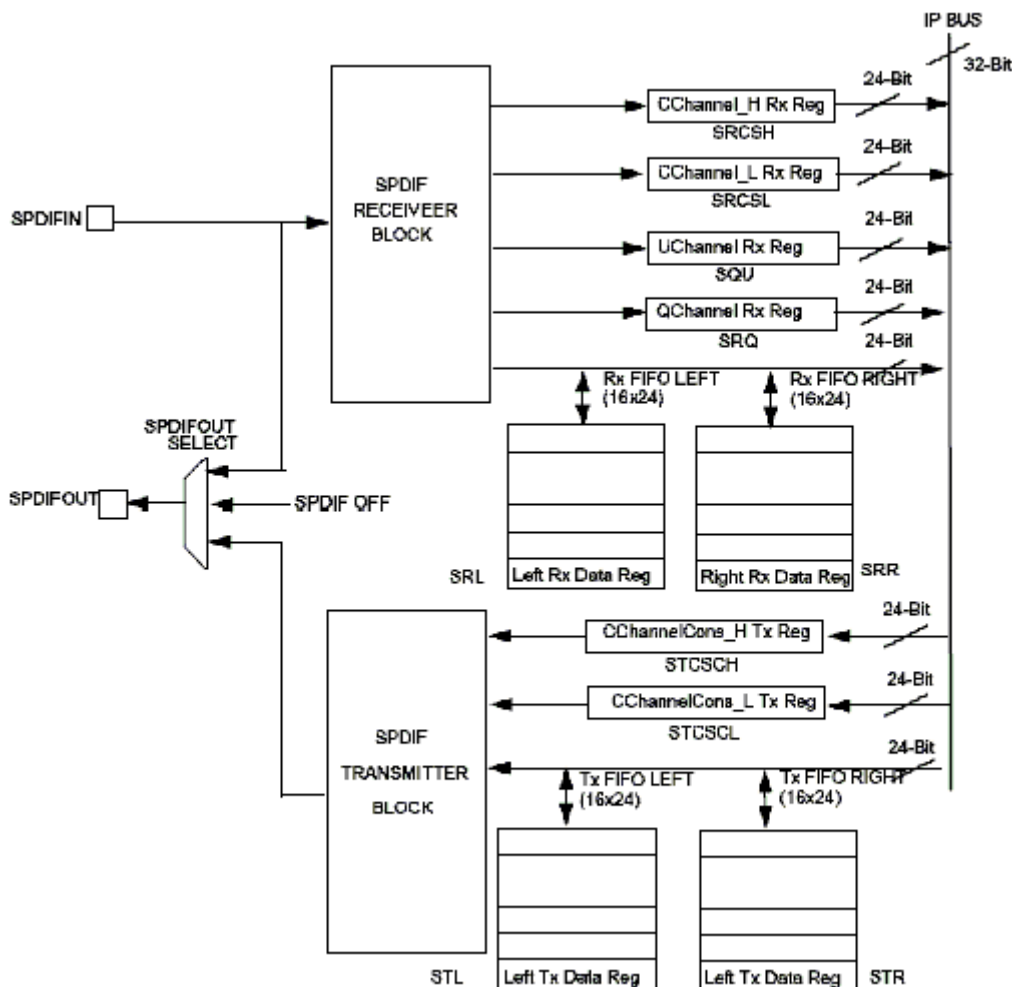
Figure 7-3. S/PDIF Transceiver Data Interface Block Diagram

7.4.1.2 Hardware Overview

The S/PDIF is composed of two parts:

- The S/PDIF receiver extracts the audio data from each S/PDIF frame and places the data in the S/PDIF Rx left and right FIFOs. The Channel Status and User Bits are also extracted from each frame and placed in the corresponding registers. The S/PDIF receiver provides a bypass option for direct transfer of the S/PDIF input signal to the S/PDIF transmitter.
- For the S/PDIF transmitter, the audio data is provided by the processor through the SPDIFTxLeft and SPDIFTxRight registers. The Channel Status bits are provided

through the corresponding registers. The S/PDIF transmitter generates a S/PDIF output bitstream in the biphase mark format (IEC958), which consists of audio data, channel status and user bits.

In the S/PDIF transmitter, the IEC958 biphase bit stream is generated on both edges of the S/PDIF Transmit clock. The S/PDIF Transmit clock is generated by the S/PDIF internal clock dividers and the sources are from outside of the S/PDIF block. The S/PDIF receiver can recover the S/PDIF Rx clock from the S/PDIF stream. Figure 7-3 shows the clock structure of the S/PDIF transceiver.

7.4.1.3 Software Overview

The S/PDIF driver is designed under ALSA System on Chip (ASoC) layer. The ASoC driver for S/PDIF provides one playback device for Tx and one capture device for Rx. The playback output audio format can be linear PCM data or compressed data with 16-bit, 20-bit, and 24-bit audio. The allowed sampling bit rates are 44.1, 48 or 32 KHz. The capture input audio format can be linear PCM data or compressed 24-bit data and the allowed sampling bit rates are from 16 to 96 KHz. The driver provides the same interface for PCM and compressed data transmission.

7.4.1.4 The ASoC layer

The ASoC layer divides audio drivers for embedded platforms into separate layers that can be reused. ASoC divides an audio driver into a codec driver, a machine layer, a DAI (digital audio interface) layer, and a platform layer. The Linux kernel documentation has some concise description of these layers in linux/Documentation/sound/alsa/soc. In the case of the S/PDIF driver, we are able to reuse the platform layer (imx-pcm-dma.c) that is used by the ssi stereo codec driver and also the generic dummy codec driver useful for DAI links creation without a real codec.

7.4.2 S/PDIF Tx Driver

The S/PDIF Tx driver supports the following features.

- 32, 44.1 and 48 KHz sample rates.

- Signed 16 and 24-bit little Endian sample format. Due to S/PDIF SDMA feature, the 24-bit output sample file must have 32-bits in each channel per frame. Only the 24 LSBs are valid.
- In the ALSA subsystem, the supported format is defined as S16_LE and S24_LE.

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

- Stereo playback.
- Information query through iecset or amixer.
- The device ID can be determined by using the 'aplay -l' utility to list out the playback audio devices.

For example:

```
root@freescale ~$ aplay -l

**** List of PLAYBACK Hardware Devices ****

card 0: imxspdif [imx-spdif], device 0: S/PDIF PCM snd-soc-dummy-dai-0 []

  Subdevices: 1/1

  Subdevice #0: subdevice #0
```

NOTE
The number at the beginning of the IMX_SPDIF line is the card ID. The string in the square brackets is the card name.

- The ALSA utility provides a common method for user spaces to operate and use ALSA drivers

```
#aplay -Dplughw:0,0 audio.wav
```

NOTE
The -D parameter of aplay indicates the PCM device with card ID and PCM device ID: hw:[card id],[pcm device id]

The "iecset" utility provides a common method to set or dump the IEC958 status bits.

```
#iecset -c 0
```

7.4.2.1 Driver Design

Before S/PDIF playback, the configuration, interrupt, clock and channel registers are initialized. During S/PDIF playback, the channel status bits are fixed. The DMA and interrupts are enabled. S/PDIF has 16 TX sample FIFOs on Left and Right channel respectively. When both FIFOs are empty, an empty interrupt is generated if the empty interrupt is enabled. If no data are refilled in the 20.8 µs (1/48000), an underrun interrupt is generated. Overrun is avoided if only 16 sample FIFOs are filled for each channel every time. If auto re-synchronization is enabled, the hardware checks if the left and right FIFO are in sync, and if not, it sets the filling pointer of the right FIFO to be equal to the filling pointer of the left FIFO and an interrupt is generated.

7.4.2.2   Provided User Interface

The S/PDIF transmitter driver provides one ALSA mixer sound control interface to the user besides the common PCM operations interface. It provides the interface for the user to write S/PDIF channel status codes into the driver so they can be sent in the S/PDIF stream. The input parameter of this interface is the IEC958 digital audio structure shown below, and only status member is used:

```
struct snd_aes_iec958 {
                unsigned char status[24];       /* AES/IEC958 channel status bits */
                unsigned char subcode[147];     /* AES/IEC958 subcode bits */
                unsigned char pad;              /* nothing */
                unsigned char dig_subframe[4];  /* AES/IEC958 subframe bits */
};
```

7.4.3   S/PDIF Rx Driver

The S/PDIF Rx driver supports the following features:

- 16, 32, 44.1, 48, 64 and 96 KHz receiving sample rate
- Signed 24-bit little endian sample format. Due to S/PDIF SDMA feature, each channel bit length in PCM recorded frame is 32 bits, and only the 24 LSBs are valid

    In ALSA subsystem, the supported format is defined as S24_LE.

- Stereo record.
- The device ID can be determined by using the 'arecord -l' to list out record devices.

    For example:

    ```
    root@freescale ~$ arecord -l

    **** List of CAPTURE Hardware Devices ****

    card 0: cs42888audio [cs42888-audio], device 0: HiFi CS42888-0 []

      Subdevices: 1/1

      Subdevice #0: subdevice #0

    card 1: imxspdif [imx-spdif], device 0: S/PDIF PCM snd-soc-dummy-dai-0 []

      Subdevices: 1/1

      Subdevice #0: subdevice #0
    ```

- The ALSA utility provides a common method for user spaces to operate and use ALSA drivers.

    ```
    #arecord -Dplughw:1,0" -c 2 -r 44100 -f S24_LE record.wav
    ```

NOTE

The -D parameter of the arecord indicates the PCM device with card ID and PCM device ID: hw:[card id],[pcm device id]

The "iecset" utility provides a common method to set or dump the IEC958 status bits.

```
#iecset -c 1
```

7.4.3.1 Driver Design

Before the driver can read a data frame from the S/PDIF receiver FIFO, it must wait for the internal DPLL to be locked. Using the high-speed system clock, the internal DPLL can extract the bit clock (advanced pulse) from the input bit stream. When this internal DPLL is locked, the LOCK bit of PhaseConfig Register is set and the driver configures the interrupt, clock and SDMA channel. After that, the driver can receive audio data, channel status, user bits and valid bits concurrently.

For channel status reception, a total of 48 channel status bits are received in two registers. The driver reads them out when a user application makes a request.

For user bits reception, there are two modes for User Channel reception: CD and non-CD. The mode is determined by the USyncMode (bit 1 of CDText_Control register). User can call the sound control interface to set the mode (see Table 7-5), but no matter what the mode is, the driver handles the user bits in the same way. For the S/PDIF Rx, the hardware block copies the Q bits from the user bits to the QChannel registers and puts the user bits in UChannel registers. The driver allocates two queue buffers for both U bits and Q bits. The U bits queue buffer is 96x2 bytes in size, the Q bits queue buffer is 12x2 bytes in size, and queue buffers are filled in the U/Q Full, Err and Sync interrupt handlers. This means that the user can get the previous ready U/Q bits while S/PDIF driver is reading new U/Q bits.

For valid bit reception, S/PDIF Rx hardware block triggers an interrupt and set interrupt status upon reception. A sound control interface is provided for the user to get the status of this valid bit.

7.4.3.2 Provided User Interface

The S/PDIF Rx driver provides interfaces for user application as shown in table below.

Table 7-5.   S/PDIF Rx Driver Interfaces

Interface	Type	Mode[1]	Parameter	Comment
Common PCM	PCM	-	-	PCM open/close prepare/trigger hw_params/sw_params
Rx Sample Rate	Sound Control[2]	r	Integer Range: [16000, 96000]	Get sample rate. It is not accurate due to DPLL frequency measure module. So the user application must do a correction to the get value.
USyncMode	Sound Control	rw	Boolean Value: 0 or 1	Set 1 for CD mode Set 0 for non-CD mode
Channel Status	Sound Control	r	struct snd_aes_iec958 Only status [6] array member is used	-
User bit	Sound Control	r	Byte array 96 bytes for U bits 12 bytes for Q bits	-
No good V bit	Sound Control	r	Boolean Value: 0 or 1	An interrupt is associated with the valid flag. (interrupt 16 - SPDIFValNoGood). This interrupt is set every time a frame is seen on the SPDIF interface with the valid bit set to invalid.

