

Linux GPU Driver Developer's Guide

Jesse Barnes, Intel Corporation <jesse.barnes@intel.com>

Laurent Pinchart, Ideas on board SPRL <laurent.pinchart@ideasonboard.com>

Daniel Vetter, Intel Corporation <daniel.vetter@ffwll.ch>

Lukas Wunner <lukas@wunner.de>

Linux GPU Driver Developer's Guide

by Jesse Barnes, Laurent Pinchart, Daniel Vetter, and Lukas Wunner

Copyright © 2008-2009, 2013-2014 Intel Corporation

Copyright © 2012 Laurent Pinchart

Copyright © 2015 Lukas Wunner

The contents of this file may be used under the terms of the GNU General Public License version 2 (the "GPL") as distributed in the kernel source COPYING file.

Table of Contents

I. DRM Core	1
1. Introduction	4
Style Guidelines	4
2. DRM Internals	5
Driver Initialization	5
Driver Information	5
Device Instance and Driver Handling	6
Driver Load	14
Bus-specific Device Registration and PCI Support	15
Memory management	21
The Translation Table Manager (TTM)	21
The Graphics Execution Manager (GEM)	22
GEM Function Reference	26
VMA Offset Manager	47
PRIME Buffer Sharing	65
PRIME Function References	66
DRM MM Range Allocator	74
DRM MM Range Allocator Function References	75
CMA Helper Functions Reference	97
Mode Setting	110
Display Modes Function Reference	110
Atomic Mode Setting Function Reference	145
Frame Buffer Abstraction	174
Dumb Buffer Objects	174
Output Polling	175
Locking	175
KMS Initialization and Cleanup	175
CRTCs (struct drm_crtc)	176
Planes (struct drm_plane)	176
Encoders (struct drm_encoder)	176
Connectors (struct drm_connector)	177
Cleanup	180
Output discovery and initialization example	180
KMS API Functions	181
KMS Data Structures	253
KMS Locking	295
Mode Setting Helper Functions	315
Atomic Modeset Helper Functions Reference	315
Modeset Helper Reference for Common Vtables	361
Legacy CRTC/Modeset Helper Functions Reference	379
Output Probing Helper Functions Reference	390
fbdev Helper Functions Reference	400
Display Port Helper Functions Reference	434
Display Port MST Helper Functions Reference	447
MIPI DSI Helper Functions Reference	471
EDID Helper Functions Reference	507
Rectangle Utilities Reference	529
Flip-work Helper Reference	546
HDMI Infoframes Helper Reference	554
Plane Helper Reference	566
Tile group	573

Bridges	573
KMS Properties	584
Existing KMS Properties	585
Vertical Blanking	594
Vertical Blanking and Interrupt Handling Functions Reference	595
Open/Close, File Operations and IOCTLs	626
Open and Close	626
File Operations	627
IOCTLs	627
Legacy Support Code	632
Legacy Suspend/Resume	632
Legacy DMA Services	632
3. Userland interfaces	633
Render nodes	633
VBlank event handling	633
II. DRM Drivers	635
4. drm/i915 Intel GFX Driver	637
Core Driver Infrastructure	637
Runtime Power Management	637
Interrupt Handling	654
Intel GVT-g Guest Support(vGPU)	657
Display Hardware Handling	660
Mode Setting Infrastructure	660
Frontbuffer Tracking	660
Display FIFO Underrun Reporting	668
Plane Configuration	674
Atomic Plane Helpers	674
Output Probing	679
Hotplug	679
High Definition Audio	683
Panel Self Refresh PSR (PSR/SRD)	691
Frame Buffer Compression (FBC)	697
Display Refresh Rate Switching (DRRS)	704
DPIO	710
CSR firmware support for DMC	711
Memory Management and Command Submission	714
Batchbuffer Parsing	714
Batchbuffer Pools	720
Logical Rings, Logical Ring Contexts and Execlists	723
Global GTT views	735
GTT Fences and Swizzling	741
Object Tiling IOCTLs	750
Buffer Object Eviction	752
Buffer Object Memory Shrinking	754
GuC	757
GuC-specific firmware loader	757
GuC-based command submission	760
GuC Firmware Layout	766
Tracing	766
i915_ppgtt_create and i915_ppgtt_release	766
i915_context_create and i915_context_free	766
switch_mm	766
III. vga_switcheroo	768
5. Modes of Use	770

Manual switching and manual power control	770
Driver power control	770
6. API	771
Public functions	771
Public structures	782
Public constants	784
Private structures	786
7. Handlers	789
apple-gmux Handler	789
Graphics mux	789
Power control	790
Backlight control	790

List of Tables

2.1. 585

Part I. DRM Core

This first part of the GPU Driver Developer's Guide documents core DRM code, helper libraries for writing drivers and generic userspace interfaces exposed by DRM drivers.

Table of Contents

1. Introduction	4
Style Guidelines	4
2. DRM Internals	5
Driver Initialization	5
Driver Information	5
Device Instance and Driver Handling	6
Driver Load	14
Bus-specific Device Registration and PCI Support	15
Memory management	21
The Translation Table Manager (TTM)	21
The Graphics Execution Manager (GEM)	22
GEM Function Reference	26
VMA Offset Manager	47
PRIME Buffer Sharing	65
PRIME Function References	66
DRM MM Range Allocator	74
DRM MM Range Allocator Function References	75
CMA Helper Functions Reference	97
Mode Setting	110
Display Modes Function Reference	110
Atomic Mode Setting Function Reference	145
Frame Buffer Abstraction	174
Dumb Buffer Objects	174
Output Polling	175
Locking	175
KMS Initialization and Cleanup	175
CRTCs (struct drm_crtc)	176
Planes (struct drm_plane)	176
Encoders (struct drm_encoder)	176
Connectors (struct drm_connector)	177
Cleanup	180
Output discovery and initialization example	180
KMS API Functions	181
KMS Data Structures	253
KMS Locking	295
Mode Setting Helper Functions	315
Atomic Modeset Helper Functions Reference	315
Modeset Helper Reference for Common Vtables	361
Legacy CRTC/Modeset Helper Functions Reference	379
Output Probing Helper Functions Reference	390
fbdev Helper Functions Reference	400
Display Port Helper Functions Reference	434
Display Port MST Helper Functions Reference	447
MIPI DSI Helper Functions Reference	471
EDID Helper Functions Reference	507
Rectangle Utilities Reference	529
Flip-work Helper Reference	546
HDMI Infoframes Helper Reference	554
Plane Helper Reference	566
Tile group	573
Bridges	573

KMS Properties	584
Existing KMS Properties	585
Vertical Blanking	594
Vertical Blanking and Interrupt Handling Functions Reference	595
Open/Close, File Operations and IOCTLs	626
Open and Close	626
File Operations	627
IOCTLs	627
Legacy Support Code	632
Legacy Suspend/Resume	632
Legacy DMA Services	632
3. Userland interfaces	633
Render nodes	633
VBlank event handling	633

Chapter 1. Introduction

The Linux DRM layer contains code intended to support the needs of complex graphics devices, usually containing programmable pipelines well suited to 3D graphics acceleration. Graphics drivers in the kernel may make use of DRM functions to make tasks like memory management, interrupt handling and DMA easier, and provide a uniform interface to applications.

A note on versions: this guide covers features found in the DRM tree, including the TTM memory manager, output configuration and mode setting, and the new vblank internals, in addition to all the regular features found in current kernels.

[Insert diagram of typical DRM stack here]

Style Guidelines

For consistency this documentation uses American English. Abbreviations are written as all-uppercase, for example: DRM, KMS, IOCTL, CRTC, and so on. To aid in reading, documentations make full use of the markup characters kerneldoc provides: @parameter for function parameters, @member for structure members, &structure to reference structures and function() for functions. These all get automatically hyperlinked if kerneldoc for the referenced objects exists. When referencing entries in function vtables please use ->vfunc(). Note that kerneldoc does not support referencing struct members directly, so please add a reference to the vtable struct somewhere in the same paragraph or at least section.

Except in special situations (to separate locked from unlocked variants) locking requirements for functions aren't documented in the kerneldoc. Instead locking should be checked at runtime using e.g. `WARN_ON(!mutex_is_locked(...))`. Since it's much easier to ignore documentation than runtime noise this provides more value. And on top of that runtime checks do need to be updated when the locking rules change, increasing the chances that they're correct. Within the documentation the locking rules should be explained in the relevant structures: Either in the comment for the lock explaining what it protects, or data fields need a note about which lock protects them, or both.

Functions which have a non-void return value should have a section called "Returns" explaining the expected return values in different cases and their meanings. Currently there's no consensus whether that section name should be all upper-case or not, and whether it should end in a colon or not. Go with the file-local style. Other common section names are "Notes" with information for dangerous or tricky corner cases, and "FIXME" where the interface could be cleaned up.

Chapter 2. DRM Internals

This chapter documents DRM internals relevant to driver authors and developers working to add support for the latest features to existing drivers.

First, we go over some typical driver initialization requirements, like setting up command buffers, creating an initial output configuration, and initializing core services. Subsequent sections cover core internals in more detail, providing implementation notes and examples.

The DRM layer provides several services to graphics drivers, many of them driven by the application interfaces it provides through libdrm, the library that wraps most of the DRM ioctls. These include vblank event handling, memory management, output management, framebuffer management, command submission & fencing, suspend/resume support, and DMA services.

Driver Initialization

At the core of every DRM driver is a `drm_driver` structure. Drivers typically statically initialize a `drm_driver` structure, and then pass it to `drm_dev_alloc()` to allocate a device instance. After the device instance is fully initialized it can be registered (which makes it accessible from userspace) using `drm_dev_register()`.

The `drm_driver` structure contains static information that describes the driver and features it supports, and pointers to methods that the DRM core will call to implement the DRM API. We will first go through the `drm_driver` static information fields, and will then describe individual operations in details as they get used in later sections.

Driver Information

Driver Features

Drivers inform the DRM core about their requirements and supported features by setting appropriate flags in the `driver_features` field. Since those flags influence the DRM core behaviour since registration time, most of them must be set to registering the `drm_driver` instance.

```
u32 driver_features;
```

Driver Feature Flags

<code>DRIVER_USE_AGP</code>	Driver uses AGP interface, the DRM core will manage AGP resources.
<code>DRIVER_REQUIRE_AGP</code>	Driver needs AGP interface to function. AGP initialization failure will become a fatal error.
<code>DRIVER_PCI_DMA</code>	Driver is capable of PCI DMA, mapping of PCI DMA buffers to userspace will be enabled. Deprecated.
<code>DRIVER_SG</code>	Driver can perform scatter/gather DMA, allocation and mapping of scatter/gather buffers will be enabled. Deprecated.
<code>DRIVER_HAVE_DMA</code>	Driver supports DMA, the userspace DMA API will be supported. Deprecated.
<code>DRIVER_HAVE_IRQ</code> , <code>DRIVER_IRQ_SHARED</code>	<code>DRIVER_HAVE_IRQ</code> indicates whether the driver has an IRQ handler managed by the DRM Core. The core will support simple IRQ handler installa-

tion when the flag is set. The installation process is described in the section called “IRQ Registration”.

DRIVER_IRQ_SHARED indicates whether the device & handler support shared IRQs (note that this is required of PCI drivers).

DRIVER_GEM	Driver use the GEM memory manager.
DRIVER_MODESET	Driver supports mode setting interfaces (KMS).
DRIVER_PRIME	Driver implements DRM PRIME buffer sharing.
DRIVER_RENDER	Driver supports dedicated render nodes.
DRIVER_ATOMIC	Driver supports atomic properties. In this case the driver must implement appropriate obj->atomic_get_property() vfuncs for any modeset objects with driver specific properties.

Major, Minor and Patchlevel

```
int major;
int minor;
int patchlevel;
```

The DRM core identifies driver versions by a major, minor and patch level triplet. The information is printed to the kernel log at initialization time and passed to userspace through the DRM_IOCTL_VERSION ioctl.

The major and minor numbers are also used to verify the requested driver API version passed to DRM_IOCTL_SET_VERSION. When the driver API changes between minor versions, applications can call DRM_IOCTL_SET_VERSION to select a specific version of the API. If the requested major isn't equal to the driver major, or the requested minor is larger than the driver minor, the DRM_IOCTL_SET_VERSION call will return an error. Otherwise the driver's set_version() method will be called with the requested version.

Name, Description and Date

```
char *name;
char *desc;
char *date;
```

The driver name is printed to the kernel log at initialization time, used for IRQ registration and passed to userspace through DRM_IOCTL_VERSION.

The driver description is a purely informative string passed to userspace through the DRM_IOCTL_VERSION ioctl and otherwise unused by the kernel.

The driver date, formatted as YYYYMMDD, is meant to identify the date of the latest modification to the driver. However, as most drivers fail to update it, its value is mostly useless. The DRM core prints it to the kernel log at initialization time and passes it to userspace through the DRM_IOCTL_VERSION ioctl.

Device Instance and Driver Handling

A device instance for a drm driver is represented by struct drm_device. This is allocated with drm_dev_alloc, usually from bus-specific ->probe callbacks implemented by the driver. The driver then needs

to initialize all the various subsystems for the drm device like memory management, vblank handling, modesetting support and initial output configuration plus obviously initialize all the corresponding hardware bits. An important part of this is also calling `drm_dev_set_unique` to set the userspace-visible unique name of this device instance. Finally when everything is up and running and ready for userspace the device instance can be published using `drm_dev_register`.

There is also deprecated support for initializing device instances using bus-specific helpers and the `->load` callback. But due to backwards-compatibility needs the device instance have to be published too early, which requires unpretty global locking to make safe and is therefore only support for existing drivers not yet converted to the new scheme.

When cleaning up a device instance everything needs to be done in reverse: First unpublish the device instance with `drm_dev_unregister`. Then clean up any other resources allocated at device initialization and drop the driver's reference to `drm_device` using `drm_dev_unref`.

Note that the lifetime rules for `drm_device` instance has still a lot of historical baggage. Hence use the reference counting provided by `drm_dev_ref` and `drm_dev_unref` only carefully.

Also note that embedding of `drm_device` is currently not (yet) supported (but it would be easy to add). Drivers can store driver-private data in the `dev_priv` field of `drm_device`.

Name

`drm_put_dev` — Unregister and release a DRM device

Synopsis

```
void drm_put_dev (struct drm_device * dev);
```

Arguments

dev DRM device

Description

Called at module unload time or when a PCI device is unplugged.

Cleans up all DRM device, calling `drm_lastclose`.

Note

Use of this function is deprecated. It will eventually go away completely. Please use `drm_dev_unregister` and `drm_dev_unref` explicitly instead to make sure that the device isn't userspace accessible any more while teardown is in progress, ensuring that userspace can't access an inconsistent state.

Name

`drm_dev_alloc` — Allocate new DRM device

Synopsis

```
struct drm_device * drm_dev_alloc (struct drm_driver * driver, struct  
device * parent);
```

Arguments

driver DRM driver to allocate device for

parent Parent device object

Description

Allocate and initialize a new DRM device. No device registration is done. Call `drm_dev_register` to advertice the device to user space and register it with other core subsystems. This should be done last in the device initialization sequence to make sure userspace can't access an inconsistent state.

The initial ref-count of the object is 1. Use `drm_dev_ref` and `drm_dev_unref` to take and drop further ref-counts.

Note that for purely virtual devices *parent* can be NULL.

RETURNS

Pointer to new DRM device, or NULL if out of memory.

Name

`drm_dev_ref` — Take reference of a DRM device

Synopsis

```
void drm_dev_ref (struct drm_device * dev);
```

Arguments

dev device to take reference of or NULL

Description

This increases the ref-count of *dev* by one. You *must* already own a reference when calling this. Use `drm_dev_unref` to drop this reference again.

This function never fails. However, this function does not provide *any* guarantee whether the device is alive or running. It only provides a reference to the object and the memory associated with it.

Name

`drm_dev_unref` — Drop reference of a DRM device

Synopsis

```
void drm_dev_unref (struct drm_device * dev);
```

Arguments

dev device to drop reference of or NULL

Description

This decreases the ref-count of *dev* by one. The device is destroyed if the ref-count drops to zero.

Name

`drm_dev_register` — Register DRM device

Synopsis

```
int drm_dev_register (struct drm_device * dev, unsigned long flags);
```

Arguments

dev Device to register

flags Flags passed to the driver's `.load` function

Description

Register the DRM device *dev* with the system, advertise device to user-space and start normal device operation. *dev* must be allocated via `drm_dev_alloc` previously.

Never call this twice on any device!

NOTE

To ensure backward compatibility with existing drivers method this function calls the `->load` method after registering the device nodes, creating race conditions. Usage of the `->load` methods is therefore deprecated, drivers must perform all initialization before calling `drm_dev_register`.

RETURNS

0 on success, negative error code on failure.

Name

`drm_dev_unregister` — Unregister DRM device

Synopsis

```
void drm_dev_unregister (struct drm_device * dev);
```

Arguments

dev Device to unregister

Description

Unregister the DRM device from the system. This does the reverse of `drm_dev_register` but does not deallocate the device. The caller must call `drm_dev_unref` to drop their final reference.

This should be called first in the device teardown code to make sure userspace can't access the device instance any more.

Name

`drm_dev_set_unique` — Set the unique name of a DRM device

Synopsis

```
int drm_dev_set_unique (struct drm_device * dev, const char * name);
```

Arguments

dev device of which to set the unique name

name unique name

Description

Sets the unique name of a DRM device using the specified string. Drivers can use this at driver probe time if the unique name of the devices they drive is static.

Return

0 on success or a negative error code on failure.

Driver Load

IRQ Registration

The DRM core tries to facilitate IRQ handler registration and unregistration by providing `drm_irq_install` and `drm_irq_uninstall` functions. Those functions only support a single interrupt per device, devices that use more than one IRQs need to be handled manually.

Managed IRQ Registration

`drm_irq_install` starts by calling the `irq_preinstall` driver operation. The operation is optional and must make sure that the interrupt will not get fired by clearing all pending interrupt flags or disabling the interrupt.

The passed-in IRQ will then be requested by a call to `request_irq`. If the `DRIVER_IRQ_SHARED` driver feature flag is set, a shared (`IRQF_SHARED`) IRQ handler will be requested.

The IRQ handler function must be provided as the mandatory `irq_handler` driver operation. It will get passed directly to `request_irq` and thus has the same prototype as all IRQ handlers. It will get called with a pointer to the DRM device as the second argument.

Finally the function calls the optional `irq_postinstall` driver operation. The operation usually enables interrupts (excluding the vblank interrupt, which is enabled separately), but drivers may choose to enable/disable interrupts at a different time.

`drm_irq_uninstall` is similarly used to uninstall an IRQ handler. It starts by waking up all processes waiting on a vblank interrupt to make sure they don't hang, and then calls the optional `irq_uninstall` driver operation. The operation must disable all hardware interrupts. Finally the function frees the IRQ by calling `free_irq`.

Manual IRQ Registration

Drivers that require multiple interrupt handlers can't use the managed IRQ registration functions. In that case IRQs must be registered and unregistered manually (usually with the `request_irq` and `free_irq` functions, or their `devm_*` equivalent).

When manually registering IRQs, drivers must not set the `DRIVER_HAVE_IRQ` driver feature flag, and must not provide the `irq_handler` driver operation. They must set the `drm_device` `irq_enabled` field to 1 upon registration of the IRQs, and clear it to 0 after unregistering the IRQs.

Memory Manager Initialization

Every DRM driver requires a memory manager which must be initialized at load time. DRM currently contains two memory managers, the Translation Table Manager (TTM) and the Graphics Execution Manager (GEM). This document describes the use of the GEM memory manager only. See the section called “Memory management” for details.

Miscellaneous Device Configuration

Another task that may be necessary for PCI devices during configuration is mapping the video BIOS. On many devices, the VBIOS describes device configuration, LCD panel timings (if any), and contains flags indicating device state. Mapping the BIOS can be done using the `pci_map_rom()` call, a convenience function that takes care of mapping the actual ROM, whether it has been shadowed into memory (typically at address `0xc0000`) or exists on the PCI device in the ROM BAR. Note that after the ROM has been mapped and any necessary information has been extracted, it should be unmapped; on many devices, the ROM address decoder is shared with other BARs, so leaving it mapped could cause undesired behaviour like hangs or memory corruption.

Bus-specific Device Registration and PCI Support

A number of functions are provided to help with device registration. The functions deal with PCI and platform devices respectively and are only provided for historical reasons. These are all deprecated and shouldn't be used in new drivers. Besides that there's a few helpers for pci drivers.

Name

`drm_pci_alloc` — Allocate a PCI consistent memory block, for DMA.

Synopsis

```
drm_dma_handle_t * drm_pci_alloc (struct drm_device * dev, size_t size,  
size_t align);
```

Arguments

dev DRM device

size size of block to allocate

align alignment of block

Return

A handle to the allocated memory block on success or NULL on failure.

Name

`drm_pci_free` — Free a PCI consistent memory block

Synopsis

```
void drm_pci_free (struct drm_device * dev, drm_dma_handle_t * dma);
```

Arguments

dev DRM device

dma handle to memory block

Name

`drm_get_pci_dev` — Register a PCI device with the DRM subsystem

Synopsis

```
int drm_get_pci_dev (struct pci_dev * pdev, const struct pci_device_id  
* ent, struct drm_driver * driver);
```

Arguments

pdev PCI device

ent entry from the PCI ID table that matches *pdev*

driver DRM device driver

Description

Attempt to get inter module “drm” information. If we are first then register the character device and inter module information. Try and register, if we fail to register, backout previous work.

NOTE

This function is deprecated, please use `drm_dev_alloc` and `drm_dev_register` instead and remove your `->load` callback.

Return

0 on success or a negative error code on failure.

Name

`drm_pci_init` — Register matching PCI devices with the DRM subsystem

Synopsis

```
int drm_pci_init (struct drm_driver * driver, struct pci_driver * pdriver);
```

Arguments

driver DRM device driver

pdriver PCI device driver

Description

Initializes a `drm_device` structures, registering the stubs and initializing the AGP device.

NOTE

This function is deprecated. Modern modesetting drm drivers should use `pci_register_driver` directly, this function only provides shadow-binding support for old legacy drivers on top of that core pci function.

Return

0 on success or a negative error code on failure.

Name

`drm_pci_exit` — Unregister matching PCI devices from the DRM subsystem

Synopsis

```
void drm_pci_exit (struct drm_driver * driver, struct pci_driver *  
pdriver);
```

Arguments

driver DRM device driver

pdriver PCI device driver

Description

Unregisters one or more devices matched by a PCI driver from the DRM subsystem.

NOTE

This function is deprecated. Modern modesetting drm drivers should use `pci_unregister_driver` directly, this function only provides shadow-binding support for old legacy drivers on top of that core pci function.

Name

`drm_platform_init` — Register a platform device with the DRM subsystem

Synopsis

```
int drm_platform_init (struct drm_driver * driver, struct platform_device * platform_device);
```

Arguments

<i>driver</i>	DRM device driver
<i>platform_device</i>	platform device to register

Description

Registers the specified DRM device driver and platform device with the DRM subsystem, initializing a `drm_device` structure and calling the driver's `.load` function.

NOTE

This function is deprecated, please use `drm_dev_alloc` and `drm_dev_register` instead and remove your `->load` callback.

Return

0 on success or a negative error code on failure.

Memory management

Modern Linux systems require large amount of graphics memory to store frame buffers, textures, vertices and other graphics-related data. Given the very dynamic nature of many of that data, managing graphics memory efficiently is thus crucial for the graphics stack and plays a central role in the DRM infrastructure.

The DRM core includes two memory managers, namely Translation Table Maps (TTM) and Graphics Execution Manager (GEM). TTM was the first DRM memory manager to be developed and tried to be a one-size-fits-them all solution. It provides a single userspace API to accommodate the need of all hardware, supporting both Unified Memory Architecture (UMA) devices and devices with dedicated video RAM (i.e. most discrete video cards). This resulted in a large, complex piece of code that turned out to be hard to use for driver development.

GEM started as an Intel-sponsored project in reaction to TTM's complexity. Its design philosophy is completely different: instead of providing a solution to every graphics memory-related problems, GEM identified common code between drivers and created a support library to share it. GEM has simpler initialization and execution requirements than TTM, but has no video RAM management capabilities and is thus limited to UMA devices.

The Translation Table Manager (TTM)

TTM design background and information belongs here.

TTM initialization

Warning

This section is outdated.

Drivers wishing to support TTM must fill out a `drm_bo_driver` structure. The structure contains several fields with function pointers for initializing the TTM, allocating and freeing memory, waiting for command completion and fence synchronization, and memory migration. See the `radeon_ttm.c` file for an example of usage.

The `ttm_global_reference` structure is made up of several fields:

```
struct ttm_global_reference {
    enum ttm_global_types global_type;
    size_t size;
    void *object;
    int (*init) (struct ttm_global_reference *);
    void (*release) (struct ttm_global_reference *);
};
```

There should be one global reference structure for your memory manager as a whole, and there will be others for each object created by the memory manager at runtime. Your global TTM should have a type of `TTM_GLOBAL_TTM_MEM`. The size field for the global object should be `sizeof(struct ttm_mem_global)`, and the init and release hooks should point at your driver-specific init and release routines, which probably eventually call `ttm_mem_global_init` and `ttm_mem_global_release`, respectively.

Once your global TTM accounting structure is set up and initialized by calling `ttm_global_item_ref()` on it, you need to create a buffer object TTM to provide a pool for buffer object allocation by clients and the kernel itself. The type of this object should be `TTM_GLOBAL_TTM_BO`, and its size should be `sizeof(struct ttm_bo_global)`. Again, driver-specific init and release functions may be provided, likely eventually calling `ttm_bo_global_init()` and `ttm_bo_global_release()`, respectively. Also, like the previous object, `ttm_global_item_ref()` is used to create an initial reference count for the TTM, which will call your initialization function.

The Graphics Execution Manager (GEM)

The GEM design approach has resulted in a memory manager that doesn't provide full coverage of all (or even all common) use cases in its userspace or kernel API. GEM exposes a set of standard memory-related operations to userspace and a set of helper functions to drivers, and let drivers implement hardware-specific operations with their own private API.

The GEM userspace API is described in the *GEM - the Graphics Execution Manager* [<http://lwn.net/Articles/283798/>] article on LWN. While slightly outdated, the document provides a good overview of the GEM API principles. Buffer allocation and read and write operations, described as part of the common GEM API, are currently implemented using driver-specific ioctls.

GEM is data-agnostic. It manages abstract buffer objects without knowing what individual buffers contain. APIs that require knowledge of buffer contents or purpose, such as buffer allocation or synchronization primitives, are thus outside of the scope of GEM and must be implemented using driver-specific ioctls.

On a fundamental level, GEM involves several operations:

- Memory allocation and freeing
- Command execution
- Aperture management at command execution time

Buffer object allocation is relatively straightforward and largely provided by Linux's shmem layer, which provides memory to back each object.

Device-specific operations, such as command execution, pinning, buffer read & write, mapping, and domain ownership transfers are left to driver-specific ioctls.

GEM Initialization

Drivers that use GEM must set the `DRIVER_GEM` bit in the struct `drm_driver` *driver_features* field. The DRM core will then automatically initialize the GEM core before calling the `load` operation. Behind the scene, this will create a DRM Memory Manager object which provides an address space pool for object allocation.

In a KMS configuration, drivers need to allocate and initialize a command ring buffer following core GEM initialization if required by the hardware. UMA devices usually have what is called a "stolen" memory region, which provides space for the initial framebuffer and large, contiguous memory regions required by the device. This space is typically not managed by GEM, and must be initialized separately into its own DRM MM object.

GEM Objects Creation

GEM splits creation of GEM objects and allocation of the memory that backs them in two distinct operations.

GEM objects are represented by an instance of struct `drm_gem_object`. Drivers usually need to extend GEM objects with private information and thus create a driver-specific GEM object structure type that embeds an instance of struct `drm_gem_object`.

To create a GEM object, a driver allocates memory for an instance of its specific GEM object type and initializes the embedded struct `drm_gem_object` with a call to `drm_gem_object_init`. The function takes a pointer to the DRM device, a pointer to the GEM object and the buffer object size in bytes.

GEM uses shmem to allocate anonymous pageable memory. `drm_gem_object_init` will create an shmf file of the requested size and store it into the struct `drm_gem_object` *filp* field. The memory is used as either main storage for the object when the graphics hardware uses system memory directly or as a backing store otherwise.

Drivers are responsible for the actual physical pages allocation by calling `shmem_read_mapping_page_gfp` for each page. Note that they can decide to allocate pages when initializing the GEM object, or to delay allocation until the memory is needed (for instance when a page fault occurs as a result of a userspace memory access or when the driver needs to start a DMA transfer involving the memory).

Anonymous pageable memory allocation is not always desired, for instance when the hardware requires physically contiguous system memory as is often the case in embedded devices. Drivers can create GEM objects with no shmf backing (called private GEM objects) by initializing them with a call to `drm_gem_private_object_init` instead of `drm_gem_object_init`. Storage for private GEM objects must be managed by drivers.

GEM Objects Lifetime

All GEM objects are reference-counted by the GEM core. References can be acquired and release by calling `drm_gem_object_reference` and `drm_gem_object_unreference` respectively. The caller must hold the `drm_device` *struct_mutex* lock when calling `drm_gem_object_reference`. As a convenience, GEM provides `drm_gem_object_unreference_unlocked` functions that can be called without holding the lock.

When the last reference to a GEM object is released the GEM core calls the `drm_driver.gem_free_object` operation. That operation is mandatory for GEM-enabled drivers and must free the GEM object and all associated resources.

```
void (*gem_free_object) (struct drm_gem_object *obj);
```

Drivers are responsible for freeing all GEM object resources. This includes the resources created by the GEM core, which need to be released with `drm_gem_object_release`.

GEM Objects Naming

Communication between userspace and the kernel refers to GEM objects using local handles, global names or, more recently, file descriptors. All of those are 32-bit integer values; the usual Linux kernel limits apply to the file descriptors.

GEM handles are local to a DRM file. Applications get a handle to a GEM object through a driver-specific `ioctl`, and can use that handle to refer to the GEM object in other standard or driver-specific `ioctls`. Closing a DRM file handle frees all its GEM handles and dereferences the associated GEM objects.

To create a handle for a GEM object drivers call `drm_gem_handle_create`. The function takes a pointer to the DRM file and the GEM object and returns a locally unique handle. When the handle is no longer needed drivers delete it with a call to `drm_gem_handle_delete`. Finally the GEM object associated with a handle can be retrieved by a call to `drm_gem_object_lookup`.

Handles don't take ownership of GEM objects, they only take a reference to the object that will be dropped when the handle is destroyed. To avoid leaking GEM objects, drivers must make sure they drop the reference(s) they own (such as the initial reference taken at object creation time) as appropriate, without any special consideration for the handle. For example, in the particular case of combined GEM object and handle creation in the implementation of the `dumb_create` operation, drivers must drop the initial reference to the GEM object before returning the handle.

GEM names are similar in purpose to handles but are not local to DRM files. They can be passed between processes to reference a GEM object globally. Names can't be used directly to refer to objects in the DRM API, applications must convert handles to names and names to handles using the `DRM_IOCTL_GEM_FLINK` and `DRM_IOCTL_GEM_OPEN` `ioctls` respectively. The conversion is handled by the DRM core without any driver-specific support.

GEM also supports buffer sharing with `dma-buf` file descriptors through `PRIME`. GEM-based drivers must use the provided helpers functions to implement the exporting and importing correctly. See the section called “`PRIME` Buffer Sharing”. Since sharing file descriptors is inherently more secure than the easily guessable and global GEM names it is the preferred buffer sharing mechanism. Sharing buffers through GEM names is only supported for legacy userspace. Furthermore `PRIME` also allows cross-device buffer sharing since it is based on `dma-bufs`.

GEM Objects Mapping

Because mapping operations are fairly heavyweight GEM favours read/write-like access to buffers, implemented through driver-specific `ioctls`, over mapping buffers to userspace. However, when random access to the buffer is needed (to perform software rendering for instance), direct access to the object can be more efficient.

The `mmap` system call can't be used directly to map GEM objects, as they don't have their own file handle. Two alternative methods currently co-exist to map GEM objects to userspace. The first method uses a driver-specific `ioctl` to perform the mapping operation, calling `do_mmap` under the hood. This is often considered dubious, seems to be discouraged for new GEM-enabled drivers, and will thus not be described here.

The second method uses the `mmap` system call on the DRM file handle.

```
void *mmap(void *addr, size_t length, int prot, int flags, int fd,
           off_t offset);
```

DRM identifies the GEM object to be mapped by a fake offset passed through the `mmap` offset argument. Prior to being mapped, a GEM object must thus be associated with a fake offset. To do so, drivers must call `drm_gem_create_mmap_offset` on the object.

Once allocated, the fake offset value must be passed to the application in a driver-specific way and can then be used as the `mmap` offset argument.

The GEM core provides a helper method `drm_gem_mmap` to handle object mapping. The method can be set directly as the `mmap` file operation handler. It will look up the GEM object based on the offset value and set the VMA operations to the `drm_driver gem_vm_ops` field. Note that `drm_gem_mmap` doesn't map memory to userspace, but relies on the driver-provided fault handler to map pages individually.

To use `drm_gem_mmap`, drivers must fill the struct `drm_driver gem_vm_ops` field with a pointer to VM operations.

```
struct vm_operations_struct *gem_vm_ops

struct vm_operations_struct {
    void (*open)(struct vm_area_struct * area);
    void (*close)(struct vm_area_struct * area);
    int (*fault)(struct vm_area_struct *vma, struct vm_fault *vmf);
};
```

The `open` and `close` operations must update the GEM object reference count. Drivers can use the `drm_gem_vm_open` and `drm_gem_vm_close` helper functions directly as `open` and `close` handlers.

The fault operation handler is responsible for mapping individual pages to userspace when a page fault occurs. Depending on the memory allocation scheme, drivers can allocate pages at fault time, or can decide to allocate memory for the GEM object at the time the object is created.

Drivers that want to map the GEM object upfront instead of handling page faults can implement their own `mmap` file operation handler.

Memory Coherency

When mapped to the device or used in a command buffer, backing pages for an object are flushed to memory and marked write combined so as to be coherent with the GPU. Likewise, if the CPU accesses an object after the GPU has finished rendering to the object, then the object must be made coherent with the CPU's view of memory, usually involving GPU cache flushing of various kinds. This core CPU<->GPU coherency management is provided by a device-specific ioctl, which evaluates an object's current domain and performs any necessary flushing or synchronization to put the object into the desired coherency domain (note that the object may be busy, i.e. an active render target; in that case, setting the domain blocks the client and waits for rendering to complete before performing any necessary flushing operations).

Command Execution

Perhaps the most important GEM function for GPU devices is providing a command execution interface to clients. Client programs construct command buffers containing references to previously allocated memory objects, and then submit them to GEM. At that point, GEM takes care to bind all the objects into the GTT, execute the buffer, and provide necessary synchronization between clients accessing the same buffers. This

often involves evicting some objects from the GTT and re-binding others (a fairly expensive operation), and providing relocation support which hides fixed GTT offsets from clients. Clients must take care not to submit command buffers that reference more objects than can fit in the GTT; otherwise, GEM will reject them and no rendering will occur. Similarly, if several objects in the buffer require fence registers to be allocated for correct rendering (e.g. 2D blits on pre-965 chips), care must be taken not to require more fence registers than are available to the client. Such resource management should be abstracted from the client in libdrm.

GEM Function Reference

Name

`drm_gem_object_init` — initialize an allocated shmem-backed GEM object

Synopsis

```
int drm_gem_object_init (struct drm_device * dev, struct drm_gem_object  
* obj, size_t size);
```

Arguments

dev `drm_device` the object should be initialized for

obj `drm_gem_object` to initialize

size object size

Description

Initialize an already allocated GEM object of the specified size with shmf's backing store.

Name

`drm_gem_private_object_init` — initialize an allocated private GEM object

Synopsis

```
void drm_gem_private_object_init (struct drm_device * dev, struct dr-  
m_gem_object * obj, size_t size);
```

Arguments

dev drm_device the object should be initialized for

obj drm_gem_object to initialize

size object size

Description

Initialize an already allocated GEM object of the specified size with no GEM provided backing store. Instead the caller is responsible for backing the object and handling it.

Name

`drm_gem_handle_delete` — deletes the given file-private handle

Synopsis

```
int drm_gem_handle_delete (struct drm_file * filp, u32 handle);
```

Arguments

filp drm file-private structure to use for the handle look up

handle userspace handle to delete

Description

Removes the GEM handle from the *filp* lookup table which has been added with `drm_gem_handle_create`. If this is the last handle also cleans up linked resources like GEM names.

Name

`drm_gem_dumb_destroy` — dumb fb callback helper for gem based drivers

Synopsis

```
int drm_gem_dumb_destroy (struct drm_file * file, struct drm_device *  
dev, uint32_t handle);
```

Arguments

file drm file-private structure to remove the dumb handle from

dev corresponding `drm_device`

handle the dumb handle to remove

Description

This implements the `->dumb_destroy` kms driver callback for drivers which use gem to manage their backing storage.

Name

`drm_gem_handle_create` — create a gem handle for an object

Synopsis

```
int drm_gem_handle_create (struct drm_file * file_priv, struct dr-  
m_gem_object * obj, u32 * handlep);
```

Arguments

file_priv drm file-private structure to register the handle for

obj object to register

handlep pointer to return the created handle to the caller

Description

Create a handle for this object. This adds a handle reference to the object, which includes a regular reference count. Callers will likely want to dereference the object afterwards.

Name

`drm_gem_free_mmap_offset` — release a fake mmap offset for an object

Synopsis

```
void drm_gem_free_mmap_offset (struct drm_gem_object * obj);
```

Arguments

obj obj in question

Description

This routine frees fake offsets allocated by `drm_gem_create_mmap_offset`.

Name

`drm_gem_create_mmap_offset_size` — create a fake mmap offset for an object

Synopsis

```
int drm_gem_create_mmap_offset_size (struct drm_gem_object * obj, size_t
size);
```

Arguments

obj obj in question

size the virtual size

Description

GEM memory mapping works by handing back to userspace a fake mmap offset it can use in a subsequent `mmap(2)` call. The DRM core code then looks up the object based on the offset and sets up the various memory mapping structures.

This routine allocates and attaches a fake offset for *obj*, in cases where the virtual size differs from the physical size (ie. `obj->size`). Otherwise just use `drm_gem_create_mmap_offset`.

Name

`drm_gem_create_mmap_offset` — create a fake mmap offset for an object

Synopsis

```
int drm_gem_create_mmap_offset (struct drm_gem_object * obj);
```

Arguments

obj obj in question

Description

GEM memory mapping works by handing back to userspace a fake mmap offset it can use in a subsequent `mmap(2)` call. The DRM core code then looks up the object based on the offset and sets up the various memory mapping structures.

This routine allocates and attaches a fake offset for *obj*.

Name

`drm_gem_get_pages` — helper to allocate backing pages for a GEM object from shmem

Synopsis

```
struct page ** drm_gem_get_pages (struct drm_gem_object * obj);
```

Arguments

obj obj in question

Description

This reads the page-array of the shmem-backing storage of the given gem object. An array of pages is returned. If a page is not allocated or swapped-out, this will allocate/swap-in the required pages. Note that the whole object is covered by the page-array and pinned in memory.

Use `drm_gem_put_pages` to release the array and unpin all pages.

This uses the GFP-mask set on the shmem-mapping (see `mapping_set_gfp_mask`). If you require other GFP-masks, you have to do those allocations yourself.

Note that you are not allowed to change gfp-zones during runtime. That is, `shmem_read_mapping_page_gfp` must be called with the same `gfp_zone(gfp)` as set during initialization. If you have special zone constraints, set them after `drm_gem_init_object` via `mapping_set_gfp_mask`. shmem-core takes care to keep pages in the required zone during swap-in.

Name

`drm_gem_put_pages` — helper to free backing pages for a GEM object

Synopsis

```
void drm_gem_put_pages (struct drm_gem_object * obj, struct page **  
pages, bool dirty, bool accessed);
```

Arguments

<i>obj</i>	obj in question
<i>pages</i>	pages to free
<i>dirty</i>	if true, pages will be marked as dirty
<i>accessed</i>	if true, the pages will be marked as accessed

Name

`drm_gem_object_lookup` — look up a GEM object from it's handle

Synopsis

```
struct drm_gem_object * drm_gem_object_lookup (struct drm_device * dev,  
struct drm_file * filp, u32 handle);
```

Arguments

dev DRM device

filp DRM file private data

handle userspace handle

Returns

A reference to the object named by the handle if such exists on *filp*, NULL otherwise.

Name

`drm_gem_object_free` — free a GEM object

Synopsis

```
void drm_gem_object_free (struct kref * kref);
```

Arguments

kref kref of the object to free

Description

Called after the last reference to the object has been lost. Must be called holding `struct_mutex`

Frees the object

Name

`drm_gem_vm_open` — `vma->ops->open` implementation for GEM

Synopsis

```
void drm_gem_vm_open (struct vm_area_struct * vma);
```

Arguments

vma VM area structure

Description

This function implements the `#vm_operations_struct` `open` callback for GEM drivers. This must be used together with `drm_gem_vm_close`.

Name

`drm_gem_vm_close` — `vma->ops->close` implementation for GEM

Synopsis

```
void drm_gem_vm_close (struct vm_area_struct * vma);
```

Arguments

vma VM area structure

Description

This function implements the `#vm_operations_struct.close` callback for GEM drivers. This must be used together with `drm_gem_vm_open`.

Name

`drm_gem_mmap_obj` — memory map a GEM object

Synopsis

```
int drm_gem_mmap_obj (struct drm_gem_object * obj, unsigned long obj_size, struct vm_area_struct * vma);
```

Arguments

obj the GEM object to map

obj_size the object size to be mapped, in bytes

vma VMA for the area to be mapped

Description

Set up the VMA to prepare mapping of the GEM object using the `gem_vm_ops` provided by the driver. Depending on their requirements, drivers can either provide a fault handler in their `gem_vm_ops` (in which case any accesses to the object will be trapped, to perform migration, GTT binding, surface register allocation, or performance monitoring), or `mmap` the buffer memory synchronously after calling `drm_gem_mmap_obj`.

This function is mainly intended to implement the DMABUF `mmap` operation, when the GEM object is not looked up based on its fake offset. To implement the DRM `mmap` operation, drivers should use the `drm_gem_mmap` function.

`drm_gem_mmap_obj` assumes the user is granted access to the buffer while `drm_gem_mmap` prevents unprivileged users from mapping random objects. So callers must verify access restrictions before calling this helper.

Return 0 on success or `-EINVAL` if the object size is smaller than the VMA size, or if no `gem_vm_ops` are provided.

Name

`drm_gem_mmap` — memory map routine for GEM objects

Synopsis

```
int drm_gem_mmap (struct file * filp, struct vm_area_struct * vma);
```

Arguments

filp DRM file pointer

vma VMA for the area to be mapped

Description

If a driver supports GEM object mapping, mmap calls on the DRM file descriptor will end up here.

Look up the GEM object based on the offset passed in (`vma->vm_pgoff` will contain the fake offset we created when the GTT map ioctl was called on the object) and map it with a call to `drm_gem_mmap_obj`.

If the caller is not granted access to the buffer object, the mmap will fail with `EACCES`. Please see the vma manager for more information.

Name

struct `drm_gem_object` — GEM buffer object

Synopsis

```
struct drm_gem_object {
    struct kref refcount;
    unsigned handle_count;
    struct drm_device * dev;
    struct file * filp;
    struct drm_vma_offset_node vma_node;
    size_t size;
    int name;
    uint32_t read_domains;
    uint32_t write_domain;
    uint32_t pending_read_domains;
    uint32_t pending_write_domain;
    struct dma_buf * dma_buf;
    struct dma_buf_attachment * import_attach;
};
```

Members

`refcount`

Reference count of this object

Please use `drm_gem_object_reference` to acquire and `drm_gem_object_unreference` or `drm_gem_object_unreference_unlocked` to release a reference to a GEM buffer object.

`handle_count`

This is the GEM `file_priv` handle count of this object.

Each handle also holds a reference. Note that when the `handle_count` drops to 0 any global names (e.g. the id in the flink namespace) will be cleared.

Protected by `dev->object_name_lock`.

`dev`

DRM dev this object belongs to.

`filp`

SHMEM file node used as backing storage for swappable buffer objects. GEM also supports driver private objects with driver-specific backing storage (contiguous CMA memory, special reserved blocks). In this case `filp` is NULL.

`vma_node`

Mapping info for this object to support mmap. Drivers are supposed to allocate the mmap offset using `drm_gem_create_mmap_offset`. The offset itself can be retrieved using `drm_vma_node_offset_addr`.

Memory mapping itself is handled by `drm_gem_mmap`, which also checks that userspace is allowed to access the object.

size	Size of the object, in bytes. Immutable over the object's lifetime.
name	Global name for this object, starts at 1. 0 means unnamed. Access is covered by dev->object_name_lock. This is used by the GEM_FLINK and GEM_OPEN ioctls.
read_domains	Read memory domains. These monitor which caches contain read/write data related to the object. When transitioning from one set of domains to another, the driver is called to ensure that caches are suitably flushed and invalidated.
write_domain	Corresponding unique write memory domain.
pending_read_domains	While validating an exec operation, the new read/write domain values are computed here. They will be transferred to the above values at the point that any cache flushing occurs
pending_write_domain	Write domain similar to <i>pending_read_domains</i> .
dma_buf	<p>dma-buf associated with this GEM object.</p> <p>Pointer to the dma-buf associated with this gem object (either through importing or exporting). We break the resulting reference loop when the last gem handle for this object is released.</p> <p>Protected by obj->object_name_lock.</p>
import_attach	<p>dma-buf attachment backing this object.</p> <p>Any foreign dma_buf imported as a gem object has this set to the attachment point for the device. This is invariant over the lifetime of a gem object.</p> <p>The driver's ->gem_free_object callback is responsible for cleaning up the dma_buf attachment and references acquired at import time.</p> <p>Note that the drm gem/prime core does not depend upon drivers setting this field any more. So for drivers where this doesn't make sense (e.g. virtual devices or a displaylink behind an usb bus) they can simply leave it as NULL.</p>

Description

This structure defines the generic parts for GEM buffer objects, which are mostly around handling mmap and userspace handles.

Buffer objects are often abbreviated to BO.

Name

`drm_gem_object_reference` — acquire a GEM BO reference

Synopsis

```
void drm_gem_object_reference (struct drm_gem_object * obj);
```

Arguments

obj GEM buffer object

Description

This acquires additional reference to *obj*. It is illegal to call this without already holding a reference. No locks required.

Name

`drm_gem_object_unreference` — release a GEM BO reference

Synopsis

```
void drm_gem_object_unreference (struct drm_gem_object * obj);
```

Arguments

obj GEM buffer object

Description

This releases a reference to *obj*. Callers must hold the `dev->struct_mutex` lock when calling this function, even when the driver doesn't use `dev->struct_mutex` for anything.

For drivers not encumbered with legacy locking use `drm_gem_object_unreference_unlocked` instead.

Name

`drm_gem_object_unreference_unlocked` — release a GEM BO reference

Synopsis

```
void drm_gem_object_unreference_unlocked (struct drm_gem_object * obj);
```

Arguments

obj GEM buffer object

Description

This releases a reference to *obj*. Callers must not hold the `dev->struct_mutex` lock when calling this function.

VMA Offset Manager

The vma-manager is responsible to map arbitrary driver-dependent memory regions into the linear user address-space. It provides offsets to the caller which can then be used on the `address_space` of the drm-device. It takes care to not overlap regions, size them appropriately and to not confuse mm-core by inconsistent fake `vm_pgoff` fields. Drivers shouldn't use this for object placement in VMEM. This manager should only be used to manage mappings into linear user-space VMs.

We use `drm_mm` as backend to manage object allocations. But it is highly optimized for alloc/free calls, not lookups. Hence, we use an rb-tree to speed up offset lookups.

You must not use multiple offset managers on a single `address_space`. Otherwise, mm-core will be unable to tear down memory mappings as the VM will no longer be linear.

This offset manager works on page-based addresses. That is, every argument and return code (with the exception of `drm_vma_node_offset_addr`) is given in number of pages, not number of bytes. That means, object sizes and offsets must always be page-aligned (as usual). If you want to get a valid byte-based user-space address for a given offset, please see `drm_vma_node_offset_addr`.

Additionally to offset management, the vma offset manager also handles access management. For every open-file context that is allowed to access a given node, you must call `drm_vma_node_allow`. Otherwise, an `mmap` call on this open-file with the offset of the node will fail with `-EACCES`. To revoke access again, use `drm_vma_node_revoke`. However, the caller is responsible for destroying already existing mappings, if required.

Name

`drm_vma_offset_manager_init` — Initialize new offset-manager

Synopsis

```
void drm_vma_offset_manager_init (struct drm_vma_offset_manager * mgr,  
unsigned long page_offset, unsigned long size);
```

Arguments

<i>mgr</i>	Manager object
<i>page_offset</i>	Offset of available memory area (page-based)
<i>size</i>	Size of available address space range (page-based)

Description

Initialize a new offset-manager. The offset and area size available for the manager are given as *page_offset* and *size*. Both are interpreted as page-numbers, not bytes.

Adding/removing nodes from the manager is locked internally and protected against concurrent access. However, node allocation and destruction is left for the caller. While calling into the vma-manager, a given node must always be guaranteed to be referenced.

Name

`drm_vma_offset_manager_destroy` — Destroy offset manager

Synopsis

```
void drm_vma_offset_manager_destroy (struct drm_vma_offset_manager *  
mgr);
```

Arguments

mgr Manager object

Description

Destroy an object manager which was previously created via `drm_vma_offset_manager_init`. The caller must remove all allocated nodes before destroying the manager. Otherwise, `drm_mm` will refuse to free the requested resources.

The manager must not be accessed after this function is called.

Name

`drm_vma_offset_lookup_locked` — Find node in offset space

Synopsis

```
struct drm_vma_offset_node * drm_vma_offset_lookup_locked (struct dr-  
m_vma_offset_manager * mgr, unsigned long start, unsigned long pages);
```

Arguments

mgr Manager object

start Start address for object (page-based)

pages Size of object (page-based)

Description

Find a node given a start address and object size. This returns the `_best_` match for the given node. That is, *start* may point somewhere into a valid region and the given node will be returned, as long as the node spans the whole requested area (given the size in number of pages as *pages*).

Note that before lookup the vma offset manager lookup lock must be acquired with `drm_vma_offset_lock_lookup`. See there for an example. This can then be used to implement weakly referenced lookups using `kref_get_unless_zero`.

Example

```
drm_vma_offset_lock_lookup(mgr);  
node = drm_vma_offset_lookup_locked(mgr);  
if (node)  
    kref_get_unless_zero(container_of(node, sth, entr));  
drm_vma_offset_unlock_lookup(mgr);
```

RETURNS

Returns NULL if no suitable node can be found. Otherwise, the best match is returned. It's the caller's responsibility to make sure the node doesn't get destroyed before the caller can access it.