1. The mode column shows the interface attribute: r (read) or w (write)
2. The sound control type of interface is called by the snd_ctl_xxx() alsa-lib function

The user application can follow the program flow from Figure 7-4 to use the S/PDIF Rx driver. First, the application opens the S/PDIF Rx PCM device, waits for the DPLL to lock the input bit stream, and gets the input sample rate. If the USyncMode needs to be set, set it before reading the U/Q bits. Next, set the hardware parameters, including channel number, format and capture sample rate which is obtained from the driver. Then, call prepare and trigger to startup S/PDIF Rx stream read. Finally, call the read function to get the data. During the reading process, applications can read the U/Q bits and channel status from the driver and valid the no good bit.
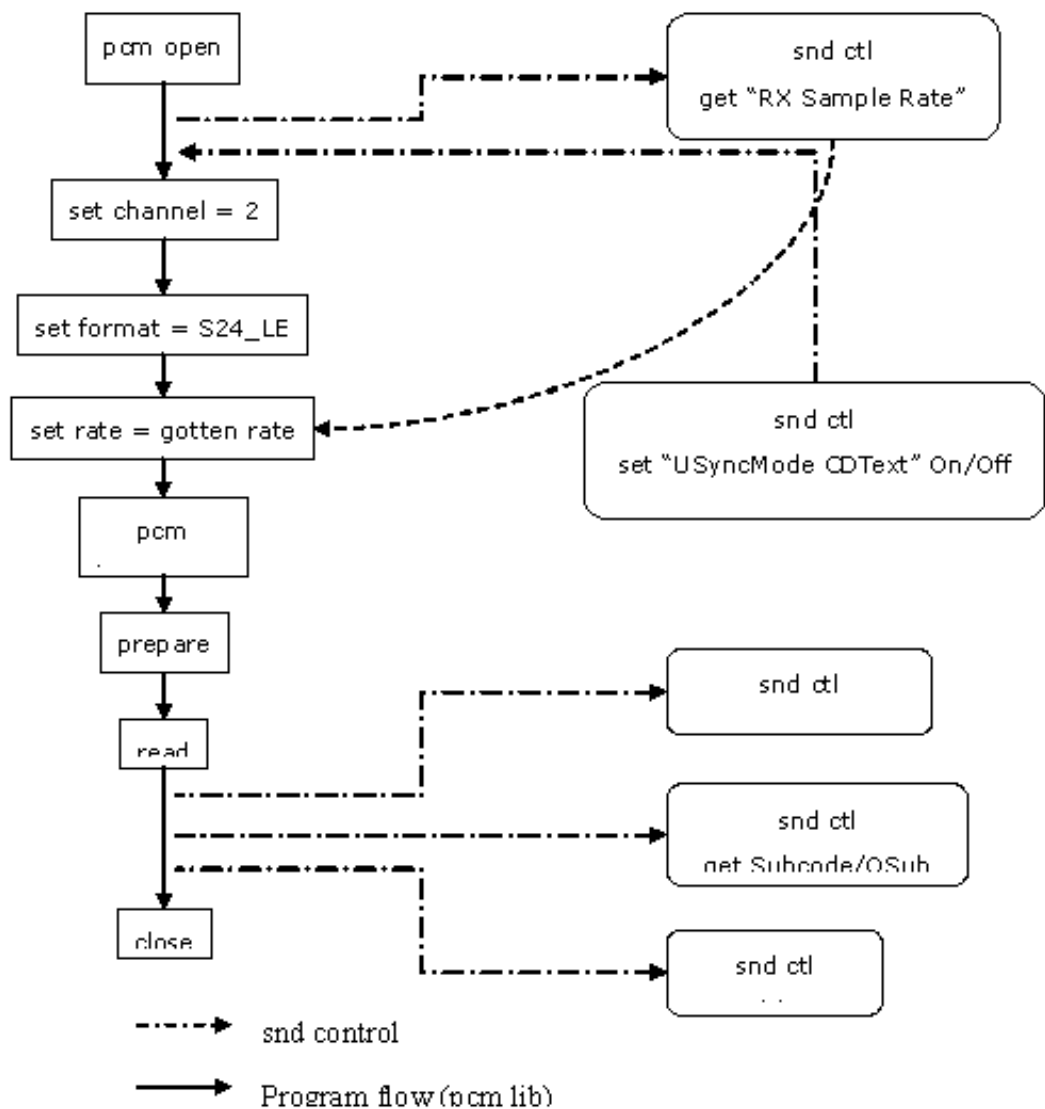
Figure 7-4. S/PDIF Rx Application Program Flow

7.4.4 Source Code Structure

Table below lists the source files for the driver.

Table 7-6. S/PDIF Driver Files

File	Description
sound/soc/soc-utils.c	Dummy ALSA SOC codec driver
sound/soc/fsl/imx-spdif.c	S/PDIF ALSA SOC machine layer
sound/soc/fsl/fsl_spdif.c	S/PDIF ALSA SOC DAI layer
sound/soc/fsl/imx-pcm-dma.c	ALSA SOC platform layer
sound/soc/fsl/imx-pcm.h	ALSA SOC platform layer header

7.4.4.1 Menu Configuration Options

The following Linux kernel configurations are provided for this module:

In menu configuration enable the following module:

- CONFIG_SND_IMX_SPDIF - Configuration option for the S/PDIF driver:
- Device Drivers -> Sound card support -> Advanced Linux Sound Architecture -> ALSA for SoC audio support -> SoC Audio for Freescale i.MX CPUs -> SoC Audio support for i.MX boards with S/PDIF

7.4.4.2 Device Tree Bindings

See the following documents:

- Documentation/devicetree/bindings/sound/fsl,spdif.txt
- Documentation/devicetree/bindings/sound/imx-audio-spdif.txt

7.4.4.3 Interrupts and Exceptions

S/PDIF Tx/Rx hardware block has many interrupts to indicate the success, exception and event.

The driver handles the following interrupts:

- DPLL Lock and Loss Lock-Saves the DPLL lock status; this is used when getting the Rx sample rate
- U/Q Channel Full and overrun/underrun-Puts the U/Q channel register data into queue buffer, and update the queue buffer write pointer
- U/Q Channel Sync-Saves the ID of the buffer whose U/Q data is ready for read out
- U/Q Channel Error-Resets the U/Q queue buffer

7.4.5 Unit Test Preparation

In order to prepare to run a unit test, perform the following actions:

- Setup M-Audio Transit USB sound card by installing M-Audio Transit driver on your PC.
- Install WaveLab tools on your PC.

7.4.5.1   Tx test step

- Plug optical line into [line|optical] port of M-Audio transit.

NOTE
Make sure the [optical out] port of M-Audio transit has no
output (red light off) after plugging the optical line.

- Start up WaveLab, press the record button on the toolbar, set up the record file name, sample rate, and channel number, and then record.
- Meanwhile, on board use following command to play one wave file:

```
#aplay -D hw:[card id],[pcm id] audioXXkYYS.wav
```

- After aplay finishes, stop recording in WaveLab.
- Play the recorded WAV file in wavelab to check.

7.4.5.2   Rx test step

- Plug optical line into [optical port] of M-Audio transit
- Startup WaveLab, open a test WAV file: audioXXkYYS.wav to play in loop
- Meanwhile, on board use the following command to record one WAV file. After finishing recording, you may playback the record WAV file on other audio card on the board or PC.

```
#arecord -D hw:[card id],[pcm id] -c 2 -d 20 -r [sample rate in Hz] -f S24_LE record.wav
```

NOTE
The sample rate argument in the arecord command must be
consistent with the WAV file playing on WaveLab.

Chapter 8
Security

8.1 Cryptographic Acceleration and Assurance Module (CAAM)

8.1.1 CAAM Device Driver Overview

This section discusses implementation specifics of the kernel driver components supporting CAAM (Cryptographic Acceleration and Assurance Module) within the Linux kernel.

CAAM's base driver packaging can be categorized on two distinct levels:

- Configuration and Job Execution Level
- API Interface Level

Configuration and Job Execution Level consists of:

- A control and configuration module which maps the main register page and writes global or system required configuration information.
- A module that feeds jobs through job rings, and reports status.

API Interface Level consists of:

- An interface to the Sctterlist Crypto API supporting asynchronous single-pass authentication-encryption operations, and common blockciphers - `caamalg`.
- An interface to the Scatterlist Crypto API supporting asynchronous hashes - `caamhash`.
- An interface to the hwrng API supporting use of the Random Number Generator - `caamrng`.

8.1.2 Configuration and Job Execution Level

This section has two parts:

- Control/Configuration Driver
- Job Ring Driver

8.1.3 Control/Configuration Driver

The control and configuration driver is responsible for initializing and setting up the master register page, initializing early-on feature initialization, providing limited debug and monitoring capability, and generally ensuring that all other dependent driver subsystems can connect to a correctly-configured device.

Step by step, it performs the following actions at startup:

- Allocates a private storage block for this level.
- Maps a virtual address to the full CAAM register page.
- Maps a virtual address for the SNVS register page.
- Maps a virtual (cache coherent) address for Secure Memory.
- Registers the security violation interrupt.
- Selects the correct DMA address size for the platform, and sets DMA address masks to match.
- Identifies other pertinent interrupt connections.
- Initializes all job ring instances.
- If the system configuration includes a DPAA Queue Interface, that interface has frame-pop enabled.

NOTE

i.MX 6 configurations do not contain this logic.

- If the instance contains a TRNG, it's oscillator/entropy configuration is set and then "kickstarted".
- Configuration information is sent to the system console to indicate that the driver is alive, and what configuration it has assumed.
- If CONFIG_DEBUG_FS is selected in the kernel configuration, then entries are added to enable debugfs views to useful registers in the performance monitor. Register views are accessible under the caam/ctl directory at the debugfs root entry.

8.1.4 Job Ring Driver

The Job Ring driver is responsible for providing job execution service to higher-level drivers. It takes care of overall management of both input and output rings and interrupt service driving the output ring.

One driver call is available for higher layers to use for queueing jobs to a ring for execution:

```
int caam_jr_enqueue(struct device *dev, u32 *desc, void (*cbk)(struct device
*dev, u32 *desc, u32 status, void *areq), void *areq);
```

Arguments:

`dev` Pointer to the struct device associated with the job ring for use. In the current configuration, one or more struct device entries exist in the controller's private data block, one for each ring.

`desc` Pointer to a CAAM job descriptor to be executed. The driver will map the descriptor prior to execution, and unmap it upon completion. However, since the driver can't reasonably know anything about the data referenced by the descriptor, it is the caller's responsibility to map/flush any of this data prior to submission, and to unmap/invalidate data after the request completes.

`cbk` Pointer to a callback function that will be called when the job has completed processing.

`areq` Pointer to metadata or context data associated with this request. Often, this can contain referenced data mapping information that request postprocessing (via the callback) can use to clean up or release resources once complete.

Callback Function Arguments:

`dev` Pointer to the struct device associated with the job ring for use.

`desc` Pointer to the original descriptor submitted for execution.

`status` Completion status received back from the CAAM DECO that executed the request. Nonzero only if an error occurred. Strings describing each error are enumerated in error.c.

`areq` Metadata/context pointer passed to the original request.

Returns:

- Zero on successful job submission
- -EBUSY if the input ring was full
- -EIO if driver could not map the job descriptor

8.1.5   API Interface Level

CAAM module provides a connection through the Scatterlist Crypto API both for common symmetric blockciphers, and for single-pass authentication-encryption services. This table lists all installed authentication-encryption algorithms by their common name, driver name, and purpose. Note that certain platforms, such as i.MX 6, contain a low-power MDHA accelerator, which cannot support SHA384 or SHA512.

Name	Driver Name	Purpose
authenc(hmac(md5),cbc(aes))	authenc-hmac-md5-cbc-aes-caam	Single-pass authentication/encryption using MD5 and AES-CBC
authenc(hmac(sha1),cbc(aes))	authenc-hmac-sha1-cbc-aes-caam	Single-pass authentication/encryption using SHA1 and AES-CBC
authenc(hmac(sha224),cbc(aes))	authenc-hmac-sha224-cbc-aes-caam	Single-pass authentication/encryption using SHA224 and AES-CBC
authenc(hmac(sha256),cbc(aes))	authenc-hmac-sha256-cbc-aes-caam	Single-pass authentication/ encryptionusing SHA256 and AES-CBC
authenc(hmac(sha384),cbc(aes))	authenc-hmac-sha384-cbc-aes-caam	Single-pass authentication/encryption using SHA384 and AES-CBC
authenc(hmac(sha512),cbc(aes))	authenc-hmac-sha512-cbc-aes-caam	Single-pass authentication/encryption using SHA512 and AES-CBC
authenc(hmac(md5),cbc(des3_ede))	authenc-hmac-md5-cbcdes3_ede-caam	Single-pass authentication/encryption using MD5 and Triple-DES-CBC
authenc(hmac(sha1),cbc(des3_ede))	authenc-hmac-sha1-cbc-des3_ede-caam	Single-pass authentication/encryption using SHA1 and Triple-DES-CBC
authenc(hmac(sha224),cbc(des3_ede))	authenc-hmac-sha224-cbc-des3_ede-caam	Single-pass authentication/encryption using SHA224 and Triple-DES-CBC
authenc(hmac(sha256),cbc(des3_ede))	authenc-hmac-sha256-cbc-des3_ede-caam	Single-pass authentication/encryption using SHA256 and Triple-DES-CBC
authenc(hmac(sha384),cbc(des3_ede))	authenc-hmac-sha384-cbc-des3_ede-caam	Single-pass authentication/encryption using SHA384 and Triple-DES-CBC
authenc(hmac(sha512),cbc(des3_ede))	authenc-hmac-sha512-cbc-des3_ede-caam	Single-pass authentication/encryption using SHA512 and Triple-DES-CBC
authenc(hmac(md5),cbc(des))	authenc-hmac-md5-cbc-des-caam	Single-pass authentication/encryption using MD5 and Single-DES-CBC
authenc(hmac(sha1),cbc(des))	authenc-hmac-sha1-cbc-des-caam	Single-pass authentication/encryption using SHA1 and Single-DES-CBC
authenc(hmac(sha224),cbc(des))	authenc-hmac-sha224-cbc-des-caam	Single-pass authentication/encryption using SHA224 and Single-DES-CBC
authenc(hmac(sha256),cbc(des))	authenc-hmac-sha256-cbc-des-caam	Single-pass authentication/encryption using SHA256 and Single-DES-CBC
authenc(hmac(sha384),cbc(des))	authenc-hmac-sha384-cbc-des-caam	Single-pass authentication/encryption using SHA384 and Single-DES-CBC
authenc(hmac(sha512),cbc(des))	authenc-hmac-sha512-cbc-des-caam	Single-pass authentication/encryption using SHA512 and Single-DES-CBC

This table lists all installed symmetric key blockcipher algorithms by their common name, driver name, and purpose.

Name	Driver Name	Purpose
cbc(aes)	cbc-aes-caam	AES with a CBC mode wrapper
cbc(des3_ede)	cbc-3des-caam	Triple DES with a CBC mode wrapper
cbc(des)	cbc-des-caam	Single DES with a CBC mode wrapper
ecb(aes)	ecb-aes-caam	AES with a ECB mode wrapper
ecb(des3_ede)	ecb-3des-caam	Triples DES with a ECB mode wrapper
ecb(des)	ecb-des-caam	Single DES with a ECB mode wrapper
ecb(arc4)	ecb-arc4-caam	ARC4 with a ECB mode wrapper
ctr(aes)	ctr-aes-caam	AES with a CTR mode wrapper

Use of these services through the API is exemplified in the common conformance/ performance testing module in the kernel's crypto subsystem, known as tcrypt, visible in the kernel source tree at crypto/tcrypt.c.

The `caamhashmodule` provides a connection through the Scatterlist Crypto API both for common asynchronous hashes.

This table lists all installed asynchronous hashes by their common name, driver name, and purpose. Note that certain platforms, such as i.MX 6, contain a low-power MDHA accelerator, which cannot support SHA384 or SHA512.

Name	Driver Name	Purpose
sha1	sha1-caam	SHA1-160 Hash Computation
sha224	sha224-caam	SHA224 Hash Computation
sha256	sha256-caam	SHA256 Hash Computation
sha384	sha384-caam	SHA384 Hash Computation
sha512	sha512-caam	SHA512 Hash Computation
md5	md5-caam	MD5 Hash Computation
hmac(sha1)	hmac-sha1-caam	SHA1-160 Hash-based Message Authentication Code
hmac(sha224)	hmac-sha224-caam	SHA224 Hash-based Message Authentication Code
hmac(sha256)	hmac-sha256-caam	SHA256 Hash-based Message Authentication Code
hmac(sha384)	hmac-sha384-caam	SHA384 Hash-based Message Authentication Code
hmac(sha512)	hmac-sha512-caam	SHA512 Hash-based Message Authentication Code
hmac(md5)	hmac-md5-caam	MD5 Hash-based Message Authentication Code

Use of these services through the API is exemplified in the common conformance/ performance testing module in the kernel's crypto subsystem, known as tcrypt, visible in the kernel source tree at crypto/tcrypt.c.

The `caamrng` module installs a mechanism to use CAAM's random number generator to feed random data into a pair of buffers that can be accessed through `/dev/hw_random`.

`/dev/hw_random` is commonly used to feed the kernel's own entropy pool, which can be used internally, as an entropy source for other random data "devices".

For more information regarding support for this service, see `rng-tools` available in sourceforge.net/projects/gkernel/files/rng-tools.

8.1.6 Driver Configuration

Configuration of the driver is controlled by the following kernel confguration parameters (found under Cryptographic API -> Hardware Crypto Devices):

CRYPTO_DEV_FSL_CAAM

Enables building the base controller driver and the job ring backend.

CRYPTO_DEV_FSL_CAAM_RINGSIZE

Selects the size (e.g., the maximum number of entries) of job rings. This is selectable as a power of 2 in the range of 2-9, allowing selection of a ring depth ranging from 4 to 512 entries.

The default selection is 9, resulting in a ring depth of 512 job entries.

CRYPTO_DEV_FSL_CAAM_INTC

Enables the use of the hardware's interrupt coalescing feature, which can reduce the amount of interrupt overhead the system incurs during periods of high utilization. Leaving this disabled forces a single interrupt for each job completion, simplifying operation, but increasing overhead.

CRYPTO_DEV_FSL_CAAM_INTC_COUNT_THLD

If coalescing is enabled, selects the number of job completions allowed to queue before an interrupt is raised. This is selectable within the range of 1 to 255. Selecting 1 effectively defeats the coalescing feature. Any selection of a size greater than the job ring size forces a situation where the interrupt times out before ever raising an interrupt.

The default selection is 255.

CRYPTO_DEV_FSL_CAAM_INTC_TIME_THLD

If coalescing is enabled, selects the count of bus clocks (divided by 64) before a coalescing timeout where, if the count threshold has not been met, an interrupt is raised at the end of the time period. The selection range is an integer from 1 to 65535.

The default selection is 2048.

CRYPTO_DEV_FSL_CAAM_CRYPTO_API

Enables Scatterlist Crypto API support for asynchronous blockciphers and for single-pass autentication-encryption operations through the API using CAAM hardware for acceleration.

CRYPTO_DEV_FSL_CAAM_AHASH_API

Enables Scatterlist Crypto API support for asynchronous hashing through the API using CAAM hardware for acceleration.

CRYPTO_DEV_FSL_CAAM_RNG_API

Enables use of the CAAM Random Number generator through the hwrng API. This can be used to generate random data to feed an entropy pool for the kernels pseudo-random number generator.

CRYPTO_DEV_FSL_CAAM_RNG_TEST

Enables a captive test to ensure that the CAAM RNG driver is operating and buffering random data.

8.1.7 Limitations

- Components of the driver do not currently build and run as modules. This may be rectified in a future version.
- Interdependencies exist between the controller and job ring backends, therefore they all must run in the same system partition. Future versions of the driver may separate out the job ring back-end as a standalone module that can run independently (and support independent API and SM instances) in its own system partition.
- The full CAAM register page is mapped by the controller driver, and derived pointers to selected subsystems are calculated and passed to higher-layer driver components. Partition-independent configurations will have to map their own subsystem pointers instead.
- Upstream variants of this driver support only Power architecture. This Arm architecture-specific port is not upstreamed at this time, although portions may be upstreamed at some point.

- TRNG kickstart may need to be moved to the bootloader in a future release, so that the RNG can be used earlier.
- The Job Ring driver has a registration and de-registration functions that are not currently necessary (and may be rewritten in future editions to provide for shutdown notifications to higher layers.
- The full CAAM function is exclusive with the Mega/Fast mix off feature in DSM. If CAAM is enabled, the Mega/Fast mix off feature needs to be disabled, and the user should "echo enabled > /sys/bus/platform/devices/2100000.aips-bus/2100000.caam/2101000.jr0/power/wakeup" after the kernel boots up, and then Mega/Fast mix will keep the power on in DSM.

8.1.8 Limitations in the Existing Implementation Overview

This chapter describes a prototype of a Keystore Management Interface intended to provide access to CAAM Secure Memory.

Secure memory provides a controlled and access-protected area where critical system security parameters can be stored and processed in a running system without bus-level exposure of clear secrets. Secrets can be imported into and exported from secure memory, but never exported from secure memory in their cleartext form. Instead, secrets may be exported from secure memory in a covered form, using keys never visible to the outside.

This driver, with its kernel-level API, exposes a basic interface to allow kernel-level services access to secure memory functionality. It is split into two pieces:

- Keystore Initialization and Maintenance Interfaces
- Keystore Access Interface

The initialization and maintenance services exist to initialize and define the instance of a keystore interface. Likewise, the access interface allows kernel-level services to use the API for management of security parameters.

8.1.9 Initialize Keystore Management Interface

Installs a set of pointers to functions that implement an underlying physical interface to the keystore subsystem.

In the present release, a default (and hidden) suite of functions implement this interface. Future implementations of this API may provide for the installation of an alternate interface. If this occurs, an alternate to this call can be provided.

```
void sm_init_keystore(struct device *dev);
```

Arguments:

`dev` points to a `struct device` established to manage resources for the secure memory subsystem.

8.1.10 Detect Available Secure Memory Storage Units

Returns the number of available units ("pages") that can be accessed by the local instance of this driver. Intended for use as a resource probe.

```
u32 sm_detect_keystore_units(struct device *dev);
```

Arguments:

`dev` Points to a `struct device` established to manage resources for the secure memory subsystem.

Returns: Number of detected units available for use, 0 through n - 1 may be used with subsequent calls to all other API functions.

8.1.11 Establish Keystore in Detected Unit

Sets up an allocation table in a detected unit that can be used for the storage of keys (or other secrets). The unit will be divided into a series of fixed-size slots, each one of which is marked available in the allocation table. The size of each slot is a build-time selectable parameter.

No calls to the keystore access interface can occur until `sm_establish_keystore()` has been called.