Name

`drm_vma_offset_add` — Add offset node to manager

Synopsis

```
int drm_vma_offset_add (struct drm_vma_offset_manager * mgr, struct
drm_vma_offset_node * node, unsigned long pages);
```

Arguments

mgr Manager object

node Node to be added

pages Allocation size visible to user-space (in number of pages)

Description

Add a node to the offset-manager. If the node was already added, this does nothing and return 0. *pages* is the size of the object given in number of pages. After this call succeeds, you can access the offset of the node until it is removed again.

If this call fails, it is safe to retry the operation or call `drm_vma_offset_remove`, anyway. However, no cleanup is required in that case.

pages is not required to be the same size as the underlying memory object that you want to map. It only limits the size that user-space can map into their address space.

RETURNS

0 on success, negative error code on failure.

Name

`drm_vma_offset_remove` — Remove offset node from manager

Synopsis

```
void drm_vma_offset_remove (struct drm_vma_offset_manager * mgr, struct  
drm_vma_offset_node * node);
```

Arguments

mgr Manager object

node Node to be removed

Description

Remove a node from the offset manager. If the node wasn't added before, this does nothing. After this call returns, the offset and size will be 0 until a new offset is allocated via `drm_vma_offset_add` again. Helper functions like `drm_vma_node_start` and `drm_vma_node_offset_addr` will return 0 if no offset is allocated.

Name

`drm_vma_node_allow` — Add open-file to list of allowed users

Synopsis

```
int drm_vma_node_allow (struct drm_vma_offset_node * node, struct file  
* filp);
```

Arguments

node Node to modify

filp Open file to add

Description

Add *filp* to the list of allowed open-files for this node. If *filp* is already on this list, the ref-count is incremented.

The list of allowed-users is preserved across `drm_vma_offset_add` and `drm_vma_offset_remove` calls. You may even call it if the node is currently not added to any offset-manager.

You must remove all open-files the same number of times as you added them before destroying the node. Otherwise, you will leak memory.

This is locked against concurrent access internally.

RETURNS

0 on success, negative error code on internal failure (out-of-mem)

Name

`drm_vma_node_revoke` — Remove open-file from list of allowed users

Synopsis

```
void drm_vma_node_revoke (struct drm_vma_offset_node * node, struct file  
* filp);
```

Arguments

node Node to modify

filp Open file to remove

Description

Decrement the ref-count of *filp* in the list of allowed open-files on *node*. If the ref-count drops to zero, remove *filp* from the list. You must call this once for every `drm_vma_node_allow` on *filp*.

This is locked against concurrent access internally.

If *filp* is not on the list, nothing is done.

Name

`drm_vma_node_is_allowed` — Check whether an open-file is granted access

Synopsis

```
bool drm_vma_node_is_allowed (struct drm_vma_offset_node * node, struct  
file * filp);
```

Arguments

node Node to check

filp Open-file to check for

Description

Search the list in *node* whether *filp* is currently on the list of allowed open-files (see `drm_vma_node_allow`).

This is locked against concurrent access internally.

RETURNS

true iff *filp* is on the list

Name

`drm_vma_offset_exact_lookup_locked` — Look up node by exact address

Synopsis

```
struct drm_vma_offset_node * drm_vma_offset_exact_lookup_locked (struct
drm_vma_offset_manager * mgr, unsigned long start, unsigned long pages);
```

Arguments

mgr Manager object

start Start address (page-based, not byte-based)

pages Size of object (page-based)

Description

Same as `drm_vma_offset_lookup_locked` but does not allow any offset into the node. It only returns the exact object with the given start address.

RETURNS

Node at exact start address *start*.

Name

`drm_vma_offset_lock_lookup` — Lock lookup for extended private use

Synopsis

```
void drm_vma_offset_lock_lookup (struct drm_vma_offset_manager * mgr);
```

Arguments

mgr Manager object

Description

Lock VMA manager for extended lookups. Only locked VMA function calls are allowed while holding this lock. All other contexts are blocked from VMA until the lock is released via `drm_vma_offset_unlock_lookup`.

Use this if you need to take a reference to the objects returned by `drm_vma_offset_lookup_locked` before releasing this lock again.

This lock must not be used for anything else than extended lookups. You must not call any other VMA helpers while holding this lock.

Note

You're in atomic-context while holding this lock!

Name

`drm_vma_offset_unlock_lookup` — Unlock lookup for extended private use

Synopsis

```
void drm_vma_offset_unlock_lookup (struct drm_vma_offset_manager *  
mgr);
```

Arguments

mgr Manager object

Description

Release lookup-lock. See `drm_vma_offset_lock_lookup` for more information.

Name

`drm_vma_node_reset` — Initialize or reset node object

Synopsis

```
void drm_vma_node_reset (struct drm_vma_offset_node * node);
```

Arguments

node Node to initialize or reset

Description

Reset a node to its initial state. This must be called before using it with any VMA offset manager.

This must not be called on an already allocated node, or you will leak memory.

Name

`drm_vma_node_start` — Return start address for page-based addressing

Synopsis

```
unsigned long drm_vma_node_start (struct drm_vma_offset_node * node);
```

Arguments

node Node to inspect

Description

Return the start address of the given node. This can be used as offset into the linear VM space that is provided by the VMA offset manager. Note that this can only be used for page-based addressing. If you need a proper offset for user-space mappings, you must apply “<< PAGE_SHIFT” or use the `drm_vma_node_offset_addr` helper instead.

RETURNS

Start address of *node* for page-based addressing. 0 if the node does not have an offset allocated.

Name

`drm_vma_node_size` — Return size (page-based)

Synopsis

```
unsigned long drm_vma_node_size (struct drm_vma_offset_node * node);
```

Arguments

node Node to inspect

Description

Return the size as number of pages for the given node. This is the same size that was passed to `drm_vma_offset_add`. If no offset is allocated for the node, this is 0.

RETURNS

Size of *node* as number of pages. 0 if the node does not have an offset allocated.

Name

`drm_vma_node_has_offset` — Check whether node is added to offset manager

Synopsis

```
bool drm_vma_node_has_offset (struct drm_vma_offset_node * node);
```

Arguments

node Node to be checked

RETURNS

true iff the node was previously allocated an offset and added to an vma offset manager.

Name

`drm_vma_node_offset_addr` — Return sanitized offset for user-space mmaps

Synopsis

```
__u64 drm_vma_node_offset_addr (struct drm_vma_offset_node * node);
```

Arguments

node Linked offset node

Description

Same as `drm_vma_node_start` but returns the address as a valid offset that can be used for user-space mappings during `mmap`. This must not be called on unlinked nodes.

RETURNS

Offset of *node* for byte-based addressing. 0 if the node does not have an object allocated.

Name

`drm_vma_node_unmap` — Unmap offset node

Synopsis

```
void drm_vma_node_unmap (struct drm_vma_offset_node * node, struct address_space * file_mapping);
```

Arguments

node Offset node

file_mapping Address space to unmap *node* from

Description

Unmap all userspace mappings for a given offset node. The mappings must be associated with the *file_mapping* address-space. If no offset exists nothing is done.

This call is unlocked. The caller must guarantee that `drm_vma_offset_remove` is not called on this node concurrently.

Name

`drm_vma_node_verify_access` — Access verification helper for TTM

Synopsis

```
int drm_vma_node_verify_access (struct drm_vma_offset_node * node,
struct file * filp);
```

Arguments

node Offset node

filp Open-file

Description

This checks whether *filp* is granted access to *node*. It is the same as `drm_vma_node_is_allowed` but suitable as drop-in helper for TTM `verify_access` callbacks.

RETURNS

0 if access is granted, -EACCES otherwise.

PRIME Buffer Sharing

PRIME is the cross device buffer sharing framework in drm, originally created for the OPTIMUS range of multi-gpu platforms. To userspace PRIME buffers are dma-buf based file descriptors.

Overview and Driver Interface

Similar to GEM global names, PRIME file descriptors are also used to share buffer objects across processes. They offer additional security: as file descriptors must be explicitly sent over UNIX domain sockets to be shared between applications, they can't be guessed like the globally unique GEM names.

Drivers that support the PRIME API must set the `DRIVER_PRIME` bit in the struct `drm_driver` *driver_features* field, and implement the `prime_handle_to_fd` and `prime_fd_to_handle` operations.

```
int (*prime_handle_to_fd)(struct drm_device *dev,
                          struct drm_file *file_priv, uint32_t handle,
                          uint32_t flags, int *prime_fd);
int (*prime_fd_to_handle)(struct drm_device *dev,
                          struct drm_file *file_priv, int prime_fd,
                          uint32_t *handle);
```

Those two operations convert a handle to a PRIME file descriptor and vice versa. Drivers must use the kernel dma-buf buffer sharing framework to manage the PRIME file descriptors. Similar to the mode setting API PRIME is agnostic to the underlying buffer object manager, as long as handles are 32bit unsigned integers.

While non-GEM drivers must implement the operations themselves, GEM drivers must use the `drm_gem_prime_handle_to_fd` and `drm_gem_prime_fd_to_handle` helper functions. Those helpers rely on the driver `gem_prime_export` and `gem_prime_import` operations to create a dma-buf instance from a GEM object (dma-buf exporter role) and to create a GEM object from a dma-buf instance (dma-buf importer role).

```
struct dma_buf * (*gem_prime_export)(struct drm_device *dev,  
                                     struct drm_gem_object *obj,  
                                     int flags);  
struct drm_gem_object * (*gem_prime_import)(struct drm_device *dev,  
                                             struct dma_buf *dma_buf);
```

These two operations are mandatory for GEM drivers that support PRIME.

PRIME Helper Functions

Drivers can implement *gem_prime_export* and *gem_prime_import* in terms of simpler APIs by using the helper functions *drm_gem_prime_export* and *drm_gem_prime_import*. These functions implement dma-buf support in terms of six lower-level driver callbacks:

Export callbacks:

* *gem_prime_pin* (optional): prepare a GEM object for exporting
* *gem_prime_get_sg_table*: provide a scatter/gather table of pinned pages
* *gem_prime_vmap*: vmap a buffer exported by your driver
* *gem_prime_vunmap*: vunmap a buffer exported by your driver
* *gem_prime_mmap* (optional): mmap a buffer exported by your driver

Import callback:

* *gem_prime_import_sg_table* (import): produce a GEM object from another driver's scatter/gather table

PRIME Function References

Name

`drm_gem_dmabuf_release` — `dma_buf` release implementation for GEM

Synopsis

```
void drm_gem_dmabuf_release (struct dma_buf * dma_buf);
```

Arguments

dma_buf buffer to be released

Description

Generic release function for `dma_bufs` exported as PRIME buffers. GEM drivers must use this in their `dma_buf` ops structure as the release callback.

Name

`drm_gem_prime_export` — helper library implementation of the export callback

Synopsis

```
struct dma_buf * drm_gem_prime_export (struct drm_device * dev, struct
drm_gem_object * obj, int flags);
```

Arguments

dev drm_device to export from

obj GEM object to export

flags flags like DRM_CLOEXEC

Description

This is the implementation of the `gem_prime_export` functions for GEM drivers using the PRIME helpers.

Name

`drm_gem_prime_handle_to_fd` — PRIME export function for GEM drivers

Synopsis

```
int drm_gem_prime_handle_to_fd (struct drm_device * dev, struct drm_file  
* file_priv, uint32_t handle, uint32_t flags, int * prime_fd);
```

Arguments

<i>dev</i>	dev to export the buffer from
<i>file_priv</i>	drm file-private structure
<i>handle</i>	buffer handle to export
<i>flags</i>	flags like <code>DRM_CLOEXEC</code>
<i>prime_fd</i>	pointer to storage for the fd id of the create dma-buf

Description

This is the PRIME export function which must be used mandatorily by GEM drivers to ensure correct lifetime management of the underlying GEM object. The actual exporting from GEM object to a dma-buf is done through the `gem_prime_export` driver callback.

Name

`drm_gem_prime_import` — helper library implementation of the import callback

Synopsis

```
struct drm_gem_object * drm_gem_prime_import (struct drm_device * dev,  
struct dma_buf * dma_buf);
```

Arguments

dev `drm_device` to import into

dma_buf `dma-buf` object to import

Description

This is the implementation of the `gem_prime_import` functions for GEM drivers using the PRIME helpers.

Name

`drm_gem_prime_fd_to_handle` — PRIME import function for GEM drivers

Synopsis

```
int drm_gem_prime_fd_to_handle (struct drm_device * dev, struct drm_file  
* file_priv, int prime_fd, uint32_t * handle);
```

Arguments

<i>dev</i>	dev to export the buffer from
<i>file_priv</i>	drm file-private structure
<i>prime_fd</i>	fd id of the dma-buf which should be imported
<i>handle</i>	pointer to storage for the handle of the imported buffer object

Description

This is the PRIME import function which must be used mandatorily by GEM drivers to ensure correct lifetime management of the underlying GEM object. The actual importing of GEM object from the dma-buf is done through the `gem_import_export` driver callback.

Name

`drm_prime_pages_to_sg` — converts a page array into an sg list

Synopsis

```
struct sg_table * drm_prime_pages_to_sg (struct page ** pages, unsigned
int nr_pages);
```

Arguments

pages pointer to the array of page pointers to convert

nr_pages length of the page vector

Description

This helper creates an sg table object from a set of pages the driver is responsible for mapping the pages into the importers address space for use with `dma_buf` itself.

Name

`drm_prime_sg_to_page_addr_arrays` — convert an sg table into a page array

Synopsis

```
int drm_prime_sg_to_page_addr_arrays (struct sg_table * sgt, struct page
** pages, dma_addr_t * addrs, int max_pages);
```

Arguments

<i>sgt</i>	scatter-gather table to convert
<i>pages</i>	array of page pointers to store the page array in
<i>addrs</i>	optional array to store the dma bus address of each page
<i>max_pages</i>	size of both the passed-in arrays

Description

Exports an sg table into an array of pages and addresses. This is currently required by the TTM driver in order to do correct fault handling.

Name

`drm_prime_gem_destroy` — helper to clean up a PRIME-imported GEM object

Synopsis

```
void drm_prime_gem_destroy (struct drm_gem_object * obj, struct sg_table  
* sg);
```

Arguments

obj GEM object which was created from a dma-buf

sg the sg-table which was pinned at import time

Description

This is the cleanup functions which GEM drivers need to call when they use `drm_gem_prime_import` to import dma-bufs.

DRM MM Range Allocator

Overview

`drm_mm` provides a simple range allocator. The drivers are free to use the resource allocator from the linux core if it suits them, the upside of `drm_mm` is that it's in the DRM core. Which means that it's easier to extend for some of the crazier special purpose needs of gpus.

The main data struct is `drm_mm`, allocations are tracked in `drm_mm_node`. Drivers are free to embed either of them into their own suitable datastructures. `drm_mm` itself will not do any allocations of its own, so if drivers choose not to embed nodes they need to still allocate them themselves.

The range allocator also supports reservation of preallocated blocks. This is useful for taking over initial mode setting configurations from the firmware, where an object needs to be created which exactly matches the firmware's scanout target. As long as the range is still free it can be inserted anytime after the allocator is initialized, which helps with avoiding looped dependencies in the driver load sequence.

`drm_mm` maintains a stack of most recently freed holes, which of all simplistic datastructures seems to be a fairly decent approach to clustering allocations and avoiding too much fragmentation. This means free space searches are $O(\text{num_holes})$. Given that all the fancy features `drm_mm` supports something better would be fairly complex and since gfx thrashing is a fairly steep cliff not a real concern. Removing a node again is $O(1)$.

`drm_mm` supports a few features: Alignment and range restrictions can be supplied. Further more every `drm_mm_node` has a color value (which is just an opaque unsigned long) which in conjunction with a driver callback can be used to implement sophisticated placement restrictions. The i915 DRM driver uses this to implement guard pages between incompatible caching domains in the graphics TT.

Two behaviors are supported for searching and allocating: bottom-up and top-down. The default is bottom-up. Top-down allocation can be used if the memory area has different restrictions, or just to reduce fragmentation.

Finally iteration helpers to walk all nodes and all holes are provided as are some basic allocator dumpers for debugging.

LRU Scan/Eviction Support

Very often GPUs need to have continuous allocations for a given object. When evicting objects to make space for a new one it is therefore not most efficient when we simply start to select all objects from the tail of an LRU until there's a suitable hole: Especially for big objects or nodes that otherwise have special allocation constraints there's a good chance we evict lots of (smaller) objects unnecessarily.

The DRM range allocator supports this use-case through the scanning interfaces. First a scan operation needs to be initialized with `drm_mm_init_scan` or `drm_mm_init_scan_with_range`. The driver adds objects to the roaster (probably by walking an LRU list, but this can be freely implemented) until a suitable hole is found or there's no further evitable object.

The driver must walk through all objects again in exactly the reverse order to restore the allocator state. Note that while the allocator is used in the scan mode no other operation is allowed.

Finally the driver evicts all objects selected in the scan. Adding and removing an object is $O(1)$, and since freeing a node is also $O(1)$ the overall complexity is $O(\text{scanned_objects})$. So like the free stack which needs to be walked before a scan operation even begins this is linear in the number of objects. It doesn't seem to hurt badly.

DRM MM Range Allocator Function References

Name

`drm_mm_reserve_node` — insert an pre-initialized node

Synopsis

```
int drm_mm_reserve_node (struct drm_mm * mm, struct drm_mm_node * node);
```

Arguments

mm `drm_mm` allocator to insert *node* into

node `drm_mm_node` to insert

Description

This functions inserts an already set-up `drm_mm_node` into the allocator, meaning that start, size and color must be set by the caller. This is useful to initialize the allocator with preallocated objects which must be set-up before the range allocator can be set-up, e.g. when taking over a firmware framebuffer.

Returns

0 on success, -ENOSPC if there's no hole where *node* is.

Name

`drm_mm_insert_node_generic` — search for space and insert *node*

Synopsis

```
int drm_mm_insert_node_generic (struct drm_mm * mm, struct drm_mm_node  
* node, u64 size, unsigned alignment, unsigned long color, enum drm_m-  
m_search_flags sflags, enum drm_mm_allocator_flags aflags);
```

Arguments

<i>mm</i>	drm_mm to allocate from
<i>node</i>	preallocate node to insert
<i>size</i>	size of the allocation
<i>alignment</i>	alignment of the allocation
<i>color</i>	opaque tag value to use for this node
<i>sflags</i>	flags to fine-tune the allocation search
<i>aflags</i>	flags to fine-tune the allocation behavior

Description

The preallocated node must be cleared to 0.

Returns

0 on success, -ENOSPC if there's no suitable hole.

Name

`drm_mm_insert_node_in_range_generic` — ranged search for space and insert *node*

Synopsis

```
int drm_mm_insert_node_in_range_generic (struct drm_mm * mm, struct dr-
m_mm_node * node, u64 size, unsigned alignment, unsigned long color,
u64 start, u64 end, enum drm_mm_search_flags sflags, enum drm_mm_allo-
cator_flags aflags);
```

Arguments

<i>mm</i>	drm_mm to allocate from
<i>node</i>	preallocate node to insert
<i>size</i>	size of the allocation
<i>alignment</i>	alignment of the allocation
<i>color</i>	opaque tag value to use for this node
<i>start</i>	start of the allowed range for this node
<i>end</i>	end of the allowed range for this node
<i>sflags</i>	flags to fine-tune the allocation search
<i>aflags</i>	flags to fine-tune the allocation behavior

Description

The preallocated node must be cleared to 0.

Returns

0 on success, -ENOSPC if there's no suitable hole.

Name

`drm_mm_remove_node` — Remove a memory node from the allocator.

Synopsis

```
void drm_mm_remove_node (struct drm_mm_node * node);
```

Arguments

node `drm_mm_node` to remove

Description

This just removes a node from its `drm_mm` allocator. The node does not need to be cleared again before it can be re-inserted into this or any other `drm_mm` allocator. It is a bug to call this function on a un-allocated node.

Name

`drm_mm_replace_node` — move an allocation from *old* to *new*

Synopsis

```
void drm_mm_replace_node (struct drm_mm_node * old, struct drm_mm_node  
* new);
```

Arguments

old `drm_mm_node` to remove from the allocator

new `drm_mm_node` which should inherit *old*'s allocation

Description

This is useful for when drivers embed the `drm_mm_node` structure and hence can't move allocations by reassigning pointers. It's a combination of remove and insert with the guarantee that the allocation start will match.

Name

`drm_mm_init_scan` — initialize lru scanning

Synopsis

```
void drm_mm_init_scan (struct drm_mm * mm, u64 size, unsigned alignment,  
unsigned long color);
```

Arguments

<i>mm</i>	drm_mm to scan
<i>size</i>	size of the allocation
<i>alignment</i>	alignment of the allocation
<i>color</i>	opaque tag value to use for the allocation

Description

This simply sets up the scanning routines with the parameters for the desired hole. Note that there's no need to specify allocation flags, since they only change the place a node is allocated from within a suitable hole.

Warning

As long as the scan list is non-empty, no other operations than adding/removing nodes to/from the scan list are allowed.

Name

`drm_mm_init_scan_with_range` — initialize range-restricted lru scanning

Synopsis

```
void drm_mm_init_scan_with_range (struct drm_mm * mm, u64 size, unsigned  
alignment, unsigned long color, u64 start, u64 end);
```

Arguments

<i>mm</i>	drm_mm to scan
<i>size</i>	size of the allocation
<i>alignment</i>	alignment of the allocation
<i>color</i>	opaque tag value to use for the allocation
<i>start</i>	start of the allowed range for the allocation
<i>end</i>	end of the allowed range for the allocation

Description

This simply sets up the scanning routines with the parameters for the desired hole. Note that there's no need to specify allocation flags, since they only change the place a node is allocated from within a suitable hole.

Warning

As long as the scan list is non-empty, no other operations than adding/removing nodes to/from the scan list are allowed.

Name

`drm_mm_scan_add_block` — add a node to the scan list

Synopsis

```
bool drm_mm_scan_add_block (struct drm_mm_node * node);
```

Arguments

node `drm_mm_node` to add

Description

Add a node to the scan list that might be freed to make space for the desired hole.

Returns

True if a hole has been found, false otherwise.

Name

`drm_mm_scan_remove_block` — remove a node from the scan list

Synopsis

```
bool drm_mm_scan_remove_block (struct drm_mm_node * node);
```

Arguments

node `drm_mm_node` to remove

Description

Nodes `_must_` be removed in the exact same order from the scan list as they have been added, otherwise the internal state of the memory manager will be corrupted.

When the scan list is empty, the selected memory nodes can be freed. An immediately following `drm_mm_search_free` with `!DRM_MM_SEARCH_BEST` will then return the just freed block (because its at the top of the `free_stack` list).

Returns

True if this block should be evicted, false otherwise. Will always return false when no hole has been found.

Name

`drm_mm_clean` — checks whether an allocator is clean

Synopsis

```
bool drm_mm_clean (struct drm_mm * mm);
```

Arguments

mm `drm_mm` allocator to check

Returns

True if the allocator is completely free, false if there's still a node allocated in it.

Name

`drm_mm_init` — initialize a drm-mm allocator

Synopsis

```
void drm_mm_init (struct drm_mm * mm, u64 start, u64 size);
```

Arguments

mm the `drm_mm` structure to initialize

start start of the range managed by *mm*

size end of the range managed by *mm*

Description

Note that *mm* must be cleared to 0 before calling this function.

Name

`drm_mm_takedown` — clean up a `drm_mm` allocator

Synopsis

```
void drm_mm_takedown (struct drm_mm * mm);
```

Arguments

mm `drm_mm` allocator to clean up

Description

Note that it is a bug to call this function on an allocator which is not clean.

Name

`drm_mm_debug_table` — dump allocator state to dmesg

Synopsis

```
void drm_mm_debug_table (struct drm_mm * mm, const char * prefix);
```

Arguments

mm `drm_mm` allocator to dump

prefix prefix to use for dumping to dmesg

Name

`drm_mm_dump_table` — dump allocator state to a `seq_file`

Synopsis

```
int drm_mm_dump_table (struct seq_file * m, struct drm_mm * mm);
```

Arguments

m `seq_file` to dump to

mm `drm_mm` allocator to dump

Name

`drm_mm_node_allocated` — checks whether a node is allocated

Synopsis

```
bool drm_mm_node_allocated (struct drm_mm_node * node);
```

Arguments

node `drm_mm_node` to check

Description

Drivers should use this helpers for proper encapsulation of `drm_mm` internals.

Returns

True if the *node* is allocated.

Name

`drm_mm_initialized` — checks whether an allocator is initialized

Synopsis

```
bool drm_mm_initialized (struct drm_mm * mm);
```

Arguments

mm `drm_mm` to check

Description

Drivers should use this helpers for proper encapsulation of `drm_mm` internals.

Returns

True if the *mm* is initialized.

Name

`drm_mm_hole_node_start` — computes the start of the hole following *node*

Synopsis

```
u64 drm_mm_hole_node_start (struct drm_mm_node * hole_node);
```

Arguments

hole_node `drm_mm_node` which implicitly tracks the following hole

Description

This is useful for driver-specific debug dumpers. Otherwise drivers should not inspect holes themselves. Drivers must check first whether a hole indeed follows by looking at `node->hole_follows`.

Returns

Start of the subsequent hole.

Name

`drm_mm_hole_node_end` — computes the end of the hole following *node*

Synopsis

```
u64 drm_mm_hole_node_end (struct drm_mm_node * hole_node);
```

Arguments

hole_node `drm_mm_node` which implicitly tracks the following hole

Description

This is useful for driver-specific debug dumpers. Otherwise drivers should not inspect holes themselves. Drivers must check first whether a hole indeed follows by looking at `node->hole_follows`.

Returns

End of the subsequent hole.

Name

`drm_mm_for_each_node` — iterator to walk over all allocated nodes

Synopsis

```
drm_mm_for_each_node ( entry, mm );
```

Arguments

entry `drm_mm_node` structure to assign to in each iteration step

mm `drm_mm` allocator to walk

Description

This iterator walks over all nodes in the range allocator. It is implemented with `list_for_each`, so not safe against removal of elements.

Name

`drm_mm_for_each_hole` — iterator to walk over all holes

Synopsis

```
drm_mm_for_each_hole ( entry, mm, hole_start, hole_end );
```

Arguments

<i>entry</i>	<code>drm_mm_node</code> used internally to track progress
<i>mm</i>	<code>drm_mm</code> allocator to walk
<i>hole_start</i>	ulong variable to assign the hole start to on each iteration
<i>hole_end</i>	ulong variable to assign the hole end to on each iteration

Description

This iterator walks over all holes in the range allocator. It is implemented with `list_for_each`, so not save against removal of elements. *entry* is used internally and will not reflect a real `drm_mm_node` for the very first hole. Hence users of this iterator may not access it.

Implementation Note

We need to inline `list_for_each_entry` in order to be able to set `hole_start` and `hole_end` on each iteration while keeping the macro sane.

The `__drm_mm_for_each_hole` version is similar, but with added support for going backwards.

Name

`drm_mm_insert_node` — search for space and insert *node*

Synopsis

```
int drm_mm_insert_node (struct drm_mm * mm, struct drm_mm_node * node,  
u64 size, unsigned alignment, enum drm_mm_search_flags flags);
```

Arguments

<i>mm</i>	drm_mm to allocate from
<i>node</i>	preallocate node to insert
<i>size</i>	size of the allocation
<i>alignment</i>	alignment of the allocation
<i>flags</i>	flags to fine-tune the allocation

Description

This is a simplified version of `drm_mm_insert_node_generic` with *color* set to 0.

The preallocated node must be cleared to 0.

Returns

0 on success, -ENOSPC if there's no suitable hole.

Name

`drm_mm_insert_node_in_range` — ranged search for space and insert *node*

Synopsis

```
int drm_mm_insert_node_in_range (struct drm_mm * mm, struct drm_mm_node  
* node, u64 size, unsigned alignment, u64 start, u64 end, enum drm_m-  
m_search_flags flags);
```

Arguments

<i>mm</i>	drm_mm to allocate from
<i>node</i>	preallocate node to insert
<i>size</i>	size of the allocation
<i>alignment</i>	alignment of the allocation
<i>start</i>	start of the allowed range for this node
<i>end</i>	end of the allowed range for this node
<i>flags</i>	flags to fine-tune the allocation

Description

This is a simplified version of `drm_mm_insert_node_in_range_generic` with *color* set to 0.

The preallocated node must be cleared to 0.

Returns

0 on success, -ENOSPC if there's no suitable hole.

CMA Helper Functions Reference

The Contiguous Memory Allocator reserves a pool of memory at early boot that is used to service requests for large blocks of contiguous memory.

The DRM GEM/CMA helpers use this allocator as a means to provide buffer objects that are physically contiguous in memory. This is useful for display drivers that are unable to map scattered buffers via an IOMMU.

Name

`drm_gem_cma_create` — allocate an object with the given size

Synopsis

```
struct drm_gem_cma_object * drm_gem_cma_create (struct drm_device * drm,  
size_t size);
```

Arguments

drm DRM device

size size of the object to allocate

Description

This function creates a CMA GEM object and allocates a contiguous chunk of memory as backing store. The backing memory has the writecombine attribute set.

Returns

A struct `drm_gem_cma_object *` on success or an `ERR_PTR`-encoded negative error code on failure.

Name

`drm_gem_cma_free_object` — free resources associated with a CMA GEM object

Synopsis

```
void drm_gem_cma_free_object (struct drm_gem_object * gem_obj);
```

Arguments

gem_obj GEM object to free

Description

This function frees the backing memory of the CMA GEM object, cleans up the GEM object state and frees the memory used to store the object itself. Drivers using the CMA helpers should set this as their DRM driver's `->gem_free_object` callback.

Name

`drm_gem_cma_dumb_create_internal` — create a dumb buffer object

Synopsis

```
int drm_gem_cma_dumb_create_internal (struct drm_file * file_priv,
struct drm_device * drm, struct drm_mode_create_dumb * args);
```

Arguments

file_priv DRM file-private structure to create the dumb buffer for

drm DRM device

args IOCTL data

Description

This aligns the pitch and size arguments to the minimum required. This is an internal helper that can be wrapped by a driver to account for hardware with more specific alignment requirements. It should not be used directly as the `->dumb_create` callback in a DRM driver.

Returns

0 on success or a negative error code on failure.

Name

`drm_gem_cma_dumb_create` — create a dumb buffer object

Synopsis

```
int drm_gem_cma_dumb_create (struct drm_file * file_priv, struct dr-  
m_device * drm, struct drm_mode_create_dumb * args);
```

Arguments

file_priv DRM file-private structure to create the dumb buffer for

drm DRM device

args IOCTL data

Description

This function computes the pitch of the dumb buffer and rounds it up to an integer number of bytes per pixel. Drivers for hardware that doesn't have any additional restrictions on the pitch can directly use this function as their `->dumb_create` callback.

For hardware with additional restrictions, drivers can adjust the fields set up by userspace and pass the IOCTL data along to the `drm_gem_cma_dumb_create_internal` function.

Returns

0 on success or a negative error code on failure.

Name

`drm_gem_cma_dumb_map_offset` — return the fake mmap offset for a CMA GEM object

Synopsis

```
int drm_gem_cma_dumb_map_offset (struct drm_file * file_priv, struct
drm_device * drm, u32 handle, u64 * offset);
```

Arguments

file_priv DRM file-private structure containing the GEM object

drm DRM device

handle GEM object handle

offset return location for the fake mmap offset

Description

This function look up an object by its handle and returns the fake mmap offset associated with it. Drivers using the CMA helpers should set this as their DRM driver's `->dumb_map_offset` callback.

Returns

0 on success or a negative error code on failure.

Name

`drm_gem_cma_mmap` — memory-map a CMA GEM object

Synopsis

```
int drm_gem_cma_mmap (struct file * filp, struct vm_area_struct * vma);
```

Arguments

filp file object

vma VMA for the area to be mapped

Description

This function implements an augmented version of the GEM DRM file mmap

operation for CMA objects

In addition to the usual GEM VMA setup it immediately faults in the entire object instead of using on-demand faulting. Drivers which employ the CMA helpers should use this function as their `->mmap` handler in the DRM device file's `file_operations` structure.

Returns

0 on success or a negative error code on failure.

Name

`drm_gem_cma_describe` — describe a CMA GEM object for debugfs

Synopsis

```
void drm_gem_cma_describe (struct drm_gem_cma_object * cma_obj, struct  
seq_file * m);
```

Arguments

cma_obj CMA GEM object

m debugfs file handle

Description

This function can be used to dump a human-readable representation of the CMA GEM object into a synthetic file.

Name

`drm_gem_cma_prime_get_sg_table` — provide a scatter/gather table of pinned pages for a CMA GEM object

Synopsis

```
struct sg_table * drm_gem_cma_prime_get_sg_table (struct drm_gem_object  
* obj);
```

Arguments

obj GEM object

Description

This function exports a scatter/gather table suitable for PRIME usage by calling the standard DMA mapping API. Drivers using the CMA helpers should set this as their DRM driver's `->gem_prime_get_sg_table` callback.

Returns

A pointer to the scatter/gather table of pinned pages or NULL on failure.

Name

`drm_gem_cma_prime_import_sg_table` — produce a CMA GEM object from another driver's scatter/gather table of pinned pages

Synopsis

```
struct drm_gem_object * drm_gem_cma_prime_import_sg_table (struct dr-  
m_device * dev, struct dma_buf_attachment * attach, struct sg_table *  
sgt);
```

Arguments

dev device to import into

attach DMA-BUF attachment

sgt scatter/gather table of pinned pages

Description

This function imports a scatter/gather table exported via DMA-BUF by another driver. Imported buffers must be physically contiguous in memory (i.e. the scatter/gather table must contain a single entry). Drivers that use the CMA helpers should set this as their DRM driver's `->gem_prime_import_sg_table` callback.

Returns

A pointer to a newly created GEM object or an `ERR_PTR`-encoded negative error code on failure.

Name

`drm_gem_cma_prime_mmap` — memory-map an exported CMA GEM object

Synopsis

```
int  drm_gem_cma_prime_mmap (struct drm_gem_object * obj, struct
vm_area_struct * vma);
```

Arguments

obj GEM object

vma VMA for the area to be mapped

Description

This function maps a buffer imported via DRM PRIME into a userspace process's address space. Drivers that use the CMA helpers should set this as their DRM driver's `->gem_prime_mmap` callback.

Returns

0 on success or a negative error code on failure.

Name

`drm_gem_cma_prime_vmap` — map a CMA GEM object into the kernel's virtual address space

Synopsis

```
void * drm_gem_cma_prime_vmap (struct drm_gem_object * obj);
```

Arguments

obj GEM object

Description

This function maps a buffer exported via DRM PRIME into the kernel's virtual address space. Since the CMA buffers are already mapped into the kernel virtual address space this simply returns the cached virtual address. Drivers using the CMA helpers should set this as their DRM driver's `->gem_prime_vmap` callback.

Returns

The kernel virtual address of the CMA GEM object's backing store.

Name

`drm_gem_cma_prime_vunmap` — unmap a CMA GEM object from the kernel's virtual address space

Synopsis

```
void drm_gem_cma_prime_vunmap (struct drm_gem_object * obj, void * vaddr);
```

Arguments

obj GEM object

vaddr kernel virtual address where the CMA GEM object was mapped

Description

This function removes a buffer exported via DRM PRIME from the kernel's virtual address space. This is a no-op because CMA buffers cannot be unmapped from kernel space. Drivers using the CMA helpers should set this as their DRM driver's `->gem_prime_vunmap` callback.

Name

struct `drm_gem_cma_object` — GEM object backed by CMA memory allocations

Synopsis

```
struct drm_gem_cma_object {  
    struct drm_gem_object base;  
    dma_addr_t paddr;  
    struct sg_table * sgt;  
    void * vaddr;  
};
```

Members

<code>base</code>	base GEM object
<code>paddr</code>	physical address of the backing memory
<code>sgt</code>	scatter/gather table for imported PRIME buffers
<code>vaddr</code>	kernel virtual address of the backing memory

Mode Setting

Drivers must initialize the mode setting core by calling `drm_mode_config_init` on the DRM device. The function initializes the `drm_device` *mode_config* field and never fails. Once done, mode configuration must be setup by initializing the following fields.

- `int min_width, min_height;`
 `int max_width, max_height;`

Minimum and maximum width and height of the frame buffers in pixel units.

- `struct drm_mode_config_funcs *funcs;`

Mode setting functions.

Display Modes Function Reference

Name

enum drm_mode_status — hardware support status of a mode

Synopsis

```
enum drm_mode_status {  
    MODE_OK,  
    MODE_HSYNC,  
    MODE_VSYNC,  
    MODE_H_ILLEGAL,  
    MODE_V_ILLEGAL,  
    MODE_BAD_WIDTH,  
    MODE_NOMODE,  
    MODE_NO_INTERLACE,  
    MODE_NO_DBLESCAN,  
    MODE_NO_VSCAN,  
    MODE_MEM,  
    MODE_VIRTUAL_X,  
    MODE_VIRTUAL_Y,  
    MODE_MEM_VIRT,  
    MODE_NOCLOCK,  
    MODE_CLOCK_HIGH,  
    MODE_CLOCK_LOW,  
    MODE_CLOCK_RANGE,  
    MODE_BAD_HVALUE,  
    MODE_BAD_VVALUE,  
    MODE_BAD_VSCAN,  
    MODE_HSYNC_NARROW,  
    MODE_HSYNC_WIDE,  
    MODE_HBLANK_NARROW,  
    MODE_HBLANK_WIDE,  
    MODE_VSYNC_NARROW,  
    MODE_VSYNC_WIDE,  
    MODE_VBLANK_NARROW,  
    MODE_VBLANK_WIDE,  
    MODE_PANEL,  
    MODE_INTERLACE_WIDTH,  
    MODE_ONE_WIDTH,  
    MODE_ONE_HEIGHT,  
    MODE_ONE_SIZE,  
    MODE_NO_REDUCED,  
    MODE_NO_STEREO,  
    MODE_STALE,  
    MODE_BAD,  
    MODE_ERROR  
};
```

Constants

MODE_OK	Mode OK
MODE_HSYNC	hsync out of range

MODE_VSYNC	vsync out of range
MODE_H_ILLEGAL	mode has illegal horizontal timings
MODE_V_ILLEGAL	mode has illegal horizontal timings
MODE_BAD_WIDTH	requires an unsupported linepitch
MODE_NOMODE	no mode with a matching name
MODE_NO_INTERLACE	interlaced mode not supported
MODE_NO_DBLESCAN	doublescan mode not supported
MODE_NO_VSCAN	multiscan mode not supported
MODE_MEM	insufficient video memory
MODE_VIRTUAL_X	mode width too large for specified virtual size
MODE_VIRTUAL_Y	mode height too large for specified virtual size
MODE_MEM_VIRT	insufficient video memory given virtual size
MODE_NOCLOCK	no fixed clock available
MODE_CLOCK_HIGH	clock required is too high
MODE_CLOCK_LOW	clock required is too low
MODE_CLOCK_RANGE	clock/mode isn't in a ClockRange
MODE_BAD_HVALUE	horizontal timing was out of range
MODE_BAD_VVALUE	vertical timing was out of range
MODE_BAD_VSCAN	VScan value out of range
MODE_HSYNC_NARROW	horizontal sync too narrow
MODE_HSYNC_WIDE	horizontal sync too wide
MODE_HBLANK_NARROW	horizontal blanking too narrow
MODE_HBLANK_WIDE	horizontal blanking too wide
MODE_VSYNC_NARROW	vertical sync too narrow
MODE_VSYNC_WIDE	vertical sync too wide
MODE_VBLANK_NARROW	vertical blanking too narrow
MODE_VBLANK_WIDE	vertical blanking too wide
MODE_PANEL	exceeds panel dimensions
MODE_INTERLACE_WIDTH	width too large for interlaced mode

MODE_ONE_WIDTH	only one width is supported
MODE_ONE_HEIGHT	only one height is supported
MODE_ONE_SIZE	only one resolution is supported
MODE_NO_REDUCED	monitor doesn't accept reduced blanking
MODE_NO_STEREO	stereo modes not supported
MODE_STALE	mode has become stale
MODE_BAD	unspecified reason
MODE_ERROR	error condition

Description

This enum is used to filter out modes not supported by the driver/hardware combination.

Name

struct drm_display_mode — DRM kernel-internal display mode structure

Synopsis

```
struct drm_display_mode {
    struct list_head head;
    struct drm_mode_object base;
    char name[DRM_DISPLAY_MODE_LEN];
    enum drm_mode_status status;
    unsigned int type;
    int clock;
    int hdisplay;
    int hsync_start;
    int hsync_end;
    int httotal;
    int hskew;
    int vdisplay;
    int vsync_start;
    int vsync_end;
    int vttotal;
    int vscan;
    unsigned int flags;
    int width_mm;
    int height_mm;
    int crtc_clock;
    int crtc_hdisplay;
    int crtc_hblank_start;
    int crtc_hblank_end;
    int crtc_hsync_start;
    int crtc_hsync_end;
    int crtc_httotal;
    int crtc_hskew;
    int crtc_vdisplay;
    int crtc_vblank_start;
    int crtc_vblank_end;
    int crtc_vsync_start;
    int crtc_vsync_end;
    int crtc_vttotal;
    int * private;
    int private_flags;
    int vrefresh;
    int hsync;
    enum hdmi_picture_aspect picture_aspect_ratio;
};
```

Members

head	struct list_head for mode lists.
base	

A display mode is a normal modeset object, possibly including public userspace id.

FIXME:

This can probably be removed since the entire concept of userspace managing modes explicitly has never landed in upstream kernel mode setting support.

name[DRM_DISPLAY_MODE_LEN]

Human-readable name of the mode, filled out with `drm_mode_set_name`.

status

Status of the mode, used to filter out modes not supported by the hardware. See enum `drm_mode_status`.

type

A bitmask of flags, mostly about the source of a mode. Possible flags are:

- `DRM_MODE_TYPE_BUILTIN`: Meant for hard-coded modes, effectively unused.
- `DRM_MODE_TYPE_PREFERRED`: Preferred mode, usually the native resolution of an LCD panel. There should only be one preferred mode per connector at any given time.
- `DRM_MODE_TYPE_DRIVER`: Mode created by the driver, which is all of them really. Drivers must set this bit for all modes they create and expose to userspace.

Plus a big list of flags which shouldn't be used at all, but are still around since these flags are also used in the userspace ABI:

- `DRM_MODE_TYPE_DEFAULT`: Again a leftover, use `DRM_MODE_TYPE_PREFERRED` instead.
- `DRM_MODE_TYPE_CLOCK_C` and `DRM_MODE_TYPE_CRTC_C`: Define leftovers which are stuck around for hysterical raisins only. No one has an idea what they were meant for. Don't use.
- `DRM_MODE_TYPE_USERDEF`: Mode defined by userspace, again a vestige from older kms designs where userspace had to first add a custom mode to the kernel's mode list before it could use it. Don't use.

clock

Pixel clock in kHz.

hdisplay

horizontal display size

hsync_start

horizontal sync start

hsync_end

horizontal sync end

htotal

horizontal total size

hskew

horizontal skew?!

vdisplay

vertical display size

vsync_start

vertical sync start

vsync_end	vertical sync end
vtotal	vertical total size
vscan	vertical scan?!
flags	<p>Sync and timing flags:</p> <ul style="list-style-type: none"> - DRM_MODE_FLAG_PHSYNC: horizontal sync is active high. - DRM_MODE_FLAG_NHSYNC: horizontal sync is active low. - DRM_MODE_FLAG_PVSYNC: vertical sync is active high. - DRM_MODE_FLAG_NVSYNC: vertical sync is active low. - DRM_MODE_FLAG_INTERLACE: mode is interlaced. - DRM_MODE_FLAG_DBLSCAN: mode uses doublescan. - DRM_MODE_FLAG_CSYNC: mode uses composite sync. - DRM_MODE_FLAG_PCSYNC: composite sync is active high. - DRM_MODE_FLAG_NCSYNC: composite sync is active low. - DRM_MODE_FLAG_HSKW: hskew provided (not used?). - DRM_MODE_FLAG_BCAST: not used? - DRM_MODE_FLAG_PIXMUX: not used? - DRM_MODE_FLAG_DBLCLK: double-clocked mode. - DRM_MODE_FLAG_CLKDIV2: half-clocked mode. <p>Additionally there's flags to specify how 3D modes are packed:</p> <ul style="list-style-type: none"> - DRM_MODE_FLAG_3D_NONE: normal, non-3D mode. - DRM_MODE_FLAG_3D_FRAME_PACKING: 2 full frames for left and right. - DRM_MODE_FLAG_3D_FIELD_ALTERNATIVE: interleaved like fields. - DRM_MODE_FLAG_3D_LINE_ALTERNATIVE: interleaved lines. - DRM_MODE_FLAG_3D_SIDE_BY_SIDE_FULL: side-by-side full frames. - DRM_MODE_FLAG_3D_L_DEPTH: ? - DRM_MODE_FLAG_3D_L_DEPTH_GFX_GFX_DEPTH: ? - DRM_MODE_FLAG_3D_TOP_AND_BOTTOM: frame split into top and bottom parts. - DRM_MODE_FLAG_3D_SIDE_BY_SIDE_HALF: frame split into left and right parts.
width_mm	Addressable size of the output in mm, projectors should set this to 0.
height_mm	Addressable size of the output in mm, projectors should set this to 0.
crtc_clock	<p>Actual pixel or dot clock in the hardware. This differs from the logical <i>clock</i> when e.g. using interlacing, double-clocking, stereo modes or other fancy stuff that changes the timings and signals actually sent over the wire.</p> <p>This is again in kHz.</p> <p>Note that with digital outputs like HDMI or DP there's usually a massive confusion between the dot clock and the signal clock at the bit encoding level. Especially when a 8b/10b encoding is used and the difference is exactly a factor of 10.</p>

<code>crtc_hdisplay</code>	hardware mode horizontal display size
<code>crtc_hblank_start</code>	hardware mode horizontal blank start
<code>crtc_hblank_end</code>	hardware mode horizontal blank end
<code>crtc_hsync_start</code>	hardware mode horizontal sync start
<code>crtc_hsync_end</code>	hardware mode horizontal sync end
<code>crtc_htotal</code>	hardware mode horizontal total size
<code>crtc_hskew</code>	hardware mode horizontal skew?!
<code>crtc_vdisplay</code>	hardware mode vertical display size
<code>crtc_vblank_start</code>	hardware mode vertical blank start
<code>crtc_vblank_end</code>	hardware mode vertical blank end
<code>crtc_vsync_start</code>	hardware mode vertical sync start
<code>crtc_vsync_end</code>	hardware mode vertical sync end
<code>crtc_vtotal</code>	hardware mode vertical total size
<code>private</code>	Pointer for driver private data. This can only be used for mode objects passed to drivers in modeset operations. It shouldn't be used by atomic drivers since they can store any additional data by subclassing state structures.
<code>private_flags</code>	Similar to <i>private</i> , but just an integer.
<code>vrefresh</code>	Vertical refresh rate, for debug output in human readable form. Not used in a functional way. This value is in Hz.
<code>hsync</code>	Horizontal refresh rate, for debug output in human readable form. Not used in a functional way. This value is in kHz.
<code>picture_aspect_ratio</code>	Field for setting the HDMI picture aspect ratio of a mode.

Description

The horizontal and vertical timings are defined per the following diagram.

```

Active Front Sync Back Region Porch Porch <-----><-----><----->
><-----> /////////////// /////////////// | /////////////// |..... ..
_____ <---- [hv]display ----> <----- [hv]sync_start -----> <----->
[hv]sync_end -----> <----- [hv]total ----->*
```

This structure contains two copies of timings. First are the plain timings, which specify the logical mode, as it would be for a progressive 1:1 scanout at the refresh rate userspace can observe through vblank timestamps. Then there's the hardware timings, which are corrected for interlacing, double-clocking and similar things. They are provided as a convenience, and can be appropriately computed using `drm_mode_set_crtcinfo`.

Name

`drm_mode_is_stereo` — check for stereo mode flags

Synopsis

```
bool drm_mode_is_stereo (const struct drm_display_mode * mode);
```

Arguments

mode `drm_display_mode` to check

Returns

True if the mode is one of the stereo modes (like side-by-side), false if not.

Name

`drm_mode_debug_printmodeline` — print a mode to dmesg

Synopsis

```
void drm_mode_debug_printmodeline (const struct drm_display_mode *  
mode);
```

Arguments

mode mode to print

Description

Describe *mode* using `DRM_DEBUG`.

Name

`drm_mode_create` — create a new display mode

Synopsis

```
struct drm_display_mode * drm_mode_create (struct drm_device * dev);
```

Arguments

dev DRM device

Description

Create a new, cleared `drm_display_mode` with `kzalloc`, allocate an ID for it and return it.

Returns

Pointer to new mode on success, `NULL` on error.

Name

`drm_mode_destroy` — remove a mode

Synopsis

```
void drm_mode_destroy (struct drm_device * dev, struct drm_display_mode  
* mode);
```

Arguments

dev DRM device

mode mode to remove

Description

Release *mode*'s unique ID, then free it *mode* structure itself using `kfree`.

Name

`drm_mode_probed_add` — add a mode to a connector's `probed_mode` list

Synopsis

```
void drm_mode_probed_add (struct drm_connector * connector, struct dr-  
m_display_mode * mode);
```

Arguments

connector connector the new mode

mode mode data

Description

Add *mode* to *connector*'s `probed_mode` list for later use. This list should then in a second step get filtered and all the modes actually supported by the hardware moved to the *connector*'s modes list.

Name

`drm_cvt_mode` — create a modeline based on the CVT algorithm

Synopsis

```
struct drm_display_mode * drm_cvt_mode (struct drm_device * dev, int
hdisplay, int vdisplay, int vrefresh, bool reduced, bool interlaced,
bool margins);
```

Arguments

<i>dev</i>	drm device
<i>hdisplay</i>	hdisplay size
<i>vdisplay</i>	vdisplay size
<i>vrefresh</i>	vrefresh rate
<i>reduced</i>	whether to use reduced blanking
<i>interlaced</i>	whether to compute an interlaced mode
<i>margins</i>	whether to add margins (borders)

Description

This function is called to generate the modeline based on CVT algorithm according to the `hdisplay`, `vdisplay`, `vrefresh`. It is based from the VESA(TM) Coordinated Video Timing Generator by Graham Loveridge April 9, 2003 available at

http

[//www.elo.utfsm.cl/~elo212/docs/CVTd6r1.xls](http://www.elo.utfsm.cl/~elo212/docs/CVTd6r1.xls)

And it is copied from `xf86CVTmode` in `xserver/hw/xfree86/modes/xf86cvt.c`. What I have done is to translate it by using integer calculation.

Returns

The modeline based on the CVT algorithm stored in a `drm_display_mode` object. The display mode object is allocated with `drm_mode_create`. Returns `NULL` when no mode could be allocated.

Name

`drm_gtf_mode_complex` — create the modeline based on the full GTF algorithm

Synopsis

```
struct drm_display_mode * drm_gtf_mode_complex (struct drm_device * dev,  
int hdisplay, int vdisplay, int vrefresh, bool interlaced, int margins,  
int GTF_M, int GTF_2C, int GTF_K, int GTF_2J);
```

Arguments

<i>dev</i>	drm device
<i>hdisplay</i>	hdisplay size
<i>vdisplay</i>	vdisplay size
<i>vrefresh</i>	vrefresh rate.
<i>interlaced</i>	whether to compute an interlaced mode
<i>margins</i>	desired margin (borders) size
<i>GTF_M</i>	extended GTF formula parameters
<i>GTF_2C</i>	extended GTF formula parameters
<i>GTF_K</i>	extended GTF formula parameters
<i>GTF_2J</i>	extended GTF formula parameters

Description

GTF feature blocks specify C and J in multiples of 0.5, so we pass them in here multiplied by two. For a C of 40, pass in 80.

Returns

The modeline based on the full GTF algorithm stored in a `drm_display_mode` object. The display mode object is allocated with `drm_mode_create`. Returns NULL when no mode could be allocated.

Name

`drm_gtf_mode` — create the modeline based on the GTF algorithm

Synopsis

```
struct drm_display_mode * drm_gtf_mode (struct drm_device * dev, int
hdisplay, int vdisplay, int vrefresh, bool interlaced, int margins);
```

Arguments

<i>dev</i>	drm device
<i>hdisplay</i>	hdisplay size
<i>vdisplay</i>	vdisplay size
<i>vrefresh</i>	vrefresh rate.
<i>interlaced</i>	whether to compute an interlaced mode
<i>margins</i>	desired margin (borders) size

Description

return the modeline based on GTF algorithm

This function is to create the modeline based on the GTF algorithm.

Generalized Timing Formula is derived from

GTF Spreadsheet by Andy Morrish (1/5/97)

available at http

[//www.vesa.org](http://www.vesa.org)

And it is copied from the file of `xserver/hw/xfree86/modes/xf86gtf.c`. What I have done is to translate it by using integer calculation. I also refer to the function of `fb_get_mode` in the file of `drivers/video/fbmon.c`

Standard GTF parameters

$M = 600$ $C = 40$ $K = 128$ $J = 20$

Returns

The modeline based on the GTF algorithm stored in a `drm_display_mode` object. The display mode object is allocated with `drm_mode_create`. Returns NULL when no mode could be allocated.

Name

`drm_display_mode_from_videomode` — fill in *dmode* using *vm*,

Synopsis

```
void drm_display_mode_from_videomode (const struct videomode * vm,  
struct drm_display_mode * dmode);
```

Arguments

vm videomode structure to use as source

dmode `drm_display_mode` structure to use as destination

Description

Fills out *dmode* using the display mode specified in *vm*.

Name

`drm_display_mode_to_videomode` — fill in *vm* using *dmode*,

Synopsis

```
void drm_display_mode_to_videomode (const struct drm_display_mode *  
dmode, struct videomode * vm);
```

Arguments

dmode `drm_display_mode` structure to use as source

vm `videomode` structure to use as destination

Description

Fills out *vm* using the display mode specified in *dmode*.

Name

`of_get_drm_display_mode` — get a `drm_display_mode` from devicetree

Synopsis

```
int of_get_drm_display_mode (struct device_node * np, struct drm_display_mode * dmode, int index);
```

Arguments

np device_node with the timing specification

dmode will be set to the return value

index index into the list of display timings in devicetree

Description

This function is expensive and should only be used, if only one mode is to be read from DT. To get multiple modes start with `of_get_display_timings` and work with that instead.

Returns

0 on success, a negative errno code when no of videomode node was found.

Name

`drm_mode_set_name` — set the name on a mode

Synopsis

```
void drm_mode_set_name (struct drm_display_mode * mode);
```

Arguments

mode name will be set in this mode

Description

Set the name of *mode* to a standard format which is <hdisplay>x<vdisplay> with an optional 'i' suffix for interlaced modes.

Name

`drm_mode_hsync` — get the hsync of a mode

Synopsis

```
int drm_mode_hsync (const struct drm_display_mode * mode);
```

Arguments

mode mode

Returns

mode's hsync rate in kHz, rounded to the nearest integer. Calculates the value first if it is not yet set.

Name

`drm_mode_vrefresh` — get the vrefresh of a mode

Synopsis

```
int drm_mode_vrefresh (const struct drm_display_mode * mode);
```

Arguments

mode mode

Returns

mode's vrefresh rate in Hz, rounded to the nearest integer. Calculates the value first if it is not yet set.

Name

`drm_mode_set_crtcinfo` — set CRTC modesetting timing parameters

Synopsis

```
void drm_mode_set_crtcinfo (struct drm_display_mode * p, int adjust_flags);
```

Arguments

p mode

adjust_flags a combination of adjustment flags

Description

Setup the CRTC modesetting timing parameters for *p*, adjusting if necessary.

- The `CRTC_INTERLACE_HALVE_V` flag can be used to halve vertical timings of interlaced modes. - The `CRTC_STEREO_DOUBLE` flag can be used to compute the timings for buffers containing two eyes (only adjust the timings when needed, eg. for “frame packing” or “side by side full”). - The `CRTC_NO_DBLSCAN` and `CRTC_NO_VSCAN` flags request that adjustment **not** be performed for double scan and `vscan > 1` modes respectively.

Name

`drm_mode_copy` — copy the mode

Synopsis

```
void drm_mode_copy (struct drm_display_mode * dst, const struct dr-  
m_display_mode * src);
```

Arguments

dst mode to overwrite

src mode to copy

Description

Copy an existing mode into another mode, preserving the object id and list head of the destination mode.

Name

`drm_mode_duplicate` — allocate and duplicate an existing mode

Synopsis

```
struct drm_display_mode * drm_mode_duplicate (struct drm_device * dev,  
const struct drm_display_mode * mode);
```

Arguments

dev `drm_device` to allocate the duplicated mode for

mode `mode` to duplicate

Description

Just allocate a new mode, copy the existing mode into it, and return a pointer to it. Used to create new instances of established modes.

Returns

Pointer to duplicated mode on success, NULL on error.

Name

`drm_mode_equal` — test modes for equality

Synopsis

```
bool drm_mode_equal (const struct drm_display_mode * mode1, const struct
drm_display_mode * mode2);
```

Arguments

mode1 first mode

mode2 second mode

Description

Check to see if *mode1* and *mode2* are equivalent.

Returns

True if the modes are equal, false otherwise.

Name

`drm_mode_equal_no_clocks` — test modes for equality

Synopsis

```
bool drm_mode_equal_no_clocks (const struct drm_display_mode * mode1,  
const struct drm_display_mode * mode2);
```

Arguments

mode1 first mode

mode2 second mode

Description

Check to see if *mode1* and *mode2* are equivalent, but don't check the pixel clocks.

Returns

True if the modes are equal, false otherwise.

Name

`drm_mode_equal_no_clocks_no_stereo` — test modes for equality

Synopsis

```
bool drm_mode_equal_no_clocks_no_stereo (const struct drm_display_mode
* mode1, const struct drm_display_mode * mode2);
```

Arguments

mode1 first mode

mode2 second mode

Description

Check to see if *mode1* and *mode2* are equivalent, but don't check the pixel clocks nor the stereo layout.

Returns

True if the modes are equal, false otherwise.

Name

`drm_mode_validate_basic` — make sure the mode is somewhat sane

Synopsis

```
enum drm_mode_status drm_mode_validate_basic (const struct drm_display_mode * mode);
```

Arguments

mode mode to check

Description

Check that the mode timings are at least somewhat reasonable. Any hardware specific limits are left up for each driver to check.

Returns

The mode status

Name

`drm_mode_validate_size` — make sure modes adhere to size constraints

Synopsis

```
enum drm_mode_status drm_mode_validate_size (const struct drm_display_mode * mode, int maxX, int maxY);
```

Arguments

mode mode to check

maxX maximum width

maxY maximum height

Description

This function is a helper which can be used to validate modes against size limitations of the DRM device/connector. If a mode is too big its status member is updated with the appropriate validation failure code. The list itself is not changed.

Returns

The mode status

Name

`drm_mode_prune_invalid` — remove invalid modes from mode list

Synopsis

```
void drm_mode_prune_invalid (struct drm_device * dev, struct list_head  
* mode_list, bool verbose);
```

Arguments

<i>dev</i>	DRM device
<i>mode_list</i>	list of modes to check
<i>verbose</i>	be verbose about it

Description

This helper function can be used to prune a display mode list after validation has been completed. All modes whose status is not `MODE_OK` will be removed from the list, and if *verbose* the status code and mode name is also printed to `dmesg`.

Name

`drm_mode_sort` — sort mode list

Synopsis

```
void drm_mode_sort (struct list_head * mode_list);
```

Arguments

mode_list list of `drm_display_mode` structures to sort

Description

Sort *mode_list* by favorability, moving good modes to the head of the list.

Name

`drm_mode_connector_list_update` — update the mode list for the connector

Synopsis

```
void drm_mode_connector_list_update (struct drm_connector * connector);
```

Arguments

connector the connector to update

Description

This moves the modes from the *connector* `probed_modes` list to the actual mode list. It compares the probed mode against the current list and only adds different/new modes.

This is just a helper functions doesn't validate any modes itself and also doesn't prune any invalid modes. Callers need to do that themselves.

Name

`drm_mode_parse_command_line_for_connector` — parse command line modeline for connector

Synopsis

```
bool drm_mode_parse_command_line_for_connector (const char * mode_option,  
struct drm_connector * connector, struct drm_cmdline_mode * mode);
```

Arguments

mode_option optional per connector mode option

connector connector to parse modeline for

mode preallocated `drm_cmdline_mode` structure to fill out

Description

This parses *mode_option* command line modeline for modes and options to configure the connector. If *mode_option* is NULL the default command line modeline in `fb_mode_option` will be parsed instead.

This uses the same parameters as the fb modedb.c, except for an extra force-enable, force-enable-digital and force-disable bit at the end:

```
<xres>x<yres>[M][R][-<bpp>][@<refresh>][i][m][eDd]
```

The intermediate `drm_cmdline_mode` structure is required to store additional options from the command line modline like the force-enable/disable flag.

Returns

True if a valid modeline has been parsed, false otherwise.

Name

`drm_mode_create_from_cmdline_mode` — convert a command line modeline into a DRM display mode

Synopsis

```
struct drm_display_mode * drm_mode_create_from_cmdline_mode (struct dr-  
m_device * dev, struct drm_cmdline_mode * cmd);
```

Arguments

dev DRM device to create the new mode for

cmd input command line modeline

Returns

Pointer to converted mode on success, NULL on error.

Atomic Mode Setting Function Reference

Name

`drm_atomic_state_default_release` — release memory initialized by `drm_atomic_state_init`

Synopsis

```
void drm_atomic_state_default_release (struct drm_atomic_state *  
state);
```

Arguments

state atomic state

Description

Free all the memory allocated by `drm_atomic_state_init`. This is useful for drivers that subclass the atomic state.

Name

`drm_atomic_state_init` — init new atomic state

Synopsis

```
int drm_atomic_state_init (struct drm_device * dev, struct drm_atomic_state * state);
```

Arguments

dev DRM device

state atomic state

Description

Default implementation for filling in a new atomic state. This is useful for drivers that subclass the atomic state.

Name

`drm_atomic_state_alloc` — allocate atomic state

Synopsis

```
struct drm_atomic_state * drm_atomic_state_alloc (struct drm_device *  
dev);
```

Arguments

dev DRM device

Description

This allocates an empty atomic state to track updates.

Name

`drm_atomic_state_default_clear` — clear base atomic state

Synopsis

```
void drm_atomic_state_default_clear (struct drm_atomic_state * state);
```

Arguments

state atomic state

Description

Default implementation for clearing atomic state. This is useful for drivers that subclass the atomic state.

Name

`drm_atomic_state_clear` — clear state object

Synopsis

```
void drm_atomic_state_clear (struct drm_atomic_state * state);
```

Arguments

state atomic state

Description

When the w/w mutex algorithm detects a deadlock we need to back off and drop all locks. So someone else could sneak in and change the current modeset configuration. Which means that all the state assembled in *state* is no longer an atomic update to the current state, but to some arbitrary earlier state. Which could break assumptions the driver's `->atomic_check` likely relies on.

Hence we must clear all cached state and completely start over, using this function.

Name

`drm_atomic_state_free` — free all memory for an atomic state

Synopsis

```
void drm_atomic_state_free (struct drm_atomic_state * state);
```

Arguments

state atomic state to deallocate

Description

This frees all memory associated with an atomic state, including all the per-object state for planes, crtcs and connectors.

Name

`drm_atomic_get_crtc_state` — get crtc state

Synopsis

```
struct drm_crtc_state * drm_atomic_get_crtc_state (struct drm_atomic_s-  
tate * state, struct drm_crtc * crtc);
```

Arguments

state global atomic state object

crtc crtc to get state object for

Description

This function returns the crtc state for the given crtc, allocating it if needed. It will also grab the relevant crtc lock to make sure that the state is consistent.

Returns

Either the allocated state or the error code encoded into the pointer. When the error is EDEADLK then the w/w mutex code has detected a deadlock and the entire atomic sequence must be restarted. All other errors are fatal.

Name

`drm_atomic_set_mode_for_crtc` — set mode for CRTC

Synopsis

```
int drm_atomic_set_mode_for_crtc (struct drm_crtc_state * state, struct
drm_display_mode * mode);
```

Arguments

state the CRTC whose incoming state to update

mode kernel-internal mode to use for the CRTC, or NULL to disable

Description

Set a mode (originating from the kernel) on the desired CRTC state. Does not change any other state properties, including enable, active, or mode_changed.

RETURNS

Zero on success, error code on failure. Cannot return -EDEADLK.

Name

`drm_atomic_set_mode_prop_for_crtc` — set mode for CRTC

Synopsis

```
int drm_atomic_set_mode_prop_for_crtc (struct drm_crtc_state * state,
    struct drm_property_blob * blob);
```

Arguments

state the CRTC whose incoming state to update

blob pointer to blob property to use for mode

Description

Set a mode (originating from a blob property) on the desired CRTC state. This function will take a reference on the blob property for the CRTC state, and release the reference held on the state's existing mode property, if any was set.

RETURNS

Zero on success, error code on failure. Cannot return -EDEADLK.

Name

`drm_atomic_crtc_set_property` — set property on CRTC

Synopsis

```
int drm_atomic_crtc_set_property (struct drm_crtc * crtc, struct dr-  
m_crtc_state * state, struct drm_property * property, uint64_t val);
```

Arguments

<i>crtc</i>	the drm CRTC to set a property on
<i>state</i>	the state object to update with the new property value
<i>property</i>	the property to set
<i>val</i>	the new property value

Description

Use this instead of calling `crtc->atomic_set_property` directly. This function handles generic/core properties and calls out to driver's `->atomic_set_property` for driver properties. To ensure consistent behavior you must call this function rather than the driver hook directly.

RETURNS

Zero on success, error code on failure

Name

`drm_atomic_get_plane_state` — get plane state

Synopsis

```
struct drm_plane_state * drm_atomic_get_plane_state (struct drm_atomic_state * state, struct drm_plane * plane);
```

Arguments

state global atomic state object

plane plane to get state object for

Description

This function returns the plane state for the given plane, allocating it if needed. It will also grab the relevant plane lock to make sure that the state is consistent.

Returns

Either the allocated state or the error code encoded into the pointer. When the error is EDEADLK then the w/w mutex code has detected a deadlock and the entire atomic sequence must be restarted. All other errors are fatal.

Name

`drm_atomic_plane_set_property` — set property on plane

Synopsis

```
int drm_atomic_plane_set_property (struct drm_plane * plane, struct
drm_plane_state * state, struct drm_property * property, uint64_t val);
```

Arguments

<i>plane</i>	the drm plane to set a property on
<i>state</i>	the state object to update with the new property value
<i>property</i>	the property to set
<i>val</i>	the new property value

Description

Use this instead of calling `plane->atomic_set_property` directly. This function handles generic/core properties and calls out to driver's `->atomic_set_property` for driver properties. To ensure consistent behavior you must call this function rather than the driver hook directly.

RETURNS

Zero on success, error code on failure

Name

`drm_atomic_get_connector_state` — get connector state

Synopsis

```
struct drm_connector_state * drm_atomic_get_connector_state (struct dr-  
m_atomic_state * state, struct drm_connector * connector);
```

Arguments

state global atomic state object

connector connector to get state object for

Description

This function returns the connector state for the given connector, allocating it if needed. It will also grab the relevant connector lock to make sure that the state is consistent.

Returns

Either the allocated state or the error code encoded into the pointer. When the error is EDEADLK then the w/w mutex code has detected a deadlock and the entire atomic sequence must be restarted. All other errors are fatal.

Name

`drm_atomic_connector_set_property` — set property on connector.

Synopsis

```
int drm_atomic_connector_set_property (struct drm_connector * connector, struct drm_connector_state * state, struct drm_property * property, uint64_t val);
```

Arguments

<i>connector</i>	the drm connector to set a property on
<i>state</i>	the state object to update with the new property value
<i>property</i>	the property to set
<i>val</i>	the new property value

Description

Use this instead of calling `connector->atomic_set_property` directly. This function handles generic/core properties and calls out to driver's `->atomic_set_property` for driver properties. To ensure consistent behavior you must call this function rather than the driver hook directly.

RETURNS

Zero on success, error code on failure

Name

`drm_atomic_set_crtc_for_plane` — set crtc for plane

Synopsis

```
int drm_atomic_set_crtc_for_plane (struct drm_plane_state * plane_s-  
tate, struct drm_crtc * crtc);
```

Arguments

plane_state the plane whose incoming state to update

crtc crtc to use for the plane

Description

Changing the assigned crtc for a plane requires us to grab the lock and state for the new crtc, as needed. This function takes care of all these details besides updating the pointer in the state object itself.

Returns

0 on success or can fail with -EDEADLK or -ENOMEM. When the error is EDEADLK then the w/w mutex code has detected a deadlock and the entire atomic sequence must be restarted. All other errors are fatal.

Name

`drm_atomic_set_fb_for_plane` — set framebuffer for plane

Synopsis

```
void drm_atomic_set_fb_for_plane (struct drm_plane_state * plane_state,  
struct drm_framebuffer * fb);
```

Arguments

plane_state atomic state object for the plane

fb fb to use for the plane

Description

Changing the assigned framebuffer for a plane requires us to grab a reference to the new fb and drop the reference to the old fb, if there is one. This function takes care of all these details besides updating the pointer in the state object itself.

Name

`drm_atomic_set_crtc_for_connector` — set crtc for connector

Synopsis

```
int drm_atomic_set_crtc_for_connector (struct drm_connector_state *  
conn_state, struct drm_crtc * crtc);
```

Arguments

conn_state atomic state object for the connector

crtc crtc to use for the connector

Description

Changing the assigned crtc for a connector requires us to grab the lock and state for the new crtc, as needed. This function takes care of all these details besides updating the pointer in the state object itself.

Returns

0 on success or can fail with -EDEADLK or -ENOMEM. When the error is EDEADLK then the w/w mutex code has detected a deadlock and the entire atomic sequence must be restarted. All other errors are fatal.

Name

`drm_atomic_add_affected_connectors` — add connectors for *crtc*

Synopsis

```
int  drm_atomic_add_affected_connectors (struct  drm_atomic_state  *  
state, struct  drm_crtc  * crtc);
```

Arguments

state atomic state

crtc DRM *crtc*

Description

This function walks the current configuration and adds all connectors currently using *crtc* to the atomic configuration *state*. Note that this function must acquire the connection mutex. This can potentially cause unneeded serialization if the update is just for the planes on one *crtc*. Hence drivers and helpers should only call this when really needed (e.g. when a full modeset needs to happen due to some change).

Returns

0 on success or can fail with -EDEADLK or -ENOMEM. When the error is EDEADLK then the w/w mutex code has detected a deadlock and the entire atomic sequence must be restarted. All other errors are fatal.

Name

`drm_atomic_add_affected_planes` — add planes for crtc

Synopsis

```
int drm_atomic_add_affected_planes (struct drm_atomic_state * state,
struct drm_crtc * crtc);
```

Arguments

state atomic state

crtc DRM crtc

Description

This function walks the current configuration and adds all planes currently used by *crtc* to the atomic configuration *state*. This is useful when an atomic commit also needs to check all currently enabled plane on *crtc*, e.g. when changing the mode. It's also useful when re-enabling a CRTC to avoid special code to force-enable all planes.

Since acquiring a plane state will always also acquire the w/w mutex of the current CRTC for that plane (if there is any) adding all the plane states for a CRTC will not reduce parallism of atomic updates.

Returns

0 on success or can fail with -EDEADLK or -ENOMEM. When the error is EDEADLK then the w/w mutex code has detected a deadlock and the entire atomic sequence must be restarted. All other errors are fatal.

Name

`drm_atomic_legacy_backoff` — locking backoff for legacy ioctls

Synopsis

```
void drm_atomic_legacy_backoff (struct drm_atomic_state * state);
```

Arguments

state atomic state

Description

This function should be used by legacy entry points which don't understand -EDEADLK semantics. For simplicity this one will grab all modeset locks after the slowpath completed.

Name

`drm_atomic_check_only` — check whether a given config would work

Synopsis

```
int drm_atomic_check_only (struct drm_atomic_state * state);
```

Arguments

state atomic configuration to check

Description

Note that this function can return `-EDEADLK` if the driver needed to acquire more locks but encountered a deadlock. The caller must then do the usual w/w backoff dance and restart. All other errors are fatal.

Returns

0 on success, negative error code on failure.

Name

`drm_atomic_commit` — commit configuration atomically

Synopsis

```
int drm_atomic_commit (struct drm_atomic_state * state);
```

Arguments

state atomic configuration to check

Description

Note that this function can return `-EDEADLK` if the driver needed to acquire more locks but encountered a deadlock. The caller must then do the usual w/w backoff dance and restart. All other errors are fatal.

Also note that on successful execution ownership of *state* is transferred from the caller of this function to the function itself. The caller must not free or in any other way access *state*. If the function fails then the caller must clean up *state* itself.

Returns

0 on success, negative error code on failure.

Name

`drm_atomic_async_commit` — atomicasync configuration commit

Synopsis

```
int drm_atomic_async_commit (struct drm_atomic_state * state);
```

Arguments

state atomic configuration to check

Description

Note that this function can return `-EDEADLK` if the driver needed to acquire more locks but encountered a deadlock. The caller must then do the usual w/w backoff dance and restart. All other errors are fatal.

Also note that on successful execution ownership of *state* is transferred from the caller of this function to the function itself. The caller must not free or in any other way access *state*. If the function fails then the caller must clean up *state* itself.

Returns

0 on success, negative error code on failure.

Name

`drm_atomic_clean_old_fb` — - Unset `old_fb` pointers and set `plane->fb` pointers.

Synopsis

```
void drm_atomic_clean_old_fb (struct drm_device * dev, unsigned
plane_mask, int ret);
```

Arguments

<i>dev</i>	drm device to check.
<i>plane_mask</i>	plane mask for planes that were updated.
<i>ret</i>	return value, can be <code>-EDEADLK</code> for a retry.

Description

Before doing an update `plane->old_fb` is set to `plane->fb`, but before dropping the locks `old_fb` needs to be set to `NULL` and `plane->fb` updated. This is a common operation for each atomic update, so this call is split off as a helper.

Name

`drm_atomic_crtc_get_property` — get property value from CRTC state

Synopsis

```
int drm_atomic_crtc_get_property (struct drm_crtc * crtc, const struct
drm_crtc_state * state, struct drm_property * property, uint64_t * val);
```

Arguments

<i>crtc</i>	the drm CRTC to set a property on
<i>state</i>	the state object to get the property value from
<i>property</i>	the property to set
<i>val</i>	return location for the property value

Description

This function handles generic/core properties and calls out to driver's `->atomic_get_property` for driver properties. To ensure consistent behavior you must call this function rather than the driver hook directly.

RETURNS

Zero on success, error code on failure

Name

`drm_atomic_crtc_check` — check crtc state

Synopsis

```
int drm_atomic_crtc_check (struct drm_crtc * crtc, struct drm_crtc_state  
* state);
```

Arguments

crtc crtc to check

state crtc state to check

Description

Provides core sanity checks for crtc state.

RETURNS

Zero on success, error code on failure

Name

`drm_atomic_plane_get_property` — get property value from plane state

Synopsis

```
int drm_atomic_plane_get_property (struct drm_plane * plane, const
struct drm_plane_state * state, struct drm_property * property, uin-
t64_t * val);
```

Arguments

<i>plane</i>	the drm plane to set a property on
<i>state</i>	the state object to get the property value from
<i>property</i>	the property to set
<i>val</i>	return location for the property value

Description

This function handles generic/core properties and calls out to driver's `->atomic_get_property` for driver properties. To ensure consistent behavior you must call this function rather than the driver hook directly.

RETURNS

Zero on success, error code on failure

Name

`drm_atomic_plane_check` — check plane state

Synopsis

```
int drm_atomic_plane_check (struct drm_plane * plane, struct dr-  
m_plane_state * state);
```

Arguments

plane plane to check

state plane state to check

Description

Provides core sanity checks for plane state.

RETURNS

Zero on success, error code on failure

Name

`drm_atomic_connector_get_property` — get property value from connector state

Synopsis

```
int drm_atomic_connector_get_property (struct drm_connector * connector,
const struct drm_connector_state * state, struct drm_property *
property, uint64_t * val);
```

Arguments

connector the drm connector to set a property on

state the state object to get the property value from

property the property to set

val return location for the property value

Description

This function handles generic/core properties and calls out to driver's `->atomic_get_property` for driver properties. To ensure consistent behavior you must call this function rather than the driver hook directly.

RETURNS

Zero on success, error code on failure

Frame Buffer Abstraction

Frame buffers are abstract memory objects that provide a source of pixels to scanout to a CRTC. Applications explicitly request the creation of frame buffers through the `DRM_IOCTL_MODE_ADDFB(2)` ioctls and receive an opaque handle that can be passed to the KMS CRTC control, plane configuration and page flip functions.

Frame buffers rely on the underneath memory manager for low-level memory operations. When creating a frame buffer applications pass a memory handle (or a list of memory handles for multi-planar formats) through the `drm_mode_fb_cmd2` argument. For drivers using GEM as their userspace buffer management interface this would be a GEM handle. Drivers are however free to use their own backing storage object handles, e.g. `vmwgfx` directly exposes special TTM handles to userspace and so expects TTM handles in the create ioctl and not GEM handles.

The lifetime of a drm framebuffer is controlled with a reference count, drivers can grab additional references with `drm_framebuffer_reference` and drop them again with `drm_framebuffer_unreference`. For driver-private framebuffers for which the last reference is never dropped (e.g. for the fbdev framebuffer when the struct `drm_framebuffer` is embedded into the fbdev helper struct) drivers can manually clean up a framebuffer at module unload time with `drm_framebuffer_unregister_private`.

Dumb Buffer Objects

The KMS API doesn't standardize backing storage object creation and leaves it to driver-specific ioctls. Furthermore actually creating a buffer object even for GEM-based drivers is done through a driver-specific

ioctl - GEM only has a common userspace interface for sharing and destroying objects. While not an issue for full-fledged graphics stacks that include device-specific userspace components (in libdrm for instance), this limit makes DRM-based early boot graphics unnecessarily complex.

Dumb objects partly alleviate the problem by providing a standard API to create dumb buffers suitable for scanout, which can then be used to create KMS frame buffers.

To support dumb objects drivers must implement the `dumb_create`, `dumb_destroy` and `dumb_map_offset` operations.

- `int (*dumb_create)(struct drm_file *file_priv, struct drm_device *dev, struct drm_mode_create_dumb *args);`

The `dumb_create` operation creates a driver object (GEM or TTM handle) suitable for scanout based on the width, height and depth from the `struct drm_mode_create_dumb` argument. It fills the argument's *handle*, *pitch* and *size* fields with a handle for the newly created object and its line pitch and size in bytes.

- `int (*dumb_destroy)(struct drm_file *file_priv, struct drm_device *dev, uint32_t handle);`

The `dumb_destroy` operation destroys a dumb object created by `dumb_create`.

- `int (*dumb_map_offset)(struct drm_file *file_priv, struct drm_device *dev, uint32_t handle, uint64_t *offset);`

The `dumb_map_offset` operation associates an mmap fake offset with the object given by the handle and returns it. Drivers must use the `drm_gem_create_mmap_offset` function to associate the fake offset as described in the section called “GEM Objects Mapping”.

Note that dumb objects may not be used for gpu acceleration, as has been attempted on some ARM embedded platforms. Such drivers really must have a hardware-specific ioctl to allocate suitable buffer objects.

Output Polling

```
void (*output_poll_changed)(struct drm_device *dev);
```

This operation notifies the driver that the status of one or more connectors has changed. Drivers that use the fb helper can just call the `drm_fb_helper_hotplug_event` function to handle this operation.

Locking

Beside some lookup structures with their own locking (which is hidden behind the interface functions) most of the modeset state is protected by the `dev-<mode_config.lock` mutex and additionally per-crtc locks to allow cursor updates, pageflips and similar operations to occur concurrently with background tasks like output detection. Operations which cross domains like a full modeset always grab all locks. Drivers there need to protect resources shared between crtcs with additional locking. They also need to be careful to always grab the relevant crtc locks if a modeset function touches crtc state, e.g. for load detection (which does only grab the `mode_config.lock` to allow concurrent screen updates on live crtcs).

KMS Initialization and Cleanup

A KMS device is abstracted and exposed as a set of planes, CRTCs, encoders and connectors. KMS drivers must thus create and initialize all those objects at load time after initializing mode setting.

CRTCs (struct `drm_crtc`)

A CRTC is an abstraction representing a part of the chip that contains a pointer to a scanout buffer. Therefore, the number of CRTCs available determines how many independent scanout buffers can be active at any given time. The CRTC structure contains several fields to support this: a pointer to some video memory (abstracted as a frame buffer object), a display mode, and an (x, y) offset into the video memory to support panning or configurations where one piece of video memory spans multiple CRTCs.

CRTC Initialization

A KMS device must create and register at least one struct `drm_crtc` instance. The instance is allocated and zeroed by the driver, possibly as part of a larger structure, and registered with a call to `drm_crtc_init` with a pointer to CRTC functions.

Planes (struct `drm_plane`)

A plane represents an image source that can be blended with or overlayed on top of a CRTC during the scanout process. Planes are associated with a frame buffer to crop a portion of the image memory (source) and optionally scale it to a destination size. The result is then blended with or overlayed on top of a CRTC.

The DRM core recognizes three types of planes:

- `DRM_PLANE_TYPE_PRIMARY` represents a "main" plane for a CRTC. Primary planes are the planes operated upon by CRTC modesetting and flipping operations described in the `page_flip` hook in `drm_crtc_funcs`.
- `DRM_PLANE_TYPE_CURSOR` represents a "cursor" plane for a CRTC. Cursor planes are the planes operated upon by the `DRM_IOCTL_MODE_CURSOR` and `DRM_IOCTL_MODE_CURSOR2` ioctls.
- `DRM_PLANE_TYPE_OVERLAY` represents all non-primary, non-cursor planes. Some drivers refer to these types of planes as "sprites" internally.

For compatibility with legacy userspace, only overlay planes are made available to userspace by default. Userspace clients may set the `DRM_CLIENT_CAP_UNIVERSAL_PLANES` client capability bit to indicate that they wish to receive a universal plane list containing all plane types.

Plane Initialization

To create a plane, a KMS driver allocates and zeroes an instance of struct `drm_plane` (possibly as part of a larger structure) and registers it with a call to `drm_universal_plane_init`. The function takes a bitmask of the CRTCs that can be associated with the plane, a pointer to the plane functions, a list of format supported formats, and the type of plane (primary, cursor, or overlay) being initialized.

Cursor and overlay planes are optional. All drivers should provide one primary plane per CRTC (although this requirement may change in the future); drivers that do not wish to provide special handling for primary planes may make use of the helper functions described in the section called "Plane Helper Reference" to create and register a primary plane with standard capabilities.

Encoders (struct `drm_encoder`)

An encoder takes pixel data from a CRTC and converts it to a format suitable for any attached connectors. On some devices, it may be possible to have a CRTC send data to more than one encoder. In that case, both encoders would receive data from the same scanout buffer, resulting in a "cloned" display configuration across the connectors attached to each encoder.

Encoder Initialization

As for CRTC, a KMS driver must create, initialize and register at least one struct `drm_encoder` instance. The instance is allocated and zeroed by the driver, possibly as part of a larger structure.

Drivers must initialize the struct `drm_encoder` *possible_crtcs* and *possible_clones* fields before registering the encoder. Both fields are bitmasks of respectively the CRTC that the encoder can be connected to, and sibling encoders candidate for cloning.

After being initialized, the encoder must be registered with a call to `drm_encoder_init`. The function takes a pointer to the encoder functions and an encoder type. Supported types are

- `DRM_MODE_ENCODER_DAC` for VGA and analog on DVI-I/DVI-A
- `DRM_MODE_ENCODER_TMDS` for DVI, HDMI and (embedded) DisplayPort
- `DRM_MODE_ENCODER_LVDS` for display panels
- `DRM_MODE_ENCODER_TVDAC` for TV output (Composite, S-Video, Component, SCART)
- `DRM_MODE_ENCODER_VIRTUAL` for virtual machine displays

Encoders must be attached to a CRTC to be used. DRM drivers leave encoders unattached at initialization time. Applications (or the fbdev compatibility layer when implemented) are responsible for attaching the encoders they want to use to a CRTC.

Connectors (struct `drm_connector`)

A connector is the final destination for pixel data on a device, and usually connects directly to an external display device like a monitor or laptop panel. A connector can only be attached to one encoder at a time. The connector is also the structure where information about the attached display is kept, so it contains fields for display data, EDID data, DPMS & connection status, and information about modes supported on the attached displays.

Connector Initialization

Finally a KMS driver must create, initialize, register and attach at least one struct `drm_connector` instance. The instance is created as other KMS objects and initialized by setting the following fields.

<i>interlace_allowed</i>	Whether the connector can handle interlaced modes.
<i>doublescan_allowed</i>	Whether the connector can handle doublescan.
<i>display_info</i>	Display information is filled from EDID information when a display is detected. For non hot-pluggable displays such as flat panels in embedded systems, the driver should initialize the <i>display_info.width_mm</i> and <i>display_info.height_mm</i> fields with the physical size of the display.
<i>polled</i>	Connector polling mode, a combination of <code>DRM_CONNECTOR_POLL_HPD</code> The connector generates hotplug events and doesn't need to be periodically polled. The <code>CONNECT</code> and

DISCONNECT flags must not be set together with the HPD flag.

DRM_CONNECTOR_POLL_CONNECT Periodically poll the connector for connection.

DRM_CONNECTOR_POLL_DISCONNECT Periodically poll the connector for disconnection.
Set to 0 for connectors that don't support connection status discovery.

The connector is then registered with a call to `drm_connector_init` with a pointer to the connector functions and a connector type, and exposed through sysfs with a call to `drm_connector_register`.

Supported connector types are

- `DRM_MODE_CONNECTOR_VGA`
- `DRM_MODE_CONNECTOR_DVII`
- `DRM_MODE_CONNECTOR_DVID`
- `DRM_MODE_CONNECTOR_DVIA`
- `DRM_MODE_CONNECTOR_Composite`
- `DRM_MODE_CONNECTOR_SVIDEO`
- `DRM_MODE_CONNECTOR_LVDS`
- `DRM_MODE_CONNECTOR_Component`
- `DRM_MODE_CONNECTOR_9PinDIN`
- `DRM_MODE_CONNECTOR_DisplayPort`
- `DRM_MODE_CONNECTOR_HDMIA`
- `DRM_MODE_CONNECTOR_HDMIB`
- `DRM_MODE_CONNECTOR_TV`
- `DRM_MODE_CONNECTOR_eDP`
- `DRM_MODE_CONNECTOR_VIRTUAL`

Connectors must be attached to an encoder to be used. For devices that map connectors to encoders 1:1, the connector should be attached at initialization time with a call to `drm_mode_connector_attach_encoder`. The driver must also set the `drm_connector` *encoder* field to point to the attached encoder.

Finally, drivers must initialize the connectors state change detection with a call to `drm_kms_helper_poll_init`. If at least one connector is pollable but can't generate hotplug interrupts (indicated by the `DRM_CONNECTOR_POLL_CONNECT` and `DRM_CONNECTOR_POLL_DISCONNECT` connector flags), a delayed work will automatically be queued to periodically poll for changes. Connectors that can generate hotplug interrupts must be marked with the `DRM_CONNECTOR_POLL_HPD`

flag instead, and their interrupt handler must call `drm_helper_hpd_irq_event`. The function will queue a delayed work to check the state of all connectors, but no periodic polling will be done.

Connector Operations

Note

Unless otherwise state, all operations are mandatory.

DPMS

```
void (*dpms)(struct drm_connector *connector, int mode);
```

The DPMS operation sets the power state of a connector. The mode argument is one of

- `DRM_MODE_DPMS_ON`
- `DRM_MODE_DPMS_STANDBY`
- `DRM_MODE_DPMS_SUSPEND`
- `DRM_MODE_DPMS_OFF`

In all but `DPMS_ON` mode the encoder to which the connector is attached should put the display in low-power mode by driving its signals appropriately. If more than one connector is attached to the encoder care should be taken not to change the power state of other displays as a side effect. Low-power mode should be propagated to the encoders and CRTC's when all related connectors are put in low-power mode.

Modes

```
int (*fill_modes)(struct drm_connector *connector, uint32_t max_width,
                  uint32_t max_height);
```

Fill the mode list with all supported modes for the connector. If the *max_width* and *max_height* arguments are non-zero, the implementation must ignore all modes wider than *max_width* or higher than *max_height*.

The connector must also fill in this operation its *display_info width_mm* and *height_mm* fields with the connected display physical size in millimeters. The fields should be set to 0 if the value isn't known or is not applicable (for instance for projector devices).

Connection Status

The connection status is updated through polling or hotplug events when supported (see *polled*). The status value is reported to userspace through `ioctl`s and must not be used inside the driver, as it only gets initialized by a call to `drm_mode_getconnector` from userspace.

```
enum drm_connector_status (*detect)(struct drm_connector *connector,
                                    bool force);
```

Check to see if anything is attached to the connector. The *force* parameter is set to false whilst polling or to true when checking the connector due to user request. *force* can be used by the driver to avoid expensive, destructive operations during automated probing.

Return `connector_status_connected` if something is connected to the connector, `connector_status_disconnected` if nothing is connected and `connector_status_unknown` if the connection state isn't known.

Drivers should only return `connector_status_connected` if the connection status has really been probed as connected. Connectors that can't detect the connection status, or failed connection status probes, should return `connector_status_unknown`.

Cleanup

The DRM core manages its objects' lifetime. When an object is not needed anymore the core calls its destroy function, which must clean up and free every resource allocated for the object. Every `drm_*_init` call must be matched with a corresponding `drm_*_cleanup` call to cleanup CRTC's (`drm_crtc_cleanup`), planes (`drm_plane_cleanup`), encoders (`drm_encoder_cleanup`) and connectors (`drm_connector_cleanup`). Furthermore, connectors that have been added to sysfs must be removed by a call to `drm_connector_unregister` before calling `drm_connector_cleanup`.

Connectors state change detection must be cleanup up with a call to `drm_kms_helper_poll_fini`.

Output discovery and initialization example

```
void intel_crt_init(struct drm_device *dev)
{
    struct drm_connector *connector;
    struct intel_output *intel_output;

    intel_output = kzalloc(sizeof(struct intel_output), GFP_KERNEL);
    if (!intel_output)
        return;

    connector = &intel_output->base;
    drm_connector_init(dev, &intel_output->base,
                      &intel_crt_connector_funcs, DRM_MODE_CONNECTOR_VGA);

    drm_encoder_init(dev, &intel_output->enc, &intel_crt_enc_funcs,
                    DRM_MODE_ENCODER_DAC);

    drm_mode_connector_attach_encoder(&intel_output->base,
                                     &intel_output->enc);

    /* Set up the DDC bus. */
    intel_output->ddc_bus = intel_i2c_create(dev, GPIOA, "CRTDDC_A");
    if (!intel_output->ddc_bus) {
        dev_printk(KERN_ERR, &dev->pdev->dev, "DDC bus registration "
                  "failed.\n");
        return;
    }

    intel_output->type = INTEL_OUTPUT_ANALOG;
    connector->interlace_allowed = 0;
    connector->doublescan_allowed = 0;

    drm_encoder_helper_add(&intel_output->enc, &intel_crt_helper_funcs);
    drm_connector_helper_add(connector, &intel_crt_connector_helper_funcs);
}
```

```
drm_connector_register(connector);  
}
```

In the example above (taken from the i915 driver), a CRTC, connector and encoder combination is created. A device-specific i2c bus is also created for fetching EDID data and performing monitor detection. Once the process is complete, the new connector is registered with sysfs to make its properties available to applications.

KMS API Functions

Name

`drm_get_connector_status_name` — return a string for connector status

Synopsis

```
const char * drm_get_connector_status_name (enum drm_connector_status
status);
```

Arguments

status connector status to compute name of

Description

In contrast to the other `drm_get_*_name` functions this one here returns a const pointer and hence is thread-safe.

Name

`drm_get_subpixel_order_name` — return a string for a given subpixel enum

Synopsis

```
const char * drm_get_subpixel_order_name (enum subpixel_order order);
```

Arguments

order enum of subpixel_order

Description

Note you could abuse this and return something out of bounds, but that would be a caller error. No unscrubbed user data should make it here.

Name

`drm_get_format_name` — return a string for drm fourcc format

Synopsis

```
const char * drm_get_format_name (uint32_t format);
```

Arguments

format format to compute name of

Description

Note that the buffer used by this function is globally shared and owned by the function itself.

FIXME

This isn't really multithreading safe.

Name

`drm_mode_object_find` — look up a drm object with static lifetime

Synopsis

```
struct drm_mode_object * drm_mode_object_find (struct drm_device * dev,  
uint32_t id, uint32_t type);
```

Arguments

dev drm device

id id of the mode object

type type of the mode object

Description

Note that framebuffers cannot be looked up with this functions - since those are reference counted, they need special treatment. Even with `DRM_MODE_OBJECT_ANY` (although that will simply return `NULL` rather than `WARN_ON`).

Name

`drm_framebuffer_init` — initialize a framebuffer

Synopsis

```
int drm_framebuffer_init (struct drm_device * dev, struct drm_frame-  
buffer * fb, const struct drm_framebuffer_funcs * funcs);
```

Arguments

dev DRM device

fb framebuffer to be initialized

funcs ... with these functions

Description

Allocates an ID for the framebuffer's parent mode object, sets its mode functions & device file and adds it to the master fd list.

IMPORTANT

This functions publishes the fb and makes it available for concurrent access by other users. Which means by this point the fb `_must_` be fully set up - since all the fb attributes are invariant over its lifetime, no further locking but only correct reference counting is required.

Returns

Zero on success, error code on failure.

Name

`drm_framebuffer_lookup` — look up a drm framebuffer and grab a reference

Synopsis

```
struct drm_framebuffer * drm_framebuffer_lookup (struct drm_device *  
dev, uint32_t id);
```

Arguments

dev drm device

id id of the fb object

Description

If successful, this grabs an additional reference to the framebuffer - callers need to make sure to eventually unreference the returned framebuffer again, using `drm_framebuffer_unreference`.

Name

`drm_framebuffer_unreference` — unref a framebuffer

Synopsis

```
void drm_framebuffer_unreference (struct drm_framebuffer * fb);
```

Arguments

fb framebuffer to unref

Description

This functions decrements the fb's refcount and frees it if it drops to zero.

Name

`drm_framebuffer_reference` — incr the fb refcnt

Synopsis

```
void drm_framebuffer_reference (struct drm_framebuffer * fb);
```

Arguments

fb framebuffer

Description

This functions increments the fb's refcount.

Name

`drm_framebuffer_unregister_private` — unregister a private fb from the lookup idr

Synopsis

```
void drm_framebuffer_unregister_private (struct drm_framebuffer * fb);
```

Arguments

fb fb to unregister

Description

Drivers need to call this when cleaning up driver-private framebuffers, e.g. those used for fbdev. Note that the caller must hold a reference of it's own, i.e. the object may not be destroyed through this call (since it'll lead to a locking inversion).

Name

`drm_framebuffer_cleanup` — remove a framebuffer object

Synopsis

```
void drm_framebuffer_cleanup (struct drm_framebuffer * fb);
```

Arguments

fb framebuffer to remove

Description

Cleanup framebuffer. This function is intended to be used from the drivers `->destroy` callback. It can also be used to clean up driver private framebuffers embedded into a larger structure.

Note that this function does not remove the fb from active usage - if it is still used anywhere, hilarity can ensue since userspace could call `getfb` on the id and get back `-EINVAL`. Obviously no concern at driver unload time.

Also, the framebuffer will not be removed from the lookup idr - for user-created framebuffers this will happen in in the `rmfb` ioctl. For driver-private objects (e.g. for fbdev) drivers need to explicitly call `drm_framebuffer_unregister_private`.

Name

`drm_framebuffer_remove` — remove and unreference a framebuffer object

Synopsis

```
void drm_framebuffer_remove (struct drm_framebuffer * fb);
```

Arguments

fb framebuffer to remove

Description

Scans all the CRTC's and planes in *dev*'s `mode_config`. If they're using *fb*, removes it, setting it to NULL. Then drops the reference to the passed-in framebuffer. Might take the modeset locks.

Note that this function optimizes the cleanup away if the caller holds the last reference to the framebuffer. It is also guaranteed to not take the modeset locks in this case.

Name

`drm_crtc_init_with_planes` — Initialise a new CRTC object with specified primary and cursor planes.

Synopsis

```
int drm_crtc_init_with_planes (struct drm_device * dev, struct drm_crtc
* crtc, struct drm_plane * primary, struct drm_plane * cursor, const
struct drm_crtc_funcs * funcs, const char * name, ...);
```

Arguments

<i>dev</i>	DRM device
<i>crtc</i>	CRTC object to init
<i>primary</i>	Primary plane for CRTC
<i>cursor</i>	Cursor plane for CRTC
<i>funcs</i>	callbacks for the new CRTC
<i>name</i>	printf style format string for the CRTC name, or NULL for default name
<i>...</i>	variable arguments

Description

Initialises a new object created as base part of a driver crtc object.

Returns

Zero on success, error code on failure.

Name

`drm_crtc_cleanup` — Clean up the core crtc usage

Synopsis

```
void drm_crtc_cleanup (struct drm_crtc * crtc);
```

Arguments

crtc CRTC to cleanup

Description

This function cleans up *crtc* and removes it from the DRM mode setting core. Note that the function does **not** free the crtc structure itself, this is the responsibility of the caller.

Name

`drm_crtc_index` — find the index of a registered CRTC

Synopsis

```
unsigned int drm_crtc_index (struct drm_crtc * crtc);
```

Arguments

crtc CRTC to find index for

Description

Given a registered CRTC, return the index of that CRTC within a DRM device's list of CRTCs.

Name

`drm_display_info_set_bus_formats` — set the supported bus formats

Synopsis

```
int drm_display_info_set_bus_formats (struct drm_display_info * info,  
const u32 * formats, unsigned int num_formats);
```

Arguments

<i>info</i>	display info to store bus formats in
<i>formats</i>	array containing the supported bus formats
<i>num_formats</i>	the number of entries in the fmts array

Description

Store the supported bus formats in display info structure. See `MEDIA_BUS_FMT_*` definitions in `include/uapi/linux/media-bus-format.h` for a full list of available formats.

Name

`drm_connector_init` — Init a preallocated connector

Synopsis

```
int drm_connector_init (struct drm_device * dev, struct drm_connector * connector, const struct drm_connector_funcs * funcs, int connector_type);
```

Arguments

<i>dev</i>	DRM device
<i>connector</i>	the connector to init
<i>funcs</i>	callbacks for this connector
<i>connector_type</i>	user visible type of the connector

Description

Initialises a preallocated connector. Connectors should be subclassed as part of driver connector objects.

Returns

Zero on success, error code on failure.

Name

`drm_connector_cleanup` — cleans up an initialised connector

Synopsis

```
void drm_connector_cleanup (struct drm_connector * connector);
```

Arguments

connector connector to cleanup

Description

Cleans up the connector but doesn't free the object.

Name

`drm_connector_index` — find the index of a registered connector

Synopsis

```
unsigned int drm_connector_index (struct drm_connector * connector);
```

Arguments

connector connector to find index for

Description

Given a registered connector, return the index of that connector within a DRM device's list of connectors.

Name

`drm_connector_register` — register a connector

Synopsis

```
int drm_connector_register (struct drm_connector * connector);
```

Arguments

connector the connector to register

Description

Register userspace interfaces for a connector

Returns

Zero on success, error code on failure.

Name

`drm_connector_unregister` — unregister a connector

Synopsis

```
void drm_connector_unregister (struct drm_connector * connector);
```

Arguments

connector the connector to unregister

Description

Unregister userspace interfaces for a connector

Name

`drm_connector_unplug_all` — unregister connector userspace interfaces

Synopsis

```
void drm_connector_unplug_all (struct drm_device * dev);
```

Arguments

dev drm device

Description

This function unregisters all connector userspace interfaces in sysfs. Should be call when the device is disconnected, e.g. from an usb driver's ->disconnect callback.

Name

`drm_encoder_init` — Init a preallocated encoder

Synopsis

```
int drm_encoder_init (struct drm_device * dev, struct drm_encoder *  
encoder, const struct drm_encoder_funcs * funcs, int encoder_type, const  
char * name, ...);
```

Arguments

<i>dev</i>	drm device
<i>encoder</i>	the encoder to init
<i>funcs</i>	callbacks for this encoder
<i>encoder_type</i>	user visible type of the encoder
<i>name</i>	printf style format string for the encoder name, or NULL for default name
...	variable arguments

Description

Initialises a preallocated encoder. Encoder should be subclassed as part of driver encoder objects.

Returns

Zero on success, error code on failure.

Name

`drm_encoder_cleanup` — cleans up an initialised encoder

Synopsis

```
void drm_encoder_cleanup (struct drm_encoder * encoder);
```

Arguments

encoder encoder to cleanup

Description

Cleans up the encoder but doesn't free the object.

Name

`drm_universal_plane_init` — Initialize a new universal plane object

Synopsis

```
int drm_universal_plane_init (struct drm_device * dev, struct drm_plane
* plane, unsigned long possible_crtcs, const struct drm_plane_funcs
* funcs, const uint32_t * formats, unsigned int format_count, enum
drm_plane_type type, const char * name, ...);
```

Arguments

<i>dev</i>	DRM device
<i>plane</i>	plane object to init
<i>possible_crtcs</i>	bitmask of possible CRTCs
<i>funcs</i>	callbacks for the new plane
<i>formats</i>	array of supported formats (<code>DRM_FORMAT_*</code>)
<i>format_count</i>	number of elements in <i>formats</i>
<i>type</i>	type of plane (overlay, primary, cursor)
<i>name</i>	printf style format string for the plane name, or NULL for default name
<i>...</i>	variable arguments

Description

Initializes a plane object of type *type*.

Returns

Zero on success, error code on failure.