`sm_establish_keystore()` should follow a call to `sm_detect_keystore_units()`.

```
int sm_establish_keystore(struct device *dev, u32 unit);
```

Arguments:

`dev` Points to a struct device established to manage resources for the secure memory subsystem.

`unit` One of the units detected with a call to sm_detect_keystore_units().

Returns:

- Zero on successful return
- -EINVAL if the keystore subsystem was not initialized
- -ENOSPC if no memory was available for the allocation table and associated context data.

8.1.12   Release Keystore

Releases all resources used by this keystore unit. No further calls to the keystore access interface can be made.

```
void sm_release_keystore(struct device *dev, u32 unit);
```

Arguments:

`dev` Points to a struct device established to manage resources for the secure memory subsystem.

`unit` One of the units detected with a call to sm_detect_keystore_units().

8.1.13   Allocate a Slot from the Keystore

Allocate a slot from the keystore for use in all other subsequent operations by the keystore access interface.

```
int sm_keystore_slot_alloc(struct device *dev, u32 unit, u32 size, u32*slot);
```

Arguments:

`dev` Points to a struct device established to manage resources for the secure memory subsystem.

`unit` One of the units detected with a call to sm_detect_keystore_units().

`size` Desired size of data for storage in the allocated slot.

`slot` Pointer to the variable to receive the allocated slot number, once known.

Returns:

- Zero for successful completion.
- -EKEYREJECTED if the requested size exceeds the selected slot size.

8.1.14   Load Data into a Keystore Slot

Load data into an allocated keystore slot so that other operations (such as encapsulation) can be carried out upon it.

```
int sm_keystore_slot_load(struct device *dev, u32 unit, u32 slot, constu8 *key_data, u32
key_length);
```

Arguments:

dev Points to a struct device established to manage resources for the secure memory subsystem.

unit One of the units detected with a call to sm_detect_keystore_units().

key_length Length (in bytes) of information to write to the slot.

key_data Pointer to buffer with the data to be loaded. Must be a contiguous buffer.

Returns:

- Zero for successful completion.
- -EFBIG if the requested size exceeds that which the slot can hold.

8.1.15   Demo Image Update

Encapsulate data written into a keystore slot as a Secure Memory Blob.

```
int sm_keystore_slot_encapsulate(struct device *dev, u32 unit, u32
inslot, u32 outslot, u16 secretlen, u8 *keymod, u16 keymodlen);
```

Arguments:

dev Points to a struct device established to manage resources for the secure memory subsystem.

unit One of the units detected with a call to sm_detect_keystore_units().

inslot Slot holding the input secret, loaded into that slot by sm_keystore_slot_load(). Note that the slot containing this secret should be overwritten or deallocated as soon as practical, since it contains cleartext at this point.

outslot Allocated slot to hold the encapsulated output as a Secure Memory Blob.

secretlen Length of the secret to be encapsulated, not including any blob storage overhead (blob key, MAC, etc.).

`keymod` Key modifier component to be used for encapsulation. The key modifier allows an extra secret to be used in the encapsulation process. The same modifier will also be required for decapsulation.

`keymodlen` Lenth of key modifier in bytes.

Returns:

- Zero on success
- CAAM job status if a failure occurs

8.1.16   Decapsulate Data in the Keystore

Decapsulate data in the keystore into a Black Key Blob for use in other cryptographic operations. A Black Key Blob allows a key to be used "covered" in main memory without exposing it as cleartext.

```
int sm_keystore_slot_decapsulate(struct device *dev, u32 unit, u32
inslot, u32 outslot, u16 secretlen, u8 *keymod, u16 keymodlen);
```

Arguments:

`dev` Points to a struct device established to manage resourcesfor the secure memory subsystem.

`unit` One of the units detected with a call to sm_detect_keystore_units().

`inslot` Slot holding the input data, processed by a prior call to `sm_keystore_slot_encapsulate()`, and containing a Secure Memory Blob.

`outslot` Allocated slot to hold the decapsulated output data in the form of a Black Key Blob.

`secretlen` Length of the secret to be decapsulated, without any blob storage overhead.

`keymod` Key modified component specified at the time of encapsulation.

`keymodlen` Lenth of key modifier in bytes.

Returns:

- Zero on success
- CAAM job status if a failure occurs

8.1.17 Read Data From a Keystore Slot

Extract data from a keystore slot back to a user buffer. Normally to be used after some other operation (e.g., decapsulation) occurs.

```
int sm_keystore_slot_read(struct device *dev, u32 unit, u32 slot, u32
key_length, u8 *key_data);
```

Arguments:

`dev` Points to a struct device established to manage resources for the secure memory subsystem.

`unit` One of the units detected with a call to sm_detect_keystore_units().

`slot` Allocated slot to read from.

`key_length` Length (in bytes) of information to read from the slot.

`key_data` Pointer to buffer to hold the extracted data. Must be a contiguous buffer.

Returns:

- Zero for successful completion.
- -EFBIG if the requested size exceeds that which the slot can hold.

8.1.18 Release a Slot back to the Keystore

Release a keystore slot back to the available pool. Information in the store is wiped clean before the deallocation occurs.

```
int sm_keystore_slot_dealloc(struct device *dev, u32 unit, u32 slot);
```

Arguments:

`dev` Points to a struct device established to manage resources for the secure memory subsystem.

`unit` One of the units detected with a call to `sm_detect_keystore_units()`.

`slot` Number of the allocated slot to be released back to the store.

Returns:

- Zero for successful completion.
- -EINVAL if an unallocated slot is specified.

Configuration of the Secure Memory Driver / Keystore API is dependent on the following kernel configuration parameters:

CRYPTO_DEV_FSL_CAAM_SM

Turns on the secure memory driver in the kernel build.

CRYPTO_DEV_FSL_CAAM_SM_SLOTSIZE

Configures the size of a secure memory "slot".

Each secure memory unit is block of internal memory, the size of which is implementation dependent. This block can be subdivided into a number of logical "slots" of a size which can be selected by this value. The size of these slots needs to be set to a value that can hold the largest secret size intended, plus the overhead of blob parameters (blob key and MAC, typically no more than 48 bytes).

The values are selectable as powers of 2, limited to a range of 32 to 512 bytes. The default value is 7, for a size of 128 bytes.

CRYPTO_DEV_FSL_CAAM_SM_TEST

Enables operation of a captive test / example module that shows how one might use the API, while verifying its functionality. The test module works along this flow:

- Creates a number of known clear keys (3 sizes).
- Allocated secure memory slots.
- Inserts those keys into secure memory slots and encapsulates.
- Decapsulates those keys into black keys.
- Enrcrypts DES, AES128, and AES256 plaintext with black keys. Since this uses symmetric ciphers, same-key encryption/decryption results will be equivalent.
- Decrypts enciphered buffers with equivalent clear keys.
- Compares decrypted results with original ciphertext and compares. If they match, the test reports OK for each key case tested.

Normal output is reported at the console as follows:

```
platform caam_sm.0: caam_sm_test: 8-byte key test match OK platform
caam_sm.0: caam_sm_test: 16-byte key test match OK platform caam_sm.0:
caam_sm_test: 32-byte key test match OK
```

- The secure memory driver is not implemented as a kernel module at this point in time.
- Implementation is presently limited to kernel-mode operations.

- One instance is possible at the present time. In the future, when job rings can run independently in different system partitions, a multiple instance secure memory driver should be considered.
- All storage requests are limited to the storage size of a single slot (which is of a build-time configurable length). It may be possible to allow a secret to span multiple slots so long as those slots can be allocated contiguously.
- Slot size is fixed across all pages/partitions.
- Encapsulation/Decapsulation interfaces could allow for authentication to be specified; the underlying interface does not request it.
- Encapsulation/Decapsulation interfaces return a job status; this status should be translated into a meaningful error from `errno.h`

8.1.19 CAAM/SNVS - Security Violation Handling Interface Overview

This chapter describes a prototype of a driver component and control interface for SNVS Security Violations. It provides a means of installing, managing, and executing application defined handlers meant to process security violation events as a response to their occurrence in a system.

SNVS allows for the continuous monitoring of a number of possible attack vectors in a running system. If the occurrence of one of these attach vectors is sensed, (e.g., a Security Violation has been detected), SNVS can, along with erasing critical security parameters and transitioning to a failure state. generate an interrupt indicating that the violation has occurred. This interrupt can dispatch an application-defined routine to take cleanup action as a consequence of the violation, such that an orderly shutdown of security services might occur.

Therefore, the purpose of this interface is to allow system-level services to install handlers for these types of events. This allows the system designer to select how he wants to respond to specific security violation causes using a simple function call written to his system-specific requirements.

8.1.20 Operation

For existing platforms, 6 security violation interrupt causes are possible within SNVS. 5 of these violation causes are normally wired for use, and these causes are defined as:

- SECVIO_CAUSE_CAAM_VIOLATION - Violation detected inside CAAM/SNVS
- SECVIO_CAUSE JTAG_ALARM - JTAG activity detected
- SECVIO_CAUSE_WATCHDOG - Watchdog expiration

- SECVIO_CAUSE_EXTERNAL_BOOT - External bootload activity
- SECVIO_CAUSE_TAMPER_DETECT - Tamper detection logic triggered

Each of these causes can be associated with an application-defined handler through the API provided with this driver. If no handler is specified, then a default handler will be called. This handler does no more than to identify the interrupt cause to the system console.

8.1.21 Configuration Interface

The following interface can be used to define or remove application-defined violation handlers from the driver's dispatch table.

8.1.22 Install a Handler

```
int caam_secvio_install_handler(struct device *dev, enum secvio_cause
cause, void (*handler)(struct device *dev, u32 cause, void *ext), u8
*cause_description, void *ext);
```

Arguments:

`dev` Points to SNVS-owning device.

`cause` Interrupt source cause from the above list of enumerated causes.

`handler` Application-defined handler, gets called with dev, source cause, and locally-defined handler argument

`cause_description` Points to a string to override the default cause name, this can be used as an alternate for error messages and such. If left NULL, the default description string is used. ext pointer to any extra data needed by the handler.

Returns:

- Zero on success.
- -EINVAL if an argument was invalid or unusable.

8.1.23 Remove an Installed Driver

```
int caam_secvio_remove_handler(struct device *dev, enum secvio_cause
cause);
```

Arguments:

dev Points to SNVS-owning device.

cause Interrupt source cause.

Returns:

- Zero on success.
- -EINVAL if an argument was invalid or unusable.

8.1.24   Driver Configuration CAAM/SNVS

CRYPTO_DEV_FSL_CAAM_SECVIO

Enables inclusion of Security Violation driver and configuration interface as part of the build configuration. The driver is not buildable as a module in its present form.

8.2   Display Content Integrity Checker (DCIC)

8.2.1   Introduction

The goal of the DCIC is to verify that a safety-critical information sent to a display is not corrupted.

The DCIC has the following features:

- Pixel clock up to 148.5 MHz
- Configurable polarity of Display Interface control signals
- 24-bit pixel data bus
- Up to 16 rectangular ROIs with a configurable location and size
- Independent CRC32 signature calculation for each ROI
- External controller mismatch indication signal

8.2.2   Source Code Structure

Table 8-1.   DCIC Driver Files

File	Description
drivers/video/fbdev/mxc/mxc_dcic.c	DCIC driver source code
include/uapi/linux/mxc_dcic.h	DCIC User Space Header

8.2.3   Menu Configuration Options

In menu configuration enable the following module:

Device Drivers -> Graphics support -> MXC DCIC

8.2.4   DTS Configuration

```
dcic_id = <0>;  /* DCIC device index 0-dcic1, i-dcic2 */
dcic_mux = "dcic-lcdif1"; /* DCIC input select */
```

Table 8-2.   DCIC Input Select

Module	i.MX 6SoloX	i.MX 6Dual/6Quad
DCIC1	dcic_lvds	dcic-ipu0-di1
	dcic_lcdif1	dcic-lvds0
		dcic-lvds1
		dcic-hdmi
DCIC2	dcic_lvds	dcic-ipu0-di0/dcic-ipu1-di0
	dcic_lcdif2	dcic-lvds0
		dcic-lvds1
		dcic-mipi_dpi

8.2.5   IOCTLs Functions

The DCIC driver supports the following IOCTLs:

- DCIC_IOC_CONFIG_DCIC: Configures the DCIC input CLK, VSYNC, HSYNC, and data signal polarity.
- DCIC_IOC_CONFIG_ROI: Configures the ROI block size and reference signature.
- DCIC_IOC_GET_RESULT: Gets the result of the ROI calculated signature.