Name

`drm_plane_init` — Initialize a legacy plane

Synopsis

```
int drm_plane_init (struct drm_device * dev, struct drm_plane * plane,
unsigned long possible_crtcs, const struct drm_plane_funcs * funcs,
const uint32_t * formats, unsigned int format_count, bool is_primary);
```

Arguments

<i>dev</i>	DRM device
<i>plane</i>	plane object to init
<i>possible_crtcs</i>	bitmask of possible CRTCs
<i>funcs</i>	callbacks for the new plane
<i>formats</i>	array of supported formats (<code>DRM_FORMAT_*</code>)
<i>format_count</i>	number of elements in <i>formats</i>
<i>is_primary</i>	plane type (primary vs overlay)

Description

Legacy API to initialize a DRM plane.

New drivers should call `drm_universal_plane_init` instead.

Returns

Zero on success, error code on failure.

Name

`drm_plane_cleanup` — Clean up the core plane usage

Synopsis

```
void drm_plane_cleanup (struct drm_plane * plane);
```

Arguments

plane plane to cleanup

Description

This function cleans up *plane* and removes it from the DRM mode setting core. Note that the function does **not** free the plane structure itself, this is the responsibility of the caller.

Name

`drm_plane_index` — find the index of a registered plane

Synopsis

```
unsigned int drm_plane_index (struct drm_plane * plane);
```

Arguments

plane plane to find index for

Description

Given a registered plane, return the index of that CRTC within a DRM device's list of planes.

Name

`drm_plane_from_index` — find the registered plane at an index

Synopsis

```
struct drm_plane * drm_plane_from_index (struct drm_device * dev, int
idx);
```

Arguments

dev DRM device

idx index of registered plane to find for

Description

Given a plane index, return the registered plane from DRM device's list of planes with matching index.

Name

`drm_plane_force_disable` — Forcibly disable a plane

Synopsis

```
void drm_plane_force_disable (struct drm_plane * plane);
```

Arguments

plane plane to disable

Description

Forces the plane to be disabled.

Used when the plane's current framebuffer is destroyed, and when restoring fbdev mode.

Name

`drm_mode_create_dvi_i_properties` — create DVI-I specific connector properties

Synopsis

```
int drm_mode_create_dvi_i_properties (struct drm_device * dev);
```

Arguments

dev DRM device

Description

Called by a driver the first time a DVI-I connector is made.

Name

`drm_mode_create_tv_properties` — create TV specific connector properties

Synopsis

```
int drm_mode_create_tv_properties (struct drm_device * dev, unsigned
int num_modes, const char *const modes[ ]);
```

Arguments

<i>dev</i>	DRM device
<i>num_modes</i>	number of different TV formats (modes) supported
<i>modes[]</i>	array of pointers to strings containing name of each format

Description

Called by a driver's TV initialization routine, this function creates the TV specific connector properties for a given device. Caller is responsible for allocating a list of format names and passing them to this routine.

Name

`drm_mode_create_scaling_mode_property` — create scaling mode property

Synopsis

```
int drm_mode_create_scaling_mode_property (struct drm_device * dev);
```

Arguments

dev DRM device

Description

Called by a driver the first time it's needed, must be attached to desired connectors.

Name

`drm_mode_create_aspect_ratio_property` — create aspect ratio property

Synopsis

```
int drm_mode_create_aspect_ratio_property (struct drm_device * dev);
```

Arguments

dev DRM device

Description

Called by a driver the first time it's needed, must be attached to desired connectors.

Returns

Zero on success, negative errno on failure.

Name

`drm_mode_create_dirty_info_property` — create dirty property

Synopsis

```
int drm_mode_create_dirty_info_property (struct drm_device * dev);
```

Arguments

dev DRM device

Description

Called by a driver the first time it's needed, must be attached to desired connectors.

Name

`drm_mode_create_suggested_offset_properties` — create suggests offset properties

Synopsis

```
int drm_mode_create_suggested_offset_properties (struct drm_device *  
dev);
```

Arguments

dev DRM device

Description

Create the the suggested x/y offset property for connectors.

Name

`drm_mode_set_config_internal` — helper to call `->set_config`

Synopsis

```
int drm_mode_set_config_internal (struct drm_mode_set * set);
```

Arguments

set modeset config to set

Description

This is a little helper to wrap internal calls to the `->set_config` driver interface. The only thing it adds is correct refcounting dance.

Returns

Zero on success, negative `errno` on failure.

Name

`drm_crtc_get_hv_timing` — Fetches hdisplay/vdisplay for given mode

Synopsis

```
void drm_crtc_get_hv_timing (const struct drm_display_mode * mode, int  
* hdisplay, int * vdisplay);
```

Arguments

mode mode to query
hdisplay hdisplay value to fill in
vdisplay vdisplay value to fill in

Description

The vdisplay value will be doubled if the specified mode is a stereo mode of the appropriate layout.

Name

`drm_crtc_check_viewport` — Checks that a framebuffer is big enough for the CRTC viewport

Synopsis

```
int drm_crtc_check_viewport (const struct drm_crtc * crtc, int x, int
                             y, const struct drm_display_mode * mode, const struct drm_framebuffer
                             * fb);
```

Arguments

crtc CRTC that framebuffer will be displayed on

x x panning

y y panning

mode mode that framebuffer will be displayed under

fb framebuffer to check size of

Name

`drm_mode_legacy_fb_format` — compute drm fourcc code from legacy description

Synopsis

```
uint32_t drm_mode_legacy_fb_format (uint32_t bpp, uint32_t depth);
```

Arguments

bpp bits per pixels

depth bit depth per pixel

Description

Computes a drm fourcc pixel format code for the given *bpp/depth* values. Useful in fbdev emulation code, since that deals in those values.

Name

`drm_property_create` — create a new property type

Synopsis

```
struct drm_property * drm_property_create (struct drm_device * dev, int
flags, const char * name, int num_values);
```

Arguments

<i>dev</i>	drm device
<i>flags</i>	flags specifying the property type
<i>name</i>	name of the property
<i>num_values</i>	number of pre-defined values

Description

This creates a new generic drm property which can then be attached to a drm object with `drm_object_attach_property`. The returned property object must be freed with `drm_property_destroy`.

Note that the DRM core keeps a per-device list of properties and that, if `drm_mode_config_cleanup` is called, it will destroy all properties created by the driver.

Returns

A pointer to the newly created property on success, NULL on failure.

Name

`drm_property_create_enum` — create a new enumeration property type

Synopsis

```
struct drm_property * drm_property_create_enum (struct drm_device * dev,  
int flags, const char * name, const struct drm_prop_enum_list * props,  
int num_values);
```

Arguments

<i>dev</i>	drm device
<i>flags</i>	flags specifying the property type
<i>name</i>	name of the property
<i>props</i>	enumeration lists with property values
<i>num_values</i>	number of pre-defined values

Description

This creates a new generic drm property which can then be attached to a drm object with `drm_object_attach_property`. The returned property object must be freed with `drm_property_destroy`.

Userspace is only allowed to set one of the predefined values for enumeration properties.

Returns

A pointer to the newly created property on success, NULL on failure.

Name

`drm_property_create_bitmask` — create a new bitmask property type

Synopsis

```
struct drm_property * drm_property_create_bitmask (struct drm_device *  
dev, int flags, const char * name, const struct drm_prop_enum_list *  
props, int num_props, uint64_t supported_bits);
```

Arguments

<i>dev</i>	drm device
<i>flags</i>	flags specifying the property type
<i>name</i>	name of the property
<i>props</i>	enumeration lists with property bitflags
<i>num_props</i>	size of the <i>props</i> array
<i>supported_bits</i>	bitmask of all supported enumeration values

Description

This creates a new bitmask drm property which can then be attached to a drm object with `drm_object_attach_property`. The returned property object must be freed with `drm_property_destroy`.

Compared to plain enumeration properties userspace is allowed to set any or'ed together combination of the predefined property bitflag values

Returns

A pointer to the newly created property on success, NULL on failure.

Name

`drm_property_create_range` — create a new unsigned ranged property type

Synopsis

```
struct drm_property * drm_property_create_range (struct drm_device *  
dev, int flags, const char * name, uint64_t min, uint64_t max);
```

Arguments

<i>dev</i>	drm device
<i>flags</i>	flags specifying the property type
<i>name</i>	name of the property
<i>min</i>	minimum value of the property
<i>max</i>	maximum value of the property

Description

This creates a new generic drm property which can then be attached to a drm object with `drm_object_attach_property`. The returned property object must be freed with `drm_property_destroy`.

Userspace is allowed to set any unsigned integer value in the (min, max) range inclusive.

Returns

A pointer to the newly created property on success, NULL on failure.

Name

`drm_property_create_signed_range` — create a new signed ranged property type

Synopsis

```
struct drm_property * drm_property_create_signed_range (struct drm_device * dev, int flags, const char * name, int64_t min, int64_t max);
```

Arguments

<i>dev</i>	drm device
<i>flags</i>	flags specifying the property type
<i>name</i>	name of the property
<i>min</i>	minimum value of the property
<i>max</i>	maximum value of the property

Description

This creates a new generic drm property which can then be attached to a drm object with `drm_object_attach_property`. The returned property object must be freed with `drm_property_destroy`.

Userspace is allowed to set any signed integer value in the (min, max) range inclusive.

Returns

A pointer to the newly created property on success, NULL on failure.

Name

`drm_property_create_object` — create a new object property type

Synopsis

```
struct drm_property * drm_property_create_object (struct drm_device *  
dev, int flags, const char * name, uint32_t type);
```

Arguments

dev drm device

flags flags specifying the property type

name name of the property

type object type from `DRM_MODE_OBJECT_*` defines

Description

This creates a new generic drm property which can then be attached to a drm object with `drm_object_attach_property`. The returned property object must be freed with `drm_property_destroy`.

Userspace is only allowed to set this to any property value of the given *type*. Only useful for atomic properties, which is enforced.

Returns

A pointer to the newly created property on success, NULL on failure.

Name

`drm_property_create_bool` — create a new boolean property type

Synopsis

```
struct drm_property * drm_property_create_bool (struct drm_device * dev,  
int flags, const char * name);
```

Arguments

dev drm device

flags flags specifying the property type

name name of the property

Description

This creates a new generic drm property which can then be attached to a drm object with `drm_object_attach_property`. The returned property object must be freed with `drm_property_destroy`.

This is implemented as a ranged property with only {0, 1} as valid values.

Returns

A pointer to the newly created property on success, NULL on failure.

Name

`drm_property_add_enum` — add a possible value to an enumeration property

Synopsis

```
int drm_property_add_enum (struct drm_property * property, int index,  
uint64_t value, const char * name);
```

Arguments

<i>property</i>	enumeration property to change
<i>index</i>	index of the new enumeration
<i>value</i>	value of the new enumeration
<i>name</i>	symbolic name of the new enumeration

Description

This functions adds enumerations to a property.

It's use is deprecated, drivers should use one of the more specific helpers to directly create the property with all enumerations already attached.

Returns

Zero on success, error code on failure.

Name

`drm_property_destroy` — destroy a drm property

Synopsis

```
void drm_property_destroy (struct drm_device * dev, struct drm_property  
* property);
```

Arguments

dev drm device

property property to destroy

Description

This function frees a property including any attached resources like enumeration values.

Name

`drm_object_attach_property` — attach a property to a modeset object

Synopsis

```
void drm_object_attach_property (struct drm_mode_object * obj, struct
drm_property * property, uint64_t init_val);
```

Arguments

obj drm modeset object

property property to attach

init_val initial value of the property

Description

This attaches the given property to the modeset object with the given initial value. Currently this function cannot fail since the properties are stored in a statically sized array.

Name

`drm_object_property_set_value` — set the value of a property

Synopsis

```
int drm_object_property_set_value (struct drm_mode_object * obj, struct
drm_property * property, uint64_t val);
```

Arguments

obj drm mode object to set property value for

property property to set

val value the property should be set to

Description

This functions sets a given property on a given object. This function only changes the software state of the property, it does not call into the driver's `->set_property` callback.

Returns

Zero on success, error code on failure.

Name

`drm_object_property_get_value` — retrieve the value of a property

Synopsis

```
int drm_object_property_get_value (struct drm_mode_object * obj, struct
drm_property * property, uint64_t * val);
```

Arguments

obj drm mode object to get property value from

property property to retrieve

val storage for the property value

Description

This function retrieves the software state of the given property for the given property. Since there is no driver callback to retrieve the current property value this might be out of sync with the hardware, depending upon the driver and property.

Returns

Zero on success, error code on failure.

Name

`drm_property_create_blob` — Create new blob property

Synopsis

```
struct drm_property_blob * drm_property_create_blob (struct drm_device  
* dev, size_t length, const void * data);
```

Arguments

dev DRM device to create property for

length Length to allocate for blob data

data If specified, copies data into blob

Description

Creates a new blob property for a specified DRM device, optionally copying data.

Returns

New blob property with a single reference on success, or an `ERR_PTR` value on failure.

Name

`drm_property_unreference_blob` — Unreference a blob property

Synopsis

```
void drm_property_unreference_blob (struct drm_property_blob * blob);
```

Arguments

blob Pointer to blob property

Description

Drop a reference on a blob property. May free the object.

Name

`drm_property_reference_blob` — Take a reference on an existing property

Synopsis

```
struct drm_property_blob * drm_property_reference_blob (struct dr-  
m_property_blob * blob);
```

Arguments

blob Pointer to blob property

Description

Take a new reference on an existing blob property.

Name

`drm_property_lookup_blob` — look up a blob property and take a reference

Synopsis

```
struct drm_property_blob * drm_property_lookup_blob (struct drm_device  
* dev, uint32_t id);
```

Arguments

dev drm device

id id of the blob property

Description

If successful, this takes an additional reference to the blob property. callers need to make sure to eventually unreference the returned property again, using `drm_property_unreference_blob`.

Name

`drm_mode_connector_set_path_property` — set tile property on connector

Synopsis

```
int drm_mode_connector_set_path_property (struct drm_connector * connector, const char * path);
```

Arguments

connector connector to set property on.

path path to use for property; must not be NULL.

Description

This creates a property to expose to userspace to specify a connector path. This is mainly used for DisplayPort MST where connectors have a topology and we want to allow userspace to give them more meaningful names.

Returns

Zero on success, negative errno on failure.

Name

`drm_mode_connector_set_tile_property` — set tile property on connector

Synopsis

```
int drm_mode_connector_set_tile_property (struct drm_connector * connector);
```

Arguments

connector connector to set property on.

Description

This looks up the tile information for a connector, and creates a property for userspace to parse if it exists. The property is of the form of 8 integers using ':' as a separator.

Returns

Zero on success, errno on failure.

Name

`drm_mode_connector_update_edid_property` — update the edid property of a connector

Synopsis

```
int drm_mode_connector_update_edid_property (struct drm_connector *  
connector, const struct edid * edid);
```

Arguments

connector drm connector

edid new value of the edid property

Description

This function creates a new blob modeset object and assigns its id to the connector's edid property.

Returns

Zero on success, negative errno on failure.

Name

`drm_mode_plane_set_obj_prop` — set the value of a property

Synopsis

```
int drm_mode_plane_set_obj_prop (struct drm_plane * plane, struct dr-  
m_property * property, uint64_t value);
```

Arguments

plane drm plane object to set property value for

property property to set

value value the property should be set to

Description

This functions sets a given property on a given plane object. This function calls the driver's `->set_property` callback and changes the software state of the property if the callback succeeds.

Returns

Zero on success, error code on failure.

Name

`drm_mode_connector_attach_encoder` — attach a connector to an encoder

Synopsis

```
int drm_mode_connector_attach_encoder (struct drm_connector * connector, struct drm_encoder * encoder);
```

Arguments

connector connector to attach

encoder encoder to attach *connector* to

Description

This function links up a connector to an encoder. Note that the routing restrictions between encoders and crtcs are exposed to userspace through the `possible_clones` and `possible_crtcs` bitmasks.

Returns

Zero on success, negative `errno` on failure.

Name

`drm_mode_crtc_set_gamma_size` — set the gamma table size

Synopsis

```
int drm_mode_crtc_set_gamma_size (struct drm_crtc * crtc, int gamma_size);
```

Arguments

crtc CRTC to set the gamma table size for

gamma_size size of the gamma table

Description

Drivers which support gamma tables should set this to the supported gamma table size when initializing the CRTC. Currently the drm core only supports a fixed gamma table size.

Returns

Zero on success, negative errno on failure.

Name

`drm_mode_config_reset` — call `->reset` callbacks

Synopsis

```
void drm_mode_config_reset (struct drm_device * dev);
```

Arguments

dev drm device

Description

This functions calls all the crtcs, encoder's and connector's `->reset` callback. Drivers can use this in e.g. their driver load or resume code to reset hardware and software state.

Name

`drm_fb_get_bpp_depth` — get the bpp/depth values for format

Synopsis

```
void drm_fb_get_bpp_depth (uint32_t format, unsigned int * depth, int  
* bpp);
```

Arguments

format pixel format (DRM_FORMAT_*)

depth storage for the depth value

bpp storage for the bpp value

Description

This only supports RGB formats here for compat with code that doesn't use pixel formats directly yet.

Name

`drm_format_num_planes` — get the number of planes for format

Synopsis

```
int drm_format_num_planes (uint32_t format);
```

Arguments

format pixel format (DRM_FORMAT_*)

Returns

The number of planes used by the specified pixel format.

Name

`drm_format_plane_cpp` — determine the bytes per pixel value

Synopsis

```
int drm_format_plane_cpp (uint32_t format, int plane);
```

Arguments

format pixel format (DRM_FORMAT_*)

plane plane index

Returns

The bytes per pixel value for the specified plane.

Name

`drm_format_horz_chroma_subsampling` — get the horizontal chroma subsampling factor

Synopsis

```
int drm_format_horz_chroma_subsampling (uint32_t format);
```

Arguments

format pixel format (DRM_FORMAT_*)

Returns

The horizontal chroma subsampling factor for the specified pixel format.

Name

`drm_format_vert_chroma_subsampling` — get the vertical chroma subsampling factor

Synopsis

```
int drm_format_vert_chroma_subsampling (uint32_t format);
```

Arguments

format pixel format (DRM_FORMAT_*)

Returns

The vertical chroma subsampling factor for the specified pixel format.

Name

`drm_rotation_simplify` — Try to simplify the rotation

Synopsis

```
unsigned int drm_rotation_simplify (unsigned int rotation, unsigned int supported_rotations);
```

Arguments

rotation Rotation to be simplified

supported_rotations Supported rotations

Description

Attempt to simplify the rotation to a form that is supported. Eg. if the hardware supports everything except `DRM_REFLECT_X`

one could call this function like this

```
drm_rotation_simplify(rotation, BIT(DRM_ROTATE_0) | BIT(DRM_ROTATE_90) | BIT(DRM_ROTATE_180) | BIT(DRM_ROTATE_270) | BIT(DRM_REFLECT_Y));
```

to eliminate the `DRM_ROTATE_X` flag. Depending on what kind of transforms the hardware supports, this function may not be able to produce a supported transform, so the caller should check the result afterwards.

Name

`drm_mode_config_init` — initialize DRM mode_configuration structure

Synopsis

```
void drm_mode_config_init (struct drm_device * dev);
```

Arguments

dev DRM device

Description

Initialize *dev*'s mode_config structure, used for tracking the graphics configuration of *dev*.

Since this initializes the modeset locks, no locking is possible. Which is no problem, since this should happen single threaded at init time. It is the driver's problem to ensure this guarantee.

Name

`drm_mode_config_cleanup` — free up DRM mode_config info

Synopsis

```
void drm_mode_config_cleanup (struct drm_device * dev);
```

Arguments

dev DRM device

Description

Free up all the connectors and CRTC's associated with this DRM device, then free up the framebuffers and associated buffer objects.

Note that since this /should/ happen single-threaded at driver/device teardown time, no locking is required. It's the driver's job to ensure that this guarantee actually holds true.

FIXME

cleanup any dangling user buffer objects too

Name

`drm_mode_get_tile_group` — get a reference to an existing tile group

Synopsis

```
struct drm_tile_group * drm_mode_get_tile_group (struct drm_device *  
dev, char topology[8]);
```

Arguments

dev DRM device

topology[8] 8-bytes unique per monitor.

Description

Use the unique bytes to get a reference to an existing tile group.

RETURNS

tile group or NULL if not found.

Name

`drm_mode_create_tile_group` — create a tile group from a displayid description

Synopsis

```
struct drm_tile_group * drm_mode_create_tile_group (struct drm_device  
* dev, char topology[8]);
```

Arguments

dev DRM device

topology[8] 8-bytes unique per monitor.

Description

Create a tile group for the unique monitor, and get a unique identifier for the tile group.

RETURNS

new tile group or error.

KMS Data Structures

Name

struct drm_framebuffer_funcs — framebuffer hooks

Synopsis

```
struct drm_framebuffer_funcs {  
    void (* destroy) (struct drm_framebuffer *framebuffer);  
    int (* create_handle) (struct drm_framebuffer *fb, struct drm_file *file_priv, unsigned int (* dirty) (struct drm_framebuffer *framebuffer, struct drm_file *file_priv, unsigned int *dirty));  
};
```

Members

destroy

Clean up framebuffer resources, specifically also unreference the backing storage. The core guarantees to call this function for every framebuffer successfully created by `fb_create` in `drm_mode_config_funcs`. Drivers must also call `drm_framebuffer_cleanup` to release DRM core resources for this framebuffer.

create_handle

Create a buffer handle in the driver-specific buffer manager (either GEM or TTM) valid for the passed-in struct `drm_file`. This is used by the core to implement the `GETFB_IOCTL`, which returns (for sufficiently privileged user) also a native buffer handle. This can be used for seamless transitions between modesetting clients by copying the current screen contents to a private buffer and blending between that and the new contents.

GEM based drivers should call `drm_gem_handle_create` to create the handle.

RETURNS:

0 on success or a negative error code on failure.

dirty

Optional callback for the dirty fb IOCTL.

Userspace can notify the driver via this callback that an area of the framebuffer has changed and should be flushed to the display hardware. This can also be used internally, e.g. by the fbdev emulation, though that's not the case currently.

See documentation in `drm_mode.h` for the struct `drm_mode_fb_dirty_cmd` for more information as all the semantics and arguments have a one to one mapping on this function.

RETURNS:

0 on success or a negative error code on failure.

Name

struct `drm_crtc_state` — mutable CRTC state

Synopsis

```
struct drm_crtc_state {
    struct drm_crtc * crtc;
    bool enable;
    bool active;
    bool planes_changed:1;
    bool mode_changed:1;
    bool active_changed:1;
    bool connectors_changed:1;
    u32 plane_mask;
    u32 connector_mask;
    u32 last_vblank_count;
    struct drm_display_mode adjusted_mode;
    struct drm_display_mode mode;
    struct drm_pending_vblank_event * event;
    struct drm_atomic_state * state;
};
```

Members

<code>crtc</code>	backpointer to the CRTC
<code>enable</code>	whether the CRTC should be enabled, gates all other state
<code>active</code>	whether the CRTC is actively displaying (used for DPMS)
<code>planes_changed</code>	planes on this crtc are updated
<code>mode_changed</code>	<code>crtc_state->mode</code> or <code>crtc_state->enable</code> has been changed
<code>active_changed</code>	<code>crtc_state->active</code> has been toggled.
<code>connectors_changed</code>	connectors to this crtc have been updated
<code>plane_mask</code>	bitmask of $(1 \ll \text{drm_plane_index(plane)})$ of attached planes
<code>connector_mask</code>	bitmask of $(1 \ll \text{drm_connector_index(connector)})$ of attached connectors
<code>last_vblank_count</code>	for helpers and drivers to capture the vblank of the update to ensure frame-buffer cleanup isn't done too early
<code>adjusted_mode</code>	for use by helpers and drivers to compute adjusted mode timings
<code>mode</code>	current mode timings
<code>event</code>	optional pointer to a DRM event to signal upon completion of the state update
<code>state</code>	backpointer to global <code>drm_atomic_state</code>

Description

Note that the distinction between *enable* and *active* is rather subtle: Flipping *active* while *enable* is set without changing anything else may never return in a failure from the `->atomic_check` callback.

Userspace assumes that a DPMS On will always succeed. In other words: *enable* controls resource assignment, *active* controls the actual hardware state.

Name

struct drm_crtc_funcs — control CRTC's for a given device

Synopsis

```
struct drm_crtc_funcs {
    void (* reset) (struct drm_crtc *crtc);
    int (* cursor_set) (struct drm_crtc *crtc, struct drm_file *file_priv, uint32_t h, uint32_t v);
    int (* cursor_set2) (struct drm_crtc *crtc, struct drm_file *file_priv, uint32_t h, uint32_t v, uint32_t w, uint32_t h2);
    int (* cursor_move) (struct drm_crtc *crtc, int x, int y);
    void (* gamma_set) (struct drm_crtc *crtc, u16 *r, u16 *g, u16 *b, uint32_t start, uint32_t end);
    void (* destroy) (struct drm_crtc *crtc);
    int (* set_config) (struct drm_mode_set *set);
    int (* page_flip) (struct drm_crtc *crtc, struct drm_framebuffer *fb, struct drm_plane *plane, uint64_t flags, uint64_t seqno, uint64_t fence);
    int (* set_property) (struct drm_crtc *crtc, struct drm_property *property, uint64_t value);
    struct drm_crtc_state *(* atomic_duplicate_state) (struct drm_crtc *crtc);
    void (* atomic_destroy_state) (struct drm_crtc *crtc, struct drm_crtc_state *state);
    int (* atomic_set_property) (struct drm_crtc *crtc, struct drm_crtc_state *state, struct drm_property *property, uint64_t value);
    int (* atomic_get_property) (struct drm_crtc *crtc, const struct drm_crtc_state *state, struct drm_property *property, uint64_t *value);
};
```

Members

reset

Reset CRTC hardware and software state to off. This function isn't called by the core directly, only through `drm_mode_config_reset`. It's not a helper hook only for historical reasons.

Atomic drivers can use `drm_atomic_helper_crtc_reset` to reset atomic state using this hook.

cursor_set

Update the cursor image. The cursor position is relative to the CRTC and can be partially or fully outside of the visible area.

Note that contrary to all other KMS functions the legacy cursor entry points don't take a framebuffer object, but instead take directly a raw buffer object id from the driver's buffer manager (which is either GEM or TTM for current drivers).

This entry point is deprecated, drivers should instead implement universal plane support and register a proper cursor plane using `drm_crtc_init_with_planes`.

This callback is optional

RETURNS:

0 on success or a negative error code on failure.

cursor_set2

Update the cursor image, including hotspot information. The hotspot must not affect the cursor position in CRTC coordinates, but is only meant as a hint for virtualized display hardware to coordinate the

guests and hosts cursor position. The cursor hotspot is relative to the cursor image. Otherwise this works exactly like *cursor_set*.

This entry point is deprecated, drivers should instead implement universal plane support and register a proper cursor plane using *drm_crtc_init_with_planes*.

This callback is optional.

RETURNS:

0 on success or a negative error code on failure.

cursor_move

Update the cursor position. The cursor does not need to be visible when this hook is called.

This entry point is deprecated, drivers should instead implement universal plane support and register a proper cursor plane using *drm_crtc_init_with_planes*.

This callback is optional.

RETURNS:

0 on success or a negative error code on failure.

gamma_set

Set gamma on the CRTC.

This callback is optional.

NOTE:

Drivers that support gamma tables and also fbdev emulation through the provided helper library need to take care to fill out the gamma hooks for both. Currently there's a bit an unfortunate duplication going on, which should eventually be unified to just one set of hooks.

destroy

Clean up plane resources. This is only called at driver unload time through *drm_mode_config_cleanup* since a CRTC cannot be hotplugged in DRM.

set_config

This is the main legacy entry point to change the modeset state on a CRTC. All the details of the desired configuration are passed in a struct *drm_mode_set* - see there for details.

Drivers implementing atomic modeset should use *drm_atomic_helper_set_config* to implement this hook.

RETURNS:

0 on success or a negative error code on failure.

page_flip

Legacy entry point to schedule a flip to the given framebuffer.

Page flipping is a synchronization mechanism that replaces the frame buffer being scanned out by the CRTC with a new frame buffer during vertical blanking, avoiding tearing (except when requested otherwise through the `DRM_MODE_PAGE_FLIP_ASYNC` flag). When an application requests a page flip the DRM core verifies that the new frame buffer is large enough to be scanned out by the CRTC in the currently configured mode and then calls the CRTC `->page_flip` operation with a pointer to the new frame buffer.

The driver must wait for any pending rendering to the new framebuffer to complete before executing the flip. It should also wait for any pending rendering from other drivers if the underlying buffer is a shared dma-buf.

An application can request to be notified when the page flip has completed. The drm core will supply a struct `drm_event` in the event parameter in this case. This can be handled by the `drm_crtc_send_vblank_event` function, which the driver should call on the provided event upon completion of the flip. Note that if the driver supports vblank signalling and timestamping the vblank counters and timestamps must agree with the ones returned from page flip events. With the current vblank helper infrastructure this can be achieved by holding a vblank reference while the page flip is pending, acquired through `drm_crtc_vblank_get` and released with `drm_crtc_vblank_put`. Drivers are free to implement their own vblank counter and timestamp tracking though, e.g. if they have accurate timestamp registers in hardware.

FIXME:

Up to that point drivers need to manage events themselves and can use `even->base.list` freely for that. Specifically they need to ensure that they don't send out page flip (or vblank) events for which the corresponding drm file has been closed already. The drm core unfortunately does not (yet) take care of that. Therefore drivers currently must clean up and release pending events in their `->preclose` driver function.

This callback is optional.

NOTE:

Very early versions of the KMS ABI mandated that the driver must block (but not reject) any rendering to the old framebuffer until the flip operation has completed and the old framebuffer is no longer visible. This requirement has been lifted, and userspace is instead expected to request delivery of an event and wait with recycling old buffers until such has been received.

RETURNS:

0 on success or a negative error code on failure. Note that if a `->page_flip` operation is already pending the callback should return `-EBUSY`. Pageflips on a disabled CRTC (either by setting a `NULL` mode or just runtime disabled through DPMS respectively the new

atomic “ACTIVE” state) should result in an -EINVAL error code. Note that `drm_atomic_helper_page_flip` checks this already for atomic drivers.

`set_property`

This is the legacy entry point to update a property attached to the CRTC.

Drivers implementing atomic modeset should use `drm_atomic_helper_crtc_set_property` to implement this hook.

This callback is optional if the driver does not support any legacy driver-private properties.

RETURNS:

0 on success or a negative error code on failure.

`atomic_duplicate_state`

Duplicate the current atomic state for this CRTC and return it. The core and helpers guarantee that any atomic state duplicated with this hook and still owned by the caller (i.e. not transferred to the driver by calling `->atomic_commit` from struct `drm_mode_config_funcs`) will be cleaned up by calling the `atomic_destroy_state` hook in this structure.

Atomic drivers which don't subclass struct `drm_crtc` should use `drm_atomic_helper_crtc_duplicate_state`. Drivers that subclass the state structure to extend it with driver-private state should use `__drm_atomic_helper_crtc_duplicate_state` to make sure shared state is duplicated in a consistent fashion across drivers.

It is an error to call this hook before `crtc->state` has been initialized correctly.

NOTE:

If the duplicate state references refcounted resources this hook must acquire a reference for each of them. The driver must release these references again in `atomic_destroy_state`.

RETURNS:

Duplicated atomic state or NULL when the allocation failed.

`atomic_destroy_state`

Destroy a state duplicated with `atomic_duplicate_state` and release or unreference all resources it references

`atomic_set_property`

Decode a driver-private property value and store the decoded value into the passed-in state structure. Since the atomic core decodes all standardized properties (even for extensions beyond the core set of properties which might not be implemented by all drivers) this requires drivers to subclass the state structure.

Such driver-private properties should really only be implemented for truly hardware/vendor specific state. Instead it is preferred to standardize atomic extension and decode the properties used to expose such an extension in the core.

Do not call this function directly, use `drm_atomic_crtc_set_property` instead.

This callback is optional if the driver does not support any driver-private atomic properties.

NOTE:

This function is called in the state assembly phase of atomic mode-sets, which can be aborted for any reason (including on userspace's request to just check whether a configuration would be possible). Drivers **MUST NOT** touch any persistent state (hardware or software) or data structures except the passed in *state* parameter.

Also since userspace controls in which order properties are set this function must not do any input validation (since the state update is incomplete and hence likely inconsistent). Instead any such input validation must be done in the various `atomic_check` callbacks.

RETURNS:

0 if the property has been found, `-EINVAL` if the property isn't implemented by the driver (which should never happen, the core only asks for properties attached to this CRTC). No other validation is allowed by the driver. The core already checks that the property value is within the range (integer, valid enum value, ...) the driver set when registering the property.

`atomic_get_property`

Reads out the decoded driver-private property. This is used to implement the `GETCRTC_IOCTL`.

Do not call this function directly, use `drm_atomic_crtc_get_property` instead.

This callback is optional if the driver does not support any driver-private atomic properties.

RETURNS:

0 on success, `-EINVAL` if the property isn't implemented by the driver (which should never happen, the core only asks for properties attached to this CRTC).

Description

The `drm_crtc_funcs` structure is the central CRTC management structure in the DRM. Each CRTC controls one or more connectors (note that the name CRTC is simply historical, a CRTC may control LVDS, VGA, DVI, TV out, etc. connectors, not just CRTs).

Each driver is responsible for filling out this structure at startup time, in addition to providing other modesetting features, like i2c and DDC bus accessors.

Name

struct drm_crtc — central CRTC control structure

Synopsis

```
struct drm_crtc {
    struct drm_device * dev;
    struct device_node * port;
    struct list_head head;
    struct drm_modeset_lock mutex;
    struct drm_mode_object base;
    struct drm_plane * primary;
    struct drm_plane * cursor;
    int cursor_x;
    int cursor_y;
    bool enabled;
    struct drm_display_mode mode;
    struct drm_display_mode hwmode;
    int x;
    int y;
    const struct drm_crtc_funcs * funcs;
    uint32_t gamma_size;
    uint16_t * gamma_store;
    const struct drm_crtc_helper_funcs * helper_private;
    struct drm_object_properties properties;
    struct drm_crtc_state * state;
    struct drm_modeset_acquire_ctx * acquire_ctx;
};
```

Members

dev	parent DRM device
port	OF node used by <code>drm_of_find_possible_crtcs</code>
head	list management
mutex	per-CRTC locking
base	base KMS object for ID tracking etc.
primary	primary plane for this CRTC
cursor	cursor plane for this CRTC
cursor_x	current x position of the cursor, used for universal cursor planes
cursor_y	current y position of the cursor, used for universal cursor planes
enabled	is this CRTC enabled?
mode	current mode timings
hwmode	mode timings as programmed to hw regs

x	x position on screen
y	y position on screen
funcs	CRTC control functions
gamma_size	size of gamma ramp
gamma_store	gamma ramp values
helper_private	mid-layer private data
properties	property tracking for this CRTC
state	current atomic state for this CRTC
acquire_ctx	per-CRTC implicit acquire context used by atomic drivers for legacy IOCTLs

Description

Each CRTC may have one or more connectors associated with it. This structure allows the CRTC to be controlled.

Name

struct drm_connector_state — mutable connector state

Synopsis

```
struct drm_connector_state {  
    struct drm_connector * connector;  
    struct drm_crtc * crtc;  
    struct drm_encoder * best_encoder;  
    struct drm_atomic_state * state;  
};
```

Members

connector	backpointer to the connector
crtc	CRTC to connect connector to, NULL if disabled
best_encoder	can be used by helpers and drivers to select the encoder
state	backpointer to global drm_atomic_state

Name

struct `drm_connector_funcs` — control connectors on a given device

Synopsis

```
struct drm_connector_funcs {
    int (* dpms) (struct drm_connector *connector, int mode);
    void (* reset) (struct drm_connector *connector);
    enum drm_connector_status (* detect) (struct drm_connector *connector, bool force);
    void (* force) (struct drm_connector *connector);
    int (* fill_modes) (struct drm_connector *connector, uint32_t max_width, uint32_t max_height);
    int (* set_property) (struct drm_connector *connector, struct drm_property *property);
    void (* destroy) (struct drm_connector *connector);
    struct drm_connector_state *(* atomic_duplicate_state) (struct drm_connector *connector);
    void (* atomic_destroy_state) (struct drm_connector *connector, struct drm_connector_state *state);
    int (* atomic_set_property) (struct drm_connector *connector, struct drm_property *property);
    int (* atomic_get_property) (struct drm_connector *connector, const struct drm_connector_state *state, struct drm_property *property);
};
```

Members

`dpms`

Legacy entry point to set the per-connector DPMS state. Legacy DPMS is exposed as a standard property on the connector, but diverted to this callback in the drm core. Note that atomic drivers don't implement the 4 level DPMS support on the connector any more, but instead only have an on/off “ACTIVE” property on the CRTC object.

Drivers implementing atomic modeset should use `drm_atomic_helper_connector_dpms` to implement this hook.

RETURNS:

0 on success or a negative error code on failure.

`reset`

Reset connector hardware and software state to off. This function isn't called by the core directly, only through `drm_mode_config_reset`. It's not a helper hook only for historical reasons.

Atomic drivers can use `drm_atomic_helper_connector_reset` to reset atomic state using this hook.

`detect`

Check to see if anything is attached to the connector. The parameter `force` is set to false whilst polling, true when checking the connector due to a user request. `force` can be used by the driver to avoid expensive, destructive operations during automated probing.

FIXME:

Note that this hook is only called by the probe helper. It's not in the helper library vtable purely for historical reasons. The only DRM core entry point to probe connector state is `fill_modes`.

RETURNS:

`drm_connector_status` indicating the connector's status.

force

This function is called to update internal encoder state when the connector is forced to a certain state by userspace, either through the sysfs interfaces or on the kernel cmdline. In that case the *detect* callback isn't called.

FIXME:

Note that this hook is only called by the probe helper. It's not in the helper library vtable purely for historical reasons. The only DRM core entry point to probe connector state is *fill_modes*.

fill_modes

Entry point for output detection and basic mode validation. The driver should reprobe the output if needed (e.g. when hotplug handling is unreliable), add all detected modes to `connector->modes` and filter out any the device can't support in any configuration. It also needs to filter out any modes wider or higher than the parameters `max_width` and `max_height` indicate.

The drivers must also prune any modes no longer valid from `connector->modes`. Furthermore it must update `connector->status` and `connector->edid`. If no EDID has been received for this output `connector->edid` must be NULL.

Drivers using the probe helpers should use `drm_helper_probe_single_connector_modes` or `drm_helper_probe_single_connector_modes_nomerge` to implement this function.

RETURNS:

The number of modes detected and filled into `connector->modes`.

set_property

This is the legacy entry point to update a property attached to the connector.

Drivers implementing atomic modeset should use `drm_atomic_helper_connector_set_property` to implement this hook.

This callback is optional if the driver does not support any legacy driver-private properties.

RETURNS:

0 on success or a negative error code on failure.

destroy

Clean up connector resources. This is called at driver unload time through `drm_mode_config_cleanup`. It can also be called at

runtime when a connector is being hot-unplugged for drivers that support connector hotplugging (e.g. DisplayPort MST).

`atomic_duplicate_state`

Duplicate the current atomic state for this connector and return it. The core and helpers guarantee that any atomic state duplicated with this hook and still owned by the caller (i.e. not transferred to the driver by calling `->atomic_commit` from struct `drm_mode_config_funcs`) will be cleaned up by calling the `atomic_destroy_state` hook in this structure.

Atomic drivers which don't subclass struct `drm_connector_state` should use `drm_atomic_helper_connector_duplicate_state`. Drivers that subclass the state structure to extend it with driver-private state should use `__drm_atomic_helper_connector_duplicate_state` to make sure shared state is duplicated in a consistent fashion across drivers.

It is an error to call this hook before `connector->state` has been initialized correctly.

NOTE:

If the duplicate state references refcounted resources this hook must acquire a reference for each of them. The driver must release these references again in `atomic_destroy_state`.

RETURNS:

Duplicated atomic state or NULL when the allocation failed.

`atomic_destroy_state`

Destroy a state duplicated with `atomic_duplicate_state` and release or unreference all resources it references

`atomic_set_property`

Decode a driver-private property value and store the decoded value into the passed-in state structure. Since the atomic core decodes all standardized properties (even for extensions beyond the core set of properties which might not be implemented by all drivers) this requires drivers to subclass the state structure.

Such driver-private properties should really only be implemented for truly hardware/vendor specific state. Instead it is preferred to standardize atomic extension and decode the properties used to expose such an extension in the core.

Do not call this function directly, use `drm_atomic_connector_set_property` instead.

This callback is optional if the driver does not support any driver-private atomic properties.

NOTE:

This function is called in the state assembly phase of atomic mode-sets, which can be aborted for any reason (including on userspace's

request to just check whether a configuration would be possible). Drivers **MUST NOT** touch any persistent state (hardware or software) or data structures except the passed in *state* parameter.

Also since userspace controls in which order properties are set this function must not do any input validation (since the state update is incomplete and hence likely inconsistent). Instead any such input validation must be done in the various *atomic_check* callbacks.

RETURNS:

0 if the property has been found, *-EINVAL* if the property isn't implemented by the driver (which shouldn't ever happen, the core only asks for properties attached to this connector). No other validation is allowed by the driver. The core already checks that the property value is within the range (integer, valid enum value, ...) the driver set when registering the property.

atomic_get_property

Reads out the decoded driver-private property. This is used to implement the *GETCONNECTOR* IOCTL.

Do not call this function directly, use *drm_atomic_connector_get_property* instead.

This callback is optional if the driver does not support any driver-private atomic properties.

RETURNS:

0 on success, *-EINVAL* if the property isn't implemented by the driver (which shouldn't ever happen, the core only asks for properties attached to this connector).

Description

Each CRTC may have one or more connectors attached to it. The functions below allow the core DRM code to control connectors, enumerate available modes, etc.

Name

struct drm_encoder_funcs — encoder controls

Synopsis

```
struct drm_encoder_funcs {  
    void (* reset) (struct drm_encoder *encoder);  
    void (* destroy) (struct drm_encoder *encoder);  
};
```

Members

- | | |
|---------|--|
| reset | Reset encoder hardware and software state to off. This function isn't called by the core directly, only through <code>drm_mode_config_reset</code> . It's not a helper hook only for historical reasons. |
| destroy | Clean up encoder resources. This is only called at driver unload time through <code>drm_mode_config_cleanup</code> since an encoder cannot be hotplugged in DRM. |

Description

Encoders sit between CRTC's and connectors.

Name

struct drm_encoder — central DRM encoder structure

Synopsis

```
struct drm_encoder {
    struct drm_device * dev;
    struct list_head head;
    struct drm_mode_object base;
    char * name;
    int encoder_type;
    uint32_t possible_crtcs;
    uint32_t possible_clones;
    struct drm_crtc * crtc;
    struct drm_bridge * bridge;
    const struct drm_encoder_funcs * funcs;
    const struct drm_encoder_helper_funcs * helper_private;
};
```

Members

dev	parent DRM device
head	list management
base	base KMS object
name	encoder name
encoder_type	one of the DRM_MODE_ENCODER_<foo> types in drm_mode.h
possible_crtcs	bitmask of potential CRTC bindings
possible_clones	bitmask of potential sibling encoders for cloning
crtc	currently bound CRTC
bridge	bridge associated to the encoder
funcs	control functions
helper_private	mid-layer private data

Description

CRTCs drive pixels to encoders, which convert them into signals appropriate for a given connector or set of connectors.

Name

struct drm_connector — central DRM connector control structure

Synopsis

```
struct drm_connector {
    struct drm_device * dev;
    struct device * kdev;
    struct device_attribute * attr;
    struct list_head head;
    struct drm_mode_object base;
    char * name;
    int connector_type;
    int connector_type_id;
    bool interlace_allowed;
    bool doublescan_allowed;
    bool stereo_allowed;
    struct list_head modes;
    enum drm_connector_status status;
    struct list_head probed_modes;
    struct drm_display_info display_info;
    const struct drm_connector_funcs * funcs;
    struct drm_property_blob * edid_blob_ptr;
    struct drm_object_properties properties;
    struct drm_property_blob * path_blob_ptr;
    uint8_t polled;
    int dpms;
    const struct drm_connector_helper_funcs * helper_private;
    struct drm_cmdline_mode cmdline_mode;
    enum drm_connector_force force;
    bool override_edid;
    uint32_t encoder_ids[DRM_CONNECTOR_MAX_ENCODER];
    struct drm_encoder * encoder;
    uint8_t eld[MAX_ELD_BYTES];
    bool dvi_dual;
    int max_tmds_clock;
    bool latency_present[2];
    int video_latency[2];
    int audio_latency[2];
    int null_edid_counter;
    unsigned bad_edid_counter;
    bool edid_corrupt;
    struct dentry * debugfs_entry;
    struct drm_connector_state * state;
    bool has_tile;
    struct drm_tile_group * tile_group;
    bool tile_is_single_monitor;
    uint8_t num_h_tile;
    uint8_t num_v_tile;
    uint8_t tile_h_loc;
    uint8_t tile_v_loc;
    uint16_t tile_h_size;
```

```
    uint16_t tile_v_size;  
};
```

Members

dev	parent DRM device
kdev	kernel device for sysfs attributes
attr	sysfs attributes
head	list management
base	base KMS object
name	connector name
connector_type	one of the <code>DRM_MODE_CONNECTOR_<foo></code> types from <code>drm_mode.h</code>
connector_type_id	index into connector type enum
interlace_allowed	can this connector handle interlaced modes?
doublescan_allowed	can this connector handle doublescan?
stereo_allowed	can this connector handle stereo modes?
modes	modes available on this connector (from <code>fill_modes</code> + user)
status	one of the <code>drm_connector_status</code> enums (connected, not, or unknown)
probed_modes	list of modes derived directly from the display
display_info	information about attached display (e.g. from EDID)
funcs	connector control functions
edid_blob_ptr	DRM property containing EDID if present
properties	property tracking for this connector
path_blob_ptr	DRM blob property data for the DP MST path property
polled	a <code>DRM_CONNECTOR_POLL_<foo></code> value for core driven polling
dpms	current dpms state
helper_private	mid-layer private data
cmdline_mode	mode line parsed from the kernel cmdline for this connector
force	a <code>DRM_FORCE_<foo></code> state for forced mode sets
override_edid	has the EDID been overwritten through debugfs for testing?
encoder_ids[DRM_CONNECTOR_MAX_ENCODER]	valid encoders for this connector

encoder	encoder driving this connector, if any
eld[MAX_ELD_BYTES]	EDID-like data, if present
dvi_dual	dual link DVI, if found
max_tmds_clock	max clock rate, if found
latency_present[2]	AV delay info from ELD, if found
video_latency[2]	video latency info from ELD, if found
audio_latency[2]	audio latency info from ELD, if found
null_edid_counter	track sinks that give us all zeros for the EDID
bad_edid_counter	track sinks that give us an EDID with invalid checksum
edid_corrupt	indicates whether the last read EDID was corrupt
debugfs_entry	debugfs directory for this connector
state	current atomic state for this connector
has_tile	is this connector connected to a tiled monitor
tile_group	tile group for the connected monitor
tile_is_single_monitor	whether the tile is one monitor housing
num_h_tile	number of horizontal tiles in the tile group
num_v_tile	number of vertical tiles in the tile group
tile_h_loc	horizontal location of this tile
tile_v_loc	vertical location of this tile
tile_h_size	horizontal size of this tile.
tile_v_size	vertical size of this tile.

Description

Each connector may be connected to one or more CRTC's, or may be clonable by another connector if they can share a CRTC. Each connector also has a specific position in the broader display (referred to as a 'screen' though it could span multiple monitors).

Name

struct `drm_plane_state` — mutable plane state

Synopsis

```
struct drm_plane_state {
    struct drm_plane * plane;
    struct drm_crtc * crtc;
    struct drm_framebuffer * fb;
    struct fence * fence;
    int32_t crtc_x;
    int32_t crtc_y;
    uint32_t crtc_w;
    uint32_t crtc_h;
    uint32_t src_x;
    uint32_t src_y;
    uint32_t src_h;
    uint32_t src_w;
    struct drm_atomic_state * state;
};
```

Members

<code>plane</code>	backpointer to the plane
<code>crtc</code>	currently bound CRTC, NULL if disabled
<code>fb</code>	currently bound framebuffer
<code>fence</code>	optional fence to wait for before scanning out <i>fb</i>
<code>crtc_x</code>	left position of visible portion of plane on crtc
<code>crtc_y</code>	upper position of visible portion of plane on crtc
<code>crtc_w</code>	width of visible portion of plane on crtc
<code>crtc_h</code>	height of visible portion of plane on crtc
<code>src_x</code>	left position of visible portion of plane within plane (in 16.16)
<code>src_y</code>	upper position of visible portion of plane within plane (in 16.16)
<code>src_h</code>	height of visible portion of plane (in 16.16)
<code>src_w</code>	width of visible portion of plane (in 16.16)
<code>state</code>	backpointer to global <code>drm_atomic_state</code>

Name

struct drm_plane_funcs — driver plane control functions

Synopsis

```
struct drm_plane_funcs {
    int (* update_plane) (struct drm_plane *plane, struct drm_crtc *crtc, struct drm_fb *fb);
    int (* disable_plane) (struct drm_plane *plane);
    void (* destroy) (struct drm_plane *plane);
    void (* reset) (struct drm_plane *plane);
    int (* set_property) (struct drm_plane *plane, struct drm_property *property, unsigned long value);
    struct drm_plane_state *(* atomic_duplicate_state) (struct drm_plane *plane);
    void (* atomic_destroy_state) (struct drm_plane *plane, struct drm_plane_state *state);
    int (* atomic_set_property) (struct drm_plane *plane, struct drm_property *property, unsigned long value);
    int (* atomic_get_property) (struct drm_plane *plane, const struct drm_plane_state *state, struct drm_property **prop);
};
```

Members

update_plane

This is the legacy entry point to enable and configure the plane for the given CRTC and framebuffer. It is never called to disable the plane, i.e. the passed-in crtc and fb parameters are never NULL.

The source rectangle in frame buffer memory coordinates is given by the src_x, src_y, src_w and src_h parameters (as 16.16 fixed point values). Devices that don't support subpixel plane coordinates can ignore the fractional part.

The destination rectangle in CRTC coordinates is given by the crtc_x, crtc_y, crtc_w and crtc_h parameters (as integer values). Devices scale the source rectangle to the destination rectangle. If scaling is not supported, and the source rectangle size doesn't match the destination rectangle size, the driver must return a `-EINVAL` error.

Drivers implementing atomic modeset should use `drm_atomic_helper_update_plane` to implement this hook.

RETURNS:

0 on success or a negative error code on failure.

disable_plane

This is the legacy entry point to disable the plane. The DRM core calls this method in response to a `DRM_IOCTL_MODE_SETPANE` IOCTL call with the frame buffer ID set to 0. Disabled planes must not be processed by the CRTC.

Drivers implementing atomic modeset should use `drm_atomic_helper_disable_plane` to implement this hook.

RETURNS:

0 on success or a negative error code on failure.

destroy

Clean up plane resources. This is only called at driver unload time through `drm_mode_config_cleanup` since a plane cannot be hotplugged in DRM.

reset

Reset plane hardware and software state to off. This function isn't called by the core directly, only through `drm_mode_config_reset`. It's not a helper hook only for historical reasons.