8.2.6   Structures

```
struct roi_params {
    unsigned int roi_n;      /* ROI index */
    unsigned int ref_sig;   /* Reference CRC32 */
    unsigned int start_y;   /* start vertical lines of ROI */
    unsigned int start_x;   /* start horizon lines of ROI */
    unsigned int end_y;    /* end vertical lines of ROI */
    unsigned int end_x;    /* end horizon lines of ROI */
```

```
    char freeze;                    /*  state of ROI */
};
```

8.2.7 DCIC CRC Calculation Functions

There are four functions in this unit test to calculate reference signature:

```
crc32_calc_18of24bit()    /* CRC calculate 18 bit of 24  */
crc32_calc_24bit()            /* CRC calculate 24  */
crc32_calc_24of16bit()  /* CRC calculate 24 bit of 16  */
crc32_calc_18of16bit() /* CRC calculate 18 bit of 16  */
```

DCIC calculates CRC according to the display bus width, but the display bus width does not always align with bytes per pixel (bpp), and the four functions above can cover different display bus widths and bpps.

8.3 Smart Card Interface - Subscriber Identification Module (SIM)

8.3.1 Introduction

The Subscriber Identification Module (SIM) is designed to facilitate communication to SIM cards or Eurochip prepaid phone cards, and compatible with ISO/IEC 7816-3 standards. The SIM module has one port that can be used to interface with various cards. The interface with the Micro Controller Unit (MCU) is a 32-bit connection as described in the reference document IP Bus Specification.

8.3.2 Modes of Operation

The SIM module I/O interface can be operated in one of the three modes of operation summarized below.
  • Two-wire interface: Both the IC pin RX and TX are used to interface to the SmartCard.
  • External one-wire interface: The IC pins RX and TX are tied together externally to the IC and routed to the SmartCard.

  • Internal one-wire interface: The IC pin TX is routed to the SmartCard. The receive pin RX is connected to the TX pin internally to the IC.

8.3.3 External Signal Description

- SIM_CLK: clock that the SIM module provides for the SmartCard. Typical frequencies are 1 MHz to 5 MHz. This clock is 372 times the data rate that is on pin SIM_TRXD.
- SIM_RST_B: reset signal from the SIM to the SmartCard.
- SIM_SVEN: SmartCard power supply enable control signal.
- SIM_TRXD: transmitted/received date from SIM module to SmartCard.
- SIM_PD: SmartCard insertion detect.

8.3.4 Source Code Structure

Table 8-3. SIM Source

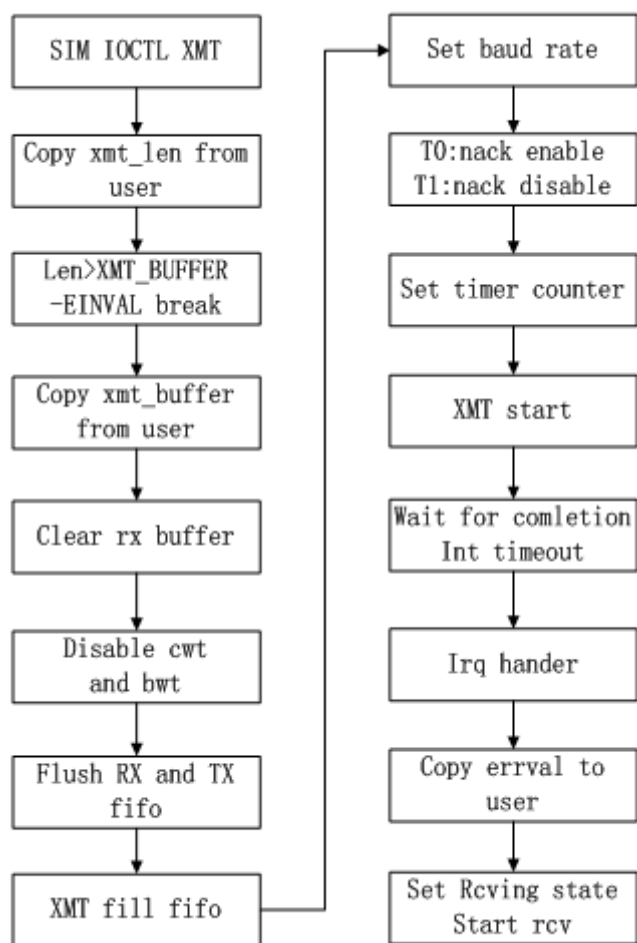| File	Description
drivers/mxc/sim/imx_sim.c	SIM Driver
drivers/mxc/sim/imx_envsim.c	SIM Env

8.3.5 Menu Configuration Options

Configure the kernel option to enable the module by menuconfig:

Device Drivers > MXC support drivers > MXC SIM Support

8.3.6 Software Framework

The following figures show the SIM TX and RX software flows.
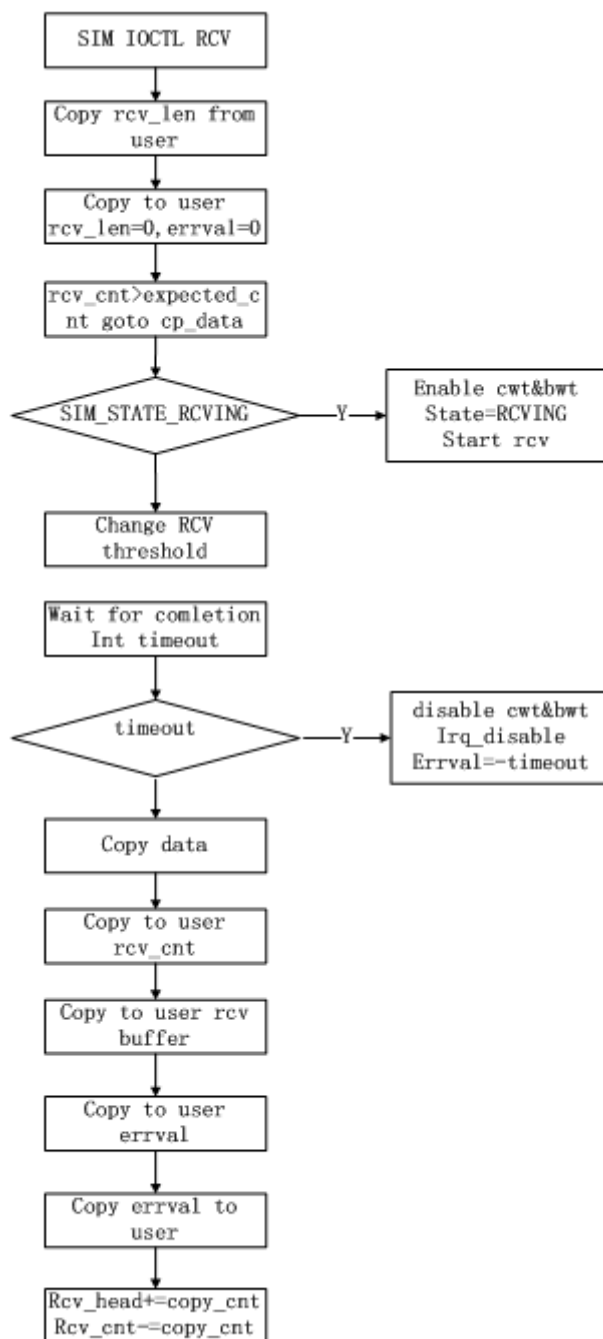
Figure 8-1. SIM transmitting flow

Figure 8-2. SIM receiving flow

8.4   Secure Non-Volatile Storage (SNVS)

8.4.1 Introduction

For more information on Secure Non-Volatile Storage (SNVS), see the i.MX Security Manual for the associated SoC.

SNVS is a block that interfaces with CAAM and SRTC.

For SNVS services related to CAAM, see Section Driver Configuration CAAM/SNVS.

For SNVS services related to srtc, see Section SRTC Introduction

8.5 SNVS Real Time Clock (SRTC)

8.5.1 Introduction

The Real Time Clock (RTC) module is used to keep the time and date. It provides a certifiable time to the user and can raise an alarm if tampering with counters is detected.

8.5.2 Hardware Operation

The RTC is a fake timer provided by the system controller firmware. It only supports basic function of read/set time, read/set alarm.

8.5.3 Software Operation

The following sections describe the software operation of the RTC driver.

8.5.4 Driver Features

The RTC driver includes the following features:

- Implements the functions required by Linux OS to provide the real time clock and alarm interrupt
- Alarm wakes up the system from low power modes

8.5.5   Source Code Structure

The RTC module is implemented in drivers/rtc.

Table 8-4.   RTC Driver Files

File	Description
drivers/rtc/rtc-imxdi.c	MX6 RTC driver
drivers/rtc/rtc-imx-sc.c	MX8 RTC System Controller driver
drivers/rtc/rtc-imx-rpmsg.c	RPMSG RTC driver

8.5.6   Menu Configuration Options

In menu configuration enable the following module:

For i.MX 6 select Device Drivers > Real Time Clock > Freescale IMX DryIce Real Time Clock

For i.MX 8 with SC select Device Drivers > Real Time Clock > NXP SC RTC support

For RPMSG select Device Drivers > Real Time Clock > NXP RPMSG RTC support

Chapter 9
Unit Tests

9.1 System

9.1.1 OProfile

9.1.1.1 Test Name

- autorun-oprofile.sh

9.1.1.1.1 Location

```
/unit_tests/OProfile/
```

9.1.1.1.2 Functionality

OProfile is a system-wide profiler capable of profiling all running code at low overhead. OProfile consists of a kernel driver, a daemon for collecting sample data, and several post-profiling tools for turning data into information.

9.1.1.1.3 Configuration

None

9.1.1.1.4 Use Case and Expected Output

```
./autorun-oprofile.sh
```

9.1.2 Owire

9.1.2.1 Test Name

- autorun-owire.sh

9.1.2.1.1 Location

`/unit_tests/OWire/`

9.1.2.1.2 Functionality

Test EEPROM functionality.

9.1.2.1.3 Configuration

None

9.1.2.1.4 Use Case and Expected Output

`./autorun-owire.sh`

9.1.3 Power Management

9.1.3.1 Test Name

- /unit_tests/Power_Management/suspend_random_auto.sh
- /unit_tests/Power_Management/suspend_quick_auto.sh

9.1.3.1.1 Location

`/unit_tests/Power_Management/`

9.1.3.1.2 Functionality

Enables low power mode on and wakes up the different cores on all i.MX boards..

9.1.3.1.3 Configuration

None

9.1.4   Remote Processor Messaging

9.1.4.1   Test Name

- mxc_mcc_tty_test.out

9.1.4.1.1   Location

```
/unit_tests/RPMsg/
```

9.1.4.1.2   Functionality

Test communication between Cortex-A and Cortex-M cores.