Atomic drivers can use `drm_atomic_helper_plane_reset` to reset atomic state using this hook.

set_property

This is the legacy entry point to update a property attached to the plane.

Drivers implementing atomic modeset should use `drm_atomic_helper_plane_set_property` to implement this hook.

This callback is optional if the driver does not support any legacy driver-private properties.

RETURNS:

0 on success or a negative error code on failure.

atomic_duplicate_state

Duplicate the current atomic state for this plane and return it. The core and helpers guarantee that any atomic state duplicated with this hook and still owned by the caller (i.e. not transferred to the driver by calling `->atomic_commit` from struct `drm_mode_config_funcs`) will be cleaned up by calling the `atomic_destroy_state` hook in this structure.

Atomic drivers which don't subclass struct `drm_plane_state` should use `drm_atomic_helper_plane_duplicate_state`. Drivers that subclass the state structure to extend it with driver-private state should use `__drm_atomic_helper_plane_duplicate_state` to make sure shared state is duplicated in a consistent fashion across drivers.

It is an error to call this hook before `plane->state` has been initialized correctly.

NOTE:

If the duplicate state references refcounted resources this hook must acquire a reference for each of them. The driver must release these references again in `atomic_destroy_state`.

RETURNS:

Duplicated atomic state or NULL when the allocation failed.

atomic_destroy_state

Destroy a state duplicated with `atomic_duplicate_state` and release or unreference all resources it references

atomic_set_property

Decode a driver-private property value and store the decoded value into the passed-in state structure. Since the atomic core decodes all standardized properties (even for extensions beyond the core set of properties which might not be implemented by all drivers) this requires drivers to subclass the state structure.

Such driver-private properties should really only be implemented for truly hardware/vendor specific state. Instead it is preferred to standardize atomic extension and decode the properties used to expose such an extension in the core.

Do not call this function directly, use `drm_atomic_plane_set_property` instead.

This callback is optional if the driver does not support any driver-private atomic properties.

NOTE:

This function is called in the state assembly phase of atomic mode-sets, which can be aborted for any reason (including on userspace's request to just check whether a configuration would be possible). Drivers **MUST NOT** touch any persistent state (hardware or software) or data structures except the passed in *state* parameter.

Also since userspace controls in which order properties are set this function must not do any input validation (since the state update is incomplete and hence likely inconsistent). Instead any such input validation must be done in the various `atomic_check` callbacks.

RETURNS:

0 if the property has been found, `-EINVAL` if the property isn't implemented by the driver (which shouldn't ever happen, the core only asks for properties attached to this plane). No other validation is allowed by the driver. The core already checks that the property value is within the range (integer, valid enum value, ...) the driver set when registering the property.

atomic_get_property

Reads out the decoded driver-private property. This is used to implement the `GETPLANE` IOCTL.

Do not call this function directly, use `drm_atomic_plane_get_property` instead.

This callback is optional if the driver does not support any driver-private atomic properties.

RETURNS:

0 on success, `-EINVAL` if the property isn't implemented by the driver (which should never happen, the core only asks for properties attached to this plane).

Name

struct `drm_plane` — central DRM plane control structure

Synopsis

```
struct drm_plane {
    struct drm_device * dev;
    struct list_head head;
    struct drm_mode_object base;
    uint32_t possible_crtcs;
    uint32_t * format_types;
    unsigned int format_count;
    bool format_default;
    struct drm_crtc * crtc;
    struct drm_framebuffer * fb;
    struct drm_framebuffer * old_fb;
    const struct drm_plane_funcs * funcs;
    struct drm_object_properties properties;
    enum drm_plane_type type;
    struct drm_plane_state * state;
};
```

Members

<code>dev</code>	DRM device this plane belongs to
<code>head</code>	for list management
<code>base</code>	base mode object
<code>possible_crtcs</code>	pipes this plane can be bound to
<code>format_types</code>	array of formats supported by this plane
<code>format_count</code>	number of formats supported
<code>format_default</code>	driver hasn't supplied supported formats for the plane
<code>crtc</code>	currently bound CRTC
<code>fb</code>	currently bound fb
<code>old_fb</code>	Temporary tracking of the old fb while a modeset is ongoing. Used by <code>drm_mode_set_config_internal</code> to implement correct refcounting.
<code>funcs</code>	helper functions
<code>properties</code>	property tracking for this plane
<code>type</code>	type of plane (overlay, primary, cursor)
<code>state</code>	current atomic state for this plane

Name

struct drm_bridge_funcs — drm_bridge control functions

Synopsis

```
struct drm_bridge_funcs {
    int (* attach) (struct drm_bridge *bridge);
    bool (* mode_fixup) (struct drm_bridge *bridge, const struct drm_display_mode *mode);
    void (* disable) (struct drm_bridge *bridge);
    void (* post_disable) (struct drm_bridge *bridge);
    void (* mode_set) (struct drm_bridge *bridge, struct drm_display_mode *mode, struct drm_encoder *encoder);
    void (* pre_enable) (struct drm_bridge *bridge);
    void (* enable) (struct drm_bridge *bridge);
};
```

Members

attach	Called during <code>drm_bridge_attach</code>
mode_fixup	<p>This callback is used to validate and adjust a mode. The parameter <code>mode</code> is the display mode that should be fed to the next element in the display chain, either the final <code>drm_connector</code> or the next <code>drm_bridge</code>. The parameter <code>adjusted_mode</code> is the input mode the bridge requires. It can be modified by this callback and does not need to match <code>mode</code>.</p> <p>This is the only hook that allows a bridge to reject a modeset. If this function passes all other callbacks must succeed for this configuration.</p> <p>NOTE:</p> <p>This function is called in the check phase of atomic modesets, which can be aborted for any reason (including on userspace's request to just check whether a configuration would be possible). Drivers MUST NOT touch any persistent state (hardware or software) or data structures except the passed in <code>state</code> parameter.</p> <p>RETURNS:</p> <p>True if an acceptable configuration is possible, false if the modeset operation should be rejected.</p>
disable	<p>This callback should disable the bridge. It is called right before the preceding element in the display pipe is disabled. If the preceding element is a bridge this means it's called before that bridge's <code>->disable</code> function. If the preceding element is a <code>drm_encoder</code> it's called right before the encoder's <code>->disable</code>, <code>->prepare</code> or <code>->dpm</code> hook from <code>struct drm_encoder_helper_funcs</code>.</p> <p>The bridge can assume that the display pipe (i.e. clocks and timing signals) feeding it is still running when this callback is called.</p>
post_disable	<p>This callback should disable the bridge. It is called right after the preceding element in the display pipe is disabled. If the preceding element is a bridge this means it's called after that bridge's <code>->post_disable</code> function. If the preceding element is a</p>

drm_encoder it's called right after the encoder's `->disable`, `->prepare` or `->dpm`s hook from struct `drm_encoder_helper_funcs`.

The bridge must assume that the display pipe (i.e. clocks and timing signals) feeding it is no longer running when this callback is called.

mode_set

This callback should set the given mode on the bridge. It is called after the `->mode_set` callback for the preceding element in the display pipeline has been called already. The display pipe (i.e. clocks and timing signals) is off when this function is called.

pre_enable

This callback should enable the bridge. It is called right before the preceding element in the display pipe is enabled. If the preceding element is a bridge this means it's called before that bridge's `->pre_enable` function. If the preceding element is a `drm_encoder` it's called right before the encoder's `->enable`, `->commit` or `->dpm`s hook from struct `drm_encoder_helper_funcs`.

The display pipe (i.e. clocks and timing signals) feeding this bridge will not yet be running when this callback is called. The bridge must not enable the display link feeding the next bridge in the chain (if there is one) when this callback is called.

enable

This callback should enable the bridge. It is called right after the preceding element in the display pipe is enabled. If the preceding element is a bridge this means it's called after that bridge's `->enable` function. If the preceding element is a `drm_encoder` it's called right after the encoder's `->enable`, `->commit` or `->dpm`s hook from struct `drm_encoder_helper_funcs`.

The bridge can assume that the display pipe (i.e. clocks and timing signals) feeding it is running when this callback is called. This callback must enable the display link feeding the next bridge in the chain if there is one.

Name

struct `drm_bridge` — central DRM bridge control structure

Synopsis

```
struct drm_bridge {
    struct drm_device * dev;
    struct drm_encoder * encoder;
    struct drm_bridge * next;
#ifdef CONFIG_OF
    struct device_node * of_node;
#endif
    struct list_head list;
    const struct drm_bridge_funcs * funcs;
    void * driver_private;
};
```

Members

<code>dev</code>	DRM device this bridge belongs to
<code>encoder</code>	encoder to which this bridge is connected
<code>next</code>	the next bridge in the encoder chain
<code>of_node</code>	device node pointer to the bridge
<code>list</code>	to keep track of all added bridges
<code>funcs</code>	control functions
<code>driver_private</code>	pointer to the bridge driver's internal context

Name

struct `drm_atomic_state` — the global state object for atomic updates

Synopsis

```
struct drm_atomic_state {
    struct drm_device * dev;
    bool allow_modeset:1;
    bool legacy_cursor_update:1;
    struct drm_plane ** planes;
    struct drm_plane_state ** plane_states;
    struct drm_crtc ** crtc;
    struct drm_crtc_state ** crtc_states;
    int num_connector;
    struct drm_connector ** connectors;
    struct drm_connector_state ** connector_states;
    struct drm_modeset_acquire_ctx * acquire_ctx;
};
```

Members

<code>dev</code>	parent DRM device
<code>allow_modeset</code>	allow full modeset
<code>legacy_cursor_update</code>	hint to enforce legacy cursor IOCTL semantics
<code>planes</code>	pointer to array of plane pointers
<code>plane_states</code>	pointer to array of plane states pointers
<code>crtc</code>	pointer to array of CRTC pointers
<code>crtc_states</code>	pointer to array of CRTC states pointers
<code>num_connector</code>	size of the <i>connectors</i> and <i>connector_states</i> arrays
<code>connectors</code>	pointer to array of connector pointers
<code>connector_states</code>	pointer to array of connector states pointers
<code>acquire_ctx</code>	acquire context for this atomic modeset state update

Name

struct `drm_mode_set` — new values for a CRTC config change

Synopsis

```
struct drm_mode_set {
    struct drm_framebuffer * fb;
    struct drm_crtc * crtc;
    struct drm_display_mode * mode;
    uint32_t x;
    uint32_t y;
    struct drm_connector ** connectors;
    size_t num_connectors;
};
```

Members

<code>fb</code>	framebuffer to use for new config
<code>crtc</code>	CRTC whose configuration we're about to change
<code>mode</code>	mode timings to use
<code>x</code>	position of this CRTC relative to <i>fb</i>
<code>y</code>	position of this CRTC relative to <i>fb</i>
<code>connectors</code>	array of connectors to drive with this CRTC if possible
<code>num_connectors</code>	size of <i>connectors</i> array

Description

Represents a single crtc the connectors that it drives with what mode and from which framebuffer it scans out from.

This is used to set modes.

Name

struct `drm_mode_config_funcs` — basic driver provided mode setting functions

Synopsis

```
struct drm_mode_config_funcs {
    struct drm_framebuffer *(* fb_create) (struct drm_device *dev, struct drm_file *f,
    void (* output_poll_changed) (struct drm_device *dev);
    int (* atomic_check) (struct drm_device *dev, struct drm_atomic_state *state);
    int (* atomic_commit) (struct drm_device *dev, struct drm_atomic_state *state, bool
    struct drm_atomic_state *(* atomic_state_alloc) (struct drm_device *dev);
    void (* atomic_state_clear) (struct drm_atomic_state *state);
    void (* atomic_state_free) (struct drm_atomic_state *state);
};
```

Members

`fb_create`

Create a new framebuffer object. The core does basic checks on the requested metadata, but most of that is left to the driver. See struct `drm_mode_fb_cmd2` for details.

If the parameters are deemed valid and the backing storage objects in the underlying memory manager all exist, then the driver allocates a new `drm_framebuffer` structure, subclassed to contain driver-specific information (like the internal native buffer object references). It also needs to fill out all relevant metadata, which should be done by calling `drm_helper_mode_fill_fb_struct`.

The initialization is finalized by calling `drm_framebuffer_init`, which registers the framebuffer and makes it accessible to other threads.

RETURNS:

A new framebuffer with an initial reference count of 1 or a negative error code encoded with `ERR_PTR`.

`output_poll_changed`

Callback used by helpers to inform the driver of output configuration changes.

Drivers implementing fbdev emulation with the helpers can call `drm_fb_helper_hotplug_changed` from this hook to inform the fbdev helper of output changes.

FIXME:

Except that there's no vtable for device-level helper callbacks there's no reason this is a core function.

`atomic_check`

This is the only hook to validate an atomic modeset update. This function must reject any modeset and state changes which the hardware or driver doesn't support. This includes but is of course not limited to:

- Checking that the modes, framebuffers, scaling and placement requirements and so on are within the limits of the hardware.
- Checking that any hidden shared resources are not oversubscribed. This can be shared PLLs, shared lanes, overall memory bandwidth, display fifo space (where shared between planes or maybe even CRTC's).
- Checking that virtualized resources exported to userspace are not oversubscribed. For various reasons it can make sense to expose more planes, crtcs or encoders than which are physically there. One example is dual-pipe operations (which generally should be hidden from userspace if when lockstepped in hardware, exposed otherwise), where a plane might need 1 hardware plane (if it's just on one pipe), 2 hardware planes (when it spans both pipes) or maybe even shared a hardware plane with a 2nd plane (if there's a compatible plane requested on the area handled by the other pipe).
- Check that any transitional state is possible and that if requested, the update can indeed be done in the vblank period without temporarily disabling some functions.
- Check any other constraints the driver or hardware might have.
- This callback also needs to correctly fill out the `drm_crtc_state` in this update to make sure that `drm_atomic_crtc_needs_modeset` reflects the nature of the possible update and returns true if and only if the update cannot be applied without tearing within one vblank on that CRTC. The core uses that information to reject updates which require a full modeset (i.e. blanking the screen, or at least pausing updates for a substantial amount of time) if userspace has disallowed that in its request.
- The driver also does not need to repeat basic input validation like done for the corresponding legacy entry points. The core does that before calling this hook.

See the documentation of `atomic_commit` for an exhaustive list of error conditions which don't have to be checked at the `->atomic_check` stage?

See the documentation for struct `drm_atomic_state` for how exactly an atomic modeset update is described.

Drivers using the atomic helpers can implement this hook using `drm_atomic_helper_check`, or one of the exported sub-functions of it.

RETURNS:

0 on success or one of the below negative error codes:

- `-EINVAL`, if any of the above constraints are violated.
- `-EDEADLK`, when returned from an attempt to acquire an additional `drm_modeset_lock` through `drm_modeset_lock`.
- `-ENOMEM`, if allocating additional state sub-structures failed due to lack of memory.

-EINTR, -EAGAIN or -ERESTARTSYS, if the IOCTL should be restarted. This can either be due to a pending signal, or because the driver needs to completely bail out to recover from an exceptional situation like a GPU hang. From a userspace point all errors are treated equally.

atomic_commit

This is the only hook to commit an atomic modeset update. The core guarantees that *atomic_check* has been called successfully before calling this function, and that nothing has been changed in the interim.

See the documentation for struct *drm_atomic_state* for how exactly an atomic modeset update is described.

Drivers using the atomic helpers can implement this hook using *drm_atomic_helper_commit*, or one of the exported sub-functions of it.

Asynchronous commits (as indicated with the *async* parameter) must do any preparatory work which might result in an unsuccessful commit in the context of this callback. The only exceptions are hardware errors resulting in -EIO. But even in that case the driver must ensure that the display pipe is at least running, to avoid compositors crashing when pageflips don't work. Anything else, specifically committing the update to the hardware, should be done without blocking the caller. For updates which do not require a modeset this must be guaranteed.

The driver must wait for any pending rendering to the new framebuffers to complete before executing the flip. It should also wait for any pending rendering from other drivers if the underlying buffer is a shared dma-buf. Asynchronous commits must not wait for rendering in the context of this callback.

An application can request to be notified when the atomic commit has completed. These events are per-CRTC and can be distinguished by the CRTC index supplied in *drm_event* to userspace.

The drm core will supply a struct *drm_event* in the event member of each CRTC's *drm_crtc_state* structure. This can be handled by the *drm_crtc_send_vblank_event* function, which the driver should call on the provided event upon completion of the atomic commit. Note that if the driver supports vblank signalling and timestamping the vblank counters and timestamps must agree with the ones returned from page flip events. With the current vblank helper infrastructure this can be achieved by holding a vblank reference while the page flip is pending, acquired through *drm_crtc_vblank_get* and released with *drm_crtc_vblank_put*. Drivers are free to implement their own vblank counter and timestamp tracking though, e.g. if they have accurate timestamp registers in hardware.

NOTE:

Drivers are not allowed to shut down any display pipe successfully enabled through an atomic commit on their own. Doing so can result in compositors crashing if a page flip is suddenly rejected because the pipe is off.

RETURNS:

0 on success or one of the below negative error codes:

- EBUSY, if an asynchronous update is requested and there is an earlier update pending. Drivers are allowed to support a queue of outstanding updates, but currently no driver supports that. Note that drivers must wait for preceding updates to complete if a synchronous update is requested, they are not allowed to fail the commit in that case.

- ENOMEM, if the driver failed to allocate memory. Specifically this can happen when trying to pin framebuffers, which must only be done when committing the state.

- ENOSPC, as a refinement of the more generic -ENOMEM to indicate that the driver has run out of vram, iommu space or similar GPU address space needed for framebuffer.

- EIO, if the hardware completely died.

- EINTR, -EAGAIN or -ERESTARTSYS, if the IOCTL should be restarted. This can either be due to a pending signal, or because the driver needs to completely bail out to recover from an exceptional situation like a GPU hang. From a userspace point of view all errors are treated equally.

This list is exhaustive. Specifically this hook is not allowed to return -EINVAL (any invalid requests should be caught in *atomic_check*) or -EDEADLK (this function must not acquire additional modeset locks).

`atomic_state_alloc`

This optional hook can be used by drivers that want to subclass struct `drm_atomic_state` to be able to track their own driver-private global state easily. If this hook is implemented, drivers must also implement *atomic_state_clear* and *atomic_state_free*.

RETURNS:

A new `drm_atomic_state` on success or NULL on failure.

`atomic_state_clear`

This hook must clear any driver private state duplicated into the passed-in `drm_atomic_state`. This hook is called when the caller encountered a `drm_modeset_lock` deadlock and needs to drop all already acquired locks as part of the deadlock avoidance dance implemented in `drm_modeset_lock_backoff`.

Any duplicated state must be invalidated since a concurrent atomic update might change it, and the `drm` atomic interfaces always apply updates as relative changes to the current state.

Drivers that implement this must call `drm_atomic_state_default_clear` to clear common state.

`atomic_state_free`

This hook needs driver private resources and the `drm_atomic_state` itself. Note that the core first calls `drm_atomic_state_clear` to avoid code duplicate between the clear and free hooks.

Drivers that implement this must call `drm_atomic_state_default_free` to release common resources.

Description

Some global (i.e. not per-CRTC, connector, etc) mode setting functions that involve drivers.

Name

struct drm_mode_config — Mode configuration control structure

Synopsis

```
struct drm_mode_config {
    struct mutex mutex;
    struct drm_modeset_lock connection_mutex;
    struct drm_modeset_acquire_ctx * acquire_ctx;
    struct mutex idr_mutex;
    struct idr crtc_idr;
    struct mutex fb_lock;
    int num_fb;
    struct list_head fb_list;
    int num_connector;
    struct list_head connector_list;
    int num_encoder;
    struct list_head encoder_list;
    int num_overlay_plane;
    int num_total_plane;
    struct list_head plane_list;
    int num_crtc;
    struct list_head crtc_list;
    struct list_head property_list;
    int min_width;
    int min_height;
    int max_width;
    int max_height;
    const struct drm_mode_config_funcs * funcs;
    resource_size_t fb_base;
    bool poll_enabled;
    bool poll_running;
    struct delayed_work output_poll_work;
    struct mutex blob_lock;
    struct list_head property_blob_list;
    uint32_t preferred_depth;
    uint32_t prefer_shadow;
    bool async_page_flip;
    uint32_t cursor_width;
    uint32_t cursor_height;
};
```

Members

mutex	mutex protecting KMS related lists and structures
connection_mutex	ww mutex protecting connector state and routing
acquire_ctx	global implicit acquire context used by atomic drivers for legacy IOCTLs
idr_mutex	mutex for KMS ID allocation and management
crtc_idr	main KMS ID tracking object

fb_lock	mutex to protect fb state and lists
num_fb	number of fbs available
fb_list	list of framebuffers available
num_connector	number of connectors on this device
connector_list	list of connector objects
num_encoder	number of encoders on this device
encoder_list	list of encoder objects
num_overlay_plane	number of overlay planes on this device
num_total_plane	number of universal (i.e. with primary/cursor) planes on this device
plane_list	list of plane objects
num_crtc	number of CRTC's on this device
crtc_list	list of CRTC objects
property_list	list of property objects
min_width	minimum pixel width on this device
min_height	minimum pixel height on this device
max_width	maximum pixel width on this device
max_height	maximum pixel height on this device
funcs	core driver provided mode setting functions
fb_base	base address of the framebuffer
poll_enabled	track polling support for this device
poll_running	track polling status for this device
output_poll_work	delayed work for polling in process context
blob_lock	mutex for blob property allocation and management
property_blob_list	list of all the blob property objects
preferred_depth	preferred RGB pixel depth, used by fb helpers
prefer_shadow	hint to userspace to prefer shadow-fb rendering
async_page_flip	does this device support async flips on the primary plane?
cursor_width	hint to userspace for max cursor width
cursor_height	hint to userspace for max cursor height

_property

core property tracking

Description

Core mode resource tracking structure. All CRTC, encoders, and connectors enumerated by the driver are added here, as are global properties. Some global restrictions are also here, e.g. dimension restrictions.

Name

`drm_for_each_plane_mask` — iterate over planes specified by bitmask

Synopsis

```
drm_for_each_plane_mask ( plane, dev, plane_mask );
```

Arguments

<i>plane</i>	the loop cursor
<i>dev</i>	the DRM device
<i>plane_mask</i>	bitmask of plane indices

Description

Iterate over all planes specified by bitmask.

Name

`drm_crtc_mask` — find the mask of a registered CRTC

Synopsis

```
uint32_t drm_crtc_mask (struct drm_crtc * crtc);
```

Arguments

crtc CRTC to find mask for

Description

Given a registered CRTC, return the mask bit of that CRTC for an encoder's `possible_crtcs` field.

Name

`drm_encoder_crtc_ok` — can a given crtc drive a given encoder?

Synopsis

```
bool drm_encoder_crtc_ok (struct drm_encoder * encoder, struct drm_crtc  
* crtc);
```

Arguments

encoder encoder to test

crtc crtc to test

Description

Return false if *encoder* can't be driven by *crtc*, true otherwise.

KMS Locking

As KMS moves toward more fine grained locking, and atomic ioctl where userspace can indirectly control locking order, it becomes necessary to use `ww_mutex` and acquire-contexts to avoid deadlocks. But because the locking is more distributed around the driver code, we want a bit of extra utility/tracking out of our acquire-ctx. This is provided by `drm_modeset_lock` / `drm_modeset_acquire_ctx`.

For basic principles of `ww_mutex`, see: [Documentation/locking/ww-mutex-design.txt](#)

The basic usage pattern is to:

```
drm_modeset_acquire_init(ctx) retry: foreach (lock in random_ordered_set_of_locks) { ret = drm_mod-  
eset_lock(lock, ctx) if (ret == -EDEADLK) { drm_modeset_backoff(ctx); goto retry; } } ... do stuff ...  
drm_modeset_drop_locks(ctx); drm_modeset_acquire_fini(ctx);
```

Name

struct drm_modeset_acquire_ctx — locking context (see ww_acquire_ctx)

Synopsis

```
struct drm_modeset_acquire_ctx {
    struct ww_acquire_ctx ww_ctx;
    struct drm_modeset_lock * contended;
    struct list_head locked;
    bool trylock_only;
};
```

Members

ww_ctx	base acquire ctx
contended	used internally for -EDEADLK handling
locked	list of held locks
trylock_only	trylock mode used in atomic contexts/panic notifiers

Description

Each thread competing for a set of locks must use one acquire ctx. And if any lock fxn returns -EDEADLK, it must backoff and retry.

Name

struct `drm_modeset_lock` — used for locking modeset resources.

Synopsis

```
struct drm_modeset_lock {  
    struct ww_mutex mutex;  
    struct list_head head;  
};
```

Members

`mutex` resource locking

`head` used to hold it's place on state->locked list when part of an atomic update

Description

Used for locking CRTC's and other modeset resources.

Name

`drm_modeset_lock_init` — initialize lock

Synopsis

```
void drm_modeset_lock_init (struct drm_modeset_lock * lock);
```

Arguments

lock lock to init

Name

`drm_modeset_lock_fini` — cleanup lock

Synopsis

```
void drm_modeset_lock_fini (struct drm_modeset_lock * lock);
```

Arguments

lock lock to cleanup

Name

`drm_modeset_is_locked` — equivalent to `mutex_is_locked`

Synopsis

```
bool drm_modeset_is_locked (struct drm_modeset_lock * lock);
```

Arguments

lock lock to check

Name

`drm_modeset_lock_all` — take all modeset locks

Synopsis

```
void drm_modeset_lock_all (struct drm_device * dev);
```

Arguments

dev DRM device

Description

This function takes all modeset locks, suitable where a more fine-grained scheme isn't (yet) implemented. Locks must be dropped by calling the `drm_modeset_unlock_all` function.

This function is deprecated. It allocates a lock acquisition context and stores it in the DRM device's `->mode_config`. This facilitate conversion of existing code because it removes the need to manually deal with the acquisition context, but it is also brittle because the context is global and care must be taken not to nest calls. New code should use the `drm_modeset_lock_all_ctx` function and pass in the context explicitly.

Name

`drm_modeset_unlock_all` — drop all modeset locks

Synopsis

```
void drm_modeset_unlock_all (struct drm_device * dev);
```

Arguments

dev DRM device

Description

This function drops all modeset locks taken by a previous call to the `drm_modeset_lock_all` function.

This function is deprecated. It uses the lock acquisition context stored in the DRM device's `->mode_config`. This facilitates conversion of existing code because it removes the need to manually deal with the acquisition context, but it is also brittle because the context is global and care must be taken not to nest calls. New code should pass the acquisition context directly to the `drm_modeset_drop_locks` function.

Name

`drm_modeset_lock_crtc` — lock crtc with hidden acquire ctx for a plane update

Synopsis

```
void drm_modeset_lock_crtc (struct drm_crtc * crtc, struct drm_plane  
* plane);
```

Arguments

crtc DRM CRTC

plane DRM plane to be updated on *crtc*

Description

This function locks the given crtc and plane (which should be either the primary or cursor plane) using a hidden acquire context. This is necessary so that drivers internally using the atomic interfaces can grab further locks with the lock acquire context.

Note that *plane* can be NULL, e.g. when the cursor support hasn't yet been converted to universal planes yet.

Name

`drm_modeset_legacy_acquire_ctx` — find acquire ctx for legacy ioctls

Synopsis

```
struct drm_modeset_acquire_ctx * drm_modeset_legacy_acquire_ctx (struct
drm_crtc * crtc);
```

Arguments

crtc drm crtc

Description

Legacy ioctl operations like cursor updates or page flips only have per-crtc locking, and store the acquire ctx in the corresponding crtc. All other legacy operations take all locks and use a global acquire context. This function grabs the right one.

Name

`drm_modeset_unlock_crtc` — drop crtc lock

Synopsis

```
void drm_modeset_unlock_crtc (struct drm_crtc * crtc);
```

Arguments

crtc drm crtc

Description

This drops the crtc lock acquire with `drm_modeset_lock_crtc` and all other locks acquired through the hidden context.

Name

`drm_warn_on_modeset_not_all_locked` — check that all modeset locks are locked

Synopsis

```
void drm_warn_on_modeset_not_all_locked (struct drm_device * dev);
```

Arguments

dev device

Description

Useful as a debug assert.

Name

`drm_modeset_acquire_init` — initialize acquire context

Synopsis

```
void drm_modeset_acquire_init (struct drm_modeset_acquire_ctx * ctx,  
uint32_t flags);
```

Arguments

ctx the acquire context

flags for future

Name

`drm_modeset_acquire_fini` — cleanup acquire context

Synopsis

```
void drm_modeset_acquire_fini (struct drm_modeset_acquire_ctx * ctx);
```

Arguments

ctx the acquire context

Name

`drm_modeset_drop_locks` — drop all locks

Synopsis

```
void drm_modeset_drop_locks (struct drm_modeset_acquire_ctx * ctx);
```

Arguments

ctx the acquire context

Description

Drop all locks currently held against this acquire context.

Name

`drm_modeset_backoff` — deadlock avoidance backoff

Synopsis

```
void drm_modeset_backoff (struct drm_modeset_acquire_ctx * ctx);
```

Arguments

ctx the acquire context

Description

If deadlock is detected (ie. `drm_modeset_lock` returns `-EDEADLK`), you must call this function to drop all currently held locks and block until the contended lock becomes available.

Name

`drm_modeset_backoff_interruptible` — deadlock avoidance backoff

Synopsis

```
int drm_modeset_backoff_interruptible (struct drm_modeset_acquire_ctx
* ctx);
```

Arguments

ctx the acquire context

Description

Interruptible version of `drm_modeset_backoff`

Name

`drm_modeset_lock` — take modeset lock

Synopsis

```
int drm_modeset_lock (struct drm_modeset_lock * lock, struct drm_modeset_acquire_ctx * ctx);
```

Arguments

lock lock to take

ctx acquire ctx

Description

If *ctx* is not NULL, then its ww acquire context is used and the lock will be tracked by the context and can be released by calling `drm_modeset_drop_locks`. If `-EDEADLK` is returned, this means a deadlock scenario has been detected and it is an error to attempt to take any more locks without first calling `drm_modeset_backoff`.

Name

`drm_modeset_lock_interruptible` — take modeset lock

Synopsis

```
int drm_modeset_lock_interruptible (struct drm_modeset_lock * lock,  
struct drm_modeset_acquire_ctx * ctx);
```

Arguments

lock lock to take

ctx acquire ctx

Description

Interruptible version of `drm_modeset_lock`

Name

`drm_modeset_unlock` — drop modeset lock

Synopsis

```
void drm_modeset_unlock (struct drm_modeset_lock * lock);
```

Arguments

lock lock to release

Name

`drm_modeset_lock_all_ctx` — take all modeset locks

Synopsis

```
int drm_modeset_lock_all_ctx (struct drm_device * dev, struct drm_modeset_acquire_ctx * ctx);
```

Arguments

dev DRM device

ctx lock acquisition context

Description

This function takes all modeset locks, suitable where a more fine-grained scheme isn't (yet) implemented.

Unlike `drm_modeset_lock_all`, it doesn't take the `dev->mode_config.mutex` since that lock isn't required for modeset state changes. Callers which need to grab that lock too need to do so outside of the acquire context *ctx*.

Locks acquired with this function should be released by calling the `drm_modeset_drop_locks` function on *ctx*.

Returns

0 on success or a negative error-code on failure.

Mode Setting Helper Functions

The plane, CRTC, encoder and connector functions provided by the drivers implement the DRM API. They're called by the DRM core and ioctl handlers to handle device state changes and configuration request. As implementing those functions often requires logic not specific to drivers, mid-layer helper functions are available to avoid duplicating boilerplate code.

The DRM core contains one mid-layer implementation. The mid-layer provides implementations of several plane, CRTC, encoder and connector functions (called from the top of the mid-layer) that pre-process requests and call lower-level functions provided by the driver (at the bottom of the mid-layer). For instance, the `drm_crtc_helper_set_config` function can be used to fill the `struct drm_crtc_funcs` *set_config* field. When called, it will split the *set_config* operation in smaller, simpler operations and call the driver to handle them.

To use the mid-layer, drivers call `drm_crtc_helper_add`, `drm_encoder_helper_add` and `drm_connector_helper_add` functions to install their mid-layer bottom operations handlers, and fill the `drm_crtc_funcs`, `drm_encoder_funcs` and `drm_connector_funcs` structures with pointers to the mid-layer top API functions. Installing the mid-layer bottom operation handlers is best done right after registering the corresponding KMS object.

The mid-layer is not split between CRTC, encoder and connector operations. To use it, a driver must provide bottom functions for all of the three KMS entities.

Atomic Modeset Helper Functions Reference

Overview

This helper library provides implementations of check and commit functions on top of the CRTC modeset helper callbacks and the plane helper callbacks. It also provides convenience implementations for the atomic state handling callbacks for drivers which don't need to subclass the drm core structures to add their own additional internal state.

This library also provides default implementations for the check callback in `drm_atomic_helper_check` and for the commit callback with `drm_atomic_helper_commit`. But the individual stages and callbacks are exposed to allow drivers to mix and match and e.g. use the plane helpers only together with a driver private modeset implementation.

This library also provides implementations for all the legacy driver interfaces on top of the atomic interface. See `drm_atomic_helper_set_config`, `drm_atomic_helper_disable_plane`, `drm_atomic_helper_disable_plane` and the various functions to implement `set_property` callbacks. New drivers must not implement these functions themselves but must use the provided helpers.

The atomic helper uses the same function table structures as all other modesetting helpers. See the documentation for struct `drm_crtc_helper_funcs`, struct `drm_encoder_helper_funcs` and struct `drm_connector_helper_funcs`. It also shares the struct `drm_plane_helper_funcs` function table with the plane helpers.

Implementing Asynchronous Atomic Commit

For now the atomic helpers don't support async commit directly. If there is real need it could be added though, using the dma-buf fence infrastructure for generic synchronization with outstanding rendering.

For now drivers have to implement async commit themselves, with the following sequence being the recommended one:

1. Run `drm_atomic_helper_prepare_planes` first. This is the only function which commit needs to call which can fail, so we want to run it first and synchronously.
2. Synchronize with any outstanding asynchronous commit worker threads which might be affected the new state update. This can be done by either cancelling or flushing the work items, depending upon whether the driver can deal with cancelled updates. Note that it is important to ensure that the framebuffer cleanup is still done when cancelling.

For sufficient parallelism it is recommended to have a work item per crtc (for updates which don't touch global state) and a global one. Then we only need to synchronize with the crtc work items for changed crtcs and the global work item, which allows nice concurrent updates on disjoint sets of crtcs.

3. The software state is updated synchronously with `drm_atomic_helper_swap_state`. Doing this under the protection of all modeset locks means concurrent callers never see inconsistent state. And doing this while it's guaranteed that no relevant async worker runs means that async workers do not need grab any locks. Actually they must not grab locks, for otherwise the work flushing will deadlock.

4. Schedule a work item to do all subsequent steps, using the split-out commit helpers: a) pre-plane commit b) plane commit c) post-plane commit and then cleaning up the framebuffers after the old framebuffer is no longer being displayed.

Atomic State Reset and Initialization

Both the drm core and the atomic helpers assume that there is always the full and correct atomic software state for all connectors, CRTCs and planes available. Which is a bit a problem on driver load and also after system suspend. One way to solve this is to have a hardware state read-out infrastructure which reconstructs the full software state (e.g. the i915 driver).

The simpler solution is to just reset the software state to everything off, which is easiest to do by calling `drm_mode_config_reset`. To facilitate this the atomic helpers provide default reset implementations for all hooks.

On the upside the precise state tracking of atomic simplifies system suspend and resume a lot. For drivers using `drm_mode_config_reset` a complete recipe is implemented in `drm_atomic_helper_suspend` and `drm_atomic_helper_resume`. For other drivers the building blocks are split out, see the documentation for these functions.

Name

`drm_atomic_crtc_for_each_plane` — iterate over planes currently attached to CRTC

Synopsis

```
drm_atomic_crtc_for_each_plane ( plane, crtc );
```

Arguments

plane the loop cursor

crtc the crtc whose planes are iterated

Description

This iterates over the current state, useful (for example) when applying atomic state after it has been checked and swapped. To iterate over the planes which **will** be attached (for `->atomic_check`) see `drm_crtc_for_each_pending_plane`

Name

`drm_atomic_crtc_state_for_each_plane` — iterate over attached planes in new state

Synopsis

```
drm_atomic_crtc_state_for_each_plane ( plane, crtc_state );
```

Arguments

plane the loop cursor

crtc_state the incoming crtc-state

Description

Similar to `drm_crtc_for_each_plane`, but iterates the planes that will be attached if the specified state is applied. Useful during (for example) `->atomic_check` operations, to validate the incoming state

Name

`drm_atomic_helper_check_modeset` — validate state object for modeset changes

Synopsis

```
int drm_atomic_helper_check_modeset (struct drm_device * dev, struct
drm_atomic_state * state);
```

Arguments

dev DRM device

state the driver state object

Description

Check the state object to see if the requested state is physically possible. This does all the crtc and connector related computations for an atomic update and adds any additional connectors needed for full modesets and calls down into `->mode_fixup` functions of the driver backend.

`crtc_state->mode_changed` is set when the input mode is changed. `crtc_state->connectors_changed` is set when a connector is added or removed from the crtc. `crtc_state->active_changed` is set when `crtc_state->active` changes, which is used for dpms.

IMPORTANT

Drivers which update `->mode_changed` (e.g. in their `->atomic_check` hooks if a plane update can't be done without a full modeset) `_must_` call this function afterwards after that change. It is permitted to call this function multiple times for the same update, e.g. when the `->atomic_check` functions depend upon the adjusted dotclock for fifo space allocation and watermark computation.

RETURNS Zero for success or `-errno`

Name

`drm_atomic_helper_check_planes` — validate state object for planes changes

Synopsis

```
int drm_atomic_helper_check_planes (struct drm_device * dev, struct
drm_atomic_state * state);
```

Arguments

dev DRM device

state the driver state object

Description

Check the state object to see if the requested state is physically possible. This does all the plane update related checks using by calling into the `->atomic_check` hooks provided by the driver.

It also sets `crtc_state->planes_changed` to indicate that a crtc has updated planes.

RETURNS Zero for success or `-errno`

Name

`drm_atomic_helper_check` — validate state object

Synopsis

```
int drm_atomic_helper_check (struct drm_device * dev, struct drm_atomic_state * state);
```

Arguments

dev DRM device

state the driver state object

Description

Check the state object to see if the requested state is physically possible. Only crtc's and planes have check callbacks, so for any additional (global) checking that a driver needs it can simply wrap that around this function. Drivers without such needs can directly use this as their `->atomic_check` callback.

This just wraps the two parts of the state checking for planes and modeset

state in the default order

First it calls `drm_atomic_helper_check_modeset` and then `drm_atomic_helper_check_planes`. The assumption is that the `->atomic_check` functions depend upon an updated `adjusted_mode.clock` to e.g. properly compute watermarks.

RETURNS Zero for success or `-errno`

Name

`drm_atomic_helper_update_legacy_modeset_state` — update legacy modeset state

Synopsis

```
void drm_atomic_helper_update_legacy_modeset_state (struct drm_device  
* dev, struct drm_atomic_state * old_state);
```

Arguments

dev DRM device

old_state atomic state object with old state structures

Description

This function updates all the various legacy modeset state pointers in connectors, encoders and crtcs. It also updates the timestamping constants used for precise vblank timestamps by calling `drm_calc_timestamping_constants`.

Drivers can use this for building their own atomic commit if they don't have a pure helper-based modeset implementation.

Name

`drm_atomic_helper_commit_modeset_disables` — modeset commit to disable outputs

Synopsis

```
void drm_atomic_helper_commit_modeset_disables (struct drm_device *  
dev, struct drm_atomic_state * old_state);
```

Arguments

dev DRM device

old_state atomic state object with old state structures

Description

This function shuts down all the outputs that need to be shut down and prepares them (if required) with the new mode.

For compatibility with legacy crtc helpers this should be called before `drm_atomic_helper_commit_planes`, which is what the default commit function does. But drivers with different needs can group the modeset commits together and do the plane commits at the end. This is useful for drivers doing runtime PM since planes updates then only happen when the CRTC is actually enabled.

Name

`drm_atomic_helper_commit_modeset_enables` — modeset commit to enable outputs

Synopsis

```
void drm_atomic_helper_commit_modeset_enables (struct drm_device * dev,  
struct drm_atomic_state * old_state);
```

Arguments

dev DRM device

old_state atomic state object with old state structures

Description

This function enables all the outputs with the new configuration which had to be turned off for the update.

For compatibility with legacy crtc helpers this should be called after `drm_atomic_helper_commit_planes`, which is what the default commit function does. But drivers with different needs can group the modeset commits together and do the plane commits at the end. This is useful for drivers doing runtime PM since planes updates then only happen when the CRTC is actually enabled.

Name

`drm_atomic_helper_wait_for_vblanks` — wait for vblank on crtcs

Synopsis

```
void drm_atomic_helper_wait_for_vblanks (struct drm_device * dev, struct  
drm_atomic_state * old_state);
```

Arguments

dev DRM device

old_state atomic state object with old state structures

Description

Helper to, after atomic commit, wait for vblanks on all effected crtcs (ie. before cleaning up old frame-buffers using `drm_atomic_helper_cleanup_planes`). It will only wait on crtcs where the frame-buffers have actually changed to optimize for the legacy cursor and plane update use-case.

Name

`drm_atomic_helper_commit` — commit validated state object

Synopsis

```
int drm_atomic_helper_commit (struct drm_device * dev, struct drm_atomic_state * state, bool async);
```

Arguments

dev DRM device

state the driver state object

async asynchronous commit

Description

This function commits a with `drm_atomic_helper_check` pre-validated state object. This can still fail when e.g. the framebuffer reservation fails. For now this doesn't implement asynchronous commits.

Note that right now this function does not support async commits, and hence driver writers must implement their own version for now. Also note that the default ordering of how the various stages are called is to match the legacy modeset helper library closest. One peculiarity of that is that it doesn't mesh well with runtime PM at all.

For drivers supporting runtime PM the recommended sequence is

```
drm_atomic_helper_commit_modeset_disables(dev, state);
```

```
drm_atomic_helper_commit_modeset_enables(dev, state);
```

```
drm_atomic_helper_commit_planes(dev, state, true);
```

See the kerneldoc entries for these three functions for more details.

RETURNS Zero for success or -errno.

Name

`drm_atomic_helper_prepare_planes` — prepare plane resources before commit

Synopsis

```
int drm_atomic_helper_prepare_planes (struct drm_device * dev, struct
drm_atomic_state * state);
```

Arguments

dev DRM device

state atomic state object with new state structures

Description

This function prepares plane state, specifically framebuffers, for the new configuration. If any failure is encountered this function will call `->cleanup_fb` on any already successfully prepared framebuffer.

Returns

0 on success, negative error code on failure.

Name

`drm_atomic_helper_commit_planes` — commit plane state

Synopsis

```
void drm_atomic_helper_commit_planes (struct drm_device * dev, struct
drm_atomic_state * old_state, bool active_only);
```

Arguments

<i>dev</i>	DRM device
<i>old_state</i>	atomic state object with old state structures
<i>active_only</i>	Only commit on active CRTC if set

Description

This function commits the new plane state using the plane and atomic helper functions for planes and crtcs. It assumes that the atomic state has already been pushed into the relevant object state pointers, since this step can no longer fail.

It still requires the global state object *old_state* to know which planes and crtcs need to be updated though.

Note that this function does all plane updates across all CRTCs in one step. If the hardware can't support this approach look at `drm_atomic_helper_commit_planes_on_crtc` instead.

Plane parameters can be updated by applications while the associated CRTC is disabled. The DRM/KMS core will store the parameters in the plane state, which will be available to the driver when the CRTC is turned on. As a result most drivers don't need to be immediately notified of plane updates for a disabled CRTC.

Unless otherwise needed, drivers are advised to set the *active_only* parameters to true in order not to receive plane update notifications related to a disabled CRTC. This avoids the need to manually ignore plane updates in driver code when the driver and/or hardware can't or just don't need to deal with updates on disabled CRTCs, for example when supporting runtime PM.

The `drm_atomic_helper_commit` default implementation only sets *active_only* to false to most closely match the behaviour of the legacy helpers. This should not be copied blindly by drivers.

Name

`drm_atomic_helper_commit_planes_on_crtc` — commit plane state for a crtc

Synopsis

```
void drm_atomic_helper_commit_planes_on_crtc (struct drm_crtc_state *  
old_crtc_state);
```

Arguments

old_crtc_state atomic state object with the old crtc state

Description

This function commits the new plane state using the plane and atomic helper functions for planes on the specific crtc. It assumes that the atomic state has already been pushed into the relevant object state pointers, since this step can no longer fail.

This function is useful when plane updates should be done crtc-by-crtc instead of one global step like `drm_atomic_helper_commit_planes` does.

This function can only be safely used when planes are not allowed to move between different CRTC's because this function doesn't handle inter-CRTC dependencies. Callers need to ensure that either no such dependencies exist, resolve them through ordering of commit calls or through some other means.

Name

`drm_atomic_helper_disable_planes_on_crtc` — helper to disable CRTC's planes

Synopsis

```
void drm_atomic_helper_disable_planes_on_crtc (struct drm_crtc * crtc,  
bool atomic);
```

Arguments

crtc CRTC

atomic if set, synchronize with CRTC's `atomic_begin/flush` hooks

Description

Disables all planes associated with the given CRTC. This can be used for instance in the CRTC helper disable callback to disable all planes before shutting down the display pipeline.

If the `atomic`-parameter is set the function calls the CRTC's `atomic_begin` hook before and `atomic_flush` hook after disabling the planes.

It is a bug to call this function without having implemented the `->atomic_disable` plane hook.

Name

`drm_atomic_helper_cleanup_planes` — cleanup plane resources after commit

Synopsis

```
void drm_atomic_helper_cleanup_planes (struct drm_device * dev, struct
drm_atomic_state * old_state);
```

Arguments

dev DRM device

old_state atomic state object with old state structures

Description

This function cleans up plane state, specifically framebuffers, from the old configuration. Hence the old configuration must be perserved in *old_state* to be able to call this function.

This function must also be called on the new state when the atomic update fails at any point after calling `drm_atomic_helper_prepare_planes`.

Name

`drm_atomic_helper_swap_state` — store atomic state into current sw state

Synopsis

```
void drm_atomic_helper_swap_state (struct drm_device * dev, struct dr-  
m_atomic_state * state);
```

Arguments

dev DRM device

state atomic state

Description

This function stores the atomic state into the current state pointers in all driver objects. It should be called after all failing steps have been done and succeeded, but before the actual hardware state is committed.

For cleanup and error recovery the current state for all changed objects will be swapped into *state*.

With that sequence it fits perfectly into the plane prepare/cleanup sequence:

1. Call `drm_atomic_helper_prepare_planes` with the staged atomic state.
2. Do any other steps that might fail.
3. Put the staged state into the current state pointers with this function.
4. Actually commit the hardware state.
5. Call `drm_atomic_helper_cleanup_planes` with *state*, which since step 3 contains the old state. Also do any other cleanup required with that state.

Name

`drm_atomic_helper_update_plane` — Helper for primary plane update using atomic

Synopsis

```
int drm_atomic_helper_update_plane (struct drm_plane * plane, struct
drm_crtc * crtc, struct drm_framebuffer * fb, int crtc_x, int crtc_y,
unsigned int crtc_w, unsigned int crtc_h, uint32_t src_x, uint32_t
src_y, uint32_t src_w, uint32_t src_h);
```

Arguments

<i>plane</i>	plane object to update
<i>crtc</i>	owning CRTC of owning plane
<i>fb</i>	framebuffer to flip onto plane
<i>crtc_x</i>	x offset of primary plane on crtc
<i>crtc_y</i>	y offset of primary plane on crtc
<i>crtc_w</i>	width of primary plane rectangle on crtc
<i>crtc_h</i>	height of primary plane rectangle on crtc
<i>src_x</i>	x offset of <i>fb</i> for panning
<i>src_y</i>	y offset of <i>fb</i> for panning
<i>src_w</i>	width of source rectangle in <i>fb</i>
<i>src_h</i>	height of source rectangle in <i>fb</i>

Description

Provides a default plane update handler using the atomic driver interface.

RETURNS

Zero on success, error code on failure

Name

`drm_atomic_helper_disable_plane` — Helper for primary plane disable using * atomic

Synopsis

```
int drm_atomic_helper_disable_plane (struct drm_plane * plane);
```

Arguments

plane plane to disable

Description

Provides a default plane disable handler using the atomic driver interface.

RETURNS

Zero on success, error code on failure

Name

`drm_atomic_helper_set_config` — set a new config from userspace

Synopsis

```
int drm_atomic_helper_set_config (struct drm_mode_set * set);
```

Arguments

set mode set configuration

Description

Provides a default crtc `set_config` handler using the atomic driver interface.

Returns

Returns 0 on success, negative `errno` numbers on failure.

Name

`drm_atomic_helper_disable_all` — disable all currently active outputs

Synopsis

```
int drm_atomic_helper_disable_all (struct drm_device * dev, struct dr-  
m_modeset_acquire_ctx * ctx);
```

Arguments

dev DRM device

ctx lock acquisition context

Description

Loops through all connectors, finding those that aren't turned off and then turns them off by setting their DPMS mode to OFF and deactivating the CRTC that they are connected to.

This is used for example in suspend/resume to disable all currently active functions when suspending.

Note that if callers haven't already acquired all modeset locks this might return `-EDEADLK`, which must be handled by calling `drm_modeset_backoff`.

Returns

0 on success or a negative error code on failure.

See also

`drm_atomic_helper_suspend`, `drm_atomic_helper_resume`

Name

`drm_atomic_helper_suspend` — subsystem-level suspend helper

Synopsis

```
struct drm_atomic_state * drm_atomic_helper_suspend (struct drm_device
* dev);
```

Arguments

dev DRM device

Description

Duplicates the current atomic state, disables all active outputs and then returns a pointer to the original atomic state to the caller. Drivers can pass this pointer to the `drm_atomic_helper_resume` helper upon resume to restore the output configuration that was active at the time the system entered suspend.

Note that it is potentially unsafe to use this. The atomic state object returned by this function is assumed to be persistent. Drivers must ensure that this holds true. Before calling this function, drivers must make sure to suspend fbdev emulation so that nothing can be using the device.

Returns

A pointer to a copy of the state before suspend on success or an `ERR_PTR`- encoded error code on failure. Drivers should store the returned atomic state object and pass it to the `drm_atomic_helper_resume` helper upon resume.

See also

`drm_atomic_helper_duplicate_state`, `drm_atomic_helper_disable_all`, `drm_atomic_helper_resume`

Name

`drm_atomic_helper_resume` — subsystem-level resume helper

Synopsis

```
int drm_atomic_helper_resume (struct drm_device * dev, struct drm_atomic_state * state);
```

Arguments

dev DRM device

state atomic state to resume to

Description

Calls `drm_mode_config_reset` to synchronize hardware and software states, grabs all modeset locks and commits the atomic state object. This can be used in conjunction with the `drm_atomic_helper_suspend` helper to implement suspend/resume for drivers that support atomic mode-setting.

Returns

0 on success or a negative error code on failure.