9.1.4.1.3   Use Case and Expected Output

Run the following command and ensure that the RPMsg TTY receiving program is running at the backend when starting RPMsg TTY tests.

```
# ./mxc_dcic_test.out -bw 18 -dev 1

Expected Output:
$ /unit_tests/mxc_mcc_tty_test.out /dev/ttyRPMSG30 115200 R 100 1000 &
Expected output
mxc_mcc_tty_test.out:
$ insmod imx_rpmsg_tty.ko
$ imx_rpmsg_tty rpmsg0: new channel: 0x400 -> 0x1!
$ Install rpmsg tty driver!
$ echo deadbeaf > /dev/ttyRPMSG30
$ imx_rpmsg_tty rpmsg0: msg(<- src 0x1) deadbeaf len 8
```

9.1.5   Watchdog (WDOG)

9.1.5.1   Test Name

- autorun-wdog.sh
- wdt_driver_test.out

9.1.5.1.1   Location

```
/unit_tests/Watchdog/
```

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

9.1.5.1.2   Functionality

Tests the Watchdog Timer module which protects against system failures by providing an escape from unexpected hang, infinite loop situations or programming errors.

9.1.5.1.3   Configuration

None

9.1.5.1.4   Use Case and Expected Output

```
Use case
./autorun-wdog.sh

or

./wdt_driver_test.out 1 2 0 &
Expected output
This should generate a reset after 3 seconds (a 1 second time-out and a 2 second sleep).

or
./wdt_driver_test.out 2 1 0

The system should keep running without being reset. This test requires the kernel to be executed
with the "jtag=on" on some platforms. Press "Ctrl+C" to terminate this test program.
```

9.2   Storage

9.2.1   Media Local Bus

9.2.1.1   Test Name
   • mxc_mlb150_test

9.2.1.1.1   Location

```
/unit_tests/Media_Local_Bus/
```

9.2.1.1.2   Functionality

MediaLB is an on-PCB or inter-chip communication bus specifically designed to standardize a common hardware interface and software API library.

9.2.1.1.3 Configuration

In menu configuration enable the following module:

```
Device Drivers > MXC support drivers > MXC Media Local Bus Driver > MLB support
```

Test only supported on i.MX6SX, i.MX6QP, i.MX6Q, i.MX6DL

9.2.1.1.4 Use Case and Expected Output

```
./mxc_mlb150_test [-v] [-h] [-b] [-f fps] [-t casetype] [-q sync quadlets] [-p isoc
packet length]\n"
-v verbose
-h help
-b block io test
-f FPS, 256/512/1024/2048/3072/4096/6144
-t CASE, CASE can be 'sync', 'ctrl', 'async', 'isoc'
-q SYNC QUADLETS, quadlets per frame in sync mode, can be 1, 2, or 3
-p Packet length, package length in isoc mode, can be 188 or 196
```

9.2.2 MMC/SD/SDIO Host

9.2.2.1 Test Name

- autorun-mmc-blockrw.sh
- autorun-mmc-fdisk.sh
- autorun-mmc-fs.sh
- autorun-mmc-mkfs.sh
- autorun-mmc.sh

9.2.2.1.1 Location

```
/unit_tests/MMC_SD_SDIO/
```

9.2.2.1.2 Functionality

The conjunction of MMC SD SDIO tests exercise the following instructions:

- MMC/SD read write test.
- MMC/SD block read write test.
- MMC/SD fdisk test.
- MMC/SD file system test.
- MMC/SD mkfs test.

9.2.2.1.3 Configuration

None

9.2.2.1.4 Use Case and Expected Output

All test return "Pass" if successful.

```
./autorun-mmc-blockrw.sh
./autorun-mmc-fdisk.sh
./autorun-mmc-fs.sh
./autorun-mmc-mkfs.sh
./autorun-mmc.sh
```

9.2.3 MMDC

9.2.3.1 Test Name

- mmdc2

9.2.3.1.1 Location

```
/unit_tests/MMDC/
```

9.2.3.1.2 Functionality

MMDC profiling utility.

9.2.3.1.3 Configuration

The following parameters allow to customize the mmcd2 test:
- export MMDC_SLEEPTIME - define profiling duration (500ms by default)
- export MMDC_LOOPCOUNT - define profiling times (1 by default, -1 means infinite loop)
- export MMDC_CUST_MADPCR1 - customize madpcr1

9.2.3.1.4 Use Case and Expected Output

The expected output will print the profiling results

```
./mmdc2 [ARM:DSP1:DSP2:GPU2D:GPU2D1:GPU2D2:GPU3D:GPU3D2:GPUVG:VPU:M4:PXP:USB:SUM]
```

9.2.4   SATA

9.2.4.1   Test Name

- autorun-ata.sh

9.2.4.1.1   Location

```
/unit_tests/SATA/
```

9.2.4.1.2   Functionality

This test writes data to the SATA drive connected to the SATA connector on the i.MX board. The data is then read back and compared to what was written.

9.2.4.1.3   Configuration

Module required: pata_fsl.ko. Hardware required: SATA drive. Only i.MX 6 Quad and QuadPlus have SATA support.

9.2.4.1.4   Use Case and Expected Output

```
./autorun-ata.sh

Expected output
Test should return "HDD test passes" if successful.
```

9.3   Connectivity

9.3.1   Enhanced Configurable Serial Peripheral Interface (ECSPI)

9.3.1.1   Test Name

- mxc_spi_test1.out

9.3.1.1.1 Location

```
/unit_tests/ECSPI/
```

9.3.1.1.2 Functionality

This test sends bytes of the last parameter to a specific SPI device. The maximum transfer bytes are 4096 bytes for bits per word less than 8(including 8), 2048 bytes for bits per word between 9 and 16, 1024 bytes for bits per word larger than 17(including 17). SPI writes data received data from the user into Tx FIFO and waits for the data in the Rx FIFO. Once the data is ready in the Rx FIFO, it is read and sent to user.

9.3.1.1.3 Configuration

For the i.MX 6QuadPlus/Quad/Dual auto boards this requires the ecspi device tree. This feature is disabled with default device tree.

9.3.1.1.4 Use Case and Expected Output

```
./mxc_spi_test1.out -D 0 -s 1000000 -b 8 E6E0
./mxc_spi_test1.out -D 1 -s 1000000 -b 8 -H -O -C E6E0E6E00001E6E00000
Usage:
./mxc_spi_test1.out [-D spi_no] [-s speed] [-b bits_per_word] [-H] [-O] [-C] $lt;value>
<spi_no> - CSPI Module number in [0, 1, 2]
<speed> - Max transfer speed
<bits_per_word> - bits per word
-H - Phase 1 operation of clock
-O - Active low polarity of clock
-C - Active high for chip select
<value> - Actual values to be sent
```

9.3.2 ENET IEEE-1588

9.3.2.1 Test Name
- autorun-firi.sh
- mxc_irlan_test.sh

9.3.2.1.1 Location

```
/unit_tests/ENET/
```

9.3.2.1.2 Functionality

The goal of the DCIC is to verify that a safety-critical information sent to a display is not corrupted.

9.3.2.1.3 Configuration

None

9.3.2.1.4 Use Case and Expected Output

```
# ./autorun-firi.sh
# ./mxc_irlan_test.sh
```

9.3.3 ETM

9.3.3.1 Test Name
- etm

9.3.3.1.1 Location

```
/unit_tests/ETM/
```

9.3.3.1.2 Functionality

Embedded Trace Macrocell, The ETM is an optional debug component that enables reconstruction of program execution. The ETM is designed as a high-speed, low-power debug tool that only supports instruction trace. This ensures that area is minimized, and that gate count is reduced.

9.3.3.1.3 Configuration

9.3.3.1.4 Use Case and Expected Output

```
# ./etm -h
Usage: ./etm [options]
Options:
--etm-3.3 ETM v3.3 trace data
--etm-3.4-alt-branch ETM v3.4 trace data with alternative branch encoding
--pft-1.1 PFT v1.1 trace data
--cycle-accurate Cycle-accurate tracing was enabled (Default 1)
```

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

```
--contextid-bytes Number of Context ID bytes (Default 4)
--formatter Enable Formatter Unpacking
--sourceid-match Enable Source ID from formatter. Also enables formatter
--print-long-waits Highlight long waits
--print-input Print input data
--print-config Print configuration data
--help Print usage information
```

9.3.4   Inter-IC (I2C)

9.3.4.1   Test Name

- autorun-mxc_i2c_slave.sh

9.3.4.1.1   Location

```
/unit_tests/I2C/
```

9.3.4.1.2   Functionality

9.3.4.1.3   Configuration

None

9.3.4.1.4   Use Case and Expected Output

9.3.5   IIM

9.3.5.1   Test Name

- mxc_iim_test.out

9.3.5.1.1   Location

```
/unit_tests/IIM_Driver/
```

9.3.5.1.2   Functionality

This test can read an iim value from a bank or fuse a value to a bank

9.3.5.1.3   Configuration

None

9.3.5.1.4   Use Case and Expected Output

For both read and fuse test input values should be in hex format.

```
Below is the usage for the read case.
./mxc_iim_test read -d <bank_addr>
<bank_addr> - bank address in fuse map file.
read - Indicate that this is a read operation.

Example:
./mxc_iim_test.out read -d 0xc30

Below is the usage for the fuse case.
./mxc_iim_test fuse -d <bank_addr> -v <value>
<bank_addr> - bank address in fuse map file.
<value> - Value to be writen to a bank.
fuse - Indicate that this is a write operation.

Example:
./mxc_iim_test.out fuse 0xc30 -v 0x40
```

9.3.6   Keyboard

9.3.6.1   Test Name

- autorun-keypad.sh
- mxc_keyb_test.sh

9.3.6.1.1   Location

```
/unit_tests/Keyboard/
```

9.3.6.1.2   Functionality

Tests keyboard input via USB.

9.3.6.1.3   Configuration

Connect Keyboard to USB OTG port.

9.3.6.1.4 Use Case and Expected Output

```
./autorun-keypad.sh
```

```
Outputs:
Print "PASS" status
```

```
./mxc_keyb_test.sh
```

```
Output:
An event will occur when a key is pressed
```

9.3.7 Low Power Universal Asynchronous Receiver/Transmitter (LPUART)

9.3.7.1 Test Name

- autorun-mxc_uart.sh
- mxc_uart_stress_test.out
- mxc_uart_test.out
- mxc_uart_xmit_test.out

9.3.7.1.1 Location

```
/unit_tests/UART/
```

9.3.7.1.2 Functionality

These tests excercise the low-level UART driver whihc is responsible for supplying information such as the UART port information and a set of control functions to the core UART driver.

9.3.7.1.3 Configuration

None

9.3.7.1.4 Use Case and Expected Output

```
./autorun-mxc_uart.sh
./mxc_uart_stress_test.out /dev/ttymxc#
./mxc_uart_test.out /dev/ttymxc#
./mxc_uart_xmit_test.out /dev/ttymxc#
```

9.3.8 USB

9.3.8.1 Test Name
- autorun-usb-gadget.sh
- autorun-usb-host.sh

9.3.8.1.1 Location

```
/unit_tests/USB/
```

9.3.8.1.2 Functionality

This tests excerise the universal serial bus (USB) driver which implements a standard Linux driver interface to the CHIPIDEA USB-HS OTG controller. The USB provides a universal link that can be used across a wide range of PC-to-peripheral interconnects. It supports plug-and-play, port expansion, and any new USB peripheral that uses the same type of port.

9.3.8.1.3 Configuration
Modules required:
- /lib/modules/$(kernel_version)/kernel/drivers/usb/gadget/g_ether.ko
- /lib/modules/$(kernel_version)/kernel/drivers/usb/gadget/arcotg_udc.ko
- /lib/modules/$(kernel_version)/kernel/drivers/usb/host/ehci-hcd.ko

9.3.8.1.4 Use Case and Expected Output

```
./autorun-usb-gadget.sh
or
./autorun-usb-host.sh
```

9.4 Graphics

9.4.1 Graphics Processing Unit (GPU)

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

9.4.1.1 Test Name

- gpu.sh
- gpuinfo.sh

9.4.1.1.1 Location

`/unit_tests/GPU`

9.4.1.1.2 Functionality

GPU function test

- tutorial3: test OpenGL ES 1.1 basic function
- tutorial4_es20: test OpenGL ES 2.0 basic function
- tiger: test OpenVG 1.1 basic function
- tvui: test Raster 2D and LibVivanteDK API

9.4.1.1.3 Configuration

For gpu.sh and gpuinfo.sh to work add the following line to the target board defconfig file:

- CONFIG_MXC_GPU_VIV=y

Hardware required: LVDS Display Panel and i.MX SoC with a GPU

9.4.1.1.4 Use Case and Expected Output

`./gpu.sh`

- Expected output are frames are drawn properly on screen

- tutorial3: a cube with texture rotating in the middle of the screen
- tutorial4_es20: draws a glass sphere inside a big sphere (enviroment mapping). The glass sphere shows both reflection and refraction effects.
- tiger: a tiger spinning on the screen.
- tvui: draws several movie clips and a tv control panel.

Example output is:

```
# ./gpu.sh
---- Running < gpu.sh > test ----
/unit_tests/GPU /unit_tests/GPU
Rendered 100 frames in 624 milliseconds: 160.26 fps
id=43, a,b,g,r=0,8,8,8, d,s=16,0, AA=0,openvgbit=71
frames:100 -- fps:58.997051
press ESC to escape...
./gpu.sh: line 28: cd: /opt/viv_samples/hal/: No such file or directory
/unit_tests/GPU
---- Test < gpu.sh > ended ----
```

```
./gpuinfo.sh
```

- Information about GPU is printed on console.

```
# ./gpuinfo.sh
----  Running < gpuinfo.sh > test  ----
GPU Info
gpu : 0
model : 2000
revision : 5108
product : 0
eco : 0
gpu : 8
model : 320
revision : 5007
product : 0
eco : 0
gpu : 9
model : 355
revision : 1215
product : 0
eco : 0
VIDEO MEMORY:
gcvPOOL_SYSTEM:
Free : 134217728 B
Used : 0 B
Total : 134217728 B
gcvPOOL_CONTIGUOUS:
Used : 0 B
gcvPOOL_VIRTUAL:
Used : 0 B
NON PAGED MEMORY:
Used : 0 B
Paged memory Info
lowMem: 0 bytes
highMem: 0 bytes
CMA memory info
cma: 138485760 bytes
>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>
Idle percentage:0.000.000.000.000.000.00%
>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>
---- Test < gpuinfo.sh > ended ----
```

9.5 Video

9.5.1 Display

9.5.1.1 Test Name
- autorun-fb.sh
- mxc_tve_test.sh
- mxc_fb_test.out
- mxc_epdc_fb_test.out
- mxc_epdc_v2_fb_test.out

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

- mxc_spdc_fb_test.out
- mxc_fb_vsync_test.out

9.5.1.1.1 Location

`/unit_tests/Display/`

9.5.1.1.2 Functionality

The tests under the display directory test some of the display options that are available with the i.MX family of boards. Some of the devices that can be tested include LVDS, HDMI and EPDC panels.

Specifically the 'mxc_fb_test.out' tests the following features:
- Basic fb operation
- Set background/foreground to 16 bpp fb
- Global alpha blending
- Color key test
- Framebuffer Panning
- Gamma test

Aditionally, the EPDC tests 'mxc_epdc_fb_test.out' and 'mxc_epdc_v2_fb_test.out' test the following features:
- Basic Updates
- Rotation Updates
- Grayscale Framebuffer Updates
- Auto-waveform Selection Updates
- Panning Updates
- Overlay Updates
- Auto-Updates
- Animation Mode Updates
- Fast Updates
- Partial to Full Update Transitions
- Test Pixel Shifting Effect
- Colormap Updates
- Collision Test Mode
- Stress Test
- RGB565, Y8 framebuffer format
- 0, 90, 180, 270 degree framebuffer rotation
- Framebuffer panning
- Use of the alternate framebuffer
- Auto-waveform mode selection

- Automatic update mode
- The force-monochrome update feature and animation mode updates
- Support for using a grayscale colormap
- Snapshot, Queue, and Queue and Merge update schemes
- The EPDC Collision Test mode

9.5.1.1.3   Configuration

In order to run some tests, changes to the defconfig file for the target board are required. These changes will add functionality in which the following tests depend upon.

For autorun-fb.sh, 'mxc_fb_test.out' and 'mxc_fb_vsync_test.out' add the following to the target board defconfig file:

```
CONFIG_FB=y
CONFIG_FB_MXC=y
CONFIG_FB_MXC_EDID=y
CONFIG_FB_MXC_SYNC_PANEL=y
CONFIG_FB_MXC_LDB=y
CONFIG_FB_MXC_HDMI=y
```

For 'mxc_epdc_fb_test.out' and 'mxc_epdc_v2_fb_test.out' add the following to the target board defconfig file:

```
CONFIG_FB=y
CONFIG_FB_MXC=y
CONFIG_FB_MXC_EINK_PANEL=y
CONFIG_MFD_MAX17135=y
CONFIG_REGULATOR_MAX17135=y
CONFIG_MXC_PXP=y
CONFIG_DMA_ENGINE=y
```

9.5.1.1.4   Use Case and Expected Output

```
# ./autorun-fb.sh
```

Expected output is:

```
---- Running < autorun-fb.sh > test ----
Checking for devnode: /dev/fb0
autorun-fb.sh: PASS devnode found: /dev/fb0
FB Blank test
Screen should be off
FB Color test
Setting FB to 16-bpp
Setting FB to 24-bpp
Setting FB to 32-bpp
FB panning test
autorun-fb.sh: Exiting PASS
Exiting PASS.
```

```
# ./mxc_tve_test.sh
```

Expected output is:

```
---- Running < mxc_tve_test.sh > test ----
Setting TV to NTSC mode
/unit_tests/Display/mxc_tve_test.sh: line 9: echo: write error: Invalid argument
/unit_tests/Display/mxc_tve_test.sh: line 11: /unit_tests/mxc_v4l2_output.out: No such
file or directory
Blank the display
Unblank the display
Setting TV to PAL mode
/unit_tests/Display/mxc_tve_test.sh: line 22: echo: write error: Invalid argument
/unit_tests/Display/mxc_tve_test.sh: line 23: /unit_tests/mxc_v4l2_output.out: No such
file or directory
Blank the display
Unblank the display
```

```
# ./mxc_fb_test.out
```

Expected Output is shown below. The test should pass without any failure messages, and the display on panel should be correct. For each test, a sequence of updates should be reflected on the screen. For almost all tests, the text printed out in the debug console describes the image that should be observed on the screen. For i.MX6Quad fb0 and fb1 are used for tests, fb0 is background framebuffer, and fb1 is foreground overlay framebuffer.

```
Opened fb: /dev/fb0 (DISP4 BG - DI1)
DISP4 BG - DI1: screen info: 1024x768 (virtual: 1024x1536) @ 32-bpp
Opened fb: /dev/fb1 (DISP4 FG)
DISP4 FG: screen info: 240x320 (virtual: 240x960) @ 16-bpp
@DISP4 BG - DI1: Set colorspace to 16-bpp
@DISP4 FG: Set colorspace to 16-bpp
Prepared DISP4 BG - DI1 (black) and DISP4 FG (white). Verify the screen and press any
key to continue!
@DISP4 BG - DI1: Succesfully changed screen to 1024x768 (virtual: 1024x768) @16-bpp
@DISP4 FG: Succesfully changed screen to 240x320 (virtual: 240x320) @16-bpp
Testing global alpha blending...
Fill the FG in black (screen is 240x320 @ 16-bpp)
Fill the BG in white (screen is 1024x768 @ 16-bpp)
Alpha is 0, FG is opaque
Alpha is 255, BG is opaque
Color key enabled
Color key disabled
Global alpha disabled
Pan test start.
@DISP4 FG: Set the colorspace to 16-bpp
Pan test done.
@DISP4 BG - DI1: Set colorspace to 16-bpp
Pan test start.
@DISP4 BG - DI1: Set the colorspace to 16-bpp
Pan test done.
Gamma test start.
Gamma 0.800000
Gamma 1.000000
Gamma 1.500000
Gamma 2.200000
Gamma 2.400000
Gamma test end.
Test bpp start
@DISP4 BG - DI1: Set colorspace to 32-bpp
@DISP4 BG - DI1: Fill the screen in red
@DISP4 BG - DI1: Set colorspace to 24-bpp
@DISP4 BG - DI1: Fill the screen in blue
```

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

```
@DISP4 BG - DI1: Set colorspace to 16-bpp
@DISP4 BG - DI1: Fill the screen in green
Test bpp end


# ./mxc_epdc_fb_test.out [-h] [-a] [-n]
EPDC framebuffer driver test program.
Usage: mxc_epdc_fb_test [-h] [-a] [-p delay] [-u s/q/m] [-n <expression>]
-h Print this message
-a Enabled animation waveforms for fast updates (tests 8-9)
-p Provide a power down delay (in ms) for the EPDC driver
0 - Immediate (default)
-1 - Never
<ms> - After <ms> milliseconds
-u Select an update scheme
s - Snapshot update scheme
q - Queue update scheme
m - Queue and merge update scheme (default)
-n Execute the tests specified in expression
Expression is a set of comma-separated numeric ranges
If not specified, all tests except Stress are run
Example:
./mxc_epdc_fb_test.out -n 1-3,5,7


EPDC tests:
1 - Basic Updates
2 - Rotation Updates
3 - Grayscale Framebuffer Updates
4 - Auto-waveform Selection Updates
5 - Panning Updates
6 - Overlay Updates
7 - Auto-Updates
8 - Animation Mode Updates
9 - Fast Updates
10 - Partial to Full Update Transitions
11 - Test Pixel Shifting Effect
12 - Colormap Updates
13 - Collision Test Mode
14 - Stress Test
15 - Dithering Y8->Y1 Test
16 - Dithering Y8->Y4 Test
17 - Hardware Dithering Test
18 - Advanced Algorithm Test
```

The full set of tests should pass without any failure messages. For each test, a sequence of updates should be reflected on the screen. For almost all tests, the text printed out in the debug console describes the image that should be observed on the screen. mxc_epdc_v2_fb_test.out: The full set of tests should pass without any failure messages. For each test, a sequence of updates should be reflected on the screen. For almost all tests, the text printed out in the debug console describes the image that should be observed on the screen.

```
# ./mxc_spdc_fb_test.out
---- Running < ./mxc_spdc_fb_test.out > test ----
Unable to open /dev/fb5


# ./mxc_fb_vsync_test.out
Usage:
/unit_tests/Display# ./mxc_fb_vsync_test.out <fb #> <count>
<fb #> the framebuffer number
<count> the frames to be rendered
```

```
Example:
/unit_tests/Display# echo 0 > /sys/class/graphics/fb0/blank
/unit_tests/Display# ./mxc_fb_vsync_test.out 0 100
```

Expected output is the following when using 100 for the < count > argument

```
total time for 100 frames = 1655674 us = 60 fps
```

9.5.2   High-Definition Multimedia Interface (HDMI) and Display Port (DP) Overview

9.5.2.1   Test Name

- mxc_cec_test.out

9.5.2.1.1   Location

```
/unit_tests/HDMI/
```

9.5.2.1.2   Functionality

Verify HDMI CEC function and send poweroff command to HDMI sink.