See also

`drm_atomic_helper_suspend`

Name

`drm_atomic_helper_crtc_set_property` — helper for crtc properties

Synopsis

```
int drm_atomic_helper_crtc_set_property (struct drm_crtc * crtc, struct
drm_property * property, uint64_t val);
```

Arguments

<i>crtc</i>	DRM crtc
<i>property</i>	DRM property
<i>val</i>	value of property

Description

Provides a default crtc `set_property` handler using the atomic driver interface.

RETURNS

Zero on success, error code on failure

Name

`drm_atomic_helper_plane_set_property` — helper for plane properties

Synopsis

```
int drm_atomic_helper_plane_set_property (struct drm_plane * plane,
struct drm_property * property, uint64_t val);
```

Arguments

<i>plane</i>	DRM plane
<i>property</i>	DRM property
<i>val</i>	value of property

Description

Provides a default plane set_property handler using the atomic driver interface.

RETURNS

Zero on success, error code on failure

Name

`drm_atomic_helper_connector_set_property` — helper for connector properties

Synopsis

```
int drm_atomic_helper_connector_set_property (struct drm_connector *  
connector, struct drm_property * property, uint64_t val);
```

Arguments

connector DRM connector

property DRM property

val value of property

Description

Provides a default connector `set_property` handler using the atomic driver interface.

RETURNS

Zero on success, error code on failure

Name

`drm_atomic_helper_page_flip` — execute a legacy page flip

Synopsis

```
int drm_atomic_helper_page_flip (struct drm_crtc * crtc, struct dr-  
m_framebuffer * fb, struct drm_pending_vblank_event * event, uint32_t  
flags);
```

Arguments

crtc DRM crtc

fb DRM framebuffer

event optional DRM event to signal upon completion

flags flip flags for non-vblank sync'ed updates

Description

Provides a default page flip implementation using the atomic driver interface.

Note that for now so called async page flips (i.e. updates which are not synchronized to vblank) are not supported, since the atomic interfaces have no provisions for this yet.

Returns

Returns 0 on success, negative errno numbers on failure.

Name

`drm_atomic_helper_connector_dpms` — connector dpms helper implementation

Synopsis

```
int drm_atomic_helper_connector_dpms (struct drm_connector * connector,  
int mode);
```

Arguments

connector affected connector

mode DPMS mode

Description

This is the main helper function provided by the atomic helper framework for implementing the legacy DPMS connector interface. It computes the new desired ->active state for the corresponding CRTC (if the connector is enabled) and updates it.

Returns

Returns 0 on success, negative errno numbers on failure.

Name

`drm_atomic_helper_crtc_reset` — default ->reset hook for CRTC's

Synopsis

```
void drm_atomic_helper_crtc_reset (struct drm_crtc * crtc);
```

Arguments

crtc drm CRTC

Description

Resets the atomic state for *crtc* by freeing the state pointer (which might be NULL, e.g. at driver load time) and allocating a new empty state object.

Name

`__drm_atomic_helper_crtc_duplicate_state` — copy atomic CRTC state

Synopsis

```
void __drm_atomic_helper_crtc_duplicate_state (struct drm_crtc * crtc,  
struct drm_crtc_state * state);
```

Arguments

crtc CRTC object

state atomic CRTC state

Description

Copies atomic state from a CRTC's current state and resets inferred values. This is useful for drivers that subclass the CRTC state.

Name

`drm_atomic_helper_crtc_duplicate_state` — default state duplicate hook

Synopsis

```
struct drm_crtc_state * drm_atomic_helper_crtc_duplicate_state (struct  
drm_crtc * crtc);
```

Arguments

crtc drm CRTC

Description

Default CRTC state duplicate hook for drivers which don't have their own subclassed CRTC state structure.

Name

`__drm_atomic_helper_crtc_destroy_state` — release CRTC state

Synopsis

```
void __drm_atomic_helper_crtc_destroy_state (struct drm_crtc * crtc,  
struct drm_crtc_state * state);
```

Arguments

crtc CRTC object

state CRTC state object to release

Description

Releases all resources stored in the CRTC state without actually freeing the memory of the CRTC state. This is useful for drivers that subclass the CRTC state.

Name

`drm_atomic_helper_crtc_destroy_state` — default state destroy hook

Synopsis

```
void drm_atomic_helper_crtc_destroy_state (struct drm_crtc * crtc,  
struct drm_crtc_state * state);
```

Arguments

crtc drm CRTC

state CRTC state object to release

Description

Default CRTC state destroy hook for drivers which don't have their own subclassed CRTC state structure.

Name

`drm_atomic_helper_plane_reset` — default ->reset hook for planes

Synopsis

```
void drm_atomic_helper_plane_reset (struct drm_plane * plane);
```

Arguments

plane drm plane

Description

Resets the atomic state for *plane* by freeing the state pointer (which might be NULL, e.g. at driver load time) and allocating a new empty state object.

Name

`__drm_atomic_helper_plane_duplicate_state` — copy atomic plane state

Synopsis

```
void __drm_atomic_helper_plane_duplicate_state (struct drm_plane *  
plane, struct drm_plane_state * state);
```

Arguments

plane plane object

state atomic plane state

Description

Copies atomic state from a plane's current state. This is useful for drivers that subclass the plane state.

Name

`drm_atomic_helper_plane_duplicate_state` — default state duplicate hook

Synopsis

```
struct   drm_plane_state   *   drm_atomic_helper_plane_duplicate_state
(struct drm_plane * plane);
```

Arguments

plane drm plane

Description

Default plane state duplicate hook for drivers which don't have their own subclassed plane state structure.

Name

`__drm_atomic_helper_plane_destroy_state` — release plane state

Synopsis

```
void __drm_atomic_helper_plane_destroy_state (struct drm_plane * plane,  
struct drm_plane_state * state);
```

Arguments

plane plane object

state plane state object to release

Description

Releases all resources stored in the plane state without actually freeing the memory of the plane state. This is useful for drivers that subclass the plane state.

Name

`drm_atomic_helper_plane_destroy_state` — default state destroy hook

Synopsis

```
void drm_atomic_helper_plane_destroy_state (struct drm_plane * plane,  
struct drm_plane_state * state);
```

Arguments

plane drm plane

state plane state object to release

Description

Default plane state destroy hook for drivers which don't have their own subclassed plane state structure.

Name

`__drm_atomic_helper_connector_reset` — reset state on connector

Synopsis

```
void __drm_atomic_helper_connector_reset (struct drm_connector * connector, struct drm_connector_state * conn_state);
```

Arguments

connector drm connector

conn_state connector state to assign

Description

Initializes the newly allocated *conn_state* and assigns it to `#connector->state`, usually required when initializing the drivers or when called from the `->reset` hook.

This is useful for drivers that subclass the connector state.

Name

`drm_atomic_helper_connector_reset` — default ->reset hook for connectors

Synopsis

```
void drm_atomic_helper_connector_reset (struct drm_connector * connector);
```

Arguments

connector drm connector

Description

Resets the atomic state for *connector* by freeing the state pointer (which might be NULL, e.g. at driver load time) and allocating a new empty state object.

Name

`__drm_atomic_helper_connector_duplicate_state` — copy atomic connector state

Synopsis

```
void __drm_atomic_helper_connector_duplicate_state (struct drm_connector * connector, struct drm_connector_state * state);
```

Arguments

connector connector object

state atomic connector state

Description

Copies atomic state from a connector's current state. This is useful for drivers that subclass the connector state.

Name

`drm_atomic_helper_connector_duplicate_state` — default state duplicate hook

Synopsis

```
struct drm_connector_state * drm_atomic_helper_connector_duplicate_s-  
tate (struct drm_connector * connector);
```

Arguments

connector drm connector

Description

Default connector state duplicate hook for drivers which don't have their own subclassed connector state structure.

Name

`drm_atomic_helper_duplicate_state` — duplicate an atomic state object

Synopsis

```
struct drm_atomic_state * drm_atomic_helper_duplicate_state (struct dr-  
m_device * dev, struct drm_modeset_acquire_ctx * ctx);
```

Arguments

dev DRM device

ctx lock acquisition context

Description

Makes a copy of the current atomic state by looping over all objects and duplicating their respective states. This is used for example by suspend/ resume support code to save the state prior to suspend such that it can be restored upon resume.

Note that this treats atomic state as persistent between save and restore. Drivers must make sure that this is possible and won't result in confusion or erroneous behaviour.

Note that if callers haven't already acquired all modeset locks this might return `-EDEADLK`, which must be handled by calling `drm_modeset_backoff`.

Returns

A pointer to the copy of the atomic state object on success or an `ERR_PTR`-encoded error code on failure.

See also

`drm_atomic_helper_suspend`, `drm_atomic_helper_resume`

Name

`__drm_atomic_helper_connector_destroy_state` — release connector state

Synopsis

```
void __drm_atomic_helper_connector_destroy_state (struct drm_connector
* connector, struct drm_connector_state * state);
```

Arguments

connector connector object

state connector state object to release

Description

Releases all resources stored in the connector state without actually freeing the memory of the connector state. This is useful for drivers that subclass the connector state.

Name

`drm_atomic_helper_connector_destroy_state` — default state destroy hook

Synopsis

```
void drm_atomic_helper_connector_destroy_state (struct drm_connector *  
connector, struct drm_connector_state * state);
```

Arguments

connector drm connector

state connector state object to release

Description

Default connector state destroy hook for drivers which don't have their own subclassed connector state structure.

Modeset Helper Reference for Common Vtables

Name

struct `drm_crtc_helper_funcs` — helper operations for CRTC's

Synopsis

```
struct drm_crtc_helper_funcs {  
    void (* dpms) (struct drm_crtc *crtc, int mode);  
    void (* prepare) (struct drm_crtc *crtc);  
    void (* commit) (struct drm_crtc *crtc);  
    bool (* mode_fixup) (struct drm_crtc *crtc, const struct drm_display_mode *mode, s  
    int (* mode_set) (struct drm_crtc *crtc, struct drm_display_mode *mode, struct dr  
    void (* mode_set_nofb) (struct drm_crtc *crtc);  
    int (* mode_set_base) (struct drm_crtc *crtc, int x, int y, struct drm_framebuffe  
    int (* mode_set_base_atomic) (struct drm_crtc *crtc, struct drm_framebuffer *fb,  
    void (* load_lut) (struct drm_crtc *crtc);  
    void (* disable) (struct drm_crtc *crtc);  
    void (* enable) (struct drm_crtc *crtc);  
    int (* atomic_check) (struct drm_crtc *crtc, struct drm_crtc_state *state);  
    void (* atomic_begin) (struct drm_crtc *crtc, struct drm_crtc_state *old_crtc_sta  
    void (* atomic_flush) (struct drm_crtc *crtc, struct drm_crtc_state *old_crtc_sta  
};
```

Members

`dpms`

Callback to control power levels on the CRTC. If the mode passed in is unsupported, the provider must use the next lowest power level. This is used by the legacy CRTC helpers to implement DPMS functionality in `drm_helper_connector_dpms`.

This callback is also used to disable a CRTC by calling it with `DRM_MODE_DPMS_OFF` if the *disable* hook isn't used.

This callback is used by the legacy CRTC helpers. Atomic helpers also support using this hook for enabling and disabling a CRTC to facilitate transitions to atomic, but it is deprecated. Instead *enable* and *disable* should be used.

`prepare`

This callback should prepare the CRTC for a subsequent modeset, which in practice means the driver should disable the CRTC if it is running. Most drivers ended up implementing this by calling their *dpms* hook with `DRM_MODE_DPMS_OFF`.

This callback is used by the legacy CRTC helpers. Atomic helpers also support using this hook for disabling a CRTC to facilitate transitions to atomic, but it is deprecated. Instead *disable* should be used.

`commit`

This callback should commit the new mode on the CRTC after a modeset, which in practice means the driver should enable the CRTC. Most drivers ended up implementing this by calling their *dpms* hook with `DRM_MODE_DPMS_ON`.

This callback is used by the legacy CRTC helpers. Atomic helpers also support using this hook for enabling a CRTC to facilitate transitions to atomic, but it is deprecated. Instead *enable* should be used.

mode_fixup

This callback is used to validate a mode. The parameter *mode* is the display mode that userspace requested, *adjusted_mode* is the mode the encoders need to be fed with. Note that this is the inverse semantics of the meaning for the *drm_encoder* and *drm_bridge* `->mode_fixup` functions. If the CRTC cannot support the requested conversion from *mode* to *adjusted_mode* it should reject the modeset.

This function is used by both legacy CRTC helpers and atomic helpers. With atomic helpers it is optional.

NOTE:

This function is called in the check phase of atomic modesets, which can be aborted for any reason (including on userspace's request to just check whether a configuration would be possible). Atomic drivers **MUST NOT** touch any persistent state (hardware or software) or data structures except the passed in *adjusted_mode* parameter.

This is in contrast to the legacy CRTC helpers where this was allowed.

Atomic drivers which need to inspect and adjust more state should instead use the *atomic_check* callback.

Also beware that neither core nor helpers filter modes before passing them to the driver: While the list of modes that is advertised to userspace is filtered using the connector's `->mode_valid` callback, neither the core nor the helpers do any filtering on modes passed in from userspace when setting a mode. It is therefore possible for userspace to pass in a mode that was previously filtered out using `->mode_valid` or add a custom mode that wasn't probed from EDID or similar to begin with. Even though this is an advanced feature and rarely used nowadays, some users rely on being able to specify modes manually so drivers must be prepared to deal with it. Specifically this means that all drivers need not only validate modes in `->mode_valid` but also in `->mode_fixup` to make sure invalid modes passed in from userspace are rejected.

RETURNS:

True if an acceptable configuration is possible, false if the modeset operation should be rejected.

mode_set

This callback is used by the legacy CRTC helpers to set a new mode, position and framebuffer. Since it ties the primary plane to every mode change it is incompatible with universal plane support. And since it can't update other planes it's incompatible with atomic modeset support.

This callback is only used by CRTC helpers and deprecated.

RETURNS:

0 on success or a negative error code on failure.

mode_set_nofb

This callback is used to update the display mode of a CRTC without changing anything of the primary plane configuration. This fits the requirement of atomic and hence is used by the atomic helpers. It is also used by the transitional plane helpers to implement a *mode_set* hook in *drm_helper_crtc_mode_set*.

Note that the display pipe is completely off when this function is called. Atomic drivers which need hardware to be running before they program the new display mode (e.g. because they implement runtime PM) should not use this hook. This is because the helper library calls this hook only once per mode change and not every time the display pipeline is suspended using either DPMS or the new “ACTIVE” property. Which means register values set in this callback might get reset when the CRTC is suspended, but not restored. Such drivers should instead move all their CRTC setup into the *enable* callback.

This callback is optional.

mode_set_base

This callback is used by the legacy CRTC helpers to set a new framebuffer and scanout position. It is optional and used as an optimized fast-path instead of a full mode set operation with all the resulting flickering. If it is not present *drm_crtc_helper_set_config* will fall back to a full modeset, using the *->mode_set* callback. Since it can't update other planes it's incompatible with atomic modeset support.

This callback is only used by the CRTC helpers and deprecated.

RETURNS:

0 on success or a negative error code on failure.

mode_set_base_atomic

This callback is used by the fbdev helpers to set a new framebuffer and scanout without sleeping, i.e. from an atomic calling context. It is only used to implement kgdb support.

This callback is optional and only needed for kgdb support in the fbdev helpers.

RETURNS:

0 on success or a negative error code on failure.

load_lut

Load a LUT prepared with the *gamma_set* functions from *drm_fb_helper_funcs*.

This callback is optional and is only used by the fbdev emulation helpers.

FIXME:

disable

This callback is functionally redundant with the core gamma table support and simply exists because the fbdev hasn't yet been refactored to use the core gamma table interfaces.

This callback should be used to disable the CRTC. With the atomic drivers it is called after all encoders connected to this CRTC have been shut off already using their own `->disable` hook. If that sequence is too simple drivers can just add their own hooks and call it from this CRTC callback here by looping over all encoders connected to it using `for_each_encoder_on_crtc`.

This hook is used both by legacy CRTC helpers and atomic helpers. Atomic drivers don't need to implement it if there's no need to disable anything at the CRTC level. To ensure that runtime PM handling (using either DPMS or the new "ACTIVE" property) works *disable* must be the inverse of *enable* for atomic drivers.

NOTE:

With legacy CRTC helpers there's a big semantic difference between *disable* and other hooks (like *prepare* or *dpms*) used to shut down a CRTC: *disable* is only called when also logically disabling the display pipeline and needs to release any resources acquired in *mode_set* (like shared PLLs, or again release pinned framebuffers).

Therefore *disable* must be the inverse of *mode_set* plus *commit* for drivers still using legacy CRTC helpers, which is different from the rules under atomic.

enable

This callback should be used to enable the CRTC. With the atomic drivers it is called before all encoders connected to this CRTC are enabled through the encoder's own `->enable` hook. If that sequence is too simple drivers can just add their own hooks and call it from this CRTC callback here by looping over all encoders connected to it using `for_each_encoder_on_crtc`.

This hook is used only by atomic helpers, for symmetry with *disable*. Atomic drivers don't need to implement it if there's no need to enable anything at the CRTC level. To ensure that runtime PM handling (using either DPMS or the new "ACTIVE" property) works *enable* must be the inverse of *disable* for atomic drivers.

atomic_check

Drivers should check plane-update related CRTC constraints in this hook. They can also check mode related limitations but need to be aware of the calling order, since this hook is used by `drm_atomic_helper_check_planes` whereas the preparations needed to check output routing and the display mode is done in `drm_atomic_helper_check_modeset`. Therefore drivers that want to check output routing and display mode constraints in this callback must ensure that `drm_atomic_helper_check_modeset` has been called beforehand. This is calling order used by the default helper implementation in `drm_atomic_helper_check`.

When using `drm_atomic_helper_check_planes` CRTC's `>atomic_check` hooks are called after the ones for planes, which allows drivers to assign shared resources requested by planes in the CRTC callback here. For more complicated dependencies the driver can call the provided check helpers multiple times until the computed state has a final configuration and everything has been checked.

This function is also allowed to inspect any other object's state and can add more state objects to the atomic commit if needed. Care must be taken though to ensure that state checkcompute functions for these added states are all called, and derived state in other objects all updated. Again the recommendation is to just call check helpers until a maximal configuration is reached.

This callback is used by the atomic modeset helpers and by the transitional plane helpers, but it is optional.

NOTE:

This function is called in the check phase of an atomic update. The driver is not allowed to change anything outside of the free-standing state objects passed-in or assembled in the overall `drm_atomic_state` update tracking structure.

RETURNS:

0 on success, `-EINVAL` if the state or the transition can't be supported, `-ENOMEM` on memory allocation failure and `-EDEADLK` if an attempt to obtain another state object ran into a `drm_modeset_lock` deadlock.

`atomic_begin`

Drivers should prepare for an atomic update of multiple planes on a CRTC in this hook. Depending upon hardware this might be vblank evasion, blocking updates by setting bits or doing preparatory work for e.g. manual update display.

This hook is called before any plane commit functions are called.

Note that the power state of the display pipe when this function is called depends upon the exact helpers and calling sequence the driver has picked. See `drm_atomic_commit_planes` for a discussion of the tradeoffs and variants of plane commit helpers.

This callback is used by the atomic modeset helpers and by the transitional plane helpers, but it is optional.

`atomic_flush`

Drivers should finalize an atomic update of multiple planes on a CRTC in this hook. Depending upon hardware this might include checking that vblank evasion was successful, unblocking updates by setting bits or setting the GO bit to flush out all updates.

Simple hardware or hardware with special requirements can commit and flush out all updates for all planes from this hook and forgo all the other commit hooks for plane updates.

This hook is called after any plane commit functions are called.

Note that the power state of the display pipe when this function is called depends upon the exact helpers and calling sequence the driver has picked. See `drm_atomic_commit_planes` for a discussion of the tradeoffs and variants of plane commit helpers.

This callback is used by the atomic modeset helpers and by the transitional plane helpers, but it is optional.

overview

These hooks are used by the legacy CRTC helpers, the transitional plane helpers and the new atomic modesetting helpers.

Name

`drm_crtc_helper_add` — sets the helper vtable for a `crtc`

Synopsis

```
void drm_crtc_helper_add (struct drm_crtc * crtc, const struct dr-  
m_crtc_helper_funcs * funcs);
```

Arguments

crtc DRM CRTC

funcs helper vtable to set for *crtc*

Name

struct drm_encoder_helper_funcs — helper operations for encoders

Synopsis

```
struct drm_encoder_helper_funcs {
    void (* dpms) (struct drm_encoder *encoder, int mode);
    bool (* mode_fixup) (struct drm_encoder *encoder, const struct drm_display_mode *
    void (* prepare) (struct drm_encoder *encoder);
    void (* commit) (struct drm_encoder *encoder);
    void (* mode_set) (struct drm_encoder *encoder, struct drm_display_mode *mode, str
    struct drm_crtc *(* get_crtc) (struct drm_encoder *encoder);
    enum drm_connector_status (* detect) (struct drm_encoder *encoder, struct drm_con
    void (* disable) (struct drm_encoder *encoder);
    void (* enable) (struct drm_encoder *encoder);
    int (* atomic_check) (struct drm_encoder *encoder, struct drm_crtc_state *crtc_st
};
```

Members

dpms

Callback to control power levels on the encoder. If the mode passed in is unsupported, the provider must use the next lowest power level. This is used by the legacy encoder helpers to implement DPMS functionality in `drm_helper_connector_dpms`.

This callback is also used to disable an encoder by calling it with `DRM_MODE_DPMS_OFF` if the *disable* hook isn't used.

This callback is used by the legacy CRTC helpers. Atomic helpers also support using this hook for enabling and disabling an encoder to facilitate transitions to atomic, but it is deprecated. Instead *enable* and *disable* should be used.

mode_fixup

This callback is used to validate and adjust a mode. The parameter *mode* is the display mode that should be fed to the next element in the display chain, either the final `drm_connector` or a `drm_bridge`. The parameter *adjusted_mode* is the input mode the encoder requires. It can be modified by this callback and does not need to match *mode*.

This function is used by both legacy CRTC helpers and atomic helpers. With atomic helpers it is optional.

NOTE:

This function is called in the check phase of atomic modesets, which can be aborted for any reason (including on userspace's request to just check whether a configuration would be possible). Atomic drivers **MUST NOT** touch any persistent state (hardware or software) or data structures except the passed in *adjusted_mode* parameter.

This is in contrast to the legacy CRTC helpers where this was allowed.

Atomic drivers which need to inspect and adjust more state should instead use the *atomic_check* callback.

Also beware that neither core nor helpers filter modes before passing them to the driver: While the list of modes that is advertised to userspace is filtered using the con-

nector's `->mode_valid` callback, neither the core nor the helpers do any filtering on modes passed in from userspace when setting a mode. It is therefore possible for userspace to pass in a mode that was previously filtered out using `->mode_valid` or add a custom mode that wasn't probed from EDID or similar to begin with. Even though this is an advanced feature and rarely used nowadays, some users rely on being able to specify modes manually so drivers must be prepared to deal with it. Specifically this means that all drivers need not only validate modes in `->mode_valid` but also in `->mode_fixup` to make sure invalid modes passed in from userspace are rejected.

RETURNS:

True if an acceptable configuration is possible, false if the modeset operation should be rejected.

prepare

This callback should prepare the encoder for a subsequent modeset, which in practice means the driver should disable the encoder if it is running. Most drivers ended up implementing this by calling their `dpms` hook with `DRM_MODE_DPMS_OFF`.

This callback is used by the legacy CRTC helpers. Atomic helpers also support using this hook for disabling an encoder to facilitate transitions to atomic, but it is deprecated. Instead `disable` should be used.

commit

This callback should commit the new mode on the encoder after a modeset, which in practice means the driver should enable the encoder. Most drivers ended up implementing this by calling their `dpms` hook with `DRM_MODE_DPMS_ON`.

This callback is used by the legacy CRTC helpers. Atomic helpers also support using this hook for enabling an encoder to facilitate transitions to atomic, but it is deprecated. Instead `enable` should be used.

mode_set

This callback is used to update the display mode of an encoder.

Note that the display pipe is completely off when this function is called. Drivers which need hardware to be running before they program the new display mode (because they implement runtime PM) should not use this hook, because the helper library calls it only once and not every time the display pipeline is suspend using either DPMS or the new "ACTIVE" property. Such drivers should instead move all their encoder setup into the `->enable` callback.

This callback is used both by the legacy CRTC helpers and the atomic modeset helpers. It is optional in the atomic helpers.

get_crtc

This callback is used by the legacy CRTC helpers to work around deficiencies in its own book-keeping.

Do not use, use atomic helpers instead, which get the book keeping right.

FIXME:

Currently only nouveau is using this, and as soon as nouveau is atomic we can ditch this hook.

detect

This callback can be used by drivers who want to do detection on the encoder object instead of in connector functions.

It is not used by any helper and therefore has purely driver-specific semantics. New drivers shouldn't use this and instead just implement their own private callbacks.

FIXME:

This should just be converted into a pile of driver vfuncs. Currently radeon, amdgpu and nouveau are using it.

disable

This callback should be used to disable the encoder. With the atomic drivers it is called before this encoder's CRTC has been shut off using the CRTC's own `->disable` hook. If that sequence is too simple drivers can just add their own driver private encoder hooks and call them from CRTC's callback by looping over all encoders connected to it using `for_each_encoder_on_crtc`.

This hook is used both by legacy CRTC helpers and atomic helpers. Atomic drivers don't need to implement it if there's no need to disable anything at the encoder level. To ensure that runtime PM handling (using either DPMS or the new "ACTIVE" property) works *disable* must be the inverse of *enable* for atomic drivers.

NOTE:

With legacy CRTC helpers there's a big semantic difference between *disable* and other hooks (like *prepare* or *dpms*) used to shut down a encoder: *disable* is only called when also logically disabling the display pipeline and needs to release any resources acquired in *mode_set* (like shared PLLs, or again release pinned frame-buffers).

Therefore *disable* must be the inverse of *mode_set* plus *commit* for drivers still using legacy CRTC helpers, which is different from the rules under atomic.

enable

This callback should be used to enable the encoder. With the atomic drivers it is called after this encoder's CRTC has been enabled using the CRTC's own `->enable` hook. If that sequence is too simple drivers can just add their own driver private encoder hooks and call them from CRTC's callback by looping over all encoders connected to it using `for_each_encoder_on_crtc`.

This hook is used only by atomic helpers, for symmetry with *disable*. Atomic drivers don't need to implement it if there's no need to enable anything at the encoder level. To ensure that runtime PM handling (using either DPMS or the new "ACTIVE" property) works *enable* must be the inverse of *disable* for atomic drivers.

atomic_check

This callback is used to validate encoder state for atomic drivers. Since the encoder is the object connecting the CRTC and connector it gets passed both states, to be able to validate interactions and update the CRTC to match what the encoder needs for the requested connector.

This function is used by the atomic helpers, but it is optional.

NOTE:

This function is called in the check phase of an atomic update. The driver is not allowed to change anything outside of the free-standing state objects passed-in or assembled in the overall `drm_atomic_state` update tracking structure.

RETURNS:

0 on success, `-EINVAL` if the state or the transition can't be supported, `-ENOMEM` on memory allocation failure and `-EDEADLK` if an attempt to obtain another state object ran into a `drm_modeset_lock` deadlock.

Description

These hooks are used by the legacy CRTC helpers, the transitional plane helpers and the new atomic modesetting helpers.

Name

`drm_encoder_helper_add` — sets the helper vtable for an encoder

Synopsis

```
void drm_encoder_helper_add (struct drm_encoder * encoder, const struct  
drm_encoder_helper_funcs * funcs);
```

Arguments

encoder DRM encoder

funcs helper vtable to set for *encoder*

Name

struct `drm_connector_helper_funcs` — helper operations for connectors

Synopsis

```
struct drm_connector_helper_funcs {
    int (* get_modes) (struct drm_connector *connector);
    enum drm_mode_status (* mode_valid) (struct drm_connector *connector, struct drm_
    struct drm_encoder *(* best_encoder) (struct drm_connector *connector);
    struct drm_encoder *(* atomic_best_encoder) (struct drm_connector *connector, str
};
```

Members

`get_modes`

This function should fill in all modes currently valid for the sink into the `connector->probed_modes` list. It should also update the EDID property by calling `drm_mode_connector_update_edid_property`.

The usual way to implement this is to cache the EDID retrieved in the probe callback somewhere in the driver-private connector structure. In this function drivers then parse the modes in the EDID and add them by calling `drm_add_edid_modes`. But connectors that driver a fixed panel can also manually add specific modes using `drm_mode_probed_add`. Drivers which manually add modes should also make sure that the *display_info*, *width_mm* and *height_mm* fields of the struct `#drm_connector` are filled in.

Virtual drivers that just want some standard VESA mode with a given resolution can call `drm_add_modes_noedid`, and mark the preferred one using `drm_set_preferred_mode`.

Finally drivers that support audio probably want to update the ELD data, too, using `drm_edid_to_eld`.

This function is only called after the `->detect` hook has indicated that a sink is connected and when the EDID isn't overridden through sysfs or the kernel commandline.

This callback is used by the probe helpers in e.g. `drm_helper_probe_single_connector_modes`.

RETURNS:

The number of modes added by calling `drm_mode_probed_add`.

`mode_valid`

Callback to validate a mode for a connector, irrespective of the specific display configuration.

This callback is used by the probe helpers to filter the mode list (which is usually derived from the EDID data block from the sink). See e.g. `drm_helper_probe_single_connector_modes`.

NOTE:

This only filters the mode list supplied to userspace in the GETCON-NECOTR IOCTL. Userspace is free to create modes of its own and ask the kernel to use them. In this case the atomic helpers or legacy CRTC helpers will not call this function. Drivers therefore must still fully validate any mode passed in in a modeset request.

RETURNS:

Either `MODE_OK` or one of the failure reasons in `enum drm_mode_status`.

`best_encoder`

This function should select the best encoder for the given connector.

This function is used by both the atomic helpers (in the `drm_atomic_helper_check_modeset` function) and in the legacy CRTC helpers.

NOTE:

In atomic drivers this function is called in the check phase of an atomic update. The driver is not allowed to change or inspect anything outside of arguments passed-in. Atomic drivers which need to inspect dynamic configuration state should instead use *atomic_best_encoder*.

RETURNS:

Encoder that should be used for the given connector and connector state, or `NULL` if no suitable encoder exists. Note that the helpers will ensure that encoders aren't used twice, drivers should not check for this.

`atomic_best_encoder`

This is the atomic version of *best_encoder* for atomic drivers which need to select the best encoder depending upon the desired configuration and can't select it statically.

This function is used by `drm_atomic_helper_check_modeset` and either this or *best_encoder* is required.

NOTE:

This function is called in the check phase of an atomic update. The driver is not allowed to change anything outside of the free-standing state objects passed-in or assembled in the overall `drm_atomic_state` update tracking structure.

RETURNS:

Encoder that should be used for the given connector and connector state, or `NULL` if no suitable encoder exists. Note that the helpers will ensure that encoders aren't used twice, drivers should not check for this.

Description

These functions are used by the atomic and legacy modeset helpers and by the probe helpers.

Name

`drm_connector_helper_add` — sets the helper vtable for a connector

Synopsis

```
void drm_connector_helper_add (struct drm_connector * connector, const  
struct drm_connector_helper_funcs * funcs);
```

Arguments

connector DRM connector

funcs helper vtable to set for *connector*

Name

struct `drm_plane_helper_funcs` — helper operations for planes

Synopsis

```
struct drm_plane_helper_funcs {  
    int (* prepare_fb) (struct drm_plane *plane, const struct drm_plane_state *new_state,  
                        void (* cleanup_fb) (struct drm_plane *plane, const struct drm_plane_state *old_state));  
    int (* atomic_check) (struct drm_plane *plane, struct drm_plane_state *state);  
    void (* atomic_update) (struct drm_plane *plane, struct drm_plane_state *old_state, struct drm_plane_state *new_state);  
    void (* atomic_disable) (struct drm_plane *plane, struct drm_plane_state *old_state);  
};
```

Members

`prepare_fb`

This hook is to prepare a framebuffer for scanout by e.g. pinning it's backing storage or relocating it into a contiguous block of VRAM. Other possible preparatory work includes flushing caches.

This function must not block for outstanding rendering, since it is called in the context of the atomic IOCTL even for async commits to be able to return any errors to userspace. Instead the recommended way is to fill out the fence member of the passed-in `drm_plane_state`. If the driver doesn't support native fences then equivalent functionality should be implemented through private members in the plane structure.

The helpers will call `cleanup_fb` with matching arguments for every successful call to this hook.

This callback is used by the atomic modeset helpers and by the transitional plane helpers, but it is optional.

RETURNS:

0 on success or one of the following negative error codes allowed by the `atomic_commit` hook in `drm_mode_config_funcs`. When using helpers this callback is the only one which can fail an atomic commit, everything else must complete successfully.

`cleanup_fb`

This hook is called to clean up any resources allocated for the given framebuffer and plane configuration in `prepare_fb`.

This callback is used by the atomic modeset helpers and by the transitional plane helpers, but it is optional.

`atomic_check`

Drivers should check plane specific constraints in this hook.

When using `drm_atomic_helper_check_planes` plane's `->atomic_check` hooks are called before the ones for CRTC's, which allows drivers to request shared resources that the CRTC controls here. For more complicated dependencies the driver can call the provided check helpers multiple times until the computed state has a final configuration and everything has been checked.

This function is also allowed to inspect any other object's state and can add more state objects to the atomic commit if needed. Care must be taken though to ensure that state checkcompute functions for these added states are all called, and derived state in other objects all updated. Again the recommendation is to just call check helpers until a maximal configuration is reached.

This callback is used by the atomic modeset helpers and by the transitional plane helpers, but it is optional.

NOTE:

This function is called in the check phase of an atomic update. The driver is not allowed to change anything outside of the free-standing state objects passed-in or assembled in the overall `drm_atomic_state` update tracking structure.

RETURNS:

0 on success, `-EINVAL` if the state or the transition can't be supported, `-ENOMEM` on memory allocation failure and `-EDEADLK` if an attempt to obtain another state object ran into a `drm_modeset_lock` deadlock.

`atomic_update`

Drivers should use this function to update the plane state. This hook is called in-between the `->atomic_begin` and `->atomic_flush` of `drm_crtc_helper_funcs`.

Note that the power state of the display pipe when this function is called depends upon the exact helpers and calling sequence the driver has picked. See `drm_atomic_commit_planes` for a discussion of the tradeoffs and variants of plane commit helpers.

This callback is used by the atomic modeset helpers and by the transitional plane helpers, but it is optional.

`atomic_disable`

Drivers should use this function to unconditionally disable a plane. This hook is called in-between the `->atomic_begin` and `->atomic_flush` of `drm_crtc_helper_funcs`. It is an alternative to `atomic_update`, which will be called for disabling planes, too, if the `atomic_disable` hook isn't implemented.

This hook is also useful to disable planes in preparation of a modeset, by calling `drm_atomic_helper_disable_planes_on_crtc` from the `->disable` hook in `drm_crtc_helper_funcs`.

Note that the power state of the display pipe when this function is called depends upon the exact helpers and calling sequence the driver has picked. See `drm_atomic_commit_planes` for a discussion of the tradeoffs and variants of plane commit helpers.

This callback is used by the atomic modeset helpers and by the transitional plane helpers, but it is optional.

Description

These functions are used by the atomic helpers and by the transitional plane helpers.

Name

`drm_plane_helper_add` — sets the helper vtable for a plane

Synopsis

```
void drm_plane_helper_add (struct drm_plane * plane, const struct dr-  
m_plane_helper_funcs * funcs);
```

Arguments

plane DRM plane

funcs helper vtable to set for *plane*

The DRM mode setting helper functions are common code for drivers to use if they wish. Drivers are not forced to use this code in their implementations but it would be useful if the code they do use at least provides a consistent interface and operation to userspace. Therefore it is highly recommended to use the provided helpers as much as possible.

Because there is only one pointer per modeset object to hold a vfunc table for helper libraries they are by necessity shared among the different helpers.

To make this clear all the helper vtables are pulled together in this location here.

Legacy CRTC/Modeset Helper Functions Reference

Name

`drm_helper_move_panel_connectors_to_head` — move panels to the front in the connector list

Synopsis

```
void drm_helper_move_panel_connectors_to_head (struct drm_device *  
dev);
```

Arguments

dev drm device to operate on

Description

Some userspace presumes that the first connected connector is the main display, where it's supposed to display e.g. the login screen. For laptops, this should be the main panel. Use this function to sort all (eDP/LVDS) panels to the front of the connector list, instead of painstakingly trying to initialize them in the right order.

Name

`drm_helper_encoder_in_use` — check if a given encoder is in use

Synopsis

```
bool drm_helper_encoder_in_use (struct drm_encoder * encoder);
```

Arguments

encoder encoder to check

Description

Checks whether *encoder* is with the current mode setting output configuration in use by any connector. This doesn't mean that it is actually enabled since the DPMS state is tracked separately.

Returns

True if *encoder* is used, false otherwise.

Name

`drm_helper_crtc_in_use` — check if a given CRTC is in a mode_config

Synopsis

```
bool drm_helper_crtc_in_use (struct drm_crtc * crtc);
```

Arguments

crtc CRTC to check

Description

Checks whether *crtc* is with the current mode setting output configuration in use by any connector. This doesn't mean that it is actually enabled since the DPMS state is tracked separately.

Returns

True if *crtc* is used, false otherwise.

Name

`drm_helper_disable_unused_functions` — disable unused objects

Synopsis

```
void drm_helper_disable_unused_functions (struct drm_device * dev);
```

Arguments

dev DRM device

Description

This function walks through the entire mode setting configuration of *dev*. It will remove any CRTC links of unused encoders and encoder links of disconnected connectors. Then it will disable all unused encoders and CRTCs either by calling their disable callback if available or by calling their dpms callback with `DRM_MODE_DPMS_OFF`.

Name

`drm_crtc_helper_set_mode` — internal helper to set a mode

Synopsis

```
bool drm_crtc_helper_set_mode (struct drm_crtc * crtc, struct drm_display_mode * mode, int x, int y, struct drm_framebuffer * old_fb);
```

Arguments

<i>crtc</i>	CRTC to program
<i>mode</i>	mode to use
<i>x</i>	horizontal offset into the surface
<i>y</i>	vertical offset into the surface
<i>old_fb</i>	old framebuffer, for cleanup

Description

Try to set *mode* on *crtc*. Give *crtc* and its associated connectors a chance to fixup or reject the mode prior to trying to set it. This is an internal helper that drivers could e.g. use to update properties that require the entire output pipe to be disabled and re-enabled in a new configuration. For example for changing whether audio is enabled on a hdmi link or for changing panel fitter or dither attributes. It is also called by the `drm_crtc_helper_set_config` helper function to drive the mode setting sequence.

Returns

True if the mode was set successfully, false otherwise.

Name

`drm_crtc_helper_set_config` — set a new config from userspace

Synopsis

```
int drm_crtc_helper_set_config (struct drm_mode_set * set);
```

Arguments

set mode set configuration

Description

The `drm_crtc_helper_set_config` helper function implements the `set_config` callback of struct `drm_crtc_funcs` for drivers using the legacy CRTC helpers.

It first tries to locate the best encoder for each connector by calling the `connector->best_encoder` (struct `drm_connector_helper_funcs`) helper operation.

After locating the appropriate encoders, the helper function will call the `mode_fixup` encoder and CRTC helper operations to adjust the requested mode, or reject it completely in which case an error will be returned to the application. If the new configuration after mode adjustment is identical to the current configuration the helper function will return without performing any other operation.

If the adjusted mode is identical to the current mode but changes to the frame buffer need to be applied, the `drm_crtc_helper_set_config` function will call the `CRTC->mode_set_base` (struct `drm_crtc_helper_funcs`) helper operation.

If the adjusted mode differs from the current mode, or if the `->mode_set_base` helper operation is not provided, the helper function performs a full mode set sequence by calling the `->prepare`, `->mode_set` and `->commit` CRTC and encoder helper operations, in that order. Alternatively it can also use the `dpms` and `disable` helper operations. For details see struct `drm_crtc_helper_funcs` and struct `drm_encoder_helper_funcs`.

This function is deprecated. New drivers must implement atomic modeset support, for which this function is unsuitable. Instead drivers should use `drm_atomic_helper_set_config`.

Returns

Returns 0 on success, negative `errno` numbers on failure.

Name

`drm_helper_connector_dpms` — connector dpms helper implementation

Synopsis

```
int drm_helper_connector_dpms (struct drm_connector * connector, int
mode);
```

Arguments

connector affected connector

mode DPMS mode

Description

The `drm_helper_connector_dpms` helper function implements the `->dpms` callback of struct `drm_connector_funcs` for drivers using the legacy CRTC helpers.

This is the main helper function provided by the CRTC helper framework for implementing the DPMS connector attribute. It computes the new desired DPMS state for all encoders and CRTCs in the output mesh and calls the `->dpms` callbacks provided by the driver in struct `drm_crtc_helper_funcs` and struct `drm_encoder_helper_funcs` appropriately.

This function is deprecated. New drivers must implement atomic modeset support, for which this function is unsuitable. Instead drivers should use `drm_atomic_helper_connector_dpms`.

Returns

Always returns 0.

Name

`drm_helper_mode_fill_fb_struct` — fill out framebuffer metadata

Synopsis

```
void drm_helper_mode_fill_fb_struct (struct drm_framebuffer * fb, const  
struct drm_mode_fb_cmd2 * mode_cmd);
```

Arguments

fb `drm_framebuffer` object to fill out

mode_cmd metadata from the userspace fb creation request

Description

This helper can be used in a drivers `fb_create` callback to pre-fill the fb's metadata fields.

Name

`drm_helper_resume_force_mode` — force-restore mode setting configuration

Synopsis

```
void drm_helper_resume_force_mode (struct drm_device * dev);
```

Arguments

dev `drm_device` which should be restored

Description

Drivers which use the mode setting helpers can use this function to force-restore the mode setting configuration e.g. on resume or when something else might have trampled over the hw state (like some overzealous old BIOSen tended to do).

This helper doesn't provide a error return value since restoring the old config should never fail due to resource allocation issues since the driver has successfully set the restored configuration already. Hence this should boil down to the equivalent of a few `dpms` on calls, which also don't provide an error code.

Drivers where simply restoring an old configuration again might fail (e.g. due to slight differences in allocating shared resources when the configuration is restored in a different order than when userspace set it up) need to use their own restore logic.

This function is deprecated. New drivers should implement atomic mode- setting and use the atomic suspend/resume helpers.

See also

`drm_atomic_helper_suspend`, `drm_atomic_helper_resume`

Name

`drm_helper_crtc_mode_set` — `mode_set` implementation for atomic plane helpers

Synopsis

```
int drm_helper_crtc_mode_set (struct drm_crtc * crtc, struct drm_display_mode * mode, struct drm_display_mode * adjusted_mode, int x, int y, struct drm_framebuffer * old_fb);
```

Arguments

<i>crtc</i>	DRM CRTC
<i>mode</i>	DRM display mode which userspace requested
<i>adjusted_mode</i>	DRM display mode adjusted by <code>->mode_fixup</code> callbacks
<i>x</i>	x offset of the CRTC scanout area on the underlying framebuffer
<i>y</i>	y offset of the CRTC scanout area on the underlying framebuffer
<i>old_fb</i>	previous framebuffer

Description

This function implements a callback useable as the `->mode_set` callback required by the CRTC helpers. Besides the atomic plane helper functions for the primary plane the driver must also provide the `->mode_set_nofb` callback to set up the CRTC.

This is a transitional helper useful for converting drivers to the atomic interfaces.

Name

`drm_helper_crtc_mode_set_base` — `mode_set_base` implementation for atomic plane helpers

Synopsis

```
int drm_helper_crtc_mode_set_base (struct drm_crtc * crtc, int x, int y, struct drm_framebuffer * old_fb);
```

Arguments

<i>crtc</i>	DRM CRTC
<i>x</i>	x offset of the CRTC scanout area on the underlying framebuffer
<i>y</i>	y offset of the CRTC scanout area on the underlying framebuffer
<i>old_fb</i>	previous framebuffer

Description

This function implements a callback useable as the `->mode_set_base` used required by the CRTC helpers. The driver must provide the atomic plane helper functions for the primary plane.

This is a transitional helper useful for converting drivers to the atomic interfaces.

The CRTC modeset helper library provides a default `set_config` implementation in `drm_crtc_helper_set_config`. Plus a few other convenience functions using the same callbacks which drivers can use to e.g. restore the modeset configuration on resume with `drm_helper_resume_force_mode`.

Note that this helper library doesn't track the current power state of CRTCs and encoders. It can call callbacks like `->dpms` even though the hardware is already in the desired state. This deficiency has been fixed in the atomic helpers.

The driver callbacks are mostly compatible with the atomic modeset helpers, except for the handling of the primary plane: Atomic helpers require that the primary plane is implemented as a real standalone plane and not directly tied to the CRTC state. For easier transition this library provides functions to implement the old semantics required by the CRTC helpers using the new plane and atomic helper callbacks.

Drivers are strongly urged to convert to the atomic helpers (by way of first converting to the plane helpers). New drivers must not use these functions but need to implement the atomic interface instead, potentially using the atomic helpers for that.

These legacy modeset helpers use the same function table structures as all other modesetting helpers. See the documentation for `struct drm_crtc_helper_funcs`, `struct drm_encoder_helper_funcs` and `struct drm_connector_helper_funcs`.

Output Probing Helper Functions Reference

This library provides some helper code for output probing. It provides an implementation of the core `connector->fill_modes` interface with `drm_helper_probe_single_connector_modes`.

It also provides support for polling connectors with a work item and for generic hotplug interrupt handling where the driver doesn't or cannot keep track of a per-connector hpd interrupt.

This helper library can be used independently of the modeset helper library. Drivers can also overwrite different parts e.g. use their own hotplug handling code to avoid probing unrelated outputs.

The probe helpers share the function table structures with other display helper libraries. See struct `drm_connector_helper_funcs` for the details.

Name

`drm_kms_helper_poll_enable_locked` — re-enable output polling.

Synopsis

```
void drm_kms_helper_poll_enable_locked (struct drm_device * dev);
```

Arguments

dev `drm_device`

Description

This function re-enables the output polling work without locking the `mode_config` mutex.

This is like `drm_kms_helper_poll_enable` however it is to be called from a context where the `mode_config` mutex is locked already.

Name

`drm_helper_probe_single_connector_modes` — get complete set of display modes

Synopsis

```
int drm_helper_probe_single_connector_modes (struct drm_connector *  
connector, uint32_t maxX, uint32_t maxY);
```

Arguments

<i>connector</i>	connector to probe
<i>maxX</i>	max width for modes
<i>maxY</i>	max height for modes

Description

Based on the helper callbacks implemented by *connector* in struct `drm_connector_helper_funcs` try to detect all valid modes. Modes will first be added to the connector's `probed_modes` list, then culled (based on validity and the *maxX*, *maxY* parameters) and put into the normal modes list.

Intended to be used as a generic implementation of the `->fill_modes connector` vfunc for drivers that use the CRTC helpers for output mode filtering and detection.

The basic procedure is as follows

1. All modes currently on the connector's modes list are marked as stale
2. New modes are added to the connector's `probed_modes` list with `drm_mode_probed_add`. New modes start their life with status as OK. Modes are added from a single source using the following priority order.

- debugfs 'override_edid' (used for testing only) - firmware EDID (`drm_load_edid_firmware`) - connector helper `->get_modes` vfunc - if the connector status is `connector_status_connected`, standard VESA DMT modes up to 1024x768 are automatically added (`drm_add_modes_noedid`)

Finally modes specified via the kernel command line (`video=...`) are added in addition to what the earlier probes produced (`drm_helper_probe_add_cmdline_mode`). These modes are generated using the VESA GTF/CVT formulas.

3. Modes are moved from the `probed_modes` list to the modes list. Potential duplicates are merged together (see `drm_mode_connector_list_update`). After this step the `probed_modes` list will be empty again.

4. Any non-stale mode on the modes list then undergoes validation

- `drm_mode_validate_basic` performs basic sanity checks - `drm_mode_validate_size` filters out modes larger than *maxX* and *maxY* (if specified) - `drm_mode_validate_flag` checks the modes against basic connector capabilities (`interlace_allowed`, `doublescan_allowed`, `stereo_allowed`) - the optional connector `->mode_valid` helper can perform driver and/or hardware specific checks

5. Any mode whose status is not OK is pruned from the connector's modes list, accompanied by a debug message indicating the reason for the mode's rejection (see `drm_mode_prune_invalid`).

Returns

The number of modes found on *connector*.

Name

`drm_kms_helper_hotplug_event` — fire off KMS hotplug events

Synopsis

```
void drm_kms_helper_hotplug_event (struct drm_device * dev);
```

Arguments

dev `drm_device` whose connector state changed

Description

This function fires off the uevent for userspace and also calls the `output_poll_changed` function, which is most commonly used to inform the fbdev emulation code and allow it to update the fbcon output configuration.

Drivers should call this from their hotplug handling code when a change is detected. Note that this function does not do any output detection of its own, like `drm_helper_hpd_irq_event` does - this is assumed to be done by the driver already.

This function must be called from process context with no mode setting locks held.

Name

`drm_kms_helper_poll_disable` — disable output polling

Synopsis

```
void drm_kms_helper_poll_disable (struct drm_device * dev);
```

Arguments

dev `drm_device`

Description

This function disables the output polling work.

Drivers can call this helper from their device suspend implementation. It is not an error to call this even when output polling isn't enabled or already disabled.

Name

`drm_kms_helper_poll_enable` — re-enable output polling.

Synopsis

```
void drm_kms_helper_poll_enable (struct drm_device * dev);
```

Arguments

dev `drm_device`

Description

This function re-enables the output polling work.

Drivers can call this helper from their device resume implementation. It is an error to call this when the output polling support has not yet been set up.

Name

`drm_kms_helper_poll_init` — initialize and enable output polling

Synopsis

```
void drm_kms_helper_poll_init (struct drm_device * dev);
```

Arguments

dev `drm_device`

Description

This function initializes and then also enables output polling support for *dev*. Drivers which do not have reliable hotplug support in hardware can use this helper infrastructure to regularly poll such connectors for changes in their connection state.

Drivers can control which connectors are polled by setting the `DRM_CONNECTOR_POLL_CONNECT` and `DRM_CONNECTOR_POLL_DISCONNECT` flags. On connectors where probing live outputs can result in visual distortion drivers should not set the `DRM_CONNECTOR_POLL_DISCONNECT` flag to avoid this. Connectors which have no flag or only `DRM_CONNECTOR_POLL_HPD` set are completely ignored by the polling logic.

Note that a connector can be both polled and probed from the hotplug handler, in case the hotplug interrupt is known to be unreliable.

Name

`drm_kms_helper_poll_fini` — disable output polling and clean it up

Synopsis

```
void drm_kms_helper_poll_fini (struct drm_device * dev);
```

Arguments

dev `drm_device`

Name

`drm_helper_hpd_irq_event` — hotplug processing

Synopsis

```
bool drm_helper_hpd_irq_event (struct drm_device * dev);
```

Arguments

dev `drm_device`

Description

Drivers can use this helper function to run a detect cycle on all connectors which have the `DRM_CONNECTOR_POLL_HPD` flag set in their polled member. All other connectors are ignored, which is useful to avoid reprobing fixed panels.