9.5.2.1.3   Configuration

For mxc_cec_test.out to work add the following line to the target board defconfig file:

```
CONFIG_MXC_HDMI_CEC=y
```

The hardware should support HDMI and TV should support HDMI CEC

9.5.2.1.4   Use Case and Expected Output

```
./mxc_cec_test.out
```

9.5.3   Video Processing Unit (VPU)

9.5.3.1   Test Name

- autorun-vpu.sh
- mxc_vpu_test.out

9.5.3.1.1   Location

```
/unit_tests/VPU/
```

9.5.3.1.2   Functionality

The VPU test exercises the following options on the Video Processing Unit (VPU):
   • Decode one stream and display on the LCD.
   • Decode a stream and save to a file.
   • Decode a stream using a config file.
   • Encode a YUV stream and save to a file.
   • Encode an image from the camera and decode it to display on the LCD.
   • Decode multiple streams with different formats simultaneously.
   • Decode and encode simultaneously.
   • Output to TV out.
   • Test VPU with VDI (HW deinterlace via IPU).

9.5.3.1.3   Configuration

This tests require libvpu.so under /usr/lib/ and LCD display. This test requires i.MX 6QuadPlus/Quad/Dual SoC.

9.5.3.1.4   Use Case and Expected Output

```
./autorun-vpu.sh
Decode one stream and display on the LCD.

To test MPEG-4 decode and display to screen:
./mxc_vpu_test.out -D "-i /usr/vectors/file.m4v -f 0"

To test H.263 decode and display to screen:
./mxc_vpu_test.out -D "-i /usr/vectors/file.263 -f 1"

To test H.264 decode and display to screen:
./mxc_vpu_test.out -D "-i /usr/vectors/file.264 -f 2"

You can get the mp4 test file from the imx-test.git server.
It is located under test/mxc_vpu_test/configs/akiyo.mp4.

Decode a stream and save to a file.
To test MPEG-4 decode and save to file:
./mxc_vpu_test.out -D "-i /usr/vectors/file.m4v -f 0 -o out.yuv"

To test H.263 decode and save to file:
./mxc_vpu_test.out -D "-i /usr/vectors/file.263 -f 1 -o out.yuv"

To test H.264 decode and save to file:
./mxc_vpu_test.out -D "-i /usr/vectors/file.264 -f 2 -o out.yuv"

Decode a stream using a config file.
Change options in config file, e.g, config_dec. Input correct input filename, output
filename, format,
./mxc_vpu_test.out -C config_dec
```

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

```
Encode a YUV stream and save to a file.
To test MPEG-4 encode and save to a file:
./mxc_vpu_test.out -E "-i file.yuv -w 240 -h 320 -f 0 -o file.mpeg4"

To test H.263 encode and save to a file:
./mxc_vpu_test.out -E "-i file.yuv -w 240 -h 320 -f 1 -o file.263"

To test H.264 encode and save to a file:
./mxc_vpu_test.out -E "-i file.yuv -w 240 -h 320 -f 2 -o file.264"

Encode an image from the camera and decode it to display on the LCD.
To encode an MPEG4 image from the camera and display on the LCD: that
./mxc_vpu_test.out -L "-f 0 -w 1280 -h 720"

To encode an H263 image from the camera and display on the LCD:
./mxc_vpu_test.out -L "-f 1 -w 1280 -h 720"

To encode an H264 image from the camera and display on the LCD:
./mxc_vpu_test.out -L "-f 2 -w 1280 -h 720"

Decode multiple streams with different formats simultaneously.
Decoder one H264 and one MPEG4 streams.
./mxc_vpu_test.out -D "-i/vectors/file.264 -f 2" -D "-i ./akiyo.mp4 -f 0 -o akiyo.yuv"

Decode and encode simultaneously.
Encode one MPEG-4 stream and decode one H.264 stream simultaneously.
./mxc_vpu_test.out -E "-w 176 -h 144 -f 0 -o enc.m4v" -D "-i/vectors/file.264 -f

Test VPU with TV out.
Decoder one stream as normal VPU test. For example, H264 video stream:
./mxc_vpu_test.out -D "-i filename -f 2"

Test VPU with VDI (HW deinterlace via IPU).
Select one stream with top and bottom fields are interlaced.
av_stress2_dsmcc4m_1_C1_V11_A6.mp4_track1.h264
To decode the stream and display on LCD:
./mxc_vpu_test.out -D "-i av_stress2_dsmcc4m_1_C1_V11_A6.mp4_track1.h264 -f2"

To decode the stream and display on LCD using high motion stream video De Interlacing
algorithm:
./mxc_vpu_test.out -D "-i av_stress2_dsmcc4m_1_C1_V11_A6.mp4_track1.h264 -v h -f2"

To decode the stream and display on LCD using low motion stream video De Interlacing
algorithm:
./mxc_vpu_test.out -D "-i av_stress2_dsmcc4m_1_C1_V11_A6.mp4_track1.h264 -v l -f2"

To decode the stream and display on LCD having input in NV12 pixel format:
./mxc_vpu_test.out -D "-i av_stress2_dsmcc4m_1_C1_V11_A6.mp4_track1.h264 -v
```

9.6 Audio

9.6.1 Advanced Linux Sound Architecture (ALSA) System on a Chip (ASoC) Sound

9.6.1.1 Test Name

- mxc_tuner_test.out

9.6.1.1.1 Location

`/unit_tests/ALSA/`

9.6.1.1.2 Functionality

Test audio capabilities using ALSA.

9.6.1.1.3 Configuration

ALSA is supported on all i.MX for test aplay, arecord and speaker-test. To use this tuner test it requires tuner hardware available only on the i.MX 6 auto reference boards

9.6.1.1.4 Use Case and Expected Output

9.6.2   Asynchronous Sample Rate Converter (ASRC)

9.6.2.1   Test Name
- mxc_asrc_test.out

9.6.2.1.1 Location

`/unit_tests/ASRC`

9.6.2.1.2 Functionality

Converts WAV to different sample rates.

9.6.2.1.3 Configuration

None

9.6.2.1.4 Use Case and Expected Output

```
#/unit_tests/ASRC/mxc_asrc_test.out -to 48000 /unit_tests/ASRC/audio8k16S.wav
audio48k16S.wav
---- Running < /unit_tests/ASRC/mxc_asrc_test.out >  test ----
Pair A requested
All tests passed with success
```

More usages for mxc_asrc_test.out can be obtained by the following command:

i.MX Linux Reference Manual, Rev. L4.9.123_2.3.0_8mm-ga, 09/2018

```
#/unit_tests/ASRC/mxc_asrc_test.out -h
---- Running  < /unit_tests/ASRC/mxc_asrc_test.out > test ----
*************************************************
* Test aplication for ASRC
* Options :
-to <output sample rate> <origin.wav$gt; <converted.wav>
<input clock source> <output clock source>
input clock source types are:
0 -- INCLK_NONE
1 -- INCLK_ESAI_RX
2 -- INCLK_SSI1_RX
3 -- INCLK_SSI2_RX
4 -- INCLK_SPDIF_RX
5 -- INCLK_MLB_CLK
6 -- INCLK_ESAI_TX
7 -- INCLK_SSI1_TX
8 -- INCLK_SSI2_TX
9 -- INCLK_SPDIF_TX
10 -- INCLK_ASRCK1_CLK
default option for output clock source is 0
output clock source types are:
0 -- OUTCLK_NONE
1 -- OUTCLK_ESAI_TX
2 -- OUTCLK_SSI1_TX
3 -- OUTCLK_SSI2_TX
4 -- OUTCLK_SPDIF_TX
5 -- OUTCLK_MLB_CLK
6 -- OUTCLK_ESAI_RX
7 -- OUTCLK_SSI1_RX
8 -- OUTCLK_SSI2_RX
9 -- OUTCLK_SPDIF_RX
10 -- OUTCLK_ASRCK1_CLK
default option for output clock source is 10
*************************************************
```

9.7 Security

9.7.1 Display Content Integrity Checker (DCIC)

9.7.1.1 Test Name

- mxc_dcic_test.out

9.7.1.1.1 Location

/unit_tests/DCIC/

9.7.1.1.2 Functionality

The goal of the DCIC is to verify that a safety-critical information sent to a display is not corrupted.

9.7.1.1.3 Configuration

None

9.7.1.1.4 Use Case and Expected Output

```
# ./mxc_dcic_test.out -bw 18 -dev 1
```

Expected output for mxc_dcic_test.out:

```
Opened fb0
open /dev/dcic1
bpp=16, bus_width=18
Config ROI=1
Config ROI=3
Config ROI=5
ROI=0,crcRS=0x0, crcCS=0x0
ROI=1,crcRS=0x6cd6b18d, crcCS=0x6cd6b18d
ROI=2,crcRS=0x0, crcCS=0x0
ROI=3,crcRS=0xc9da7ae6, crcCS=0xc9da7ae6
ROI=4,crcRS=0x0, crcCS=0x0
ROI=5,crcRS=0xb5ba1453, crcCS=0xb5ba1453
ROI=6,crcRS=0x0, crcCS=0x0
ROI=7,crcRS=0x0, crcCS=0x0
ROI=8,crcRS=0x0, crcCS=0x0
ROI=9,crcRS=0x0, crcCS=0x0
ROI=10,crcRS=0x0, crcCS=0x0
ROI=11,crcRS=0x0, crcCS=0x0
ROI=12,crcRS=0x0, crcCS=0x0
ROI=13,crcRS=0x0, crcCS=0x0
ROI=14,crcRS=0x0, crcCS=0x0
ROI=15,crcRS=0x0, crcCS=0x0
All ROI CRC check success!
```

9.7.2 SIM

9.7.2.1 Test Name

- mxc_sim_test.out

9.7.2.1.1 Location

```
/unit_tests/SIM/
```

9.7.2.1.2 Functionality

Basic testing of SIM card interface.

9.7.2.1.3   Configuration

None

9.7.2.1.4   Use Case and Expected Output

```
/unit_tests/mxc_sim_test.out

Expected output
atr[0]= 0x3b atr[1]= 0x68 atr[2]= 0x0 atr[3]= 0x0 atr[4]= 0x0 atr[5]= 0x73 atr[6]=
0xc8
atr[7]= 0x40 atr[8]= 0x0 atr[9]= 0x0 atr[10]= 0x90 atr[11]= 0x0
rx[0] = 0x6e rx[1] = 0x0
rx[0] = 0x6d rx[1] = 0x0
rx[0] = 0x6e rx[1] = 0x0
```

9.7.3   SNVS Real Time Clock (SRTC)

9.7.3.1   Test Name
- autorun-rtc.sh
- rtctest.out
- rtcwakeup.out

9.7.3.1.1   Location

```
/unit_tests/SRTC/
```

9.7.3.1.2   Functionality

These tests check the Real Time Clock (RTC) module which is used to keep the time and date. It provides a certifiable time to the user and can raise an alarm if tampering with counters is detected.

9.7.3.1.3   Configuration

For autorun-rtc.sh, rtctest.out and rtcwakeup.out to work add the following line to the target board defconfig file:

```
CONFIG_RTC_DRV_SNVS=y
```

Chapter 10
Revision History

10.1 Revision History

This table provides the revision history.

Table 10-1. Revision history

Revision number	Date	Substantive changes
L4.9.51_imx8qxp-alpha	11/2017	Initial release
L4.9.51_imx8qm-beta1	12/2017	Added i.MX 8QuadMax
L4.9.51_imx8mq-beta	12/2017	Added i.MX 8M Quad
L4.9.51_8qm-beta2/8qxp-beta	02/2018	Added i.MX 8QuadMax Beta2 and i.MX 8QuadXPlus Beta
L4.9.51_imx8mq-ga	03/2018	Added i.MX 8M Quad GA
L4.9.88_2.0.0-ga	05/2018	i.MX 7ULP and i.MX 8M Quad GA release
L4.9.88_2.1.0_8mm-alpha	06/2018	i.MX 8M Mini Alpha release
L4.9.88_2.2.0_8qxp-beta2	07/2018	i.MX 8QuadXPlus Beta2 release
L4.9.123_2.3.0_8mm	09/2018	i.MX 8M Mini GA release