This helper function is useful for drivers which can't or don't track hotplug interrupts for each connector.

Drivers which support hotplug interrupts for each connector individually and which have a more fine-grained detect logic should bypass this code and directly call `drm_kms_helper_hotplug_event` in case the connector state changed.

This function must be called from process context with no mode setting locks held.

Note that a connector can be both polled and probed from the hotplug handler, in case the hotplug interrupt is known to be unreliable.

fbdev Helper Functions Reference

The fb helper functions are useful to provide an fbdev on top of a drm kernel mode setting driver. They can be used mostly independently from the crtc helper functions used by many drivers to implement the kernel mode setting interfaces.

Initialization is done as a four-step process with `drm_fb_helper_prepare`, `drm_fb_helper_init`, `drm_fb_helper_single_add_all_connectors` and `drm_fb_helper_initial_config`. Drivers with fancier requirements than the default behaviour can override the third step with their own code. Teardown is done with `drm_fb_helper_fini`.

At runtime drivers should restore the fbdev console by calling `drm_fb_helper_restore_fbdev_mode_unlocked` from their `->lastclose` callback. They should also notify the fb helper code from updates to the output configuration by calling `drm_fb_helper_hotplug_event`. For easier integration with the output polling code in `drm_crtc_helper.c` the modeset code provides a `->output_poll_changed` callback.

All other functions exported by the fb helper library can be used to implement the fbdev driver interface by the driver.

It is possible, though perhaps somewhat tricky, to implement race-free hotplug detection using the fbdev helpers. The `drm_fb_helper_prepare` helper must be called first to initialize the minimum required to make hotplug detection work. Drivers also need to make sure to properly set up the `dev->mode_config.funcs` member. After calling `drm_kms_helper_poll_init` it is safe to enable interrupts and start processing hotplug events. At the same time, drivers should initialize all modeset objects such as CRTC's, encoders and connectors. To finish up the fbdev helper initialization, the `drm_fb_helper_init` func-

tion is called. To probe for all attached displays and set up an initial configuration using the detected hardware, drivers should call `drm_fb_helper_single_add_all_connectors` followed by `drm_fb_helper_initial_config`.

Name

`drm_fb_helper_single_add_all_connectors` — add all connectors to fbdev emulation helper

Synopsis

```
int drm_fb_helper_single_add_all_connectors (struct drm_fb_helper *  
fb_helper);
```

Arguments

fb_helper fbdev initialized with `drm_fb_helper_init`

Description

This functions adds all the available connectors for use with the given `fb_helper`. This is a separate step to allow drivers to freely assign connectors to the fbdev, e.g. if some are reserved for special purposes or not adequate to be used for the fbcon.

This function is protected against concurrent connector hotadds/removals using `drm_fb_helper_add_one_connector` and `drm_fb_helper_remove_one_connector`.

Name

`drm_fb_helper_debug_enter` — implementation for `->fb_debug_enter`

Synopsis

```
int drm_fb_helper_debug_enter (struct fb_info * info);
```

Arguments

info fbdev registered by the helper

Name

`drm_fb_helper_debug_leave` — implementation for `->fb_debug_leave`

Synopsis

```
int drm_fb_helper_debug_leave (struct fb_info * info);
```

Arguments

info fbdev registered by the helper

Name

`drm_fb_helper_restore_fbdev_mode_unlocked` — restore fbdev configuration

Synopsis

```
int drm_fb_helper_restore_fbdev_mode_unlocked (struct drm_fb_helper *  
fb_helper);
```

Arguments

fb_helper fbcon to restore

Description

This should be called from driver's `drm ->lastclose` callback when implementing an fbcon on top of kms using this helper. This ensures that the user isn't greeted with a black screen when e.g. X dies.

RETURNS

Zero if everything went ok, negative error code otherwise.

Name

`drm_fb_helper_blank` — implementation for `->fb_blank`

Synopsis

```
int drm_fb_helper_blank (int blank, struct fb_info * info);
```

Arguments

blank desired blanking state

info fbdev registered by the helper

Name

`drm_fb_helper_prepare` — setup a `drm_fb_helper` structure

Synopsis

```
void drm_fb_helper_prepare (struct drm_device * dev, struct drm_fb_helper * helper, const struct drm_fb_helper_funcs * funcs);
```

Arguments

dev DRM device

helper driver-allocated fbdev helper structure to set up

funcs pointer to structure of functions associate with this helper

Description

Sets up the bare minimum to make the framebuffer helper usable. This is useful to implement race-free initialization of the polling helpers.

Name

`drm_fb_helper_init` — initialize a `drm_fb_helper` structure

Synopsis

```
int drm_fb_helper_init (struct drm_device * dev, struct drm_fb_helper
* fb_helper, int crtc_count, int max_conn_count);
```

Arguments

<i>dev</i>	drm device
<i>fb_helper</i>	driver-allocated fbdev helper structure to initialize
<i>crtc_count</i>	maximum number of crtc's to support in this fbdev emulation
<i>max_conn_count</i>	max connector count

Description

This allocates the structures for the fbdev helper with the given limits. Note that this won't yet touch the hardware (through the driver interfaces) nor register the fbdev. This is only done in `drm_fb_helper_initial_config` to allow driver writes more control over the exact init sequence.

Drivers must call `drm_fb_helper_prepare` before calling this function.

RETURNS

Zero if everything went ok, nonzero otherwise.

Name

`drm_fb_helper_alloc_fbi` — allocate `fb_info` and some of its members

Synopsis

```
struct fb_info * drm_fb_helper_alloc_fbi (struct drm_fb_helper *  
fb_helper);
```

Arguments

fb_helper driver-allocated fbdev helper

Description

A helper to alloc `fb_info` and the members `cmap` and `apertures`. Called by the driver within the `fb_probe` `fb_helper` callback function.

RETURNS

`fb_info` pointer if things went okay, pointer containing error code otherwise

Name

`drm_fb_helper_unregister_fbi` — unregister fb_info framebuffer device

Synopsis

```
void drm_fb_helper_unregister_fbi (struct drm_fb_helper * fb_helper);
```

Arguments

fb_helper driver-allocated fbdev helper

Description

A wrapper around `unregister_framebuffer`, to release the fb_info framebuffer device

Name

`drm_fb_helper_release_fbi` — dealloc `fb_info` and its members

Synopsis

```
void drm_fb_helper_release_fbi (struct drm_fb_helper * fb_helper);
```

Arguments

fb_helper driver-allocated fbdev helper

Description

A helper to free memory taken by `fb_info` and the members `cmap` and `apertures`

Name

`drm_fb_helper_unlink_fbi` — wrapper around `unlink_framebuffer`

Synopsis

```
void drm_fb_helper_unlink_fbi (struct drm_fb_helper * fb_helper);
```

Arguments

fb_helper driver-allocated fbdev helper

Description

A wrapper around `unlink_framebuffer` implemented by fbdev core

Name

`drm_fb_helper_sys_read` — wrapper around `fb_sys_read`

Synopsis

```
ssize_t drm_fb_helper_sys_read (struct fb_info * info, char __user *  
buf, size_t count, loff_t * ppos);
```

Arguments

info fb_info struct pointer

buf userspace buffer to read from framebuffer memory

count number of bytes to read from framebuffer memory

ppos read offset within framebuffer memory

Description

A wrapper around `fb_sys_read` implemented by fbdev core

Name

`drm_fb_helper_sys_write` — wrapper around `fb_sys_write`

Synopsis

```
ssize_t drm_fb_helper_sys_write (struct fb_info * info, const char
__user * buf, size_t count, loff_t * ppos);
```

Arguments

info fb_info struct pointer

buf userspace buffer to write to framebuffer memory

count number of bytes to write to framebuffer memory

ppos write offset within framebuffer memory

Description

A wrapper around `fb_sys_write` implemented by fbdev core

Name

`drm_fb_helper_sys_fillrect` — wrapper around `sys_fillrect`

Synopsis

```
void drm_fb_helper_sys_fillrect (struct fb_info * info, const struct  
fb_fillrect * rect);
```

Arguments

info fbdev registered by the helper

rect info about rectangle to fill

Description

A wrapper around `sys_fillrect` implemented by fbdev core

Name

`drm_fb_helper_sys_copyarea` — wrapper around `sys_copyarea`

Synopsis

```
void drm_fb_helper_sys_copyarea (struct fb_info * info, const struct  
fb_copyarea * area);
```

Arguments

info fbdev registered by the helper

area info about area to copy

Description

A wrapper around `sys_copyarea` implemented by fbdev core

Name

`drm_fb_helper_sys_imageblit` — wrapper around `sys_imageblit`

Synopsis

```
void drm_fb_helper_sys_imageblit (struct fb_info * info, const struct  
fb_image * image);
```

Arguments

info fbdev registered by the helper

image info about image to blit

Description

A wrapper around `sys_imageblit` implemented by fbdev core

Name

`drm_fb_helper_cfb_fillrect` — wrapper around `cfb_fillrect`

Synopsis

```
void drm_fb_helper_cfb_fillrect (struct fb_info * info, const struct  
fb_fillrect * rect);
```

Arguments

info fbdev registered by the helper

rect info about rectangle to fill

Description

A wrapper around `cfb_imageblit` implemented by fbdev core

Name

`drm_fb_helper_cfb_copyarea` — wrapper around `cfb_copyarea`

Synopsis

```
void drm_fb_helper_cfb_copyarea (struct fb_info * info, const struct  
fb_copyarea * area);
```

Arguments

info fbdev registered by the helper

area info about area to copy

Description

A wrapper around `cfb_copyarea` implemented by fbdev core

Name

`drm_fb_helper_cfb_imageblit` — wrapper around `cfb_imageblit`

Synopsis

```
void drm_fb_helper_cfb_imageblit (struct fb_info * info, const struct  
fb_image * image);
```

Arguments

info fbdev registered by the helper

image info about image to blit

Description

A wrapper around `cfb_imageblit` implemented by fbdev core

Name

`drm_fb_helper_set_suspend` — wrapper around `fb_set_suspend`

Synopsis

```
void drm_fb_helper_set_suspend (struct drm_fb_helper * fb_helper, int
state);
```

Arguments

fb_helper driver-allocated fbdev helper

state desired state, zero to resume, non-zero to suspend

Description

A wrapper around `fb_set_suspend` implemented by fbdev core

Name

`drm_fb_helper_setcmap` — implementation for `->fb_setcmap`

Synopsis

```
int drm_fb_helper_setcmap (struct fb_cmap * cmap, struct fb_info * info);
```

Arguments

cmap cmap to set

info fbdev registered by the helper

Name

`drm_fb_helper_check_var` — implementation for `->fb_check_var`

Synopsis

```
int drm_fb_helper_check_var (struct fb_var_screeninfo * var, struct
fb_info * info);
```

Arguments

var screeninfo to check

info fbdev registered by the helper

Name

`drm_fb_helper_set_par` — implementation for `->fb_set_par`

Synopsis

```
int drm_fb_helper_set_par (struct fb_info * info);
```

Arguments

info fbdev registered by the helper

Description

This will let fbcon do the mode init and is called at initialization time by the fbdev core when registering the driver, and later on through the hotplug callback.

Name

`drm_fb_helper_pan_display` — implementation for `->fb_pan_display`

Synopsis

```
int drm_fb_helper_pan_display (struct fb_var_screeninfo * var, struct
fb_info * info);
```

Arguments

var updated screen information

info fbdev registered by the helper

Name

`drm_fb_helper_fill_fix` — initializes fixed fbdev information

Synopsis

```
void drm_fb_helper_fill_fix (struct fb_info * info, uint32_t pitch,  
uint32_t depth);
```

Arguments

info fbdev registered by the helper

pitch desired pitch

depth desired depth

Description

Helper to fill in the fixed fbdev information useful for a non-accelerated fbdev emulations. Drivers which support acceleration methods which impose additional constraints need to set up their own limits.

Drivers should call this (or their equivalent setup code) from their `->fb_probe` callback.

Name

`drm_fb_helper_fill_var` — initializes variable fbdev information

Synopsis

```
void drm_fb_helper_fill_var (struct fb_info * info, struct drm_fb_helper  
* fb_helper, uint32_t fb_width, uint32_t fb_height);
```

Arguments

info fbdev instance to set up

fb_helper fb helper instance to use as template

fb_width desired fb width

fb_height desired fb height

Description

Sets up the variable fbdev metainformation from the given fb helper instance and the drm framebuffer allocated in `fb_helper->fb`.

Drivers should call this (or their equivalent setup code) from their `->fb_probe` callback after having allocated the fbdev backing storage framebuffer.

Name

`drm_fb_helper_initial_config` — setup a sane initial connector configuration

Synopsis

```
int drm_fb_helper_initial_config (struct drm_fb_helper * fb_helper, int
bpp_sel);
```

Arguments

fb_helper fb_helper device struct

bpp_sel bpp value to use for the framebuffer configuration

Description

Scans the CRTC's and connectors and tries to put together an initial setup. At the moment, this is a cloned configuration across all heads with a new framebuffer object as the backing store.

Note that this also registers the fbdev and so allows userspace to call into the driver through the fbdev interfaces.

This function will call down into the `->fb_probe` callback to let the driver allocate and initialize the fbdev info structure and the drm framebuffer used to back the fbdev. `drm_fb_helper_fill_var` and `drm_fb_helper_fill_fix` are provided as helpers to setup simple default values for the fbdev info structure.

RETURNS

Zero if everything went ok, nonzero otherwise.

Name

`drm_fb_helper_hotplug_event` — respond to a hotplug notification by probing all the outputs attached to the fb

Synopsis

```
int drm_fb_helper_hotplug_event (struct drm_fb_helper * fb_helper);
```

Arguments

fb_helper the `drm_fb_helper`

Description

Scan the connectors attached to the `fb_helper` and try to put together a setup after *notification of a change in output configuration.

Called at runtime, takes the mode config locks to be able to check/change the modeset configuration. Must be run from process context (which usually means either the output polling work or a work item launched from the driver's hotplug interrupt).

Note that drivers may call this even before calling `drm_fb_helper_initial_config` but only after `drm_fb_helper_init`. This allows for a race-free fbcon setup and will make sure that the fbdev emulation will not miss any hotplug events.

RETURNS

0 on success and a non-zero error code otherwise.

Name

struct `drm_fb_helper_surface_size` — describes fbdev size and scanout surface size

Synopsis

```
struct drm_fb_helper_surface_size {  
    u32 fb_width;  
    u32 fb_height;  
    u32 surface_width;  
    u32 surface_height;  
    u32 surface_bpp;  
    u32 surface_depth;  
};
```

Members

<code>fb_width</code>	fbdev width
<code>fb_height</code>	fbdev height
<code>surface_width</code>	scanout buffer width
<code>surface_height</code>	scanout buffer height
<code>surface_bpp</code>	scanout buffer bpp
<code>surface_depth</code>	scanout buffer depth

Description

Note that the scanout surface width/height may be larger than the fbdev width/height. In case of multiple displays, the scanout surface is sized according to the largest width/height (so it is large enough for all CRTC's to scanout). But the fbdev width/height is sized to the minimum width/ height of all the displays. This ensures that fbcon fits on the smallest of the attached displays.

So what is passed to `drm_fb_helper_fill_var` should be `fb_width/fb_height`, rather than the surface size.

Name

struct drm_fb_helper_funcs — driver callbacks for the fbdev emulation library

Synopsis

```
struct drm_fb_helper_funcs {  
    void (* gamma_set) (struct drm_crtc *crtc, u16 red, u16 green,u16 blue, int regn  
    void (* gamma_get) (struct drm_crtc *crtc, u16 *red, u16 *green,u16 *blue, int r  
    int (* fb_probe) (struct drm_fb_helper *helper,struct drm_fb_helper_surface_size  
    bool (* initial_config) (struct drm_fb_helper *fb_helper,struct drm_fb_helper_cr  
};
```

Members

gamma_set

Set the given gamma LUT register on the given CRTC.

This callback is optional.

FIXME:

This callback is functionally redundant with the core gamma table support and simply exists because the fbdev hasn't yet been refactored to use the core gamma table interfaces.

gamma_get

Read the given gamma LUT register on the given CRTC, used to save the current LUT when force-restoring the fbdev for e.g. kdbg.

This callback is optional.

FIXME:

This callback is functionally redundant with the core gamma table support and simply exists because the fbdev hasn't yet been refactored to use the core gamma table interfaces.

fb_probe

Driver callback to allocate and initialize the fbdev info structure. Furthermore it also needs to allocate the DRM framebuffer used to back the fbdev.

This callback is mandatory.

RETURNS:

The driver should return 0 on success and a negative error code on failure.

initial_config

Driver callback to setup an initial fbdev display configuration. Drivers can use this callback to tell the fbdev emulation what the preferred initial configuration is. This is useful to implement smooth booting where the fbdev (and subsequently all userspace) never changes the mode, but always inherits the existing configuration.

This callback is optional.

RETURNS:

The driver should return true if a suitable initial configuration has been filled out and false when the fbdev helper should fall back to the default probing logic.

Description

Driver callbacks used by the fbdev emulation helper library.

Name

struct `drm_fb_helper` — main structure to emulate fbdev on top of KMS

Synopsis

```
struct drm_fb_helper {
    struct drm_framebuffer * fb;
    struct drm_device * dev;
    int crtc_count;
    struct drm_fb_helper_crtc * crtc_info;
    int connector_count;
    int connector_info_alloc_count;
    struct drm_fb_helper_connector ** connector_info;
    const struct drm_fb_helper_funcs * funcs;
    struct fb_info * fbdev;
    u32 pseudo_palette[17];
    struct list_head kernel_fb_list;
    bool delayed_hotplug;
    bool atomic;
};
```

Members

<code>fb</code>	Scanout framebuffer object
<code>dev</code>	DRM device
<code>crtc_count</code>	number of possible CRTC's
<code>crtc_info</code>	per-CRTC helper state (mode, x/y offset, etc)
<code>connector_count</code>	number of connected connectors
<code>connector_info_alloc_count</code>	size of <code>connector_info</code>
<code>connector_info</code>	array of per-connector information
<code>funcs</code>	driver callbacks for fb helper
<code>fbdev</code>	emulated fbdev device info struct
<code>pseudo_palette[17]</code>	fake palette of 16 colors
<code>kernel_fb_list</code>	Entry on the global <code>kernel_fb_helper_list</code> , used for kgdb entry/exit.
<code>delayed_hotplug</code>	A hotplug was received while fbdev wasn't in control of the DRM device, i.e. another KMS master was active. The output configuration needs to be reprobe when fbdev is in control again.
<code>atomic</code>	Use atomic updates for <code>restore_fbdev_mode</code> , etc. This defaults to true if driver has <code>DRIVER_ATOMIC</code> feature flag, but drivers can override it to true after <code>drm_fb_helper_init</code> if they

support atomic modeset but do not yet advertise `DRIVER_ATOMIC` (note that fb-helper does not require `ASYNC` commits).

Description

This is the main structure used by the fbdev helpers. Drivers supporting fbdev emulation should embed this into their overall driver structure. Drivers must also fill out a struct `drm_fb_helper_funcs` with a few operations.

Display Port Helper Functions Reference

These functions contain some common logic and helpers at various abstraction levels to deal with Display Port sink devices and related things like DP aux channel transfers, EDID reading over DP aux channels, decoding certain DPCD blocks, ...

The DisplayPort AUX channel is an abstraction to allow generic, driver- independent access to AUX functionality. Drivers can take advantage of this by filling in the fields of the `drm_dp_aux` structure.

Transactions are described using a hardware-independent `drm_dp_aux_msg` structure, which is passed into a driver's `.transfer` implementation. Both native and I2C-over-AUX transactions are supported.

Name

struct `drm_dp_aux_msg` — DisplayPort AUX channel transaction

Synopsis

```
struct drm_dp_aux_msg {  
    unsigned int address;  
    u8 request;  
    u8 reply;  
    void * buffer;  
    size_t size;  
};
```

Members

<code>address</code>	address of the (first) register to access
<code>request</code>	contains the type of transaction (see <code>DP_AUX_*</code> macros)
<code>reply</code>	upon completion, contains the reply type of the transaction
<code>buffer</code>	pointer to a transmission or reception buffer
<code>size</code>	size of <i>buffer</i>

Name

struct drm_dp_aux — DisplayPort AUX channel

Synopsis

```
struct drm_dp_aux {
    const char * name;
    struct i2c_adapter ddc;
    struct device * dev;
    struct mutex hw_mutex;
    ssize_t (* transfer) (struct drm_dp_aux *aux, struct drm_dp_aux_msg *msg);
};
```

Members

name	user-visible name of this AUX channel and the I2C-over-AUX adapter
ddc	I2C adapter that can be used for I2C-over-AUX communication
dev	pointer to struct device that is the parent for this AUX channel
hw_mutex	internal mutex used for locking transfers
transfer	transfers a message representing a single AUX transaction

Description

The `.dev` field should be set to a pointer to the device that implements the AUX channel.

The `.name` field may be used to specify the name of the I2C adapter. If set to NULL, `dev_name` of `.dev` will be used.

Drivers provide a hardware-specific implementation of how transactions are executed via the `.transfer` function. A pointer to a `drm_dp_aux_msg` structure describing the transaction is passed into this function. Upon success, the implementation should return the number of payload bytes that were transferred, or a negative error-code on failure. Helpers propagate errors from the `.transfer` function, with the exception of the `-EBUSY` error, which causes a transaction to be retried. On a short, helpers will return `-EPROTO` to make it simpler to check for failure.

An AUX channel can also be used to transport I2C messages to a sink. A typical application of that is to access an EDID that's present in the sink device. The `.transfer` function can also be used to execute such transactions. The `drm_dp_aux_register` function registers an I2C adapter that can be passed to `drm_probe_ddc`. Upon removal, drivers should call `drm_dp_aux_unregister` to remove the I2C adapter. The I2C adapter uses long transfers by default; if a partial response is received, the adapter will drop down to the size given by the partial response for this transaction only.

Note that the aux helper code assumes that the `.transfer` function only modifies the reply field of the `drm_dp_aux_msg` structure. The retry logic and i2c helpers assume this is the case.

Name

`drm_dp_dpcd_readb` — read a single byte from the DPCD

Synopsis

```
ssize_t drm_dp_dpcd_readb (struct drm_dp_aux * aux, unsigned int offset,  
u8 * valuep);
```

Arguments

aux DisplayPort AUX channel

offset address of the register to read

valuep location where the value of the register will be stored

Description

Returns the number of bytes transferred (1) on success, or a negative error code on failure.

Name

`drm_dp_dpcd_writeb` — write a single byte to the DPCD

Synopsis

```
ssize_t drm_dp_dpcd_writeb (struct drm_dp_aux * aux, unsigned int offset,  
                             u8 value);
```

Arguments

aux DisplayPort AUX channel

offset address of the register to write

value value to write to the register

Description

Returns the number of bytes transferred (1) on success, or a negative error code on failure.

Name

`drm_dp_dpcd_read` — read a series of bytes from the DPCD

Synopsis

```
ssize_t drm_dp_dpcd_read (struct drm_dp_aux * aux, unsigned int offset,  
void * buffer, size_t size);
```

Arguments

<i>aux</i>	DisplayPort AUX channel
<i>offset</i>	address of the (first) register to read
<i>buffer</i>	buffer to store the register values
<i>size</i>	number of bytes in <i>buffer</i>

Description

Returns the number of bytes transferred on success, or a negative error code on failure. -EIO is returned if the request was NAKed by the sink or if the retry count was exceeded. If not all bytes were transferred, this function returns -EPROTO. Errors from the underlying AUX channel transfer function, with the exception of -EBUSY (which causes the transaction to be retried), are propagated to the caller.

Name

`drm_dp_dpcd_write` — write a series of bytes to the DPCD

Synopsis

```
ssize_t drm_dp_dpcd_write (struct drm_dp_aux * aux, unsigned int offset,  
void * buffer, size_t size);
```

Arguments

<i>aux</i>	DisplayPort AUX channel
<i>offset</i>	address of the (first) register to write
<i>buffer</i>	buffer containing the values to write
<i>size</i>	number of bytes in <i>buffer</i>

Description

Returns the number of bytes transferred on success, or a negative error code on failure. -EIO is returned if the request was NAKed by the sink or if the retry count was exceeded. If not all bytes were transferred, this function returns -EPROTO. Errors from the underlying AUX channel transfer function, with the exception of -EBUSY (which causes the transaction to be retried), are propagated to the caller.

Name

`drm_dp_dpcd_read_link_status` — read DPCD link status (bytes 0x202-0x207)

Synopsis

```
int drm_dp_dpcd_read_link_status (struct drm_dp_aux * aux, u8 status[D-  
P_LINK_STATUS_SIZE]);
```

Arguments

<i>aux</i>	DisplayPort AUX channel
<i>status</i> [<i>DP_LINK_S- TATUS_SIZE</i>]	buffer to store the link status in (must be at least 6 bytes)

Description

Returns the number of bytes transferred on success or a negative error code on failure.

Name

`drm_dp_link_probe` — probe a DisplayPort link for capabilities

Synopsis

```
int drm_dp_link_probe (struct drm_dp_aux * aux, struct drm_dp_link *  
link);
```

Arguments

aux DisplayPort AUX channel

link pointer to structure in which to return link capabilities

Description

The structure filled in by this function can usually be passed directly into `drm_dp_link_power_up` and `drm_dp_link_configure` to power up and configure the link based on the link's capabilities.

Returns 0 on success or a negative error code on failure.

Name

`drm_dp_link_power_up` — power up a DisplayPort link

Synopsis

```
int drm_dp_link_power_up (struct drm_dp_aux * aux, struct drm_dp_link  
* link);
```

Arguments

aux DisplayPort AUX channel

link pointer to a structure containing the link configuration

Description

Returns 0 on success or a negative error code on failure.

Name

`drm_dp_link_power_down` — power down a DisplayPort link

Synopsis

```
int drm_dp_link_power_down (struct drm_dp_aux * aux, struct drm_dp_link  
* link);
```

Arguments

aux DisplayPort AUX channel

link pointer to a structure containing the link configuration

Description

Returns 0 on success or a negative error code on failure.

Name

`drm_dp_link_configure` — configure a DisplayPort link

Synopsis

```
int drm_dp_link_configure (struct drm_dp_aux * aux, struct drm_dp_link  
* link);
```

Arguments

aux DisplayPort AUX channel

link pointer to a structure containing the link configuration

Description

Returns 0 on success or a negative error code on failure.

Name

`drm_dp_aux_register` — initialise and register aux channel

Synopsis

```
int drm_dp_aux_register (struct drm_dp_aux * aux);
```

Arguments

aux DisplayPort AUX channel

Description

Returns 0 on success or a negative error code on failure.

Name

`drm_dp_aux_unregister` — unregister an AUX adapter

Synopsis

```
void drm_dp_aux_unregister (struct drm_dp_aux * aux);
```

Arguments

aux DisplayPort AUX channel

Display Port MST Helper Functions Reference

These functions contain parts of the DisplayPort 1.2a MultiStream Transport protocol. The helpers contain a topology manager and bandwidth manager. The helpers encapsulate the sending and received of sideband msgs.

Name

struct drm_dp_vcpi — Virtual Channel Payload Identifier

Synopsis

```
struct drm_dp_vcpi {  
    int vcpi;  
    int pbn;  
    int aligned_pbn;  
    int num_slots;  
};
```

Members

vcpi	Virtual channel ID.
pbn	Payload Bandwidth Number for this channel
aligned_pbn	PBN aligned with slot size
num_slots	number of slots for this PBN

Name

struct drm_dp_mst_port — MST port

Synopsis

```
struct drm_dp_mst_port {
    struct kref kref;
    bool guid_valid;
    u8 guid[16];
    u8 port_num;
    bool input;
    bool mcs;
    bool ddps;
    u8 pdt;
    bool ldps;
    u8 dpcd_rev;
    u8 num_sdp_streams;
    u8 num_sdp_stream_sinks;
    uint16_t available_pbn;
    struct list_head next;
    struct drm_dp_mst_branch * mstb;
    struct drm_dp_aux aux;
    struct drm_dp_mst_branch * parent;
    struct drm_dp_vcpi vcpi;
    struct drm_connector * connector;
    struct drm_dp_mst_topology_mgr * mgr;
};
```

Members

kref	reference count for this port.
guid_valid	for DP 1.2 devices if we have validated the GUID.
guid[16]	guid for DP 1.2 device on this port.
port_num	port number
input	if this port is an input port.
mcs	message capability status - DP 1.2 spec.
ddps	DisplayPort Device Plug Status - DP 1.2
pdt	Peer Device Type
ldps	Legacy Device Plug Status
dpcd_rev	DPCD revision of device on this port
num_sdp_streams	Number of simultaneous streams
num_sdp_stream_sinks	Number of stream sinks
available_pbn	Available bandwidth for this port.

next	link to next port on this branch device
mstb	branch device attach below this port
aux	i2c aux transport to talk to device connected to this port.
parent	branch device parent of this port
vcpi	Virtual Channel Payload info for this port.
connector	DRM connector this port is connected to.
mgr	topology manager this port lives under.

Description

This structure represents an MST port endpoint on a device somewhere in the MST topology.

Name

struct drm_dp_mst_branch — MST branch device.

Synopsis

```
struct drm_dp_mst_branch {
    struct kref kref;
    u8 rad[8];
    u8 lct;
    int num_ports;
    int msg_slots;
    struct list_head ports;
    struct drm_dp_mst_port * port_parent;
    struct drm_dp_mst_topology_mgr * mgr;
    struct drm_dp_sideband_msg_tx * tx_slots[2];
    int last_seqno;
    bool link_address_sent;
};
```

Members

kref	reference count for this port.
rad[8]	Relative Address to talk to this branch device.
lct	Link count total to talk to this branch device.
num_ports	number of ports on the branch.
msg_slots	one bit per transmitted msg slot.
ports	linked list of ports on this branch.
port_parent	pointer to the port parent, NULL if toplevel.
mgr	topology manager for this branch device.
tx_slots[2]	transmission slots for this device.
last_seqno	last sequence number used to talk to this.
link_address_sent	if a link address message has been sent to this device yet.

Description

This structure represents an MST branch device, there is one primary branch device at the root, along with any others connected to downstream ports

Name

struct drm_dp_mst_topology_mgr — DisplayPort MST manager

Synopsis

```
struct drm_dp_mst_topology_mgr {
    struct device * dev;
    const struct drm_dp_mst_topology_cbs * cbs;
    struct drm_dp_aux * aux;
    int max_payloads;
    int conn_base_id;
    struct drm_dp_sideband_msg_rx down_rep_recv;
    struct drm_dp_sideband_msg_rx up_req_recv;
    struct mutex lock;
    bool mst_state;
    struct drm_dp_mst_branch * mst_primary;
    bool guid_valid;
    u8 guid[16];
    u8 dpcd[DP_RECEIVER_CAP_SIZE];
    int pbn_div;
};
```

Members

dev	device pointer for adding i2c devices etc.
cbs	callbacks for connector addition and destruction. <i>max_dpcd_transaction_bytes</i> - maximum number of bytes to read/write in one go.
aux	aux channel for the DP connector.
max_payloads	maximum number of payloads the GPU can generate.
conn_base_id	DRM connector ID this mgr is connected to.
down_rep_recv	msg receiver state for down replies.
up_req_recv	msg receiver state for up requests.
lock	protects mst state, primary, guid, dpcd.
mst_state	if this manager is enabled for an MST capable port.
mst_primary	pointer to the primary branch device.
guid_valid	GUID valid for the primary branch device.
guid[16]	GUID for primary port.
dpcd[DP_RECEIVER_CAP_SIZE]	cache of DPCD for primary port.
pbn_div	PBN to slots divisor.

Description

This struct represents the toplevel displayport MST topology manager. There should be one instance of this for every MST capable DP connector on the GPU.

Name

`drm_dp_update_payload_part1` — Execute payload update part 1

Synopsis

```
int drm_dp_update_payload_part1 (struct drm_dp_mst_topology_mgr * mgr);
```

Arguments

mgr manager to use.

Description

This iterates over all proposed virtual channels, and tries to allocate space in the link for them. For 0->slots transitions, this step just writes the VCPI to the MST device. For slots->0 transitions, this writes the updated VCPIs and removes the remote VC payloads.

after calling this the driver should generate ACT and payload packets.

Name

`drm_dp_update_payload_part2` — Execute payload update part 2

Synopsis

```
int drm_dp_update_payload_part2 (struct drm_dp_mst_topology_mgr * mgr);
```

Arguments

mgr manager to use.

Description

This iterates over all proposed virtual channels, and tries to allocate space in the link for them. For 0->slots transitions, this step writes the remote VC payload commands. For slots->0 this just resets some internal state.

Name

`drm_dp_mst_topology_mgr_set_mst` — Set the MST state for a topology manager

Synopsis

```
int drm_dp_mst_topology_mgr_set_mst (struct drm_dp_mst_topology_mgr *  
mgr, bool mst_state);
```

Arguments

mgr manager to set state for

mst_state true to enable MST on this connector - false to disable.

Description

This is called by the driver when it detects an MST capable device plugged into a DP MST capable port, or when a DP MST capable device is unplugged.

Name

`drm_dp_mst_topology_mgr_suspend` — suspend the MST manager

Synopsis

```
void drm_dp_mst_topology_mgr_suspend (struct drm_dp_mst_topology_mgr *  
mgr);
```

Arguments

mgr manager to suspend

Description

This function tells the MST device that we can't handle UP messages anymore. This should stop it from sending any since we are suspended.

Name

`drm_dp_mst_topology_mgr_resume` — resume the MST manager

Synopsis

```
int drm_dp_mst_topology_mgr_resume (struct drm_dp_mst_topology_mgr *  
mgr);
```

Arguments

mgr manager to resume

Description

This will fetch DPCD and see if the device is still there, if it is, it will rewrite the MSTM control bits, and return.

if the device fails this returns -1, and the driver should do a full MST reprobe, in case we were undocked.

Name

`drm_dp_mst_hpd_irq` — MST hotplug IRQ notify

Synopsis

```
int drm_dp_mst_hpd_irq (struct drm_dp_mst_topology_mgr * mgr, u8 * esi,  
bool * handled);
```

Arguments

mgr manager to notify irq for.

esi 4 bytes from SINK_COUNT_ESI

handled whether the hpd interrupt was consumed or not

Description

This should be called from the driver when it detects a short IRQ, along with the value of the `DEVICE_SERVICE_IRQ_VECTOR_ESI0`. The topology manager will process the sideband messages received as a result of this.

Name

`drm_dp_mst_detect_port` — get connection status for an MST port

Synopsis

```
enum drm_connector_status drm_dp_mst_detect_port (struct drm_connector * connector, struct drm_dp_mst_topology_mgr * mgr, struct drm_dp_mst_port * port);
```

Arguments

connector -- undecribed --

mgr manager for this port

port unverified pointer to a port

Description

This returns the current connection state for a port. It validates the port pointer still exists so the caller doesn't require a reference

Name

`drm_dp_mst_port_has_audio` — Check whether port has audio capability or not

Synopsis

```
bool drm_dp_mst_port_has_audio (struct drm_dp_mst_topology_mgr * mgr,  
struct drm_dp_mst_port * port);
```

Arguments

mgr manager for this port

port unverified pointer to a port.

Description

This returns whether the port supports audio or not.

Name

`drm_dp_mst_get_edid` — get EDID for an MST port

Synopsis

```
struct edid * drm_dp_mst_get_edid (struct drm_connector * connector,  
struct drm_dp_mst_topology_mgr * mgr, struct drm_dp_mst_port * port);
```

Arguments

connector toplevel connector to get EDID for

mgr manager for this port

port unverified pointer to a port.

Description

This returns an EDID for the port connected to a connector, It validates the pointer still exists so the caller doesn't require a reference.

Name

`drm_dp_find_vcpi_slots` — find slots for this PBN value

Synopsis

```
int drm_dp_find_vcpi_slots (struct drm_dp_mst_topology_mgr * mgr, int  
pbn);
```

Arguments

mgr manager to use

pbn payload bandwidth to convert into slots.

Name

`drm_dp_mst_allocate_vcpi` — Allocate a virtual channel

Synopsis

```
bool drm_dp_mst_allocate_vcpi (struct drm_dp_mst_topology_mgr * mgr,  
struct drm_dp_mst_port * port, int pbn, int * slots);
```

Arguments

mgr manager for this port

port port to allocate a virtual channel for.

pbn payload bandwidth number to request

slots returned number of slots for this PBN.

Name

`drm_dp_mst_reset_vcpi_slots` — Reset number of slots to 0 for VCPI

Synopsis

```
void drm_dp_mst_reset_vcpi_slots (struct drm_dp_mst_topology_mgr * mgr,  
struct drm_dp_mst_port * port);
```

Arguments

mgr manager for this port

port unverified pointer to a port.

Description

This just resets the number of slots for the ports VCPI for later programming.

Name

`drm_dp_mst_deallocate_vcpi` — deallocate a VCPI

Synopsis

```
void drm_dp_mst_deallocate_vcpi (struct drm_dp_mst_topology_mgr * mgr,  
struct drm_dp_mst_port * port);
```

Arguments

mgr manager for this port

port unverified port to deallocate vcpi for

Name

`drm_dp_check_act_status` — Check ACT handled status.

Synopsis

```
int drm_dp_check_act_status (struct drm_dp_mst_topology_mgr * mgr);
```

Arguments

mgr manager to use

Description

Check the payload status bits in the DPCD for ACT handled completion.

Name

`drm_dp_calc_pbn_mode` — Calculate the PBN for a mode.

Synopsis

```
int drm_dp_calc_pbn_mode (int clock, int bpp);
```

Arguments

clock dot clock for the mode

bpp bpp for the mode.

Description

This uses the formula in the spec to calculate the PBN value for a mode.

Name

`drm_dp_mst_dump_topology` —

Synopsis

```
void drm_dp_mst_dump_topology (struct seq_file * m, struct drm_dp_mst_topology_mgr * mgr);
```

Arguments

m seq_file to dump output to

mgr manager to dump current topology for.

Description

helper to dump MST topology to a seq file for debugfs.

Name

`drm_dp_mst_topology_mgr_init` — initialise a topology manager

Synopsis

```
int drm_dp_mst_topology_mgr_init (struct drm_dp_mst_topology_mgr * mgr,
struct device * dev, struct drm_dp_aux * aux, int max_dpcd_transaction_bytes, int max_payloads, int conn_base_id);
```

Arguments

<i>mgr</i>	manager struct to initialise
<i>dev</i>	device providing this structure - for i2c addition.
<i>aux</i>	DP helper aux channel to talk to this device
<i>max_dpcd_transaction_bytes</i>	hw specific DPCD transaction limit
<i>max_payloads</i>	maximum number of payloads this GPU can source
<i>conn_base_id</i>	the connector object ID the MST device is connected to.

Description

Return 0 for success, or negative error code on failure

Name

`drm_dp_mst_topology_mgr_destroy` — destroy topology manager.

Synopsis

```
void drm_dp_mst_topology_mgr_destroy (struct drm_dp_mst_topology_mgr *  
mgr);
```

Arguments

mgr manager to destroy

MIPI DSI Helper Functions Reference

These functions contain some common logic and helpers to deal with MIPI DSI peripherals.

Helpers are provided for a number of standard MIPI DSI command as well as a subset of the MIPI DCS command set.

Name

struct mipi_dsi_msg — read/write DSI buffer

Synopsis

```
struct mipi_dsi_msg {  
    u8 channel;  
    u8 type;  
    u16 flags;  
    size_t tx_len;  
    const void * tx_buf;  
    size_t rx_len;  
    void * rx_buf;  
};
```

Members

channel	virtual channel id
type	payload data type
flags	flags controlling this message transmission
tx_len	length of <i>tx_buf</i>
tx_buf	data to be written
rx_len	length of <i>rx_buf</i>
rx_buf	data to be read, or NULL

Name

struct mipi_dsi_packet — represents a MIPI DSI packet in protocol format

Synopsis

```
struct mipi_dsi_packet {  
    size_t size;  
    u8 header[4];  
    size_t payload_length;  
    const u8 * payload;  
};
```

Members

size	size (in bytes) of the packet
header[4]	the four bytes that make up the header (Data ID, Word Count or Packet Data, and ECC)
payload_length	number of bytes in the payload
payload	a pointer to a buffer containing the payload, if any

Name

struct mipi_dsi_host_ops — DSI bus operations

Synopsis

```
struct mipi_dsi_host_ops {  
    int (* attach) (struct mipi_dsi_host *host, struct mipi_dsi_device *dsi);  
    int (* detach) (struct mipi_dsi_host *host, struct mipi_dsi_device *dsi);  
    ssize_t (* transfer) (struct mipi_dsi_host *host, const struct mipi_dsi_msg *msg)  
};
```

Members

attach	attach DSI device to DSI host
detach	detach DSI device from DSI host
transfer	transmit a DSI packet

Description

DSI packets transmitted by `.transfer` are passed in as `mipi_dsi_msg` structures. This structure contains information about the type of packet being transmitted as well as the transmit and receive buffers. When an error is encountered during transmission, this function will return a negative error code. On success it shall return the number of bytes transmitted for write packets or the number of bytes received for read packets.

Note that typically DSI packet transmission is atomic, so the `.transfer` function will seldomly return anything other than the number of bytes contained in the transmit buffer on success.

Name

struct mipi_dsi_host — DSI host device

Synopsis

```
struct mipi_dsi_host {  
    struct device * dev;  
    const struct mipi_dsi_host_ops * ops;  
};
```

Members

dev driver model device node for this DSI host

ops DSI host operations

Name

struct mipi_dsi_device — DSI peripheral device

Synopsis

```
struct mipi_dsi_device {  
    struct mipi_dsi_host * host;  
    struct device dev;  
    unsigned int channel;  
    unsigned int lanes;  
    enum mipi_dsi_pixel_format format;  
    unsigned long mode_flags;  
};
```

Members

host	DSI host for this peripheral
dev	driver model device node for this peripheral
channel	virtual channel assigned to the peripheral
lanes	number of active data lanes
format	pixel format for video mode
mode_flags	DSI operation mode related flags

Name

`mipi_dsi_pixel_format_to_bpp` — obtain the number of bits per pixel for any given pixel format defined by the MIPI DSI specification

Synopsis

```
int mipi_dsi_pixel_format_to_bpp (enum mipi_dsi_pixel_format fmt);
```

Arguments

fmt MIPI DSI pixel format

Returns

The number of bits per pixel of the given pixel format.

Name

enum mipi_dsi_dcs_tear_mode — Tearing Effect Output Line mode

Synopsis

```
enum mipi_dsi_dcs_tear_mode {  
    MIPI_DSI_DCS_TEAR_MODE_VBLANK,  
    MIPI_DSI_DCS_TEAR_MODE_VHBLANK  
};
```

Constants

MIPI_DSI_DCS_TEAR_MODE_VBLANK	the TE output line consists of V-Blanking information only
MIPI_DSI_DCS_TEAR_MODE_VHBLANK	the TE output line consists of both V-Blanking and H-Blanking information

Name

struct mipi_dsi_driver — DSI driver

Synopsis

```
struct mipi_dsi_driver {  
    struct device_driver driver;  
    int(* probe) (struct mipi_dsi_device *dsi);  
    int(* remove) (struct mipi_dsi_device *dsi);  
    void (* shutdown) (struct mipi_dsi_device *dsi);  
};
```

Members

driver	device driver model driver
probe	callback for device binding
remove	callback for device unbinding
shutdown	called at shutdown time to quiesce the device

Name

`of_find_mipi_dsi_device_by_node` — find the MIPI DSI device matching a device tree node

Synopsis

```
struct mipi_dsi_device * of_find_mipi_dsi_device_by_node (struct de-  
vice_node * np);
```

Arguments

np device tree node

Return

A pointer to the MIPI DSI device corresponding to *np* or NULL if no such device exists (or has not been registered yet).

Name

`mipi_dsi_attach` — attach a DSI device to its DSI host

Synopsis

```
int mipi_dsi_attach (struct mipi_dsi_device * dsi);
```

Arguments

dsi DSI peripheral

Name

`mipi_dsi_detach` — detach a DSI device from its DSI host

Synopsis

```
int mipi_dsi_detach (struct mipi_dsi_device * dsi);
```

Arguments

dsi DSI peripheral

Name

`mipi_dsi_packet_format_is_short` — check if a packet is of the short format

Synopsis

```
bool mipi_dsi_packet_format_is_short (u8 type);
```

Arguments

type MIPI DSI data type of the packet

Return

true if the packet for the given data type is a short packet, false otherwise.

Name

`mipi_dsi_packet_format_is_long` — check if a packet is of the long format

Synopsis

```
bool mipi_dsi_packet_format_is_long (u8 type);
```

Arguments

type MIPI DSI data type of the packet

Return

true if the packet for the given data type is a long packet, false otherwise.

Name

`mipi_dsi_create_packet` — create a packet from a message according to the DSI protocol

Synopsis

```
int mipi_dsi_create_packet (struct mipi_dsi_packet * packet, const
struct mipi_dsi_msg * msg);
```

Arguments

packet pointer to a DSI packet structure

msg message to translate into a packet

Return

0 on success or a negative error code on failure.

Name

`mipi_dsi_shutdown_peripheral` — sends a Shutdown Peripheral command

Synopsis

```
int mipi_dsi_shutdown_peripheral (struct mipi_dsi_device * dsi);
```

Arguments

dsi DSI peripheral device

Return

0 on success or a negative error code on failure.

Name

`mipi_dsi_turn_on_peripheral` — sends a Turn On Peripheral command

Synopsis

```
int mipi_dsi_turn_on_peripheral (struct mipi_dsi_device * dsi);
```

Arguments

dsi DSI peripheral device

Return

0 on success or a negative error code on failure.

Name

`mipi_dsi_generic_write` — transmit data using a generic write packet

Synopsis

```
ssize_t mipi_dsi_generic_write (struct mipi_dsi_device * dsi, const void  
* payload, size_t size);
```

Arguments

<i>dsi</i>	DSI peripheral device
<i>payload</i>	buffer containing the payload
<i>size</i>	size of payload buffer

Description

This function will automatically choose the right data type depending on the payload length.

Return

The number of bytes transmitted on success or a negative error code on failure.

Name

`mipi_dsi_generic_read` — receive data using a generic read packet

Synopsis

```
ssize_t mipi_dsi_generic_read (struct mipi_dsi_device * dsi, const void  
* params, size_t num_params, void * data, size_t size);
```

Arguments

<i>dsi</i>	DSI peripheral device
<i>params</i>	buffer containing the request parameters
<i>num_params</i>	number of request parameters
<i>data</i>	buffer in which to return the received data
<i>size</i>	size of receive buffer

Description

This function will automatically choose the right data type depending on the number of parameters passed in.

Return

The number of bytes successfully read or a negative error code on failure.

Name

`mipi_dsi_dcs_write_buffer` — transmit a DCS command with payload

Synopsis

```
ssize_t mipi_dsi_dcs_write_buffer (struct mipi_dsi_device * dsi, const  
void * data, size_t len);
```

Arguments

dsi DSI peripheral device

data buffer containing data to be transmitted

len size of transmission buffer

Description

This function will automatically choose the right data type depending on the command payload length.

Return

The number of bytes successfully transmitted or a negative error code on failure.

Name

`mipi_dsi_dcs_write` — send DCS write command

Synopsis

```
ssize_t mipi_dsi_dcs_write (struct mipi_dsi_device * dsi, u8 cmd, const  
void * data, size_t len);
```

Arguments

dsi DSI peripheral device

cmd DCS command

data buffer containing the command payload

len command payload length

Description

This function will automatically choose the right data type depending on the command payload length.

Return

The number of bytes successfully transmitted or a negative error code on failure.

Name

`mipi_dsi_dcs_read` — send DCS read request command

Synopsis

```
ssize_t mipi_dsi_dcs_read (struct mipi_dsi_device * dsi, u8 cmd, void  
* data, size_t len);
```

Arguments

dsi DSI peripheral device

cmd DCS command

data buffer in which to receive data

len size of receive buffer

Return

The number of bytes read or a negative error code on failure.

Name

`mipi_dsi_dcs_nop` — send DCS nop packet

Synopsis

```
int mipi_dsi_dcs_nop (struct mipi_dsi_device * dsi);
```

Arguments

dsi DSI peripheral device

Return

0 on success or a negative error code on failure.

Name

`mipi_dsi_dcs_soft_reset` — perform a software reset of the display module

Synopsis

```
int mipi_dsi_dcs_soft_reset (struct mipi_dsi_device * dsi);
```

Arguments

dsi DSI peripheral device

Return

0 on success or a negative error code on failure.

Name

`mipi_dsi_dcs_get_power_mode` — query the display module's current power mode

Synopsis

```
int mipi_dsi_dcs_get_power_mode (struct mipi_dsi_device * dsi, u8 *  
mode);
```

Arguments

dsi DSI peripheral device

mode return location for the current power mode

Return

0 on success or a negative error code on failure.

Name

`mipi_dsi_dcs_get_pixel_format` — gets the pixel format for the RGB image data used by the interface

Synopsis

```
int mipi_dsi_dcs_get_pixel_format (struct mipi_dsi_device * dsi, u8 *  
format);
```

Arguments

dsi DSI peripheral device

format return location for the pixel format

Return

0 on success or a negative error code on failure.

Name

`mipi_dsi_dcs_enter_sleep_mode` — disable all unnecessary blocks inside the display module except interface communication

Synopsis

```
int mipi_dsi_dcs_enter_sleep_mode (struct mipi_dsi_device * dsi);
```

Arguments

dsi DSI peripheral device

Return

0 on success or a negative error code on failure.

Name

`mipi_dsi_dcs_exit_sleep_mode` — enable all blocks inside the display module

Synopsis

```
int mipi_dsi_dcs_exit_sleep_mode (struct mipi_dsi_device * dsi);
```

Arguments

dsi DSI peripheral device

Return

0 on success or a negative error code on failure.

Name

`mipi_dsi_dcs_set_display_off` — stop displaying the image data on the display device

Synopsis

```
int mipi_dsi_dcs_set_display_off (struct mipi_dsi_device * dsi);
```

Arguments

dsi DSI peripheral device

Return

0 on success or a negative error code on failure.

Name

`mipi_dsi_dcs_set_display_on` — start displaying the image data on the display device

Synopsis

```
int mipi_dsi_dcs_set_display_on (struct mipi_dsi_device * dsi);
```

Arguments

dsi DSI peripheral device

Return

0 on success or a negative error code on failure

Name

`mipi_dsi_dcs_set_column_address` — define the column extent of the frame memory accessed by the host processor

Synopsis

```
int mipi_dsi_dcs_set_column_address (struct mipi_dsi_device * dsi, u16
start, u16 end);
```

Arguments

dsi DSI peripheral device

start first column of frame memory

end last column of frame memory

Return

0 on success or a negative error code on failure.

Name

`mipi_dsi_dcs_set_page_address` — define the page extent of the frame memory accessed by the host processor

Synopsis

```
int mipi_dsi_dcs_set_page_address (struct mipi_dsi_device * dsi, u16
start, u16 end);
```

Arguments

dsi DSI peripheral device

start first page of frame memory

end last page of frame memory

Return

0 on success or a negative error code on failure.

Name

`mipi_dsi_dcs_set_tear_off` — turn off the display module's Tearing Effect output signal on the TE signal line

Synopsis

```
int mipi_dsi_dcs_set_tear_off (struct mipi_dsi_device * dsi);
```

Arguments

dsi DSI peripheral device

Return

0 on success or a negative error code on failure

Name

`mipi_dsi_dcs_set_tear_on` — turn on the display module's Tearing Effect output signal on the TE signal line.

Synopsis

```
int mipi_dsi_dcs_set_tear_on (struct mipi_dsi_device * dsi, enum mipi_dsi_dcs_tear_mode mode);
```

Arguments

dsi DSI peripheral device

mode the Tearing Effect Output Line mode

Return

0 on success or a negative error code on failure

Name

`mipi_dsi_dcs_set_pixel_format` — sets the pixel format for the RGB image data used by the interface

Synopsis

```
int mipi_dsi_dcs_set_pixel_format (struct mipi_dsi_device * dsi, u8
format);
```

Arguments

dsi DSI peripheral device

format pixel format

Return

0 on success or a negative error code on failure.

Name

`mipi_dsi_driver_register_full` — register a driver for DSI devices

Synopsis

```
int mipi_dsi_driver_register_full (struct mipi_dsi_driver * drv, struct  
module * owner);
```

Arguments

drv DSI driver structure

owner owner module

Return

0 on success or a negative error code on failure.

Name

`mipi_dsi_driver_unregister` — unregister a driver for DSI devices

Synopsis

```
void mipi_dsi_driver_unregister (struct mipi_dsi_driver * drv);
```

Arguments

drv DSI driver structure

Return

0 on success or a negative error code on failure.

EDID Helper Functions Reference

Name

`drm_edid_header_is_valid` — sanity check the header of the base EDID block

Synopsis

```
int drm_edid_header_is_valid (const u8 * raw_edid);
```

Arguments

raw_edid pointer to raw base EDID block

Description

Sanity check the header of the base EDID block.

Return

8 if the header is perfect, down to 0 if it's totally wrong.

Name

`drm_edid_block_valid` — Sanity check the EDID block (base or extension)

Synopsis

```
bool drm_edid_block_valid (u8 * raw_edid, int block, bool print_bad_e-  
did, bool * edid_corrupt);
```

Arguments

<i>raw_edid</i>	pointer to raw EDID block
<i>block</i>	type of block to validate (0 for base, extension otherwise)
<i>print_bad_edid</i>	if true, dump bad EDID blocks to the console
<i>edid_corrupt</i>	if true, the header or checksum is invalid

Description

Validate a base or extension EDID block and optionally dump bad blocks to the console.

Return

True if the block is valid, false otherwise.

Name

`drm_edid_is_valid` — sanity check EDID data

Synopsis

```
bool drm_edid_is_valid (struct edid * edid);
```

Arguments

edid EDID data

Description

Sanity-check an entire EDID record (including extensions)

Return

True if the EDID data is valid, false otherwise.

Name

`drm_do_get_edid` — get EDID data using a custom EDID block read function

Synopsis

```
struct edid * drm_do_get_edid (struct drm_connector * connector, int
(*get_edid_block) (void *data, u8 *buf, unsigned int block, size_t len),
void * data);
```

Arguments

<i>connector</i>	connector we're probing
<i>get_edid_block</i>	EDID block read function
<i>data</i>	private data passed to the block read function

Description

When the I2C adapter connected to the DDC bus is hidden behind a device that exposes a different interface to read EDID blocks this function can be used to get EDID data using a custom block read function.

As in the general case the DDC bus is accessible by the kernel at the I2C level, drivers must make all reasonable efforts to expose it as an I2C adapter and use `drm_get_edid` instead of abusing this function.

Return

Pointer to valid EDID or NULL if we couldn't find any.

Name

`drm_probe_ddc` — probe DDC presence

Synopsis

```
bool drm_probe_ddc (struct i2c_adapter * adapter);
```

Arguments

adapter I2C adapter to probe

Return

True on success, false on failure.

Name

`drm_get_edid` — get EDID data, if available

Synopsis

```
struct edid * drm_get_edid (struct drm_connector * connector, struct  
i2c_adapter * adapter);
```

Arguments

connector connector we're probing

adapter I2C adapter to use for DDC

Description

Poke the given I2C channel to grab EDID data if possible. If found, attach it to the connector.

Return

Pointer to valid EDID or NULL if we couldn't find any.

Name

`drm_edid_duplicate` — duplicate an EDID and the extensions

Synopsis

```
struct edid * drm_edid_duplicate (const struct edid * edid);
```

Arguments

edid EDID to duplicate

Return

Pointer to duplicated EDID or NULL on allocation failure.

Name

`drm_match_cea_mode` — look for a CEA mode matching given mode

Synopsis

```
u8 drm_match_cea_mode (const struct drm_display_mode * to_match);
```

Arguments

to_match display mode

Return

The CEA Video ID (VIC) of the mode or 0 if it isn't a CEA-861 mode.

Name

`drm_get_cea_aspect_ratio` — get the picture aspect ratio corresponding to the input VIC from the CEA mode list

Synopsis

```
enum    hdmi_picture_aspect    drm_get_cea_aspect_ratio    (const    u8  
video_code);
```

Arguments

video_code ID given to each of the CEA modes

Description

Returns picture aspect ratio

Name

`drm_edid_to_eld` — build ELD from EDID

Synopsis

```
void drm_edid_to_eld (struct drm_connector * connector, struct edid *  
edid);
```

Arguments

connector connector corresponding to the HDMI/DP sink

edid EDID to parse

Description

Fill the ELD (EDID-Like Data) buffer for passing to the audio driver. The `Conn_Type`, `HDCP` and `Port_ID` ELD fields are left for the graphics driver to fill in.

Name

`drm_edid_to_sad` — extracts SADs from EDID

Synopsis

```
int drm_edid_to_sad (struct edid * edid, struct cea_sad ** sads);
```

Arguments

edid EDID to parse

sads pointer that will be set to the extracted SADs

Description

Looks for CEA EDID block and extracts SADs (Short Audio Descriptors) from it.

Note

The returned pointer needs to be freed using `kfree`.

Return

The number of found SADs or negative number on error.

Name

`drm_edid_to_speaker_allocation` — extracts Speaker Allocation Data Blocks from EDID

Synopsis

```
int drm_edid_to_speaker_allocation (struct edid * edid, u8 ** sadb);
```

Arguments

edid EDID to parse

sadb pointer to the speaker block

Description

Looks for CEA EDID block and extracts the Speaker Allocation Data Block from it.

Note

The returned pointer needs to be freed using `kfree`.

Return

The number of found Speaker Allocation Blocks or negative number on error.

Name

`drm_av_sync_delay` — compute the HDMI/DP sink audio-video sync delay

Synopsis

```
int drm_av_sync_delay (struct drm_connector * connector, const struct
drm_display_mode * mode);
```

Arguments

connector connector associated with the HDMI/DP sink

mode the display mode

Return

The HDMI/DP sink's audio-video sync delay in milliseconds or 0 if the sink doesn't support audio or video.

Name

`drm_select_eld` — select one ELD from multiple HDMI/DP sinks

Synopsis

```
struct drm_connector * drm_select_eld (struct drm_encoder * encoder);
```

Arguments

encoder the encoder just changed display mode

Description

It's possible for one encoder to be associated with multiple HDMI/DP sinks. The policy is now hard coded to simply use the first HDMI/DP sink's ELD.

Return

The connector associated with the first HDMI/DP sink that has ELD attached to it.

Name

`drm_detect_hdmi_monitor` — detect whether monitor is HDMI

Synopsis

```
bool drm_detect_hdmi_monitor (struct edid * edid);
```

Arguments

edid monitor EDID information

Description

Parse the CEA extension according to CEA-861-B.

Return

True if the monitor is HDMI, false if not or unknown.

Name

`drm_detect_monitor_audio` — check monitor audio capability

Synopsis

```
bool drm_detect_monitor_audio (struct edid * edid);
```

Arguments

edid EDID block to scan

Description

Monitor should have CEA extension block. If monitor has 'basic audio', but no CEA audio blocks, it's 'basic audio' only. If there is any audio extension block and supported audio format, assume at least 'basic audio' support, even if 'basic audio' is not defined in EDID.

Return

True if the monitor supports audio, false otherwise.

Name

`drm_rgb_quant_range_selectable` — is RGB quantization range selectable?

Synopsis

```
bool drm_rgb_quant_range_selectable (struct edid * edid);
```

Arguments

edid EDID block to scan

Description

Check whether the monitor reports the RGB quantization range selection as supported. The AVI infoframe can then be used to inform the monitor which quantization range (full or limited) is used.

Return

True if the RGB quantization range is selectable, false otherwise.

Name

`drm_add_edid_modes` — add modes from EDID data, if available

Synopsis

```
int drm_add_edid_modes (struct drm_connector * connector, struct edid
* edid);
```

Arguments

connector connector we're probing

edid EDID data

Description

Add the specified modes to the connector's mode list.

Return

The number of modes added or 0 if we couldn't find any.

Name

`drm_add_modes_noedid` — add modes for the connectors without EDID

Synopsis

```
int drm_add_modes_noedid (struct drm_connector * connector, int hdisplay, int vdisplay);
```

Arguments

connector connector we're probing

hdisplay the horizontal display limit

vdisplay the vertical display limit

Description

Add the specified modes to the connector's mode list. Only when the `hdisplay/vdisplay` is not beyond the given limit, it will be added.

Return

The number of modes added or 0 if we couldn't find any.

Name

`drm_set_preferred_mode` — Sets the preferred mode of a connector

Synopsis

```
void drm_set_preferred_mode (struct drm_connector * connector, int  
hpref, int vpref);
```

Arguments

connector connector whose mode list should be processed

hpref horizontal resolution of preferred mode

vpref vertical resolution of preferred mode

Description

Marks a mode as preferred if it matches the resolution specified by *hpref* and *vpref*.

Name

`drm_hdmi_avi_infoframe_from_display_mode` — fill an HDMI AVI infoframe with data from a DRM display mode

Synopsis

```
int drm_hdmi_avi_infoframe_from_display_mode (struct hdmi_avi_infoframe
* frame, const struct drm_display_mode * mode);
```

Arguments

frame HDMI AVI infoframe

mode DRM display mode

Return

0 on success or a negative error code on failure.

Name

`drm_hdmi_vendor_infoframe_from_display_mode` — fill an HDMI infoframe with data from a DRM display mode

Synopsis

```
int drm_hdmi_vendor_infoframe_from_display_mode (struct hdmi_vendor_infoframe * frame, const struct drm_display_mode * mode);
```

Arguments

frame HDMI vendor infoframe

mode DRM display mode

Description

Note that there's is a need to send HDMI vendor infoframes only when using a 4k or stereoscopic 3D mode. So when giving any other mode as input this function will return `-EINVAL`, error that can be safely ignored.

Return

0 on success or a negative error code on failure.

Rectangle Utilities Reference

Utility functions to help manage rectangular areas for clipping, scaling, etc. calculations.

Name

struct `drm_rect` — two dimensional rectangle

Synopsis

```
struct drm_rect {  
    int x1;  
    int y1;  
    int x2;  
    int y2;  
};
```

Members

`x1` horizontal starting coordinate (inclusive)
`y1` vertical starting coordinate (inclusive)
`x2` horizontal ending coordinate (exclusive)
`y2` vertical ending coordinate (exclusive)

Name

`drm_rect_adjust_size` — adjust the size of the rectangle

Synopsis

```
void drm_rect_adjust_size (struct drm_rect * r, int dw, int dh);
```

Arguments

r rectangle to be adjusted

dw horizontal adjustment

dh vertical adjustment

Description

Change the size of rectangle *r* by *dw* in the horizontal direction, and by *dh* in the vertical direction, while keeping the center of *r* stationary.

Positive *dw* and *dh* increase the size, negative values decrease it.

Name

`drm_rect_translate` — translate the rectangle

Synopsis

```
void drm_rect_translate (struct drm_rect * r, int dx, int dy);
```

Arguments

r rectangle to be translated

dx horizontal translation

dy vertical translation

Description

Move rectangle *r* by *dx* in the horizontal direction, and by *dy* in the vertical direction.

Name

`drm_rect_downscale` — downscale a rectangle

Synopsis

```
void drm_rect_downscale (struct drm_rect * r, int horz, int vert);
```

Arguments

r rectangle to be downscaled

horz horizontal downscale factor

vert vertical downscale factor

Description

Divide the coordinates of rectangle *r* by *horz* and *vert*.

Name

`drm_rect_width` — determine the rectangle width

Synopsis

```
int drm_rect_width (const struct drm_rect * r);
```

Arguments

r rectangle whose width is returned

RETURNS

The width of the rectangle.

Name

`drm_rect_height` — determine the rectangle height

Synopsis

```
int drm_rect_height (const struct drm_rect * r);
```

Arguments

r rectangle whose height is returned

RETURNS

The height of the rectangle.

Name

`drm_rect_visible` — determine if the the rectangle is visible

Synopsis

```
bool drm_rect_visible (const struct drm_rect * r);
```

Arguments

r rectangle whose visibility is returned

RETURNS

`true` if the rectangle is visible, `false` otherwise.

Name

`drm_rect_equals` — determine if two rectangles are equal

Synopsis

```
bool drm_rect_equals (const struct drm_rect * r1, const struct drm_rect  
* r2);
```

Arguments

r1 first rectangle

r2 second rectangle

RETURNS

true if the rectangles are equal, false otherwise.

Name

`drm_rect_intersect` — intersect two rectangles

Synopsis

```
bool drm_rect_intersect (struct drm_rect * r1, const struct drm_rect  
* r2);
```

Arguments

r1 first rectangle

r2 second rectangle

Description

Calculate the intersection of rectangles *r1* and *r2*. *r1* will be overwritten with the intersection.

RETURNS

`true` if rectangle *r1* is still visible after the operation, `false` otherwise.

Name

`drm_rect_clip_scaled` — perform a scaled clip operation

Synopsis

```
bool drm_rect_clip_scaled (struct drm_rect * src, struct drm_rect * dst,  
const struct drm_rect * clip, int hscale, int vscale);
```

Arguments

<i>src</i>	source window rectangle
<i>dst</i>	destination window rectangle
<i>clip</i>	clip rectangle
<i>hscale</i>	horizontal scaling factor
<i>vscale</i>	vertical scaling factor

Description

Clip rectangle *dst* by rectangle *clip*. Clip rectangle *src* by the same amounts multiplied by *hscale* and *vscale*.

RETURNS

true if rectangle *dst* is still visible after being clipped, false otherwise

Name

`drm_rect_calc_hscale` — calculate the horizontal scaling factor

Synopsis

```
int drm_rect_calc_hscale (const struct drm_rect * src, const struct
drm_rect * dst, int min_hscale, int max_hscale);
```

Arguments

src source window rectangle

dst destination window rectangle

min_hscale minimum allowed horizontal scaling factor

max_hscale maximum allowed horizontal scaling factor

Description

Calculate the horizontal scaling factor as $(src \text{ width}) / (dst \text{ width})$.

RETURNS

The horizontal scaling factor, or `errno` of out of limits.

Name

`drm_rect_calc_vscale` — calculate the vertical scaling factor

Synopsis

```
int drm_rect_calc_vscale (const struct drm_rect * src, const struct
drm_rect * dst, int min_vscale, int max_vscale);
```

Arguments

<i>src</i>	source window rectangle
<i>dst</i>	destination window rectangle
<i>min_vscale</i>	minimum allowed vertical scaling factor
<i>max_vscale</i>	maximum allowed vertical scaling factor

Description

Calculate the vertical scaling factor as $(src \text{ height}) / (dst \text{ height})$.

RETURNS

The vertical scaling factor, or `errno` of out of limits.

Name

`drm_rect_calc_hscale_relaxed` — calculate the horizontal scaling factor

Synopsis

```
int drm_rect_calc_hscale_relaxed (struct drm_rect * src, struct drm_rect  
* dst, int min_hscale, int max_hscale);
```

Arguments

<i>src</i>	source window rectangle
<i>dst</i>	destination window rectangle
<i>min_hscale</i>	minimum allowed horizontal scaling factor
<i>max_hscale</i>	maximum allowed horizontal scaling factor

Description

Calculate the horizontal scaling factor as $(src \text{ width}) / (dst \text{ width})$.

If the calculated scaling factor is below *min_vscale*, decrease the height of rectangle *dst* to compensate.

If the calculated scaling factor is above *max_vscale*, decrease the height of rectangle *src* to compensate.

RETURNS

The horizontal scaling factor.

Name

`drm_rect_calc_vscale_relaxed` — calculate the vertical scaling factor

Synopsis

```
int drm_rect_calc_vscale_relaxed (struct drm_rect * src, struct drm_rect  
* dst, int min_vscale, int max_vscale);
```

Arguments

<i>src</i>	source window rectangle
<i>dst</i>	destination window rectangle
<i>min_vscale</i>	minimum allowed vertical scaling factor
<i>max_vscale</i>	maximum allowed vertical scaling factor

Description

Calculate the vertical scaling factor as $(src \text{ height}) / (dst \text{ height})$.

If the calculated scaling factor is below *min_vscale*, decrease the height of rectangle *dst* to compensate.

If the calculated scaling factor is above *max_vscale*, decrease the height of rectangle *src* to compensate.

RETURNS

The vertical scaling factor.

Name

`drm_rect_debug_print` — print the rectangle information

Synopsis

```
void drm_rect_debug_print (const char * prefix, const struct drm_rect  
* r, bool fixed_point);
```

Arguments

<i>prefix</i>	prefix string
<i>r</i>	rectangle to print
<i>fixed_point</i>	rectangle is in 16.16 fixed point format

Name

`drm_rect_rotate` — Rotate the rectangle

Synopsis

```
void drm_rect_rotate (struct drm_rect * r, int width, int height, unsigned int rotation);
```

Arguments

<i>r</i>	rectangle to be rotated
<i>width</i>	Width of the coordinate space
<i>height</i>	Height of the coordinate space
<i>rotation</i>	Transformation to be applied

Description

Apply *rotation* to the coordinates of rectangle *r*.

width and *height* combined with *rotation* define the location of the new origin.

width corresponds to the horizontal and *height* to the vertical axis of the untransformed coordinate space.

Name

`drm_rect_rotate_inv` — Inverse rotate the rectangle

Synopsis

```
void drm_rect_rotate_inv (struct drm_rect * r, int width, int height,  
unsigned int rotation);
```

Arguments

<i>r</i>	rectangle to be rotated
<i>width</i>	Width of the coordinate space
<i>height</i>	Height of the coordinate space
<i>rotation</i>	Transformation whose inverse is to be applied

Description

Apply the inverse of *rotation* to the coordinates of rectangle *r*.

width and *height* combined with *rotation* define the location of the new origin.

width corresponds to the horizontal and *height* to the vertical axis of the original untransformed coordinate space, so that you never have to flip them when doing a rotation and its inverse. That is, if you do:

```
drm_rotate(r, width, height, rotation);  
drm_rotate_inv(r, width, height, rotation);
```

you will always get back the original rectangle.

Flip-work Helper Reference

Util to queue up work to run from work-queue context after flip/vblank. Typically this can be used to defer unref of framebuffer's, cursor bo's, etc until after vblank. The APIs are all thread-safe. Moreover, `drm_flip_work_queue_task` and `drm_flip_work_queue` can be called in atomic context.

Name

struct drm_flip_task — flip work task

Synopsis

```
struct drm_flip_task {  
    struct list_head node;  
    void * data;  
};
```

Members

node list entry element

data data to pass to work->func

Name

struct `drm_flip_work` — flip work queue

Synopsis

```
struct drm_flip_work {  
    const char * name;  
    drm_flip_func_t func;  
    struct work_struct worker;  
    struct list_head queued;  
    struct list_head committed;  
    spinlock_t lock;  
};
```

Members

<code>name</code>	debug name
<code>func</code>	callback fxn called for each committed item
<code>worker</code>	worker which calls <i>func</i>
<code>queued</code>	queued tasks
<code>committed</code>	committed tasks
<code>lock</code>	lock to access queued and committed lists

Name

`drm_flip_work_allocate_task` — allocate a flip-work task

Synopsis

```
struct drm_flip_task * drm_flip_work_allocate_task (void * data, gfp_t  
flags);
```

Arguments

data data associated to the task

flags allocator flags

Description

Allocate a `drm_flip_task` object and attach private data to it.

Name

`drm_flip_work_queue_task` — queue a specific task

Synopsis

```
void drm_flip_work_queue_task (struct drm_flip_work * work, struct dr-  
m_flip_task * task);
```

Arguments

work the flip-work

task the task to handle

Description

Queues task, that will later be run (passed back to `drm_flip_func_t` func) on a work queue after `drm_flip_work_commit` is called.

Name

`drm_flip_work_queue` — queue work

Synopsis

```
void drm_flip_work_queue (struct drm_flip_work * work, void * val);
```

Arguments

work the flip-work

val the value to queue

Description

Queues work, that will later be run (passed back to `drm_flip_func_t` func) on a work queue after `drm_flip_work_commit` is called.

Name

`drm_flip_work_commit` — commit queued work

Synopsis

```
void  drm_flip_work_commit (struct  drm_flip_work  *  work,  struct
workqueue_struct  *  wq);
```

Arguments

work the flip-work

wq the work-queue to run the queued work on

Description

Trigger work previously queued by `drm_flip_work_queue` to run on a workqueue. The typical usage would be to queue work (via `drm_flip_work_queue`) at any point (from vblank irq and/or prior), and then from vblank irq commit the queued work.

Name

`drm_flip_work_init` — initialize flip-work

Synopsis

```
void drm_flip_work_init (struct drm_flip_work * work, const char * name,  
drm_flip_func_t func);
```

Arguments

work the flip-work to initialize

name debug name

func the callback work function

Description

Initializes/allocates resources for the flip-work

Name

`drm_flip_work_cleanup` — cleans up flip-work

Synopsis

```
void drm_flip_work_cleanup (struct drm_flip_work * work);
```

Arguments

work the flip-work to cleanup

Description

Destroy resources allocated for the flip-work

HDMI Infoframes Helper Reference

Strictly speaking this is not a DRM helper library but generally useable by any driver interfacing with HDMI outputs like v4l or alsa drivers. But it nicely fits into the overall topic of mode setting helper libraries and hence is also included here.

Name

union hdmi_infoframe — overall union of all abstract infoframe representations

Synopsis

```
union hdmi_infoframe {
    struct hdmi_any_infoframe any;
    struct hdmi_avi_infoframe avi;
    struct hdmi_spd_infoframe spd;
    union hdmi_vendor_any_infoframe vendor;
    struct hdmi_audio_infoframe audio;
};
```

Members

any	generic infoframe
avi	avi infoframe
spd	spd infoframe
vendor	union of all vendor infoframes
audio	audio infoframe

Description

This is used by the generic pack function. This works since all infoframes have the same header which also indicates which type of infoframe should be packed.

Name

`hdmi_avi_infoframe_init` — initialize an HDMI AVI infoframe

Synopsis

```
int hdmi_avi_infoframe_init (struct hdmi_avi_infoframe * frame);
```

Arguments

frame HDMI AVI infoframe

Description

Returns 0 on success or a negative error code on failure.

Name

`hdmi_avi_infoframe_pack` — write HDMI AVI infoframe to binary buffer

Synopsis

```
ssize_t hdmi_avi_infoframe_pack (struct hdmi_avi_infoframe * frame, void  
* buffer, size_t size);
```

Arguments

frame HDMI AVI infoframe

buffer destination buffer

size size of buffer

Description

Packs the information contained in the *frame* structure into a binary representation that can be written into the corresponding controller registers. Also computes the checksum as required by section 5.3.5 of the HDMI 1.4 specification.

Returns the number of bytes packed into the binary buffer or a negative error code on failure.

Name

`hdmi_spd_infoframe_init` — initialize an HDMI SPD infoframe

Synopsis

```
int hdmi_spd_infoframe_init (struct hdmi_spd_infoframe * frame, const  
char * vendor, const char * product);
```

Arguments

frame HDMI SPD infoframe

vendor vendor string

product product string

Description

Returns 0 on success or a negative error code on failure.

Name

`hdmi_spd_infoframe_pack` — write HDMI SPD infoframe to binary buffer

Synopsis

```
ssize_t hdmi_spd_infoframe_pack (struct hdmi_spd_infoframe * frame, void  
* buffer, size_t size);
```

Arguments

frame HDMI SPD infoframe

buffer destination buffer

size size of buffer

Description

Packs the information contained in the *frame* structure into a binary representation that can be written into the corresponding controller registers. Also computes the checksum as required by section 5.3.5 of the HDMI 1.4 specification.

Returns the number of bytes packed into the binary buffer or a negative error code on failure.

Name

`hdmi_audio_infoframe_init` — initialize an HDMI audio infoframe

Synopsis

```
int hdmi_audio_infoframe_init (struct hdmi_audio_infoframe * frame);
```

Arguments

frame HDMI audio infoframe

Description

Returns 0 on success or a negative error code on failure.

Name

`hdmi_audio_infoframe_pack` — write HDMI audio infoframe to binary buffer

Synopsis

```
ssize_t hdmi_audio_infoframe_pack (struct hdmi_audio_infoframe * frame,  
void * buffer, size_t size);
```

Arguments

frame HDMI audio infoframe

buffer destination buffer

size size of buffer

Description

Packs the information contained in the *frame* structure into a binary representation that can be written into the corresponding controller registers. Also computes the checksum as required by section 5.3.5 of the HDMI 1.4 specification.

Returns the number of bytes packed into the binary buffer or a negative error code on failure.

Name

`hdmi_vendor_infoframe_init` — initialize an HDMI vendor infoframe

Synopsis

```
int hdmi_vendor_infoframe_init (struct hdmi_vendor_infoframe * frame);
```

Arguments

frame HDMI vendor infoframe

Description

Returns 0 on success or a negative error code on failure.

Name

`hdmi_vendor_infoframe_pack` — write a HDMI vendor infoframe to binary buffer

Synopsis

```
ssize_t hdmi_vendor_infoframe_pack (struct hdmi_vendor_infoframe *  
frame, void * buffer, size_t size);
```

Arguments

frame HDMI infoframe

buffer destination buffer

size size of buffer

Description

Packs the information contained in the *frame* structure into a binary representation that can be written into the corresponding controller registers. Also computes the checksum as required by section 5.3.5 of the HDMI 1.4 specification.

Returns the number of bytes packed into the binary buffer or a negative error code on failure.

Name

`hdmi_infoframe_pack` — write a HDMI infoframe to binary buffer

Synopsis

```
ssize_t hdmi_infoframe_pack (union hdmi_infoframe * frame, void *  
buffer, size_t size);
```

Arguments

frame HDMI infoframe

buffer destination buffer

size size of buffer

Description

Packs the information contained in the *frame* structure into a binary representation that can be written into the corresponding controller registers. Also computes the checksum as required by section 5.3.5 of the HDMI 1.4 specification.

Returns the number of bytes packed into the binary buffer or a negative error code on failure.

Name

`hdmi_infoframe_log` — log info of HDMI infoframe

Synopsis

```
void hdmi_infoframe_log (const char * level, struct device * dev, union  
hdmi_infoframe * frame);
```

Arguments

level logging level

dev device

frame HDMI infoframe

Name

`hdmi_infoframe_unpack` — unpack binary buffer to a HDMI infoframe

Synopsis

```
int hdmi_infoframe_unpack (union hdmi_infoframe * frame, void * buffer);
```

Arguments

frame HDMI infoframe

buffer source buffer

Description

Unpacks the information contained in binary buffer *buffer* into a structured *frame* of a HDMI infoframe. Also verifies the checksum as required by section 5.3.5 of the HDMI 1.4 specification.

Returns 0 on success or a negative error code on failure.

Plane Helper Reference

Name

`drm_plane_helper_check_update` — Check plane update for validity

Synopsis

```
int drm_plane_helper_check_update (struct drm_plane * plane, struct
drm_crtc * crtc, struct drm_framebuffer * fb, struct drm_rect * src,
struct drm_rect * dest, const struct drm_rect * clip, int min_scale, int
max_scale, bool can_position, bool can_update_disabled, bool * visible);
```

Arguments

<i>plane</i>	plane object to update
<i>crtc</i>	owning CRTC of owning plane
<i>fb</i>	framebuffer to flip onto plane
<i>src</i>	source coordinates in 16.16 fixed point
<i>dest</i>	integer destination coordinates
<i>clip</i>	integer clipping coordinates
<i>min_scale</i>	minimum <i>src:dest</i> scaling factor in 16.16 fixed point
<i>max_scale</i>	maximum <i>src:dest</i> scaling factor in 16.16 fixed point
<i>can_position</i>	is it legal to position the plane such that it doesn't cover the entire crtc? This will generally only be false for primary planes.
<i>can_update_disabled</i>	can the plane be updated while the crtc is disabled?
<i>visible</i>	output parameter indicating whether plane is still visible after clipping

Description

Checks that a desired plane update is valid. Drivers that provide their own plane handling rather than helper-provided implementations may still wish to call this function to avoid duplication of error checking code.

RETURNS

Zero if update appears valid, error code on failure

Name

`drm_primary_helper_update` — Helper for primary plane update

Synopsis

```
int drm_primary_helper_update (struct drm_plane * plane, struct drm_crtc
* crtc, struct drm_framebuffer * fb, int crtc_x, int crtc_y, unsigned int
crtc_w, unsigned int crtc_h, uint32_t src_x, uint32_t src_y, uint32_t
src_w, uint32_t src_h);
```

Arguments

<i>plane</i>	plane object to update
<i>crtc</i>	owning CRTC of owning plane
<i>fb</i>	framebuffer to flip onto plane
<i>crtc_x</i>	x offset of primary plane on crtc
<i>crtc_y</i>	y offset of primary plane on crtc
<i>crtc_w</i>	width of primary plane rectangle on crtc
<i>crtc_h</i>	height of primary plane rectangle on crtc
<i>src_x</i>	x offset of <i>fb</i> for panning
<i>src_y</i>	y offset of <i>fb</i> for panning
<i>src_w</i>	width of source rectangle in <i>fb</i>
<i>src_h</i>	height of source rectangle in <i>fb</i>

Description

Provides a default plane update handler for primary planes. This handler is called in response to a userspace SetPlane operation on the plane with a non-NULL framebuffer. We call the driver's modeset handler to update the framebuffer.

SetPlane on a primary plane of a disabled CRTC is not supported, and will return an error.

Note that we make some assumptions about hardware limitations that may not be true for all hardware -- 1) Primary plane cannot be repositioned. 2) Primary plane cannot be scaled. 3) Primary plane must cover the entire CRTC. 4) Subpixel positioning is not supported. Drivers for hardware that don't have these restrictions can provide their own implementation rather than using this helper.

RETURNS

Zero on success, error code on failure

Name

`drm_primary_helper_disable` — Helper for primary plane disable

Synopsis

```
int drm_primary_helper_disable (struct drm_plane * plane);
```

Arguments

plane plane to disable

Description

Provides a default plane disable handler for primary planes. This handler is called in response to a userspace SetPlane operation on the plane with a NULL framebuffer parameter. It unconditionally fails the disable call with -EINVAL the only way to disable the primary plane without driver support is to disable the entire CRTC. Which does not match the plane ->disable hook.

Note that some hardware may be able to disable the primary plane without disabling the whole CRTC. Drivers for such hardware should provide their own disable handler that disables just the primary plane (and they'll likely need to provide their own update handler as well to properly re-enable a disabled primary plane).

RETURNS

Unconditionally returns -EINVAL.

Name

`drm_primary_helper_destroy` — Helper for primary plane destruction

Synopsis

```
void drm_primary_helper_destroy (struct drm_plane * plane);
```

Arguments

plane plane to destroy

Description

Provides a default plane destroy handler for primary planes. This handler is called during CRTC destruction. We disable the primary plane, remove it from the DRM plane list, and deallocate the plane structure.

Name

`drm_crtc_init` — Legacy CRTC initialization function

Synopsis

```
int drm_crtc_init (struct drm_device * dev, struct drm_crtc * crtc,  
const struct drm_crtc_funcs * funcs);
```

Arguments

dev DRM device

crtc CRTC object to init

funcs callbacks for the new CRTC

Description

Initialize a CRTC object with a default helper-provided primary plane and no cursor plane.

Returns

Zero on success, error code on failure.

Name

`drm_plane_helper_update` — Transitional helper for plane update

Synopsis

```
int drm_plane_helper_update (struct drm_plane * plane, struct drm_crtc *  
crtc, struct drm_framebuffer * fb, int crtc_x, int crtc_y, unsigned int  
crtc_w, unsigned int crtc_h, uint32_t src_x, uint32_t src_y, uint32_t  
src_w, uint32_t src_h);
```

Arguments

<i>plane</i>	plane object to update
<i>crtc</i>	owning CRTC of owning plane
<i>fb</i>	framebuffer to flip onto plane
<i>crtc_x</i>	x offset of primary plane on crtc
<i>crtc_y</i>	y offset of primary plane on crtc
<i>crtc_w</i>	width of primary plane rectangle on crtc
<i>crtc_h</i>	height of primary plane rectangle on crtc
<i>src_x</i>	x offset of <i>fb</i> for panning
<i>src_y</i>	y offset of <i>fb</i> for panning
<i>src_w</i>	width of source rectangle in <i>fb</i>
<i>src_h</i>	height of source rectangle in <i>fb</i>

Description

Provides a default plane update handler using the atomic plane update functions. It is fully left to the driver to check plane constraints and handle corner-cases like a fully occluded or otherwise invisible plane.

This is useful for piecewise transitioning of a driver to the atomic helpers.

RETURNS

Zero on success, error code on failure

Name

`drm_plane_helper_disable` — Transitional helper for plane disable

Synopsis

```
int drm_plane_helper_disable (struct drm_plane * plane);
```

Arguments

plane plane to disable

Description

Provides a default plane disable handler using the atomic plane update functions. It is fully left to the driver to check plane constraints and handle corner-cases like a fully occluded or otherwise invisible plane.

This is useful for piecewise transitioning of a driver to the atomic helpers.

RETURNS

Zero on success, error code on failure

This helper library has two parts. The first part has support to implement primary plane support on top of the normal CRTC configuration interface. Since the legacy `->set_config` interface ties the primary plane together with the CRTC state this does not allow userspace to disable the primary plane itself. To avoid too much duplicated code use `drm_plane_helper_check_update` which can be used to enforce the same restrictions as primary planes had thus. The default primary plane only expose XRGB8888 and ARGB8888 as valid pixel formats for the attached framebuffer.

Drivers are highly recommended to implement proper support for primary planes, and newly merged drivers must not rely upon these transitional helpers.

The second part also implements transitional helpers which allow drivers to gradually switch to the atomic helper infrastructure for plane updates. Once that switch is complete drivers shouldn't use these any longer, instead using the proper legacy implementations for update and disable plane hooks provided by the atomic helpers.

Again drivers are strongly urged to switch to the new interfaces.

The plane helpers share the function table structures with other helpers, specifically also the atomic helpers. See `struct drm_plane_helper_funcs` for the details.

Tile group

Tile groups are used to represent tiled monitors with a unique integer identifier. Tiled monitors using DisplayID v1.3 have a unique 8-byte handle, we store this in a tile group, so we have a common identifier for all tiles in a monitor group.

Bridges

Overview

struct `drm_bridge` represents a device that hangs on to an encoder. These are handy when a regular `drm_encoder` entity isn't enough to represent the entire encoder chain.

A bridge is always attached to a single `drm_encoder` at a time, but can be either connected to it directly, or through an intermediate bridge:

encoder ---> bridge B ---> bridge A

Here, the output of the encoder feeds to bridge B, and that further feeds to bridge A.

The driver using the bridge is responsible to make the associations between the encoder and bridges. Once these links are made, the bridges will participate along with encoder functions to perform `mode_set/enable/disable` through the ops provided in `drm_bridge_funcs`.

`drm_bridge`, like `drm_panel`, aren't `drm_mode_object` entities like planes, CRTC's, encoders or connectors and hence are not visible to userspace. They just provide additional hooks to get the desired output at the end of the encoder chain.

Bridges can also be chained up using the next pointer in struct `drm_bridge`.

Both legacy CRTC helpers and the new atomic modeset helpers support bridges.

Default bridge callback sequence

The `drm_bridge_funcs` ops are populated by the bridge driver. The DRM internals (atomic and CRTC helpers) use the helpers defined in `drm_bridge.c`. These helpers call a specific `drm_bridge_funcs` op for all the bridges during encoder configuration.

For detailed specification of the bridge callbacks see `drm_bridge_funcs`.

Name

`drm_bridge_add` — add the given bridge to the global bridge list

Synopsis

```
int drm_bridge_add (struct drm_bridge * bridge);
```

Arguments

bridge bridge control structure

RETURNS

Unconditionally returns Zero.

Name

`drm_bridge_remove` — remove the given bridge from the global bridge list

Synopsis

```
void drm_bridge_remove (struct drm_bridge * bridge);
```

Arguments

bridge bridge control structure

Name

`drm_bridge_attach` — associate given bridge to our DRM device

Synopsis

```
int drm_bridge_attach (struct drm_device * dev, struct drm_bridge *  
bridge);
```

Arguments

dev DRM device

bridge bridge control structure

Description

called by a kms driver to link one of our encoder/bridge to the given bridge.

Note that setting up links between the bridge and our encoder/bridge objects needs to be handled by the kms driver itself

RETURNS

Zero on success, error code on failure

Name

`drm_bridge_mode_fixup` — fixup proposed mode for all bridges in the encoder chain

Synopsis

```
bool drm_bridge_mode_fixup (struct drm_bridge * bridge, const struct  
drm_display_mode * mode, struct drm_display_mode * adjusted_mode);
```

Arguments

<i>bridge</i>	bridge control structure
<i>mode</i>	desired mode to be set for the bridge
<i>adjusted_mode</i>	updated mode that works for this bridge

Description

Calls `->mode_fixup` `drm_bridge_funcs` op for all the bridges in the encoder chain, starting from the first bridge to the last.

Note

the bridge passed should be the one closest to the encoder

RETURNS

true on success, false on failure

Name

`drm_bridge_disable` — calls `->disable` `drm_bridge_funcs` op for all bridges in the encoder chain.

Synopsis

```
void drm_bridge_disable (struct drm_bridge * bridge);
```

Arguments

bridge bridge control structure

Description

Calls `->disable` `drm_bridge_funcs` op for all the bridges in the encoder chain, starting from the last bridge to the first. These are called before calling the encoder's prepare op.

Note

the bridge passed should be the one closest to the encoder

Name

`drm_bridge_post_disable` — calls `->post_disable` `drm_bridge_funcs` op for all bridges in the encoder chain.

Synopsis

```
void drm_bridge_post_disable (struct drm_bridge * bridge);
```

Arguments

bridge bridge control structure

Description

Calls `->post_disable` `drm_bridge_funcs` op for all the bridges in the encoder chain, starting from the first bridge to the last. These are called after completing the encoder's prepare op.

Note

the bridge passed should be the one closest to the encoder

Name

`drm_bridge_mode_set` — set proposed mode for all bridges in the encoder chain

Synopsis

```
void drm_bridge_mode_set (struct drm_bridge * bridge, struct drm_display_mode * mode, struct drm_display_mode * adjusted_mode);
```

Arguments

<i>bridge</i>	bridge control structure
<i>mode</i>	desired mode to be set for the bridge
<i>adjusted_mode</i>	updated mode that works for this bridge

Description

Calls `->mode_set` `drm_bridge_funcs` op for all the bridges in the encoder chain, starting from the first bridge to the last.

Note

the bridge passed should be the one closest to the encoder

Name

`drm_bridge_pre_enable` — calls `->pre_enable` `drm_bridge_funcs` op for all bridges in the encoder chain.

Synopsis

```
void drm_bridge_pre_enable (struct drm_bridge * bridge);
```

Arguments

bridge bridge control structure

Description

Calls `->pre_enable` `drm_bridge_funcs` op for all the bridges in the encoder chain, starting from the last bridge to the first. These are called before calling the encoder's commit op.

Note

the bridge passed should be the one closest to the encoder

Name

`drm_bridge_enable` — calls `->enable` `drm_bridge_funcs` op for all bridges in the encoder chain.

Synopsis

```
void drm_bridge_enable (struct drm_bridge * bridge);
```

Arguments

bridge bridge control structure

Description

Calls `->enable` `drm_bridge_funcs` op for all the bridges in the encoder chain, starting from the first bridge to the last. These are called after completing the encoder's commit op.

Note that the bridge passed should be the one closest to the encoder

Name

`of_drm_find_bridge` — find the bridge corresponding to the device node in the global bridge list

Synopsis

```
struct drm_bridge * of_drm_find_bridge (struct device_node * np);
```

Arguments

np device node

RETURNS

`drm_bridge` control struct on success, NULL on failure

KMS Properties

Drivers may need to expose additional parameters to applications than those described in the previous sections. KMS supports attaching properties to CRTC's, connectors and planes and offers a userspace API to list, get and set the property values.

Properties are identified by a name that uniquely defines the property purpose, and store an associated value. For all property types except blob properties the value is a 64-bit unsigned integer.

KMS differentiates between properties and property instances. Drivers first create properties and then create and associate individual instances of those properties to objects. A property can be instantiated multiple times and associated with different objects. Values are stored in property instances, and all other property information are stored in the property and shared between all instances of the property.

Every property is created with a type that influences how the KMS core handles the property. Supported property types are

DRM_MODE_PROP_RANGE	Range properties report their minimum and maximum admissible values. The KMS core verifies that values set by application fit in that range.
DRM_MODE_PROP_ENUM	Enumerated properties take a numerical value that ranges from 0 to the number of enumerated values defined by the property minus one, and associate a free-formed string name to each value. Applications can retrieve the list of defined value-name pairs and use the numerical value to get and set property instance values.
DRM_MODE_PROP_BITMASK	Bitmask properties are enumeration properties that additionally restrict all enumerated values to the 0..63 range. Bitmask property instance values combine one or more of the enumerated bits defined by the property.
DRM_MODE_PROP_BLOB	<p>Blob properties store a binary blob without any format restriction. The binary blobs are created as KMS standalone objects, and blob property instance values store the ID of their associated blob object.</p> <p>Blob properties are only used for the connector EDID property and cannot be created by drivers.</p>

To create a property drivers call one of the following functions depending on the property type. All property creation functions take property flags and name, as well as type-specific arguments.

- `struct drm_property *drm_property_create_range(struct drm_device *dev, int flags`

```
const char *name,
uint64_t min, uint64_t max);
```

Create a range property with the given minimum and maximum values.

- `struct drm_property *drm_property_create_enum(struct drm_device *dev, int flags, const char *name, const struct drm_prop_enum_list *props, int num_values);`

Create an enumerated property. The *props* argument points to an array of *num_values* value-name pairs.

- `struct drm_property *drm_property_create_bitmask(struct drm_device *dev, int flags, const char *name, const struct drm_prop_enum_list *props, int num_values);`

Create a bitmask property. The *props* argument points to an array of *num_values* value-name pairs.

Properties can additionally be created as immutable, in which case they will be read-only for applications but can be modified by the driver. To create an immutable property drivers must set the `DRM_MODE_PROP_IMMUTABLE` flag at property creation time.

When no array of value-name pairs is readily available at property creation time for enumerated or range properties, drivers can create the property using the `drm_property_create` function and manually add enumeration value-name pairs by calling the `drm_property_add_enum` function. Care must be taken to properly specify the property type through the *flags* argument.

After creating properties drivers can attach property instances to CRTC, connector and plane objects by calling the `drm_object_attach_property`. The function takes a pointer to the target object, a pointer to the previously created property and an initial instance value.

Existing KMS Properties

The following table gives description of drm properties exposed by various modules/drivers.

Table 2.1.

Owner Module/Drivers	Group	Property Name	Type	Property Values	Object attached	Description/Restrictions
DRM	Generic	“rotation”	BITMASK	{ 0, "rotate-0" }, { 1, "rotate-90" }, { 2, "rotate-180" }, { 3, "rotate-270" }, { 4, "reflect-x" }, { 5, "reflect-y" }	CRTC, Plane	rotate-(degrees) rotates the image by the specified amount in degrees in counter clockwise direction. reflect-x and reflect-y reflects the image along the specified axis

					prior to rotation
Connector	“EDID”	BLOB IMMUTABLE	0	Connector	Contains id of edid blob ptr object.
	“DPMS”	ENUM	{ “On”, “Standby”, “Suspend”, “Off” }	Connector	Contains DPMS operation mode value.
	“PATH”	BLOB IMMUTABLE	0	Connector	Contains topology path to a connector.
	“TILE”	BLOB IMMUTABLE	0	Connector	Contains tiling information for a connector.
	“CRTC_ID”	OBJECT	DRM_MODE_OBJECT_CRTC	Connector	CRTC that connector is attached to (atomic)
Plane	“type”	ENUM IMMUTABLE	{ “Overlay”, “Primary”, “Cursor” }	Plane	Plane type
	“SRC_X”	RANGE	Min=0, Max=UINT_MAX	Plane	Scanout source x coordinate in 16.16 fixed point (atomic)
	“SRC_Y”	RANGE	Min=0, Max=UINT_MAX	Plane	Scanout source y coordinate in 16.16 fixed point (atomic)
	“SRC_W”	RANGE	Min=0, Max=UINT_MAX	Plane	Scanout source width in 16.16 fixed point (atomic)
	“SRC_H”	RANGE	Min=0, Max=UINT_MAX	Plane	Scanout source height in 16.16 fixed point (atomic)
	“CRTC_X”	SIGNED_RANGE	Min=INT_MIN, Max=INT_MAX	Plane	Scanout CRTC (destination)

				Max=IN-T_MAX		nation) x co-ordinate (atomic)
	“CRTC_Y”	SIGNED_RANGE	Min=INT_MIN, Max=INT_MAX	Plane		Scanout CRTC (destination) y coordinate (atomic)
	“CRTC_W”	RANGE	Min=0, Max=UINT_MAX	Plane		Scanout CRTC (destination) width (atomic)
	“CRTC_H”	RANGE	Min=0, Max=UINT_MAX	Plane		Scanout CRTC (destination) height (atomic)
	“FB_ID”	OBJECT	DRM_MODE_OBJECT_FB	Plane		Scanout framebuffer (atomic)
	“CRTC_ID”	OBJECT	DRM_MODE_OBJECT_CRTC	Plane		CRTC that plane is attached to (atomic)
DVI-I	“subconnector”	ENUM	{ “Unknown”, “DVI-D”, “DVI-A” }	Connector		TBD
	“select subconnector”	ENUM	{ “Automatic”, “DVI-D”, “DVI-A” }	Connector		TBD
TV	“subconnector”	ENUM	{ "Unknown", "Composite", "SVIDEO", "Component", "SCART" }	Connector		TBD
	“select subconnector”	ENUM	{ "Automatic", "Composite", "SVIDEO", "Component", "SCART" }	Connector		TBD
	“mode”	ENUM	{ "NTSC_M", "NTSC_J", "NTSC_443", "PAL_B" } etc.	Connector		TBD

		"left margin"	RANGE	Min=0, Max=100	Connector	TBD
		"right margin"	RANGE	Min=0, Max=100	Connector	TBD
		"top margin"	RANGE	Min=0, Max=100	Connector	TBD
		"bottom margin"	RANGE	Min=0, Max=100	Connector	TBD
		"brightness"	RANGE	Min=0, Max=100	Connector	TBD
		"contrast"	RANGE	Min=0, Max=100	Connector	TBD
		"flicker reduction"	RANGE	Min=0, Max=100	Connector	TBD
		"overscan"	RANGE	Min=0, Max=100	Connector	TBD
		"saturation"	RANGE	Min=0, Max=100	Connector	TBD
		"hue"	RANGE	Min=0, Max=100	Connector	TBD
	Virtual GPU	"suggested X"	RANGE	Min=0, Max=0xffffffff	Connector	property to suggest an X offset for a connector
		"suggested Y"	RANGE	Min=0, Max=0xffffffff	Connector	property to suggest an Y offset for a connector
	Optional	"scaling mode"	ENUM	{ "None", "Full", "Center", "Full aspect" }	Connector	TBD
		"aspect ratio"	ENUM	{ "None", "4:3", "16:9" }	Connector	DRM property to set aspect ratio from user space app. This enum is made generic to allow addition of custom aspect ratios.
		"dirty"	ENUM IMMUTABLE	{ "Off", "On", "Annotate" }	Connector	TBD
i915	Generic	"Broadcast RGB"	ENUM	{ "Automatic", "Full",	Connector	TBD

			"Limited 16:235" }		
	"audio"	ENUM	{ "force-dvi", "off", "auto", "on" }	Connector	TBD
SDVO-TV	"mode"	ENUM	{ "NTSC_M", "NTSC_J", "NTSC_443", "PAL_B" } etc.	Connector	TBD
	"left_margin"	RANGE	Min=0, Max=SDVO dependent	Connector	TBD
	"right_margin"	RANGE	Min=0, Max=SDVO dependent	Connector	TBD
	"top_margin"	RANGE	Min=0, Max=SDVO dependent	Connector	TBD
	"bottom_margin"	RANGE	Min=0, Max=SDVO dependent	Connector	TBD
	"hpos"	RANGE	Min=0, Max=SDVO dependent	Connector	TBD
	"vpos"	RANGE	Min=0, Max=SDVO dependent	Connector	TBD
	"contrast"	RANGE	Min=0, Max=SDVO dependent	Connector	TBD
	"saturation"	RANGE	Min=0, Max=SDVO dependent	Connector	TBD
	"hue"	RANGE	Min=0, Max=SDVO dependent	Connector	TBD
	"sharpness"	RANGE	Min=0, Max=SDVO dependent	Connector	TBD
	"flicker_filter"	RANGE	Min=0, Max=SDVO dependent	Connector	TBD
	"flicker_filter_adaptive"	RANGE	Min=0, Max=SDVO dependent	Connector	TBD

		"flicker_filter_2d"	RANGE	Min=0, Max=SDVO dependent	Connector	TBD
		"tv_chroma_filter"	RANGE	Min=0, Max=SDVO dependent	Connector	TBD
		"tv_luma_filter"	RANGE	Min=0, Max=SDVO dependent	Connector	TBD
		"dot_crawl"	RANGE	Min=0, Max=1	Connector	TBD
	SDVO-TV/LVDS	"brightness"	RANGE	Min=0, Max=SDVO dependent	Connector	TBD
CDV gma-500	Generic	"Broadcast RGB"	ENUM	{ "Full", "Limited 16:235" }	Connector	TBD
		"Broadcast RGB"	ENUM	{ "off", "auto", "on" }	Connector	TBD
Poulsbo	Generic	"backlight"	RANGE	Min=0, Max=100	Connector	TBD
	SDVO-TV	"mode"	ENUM	{ "NTSC_M", "NTSC_J", "NTSC_443", "PAL_B" } etc.	Connector	TBD
		"left_margin"	RANGE	Min=0, Max=SDVO dependent	Connector	TBD
		"right_margin"	RANGE	Min=0, Max=SDVO dependent	Connector	TBD
		"top_margin"	RANGE	Min=0, Max=SDVO dependent	Connector	TBD
		"bottom_margin"	RANGE	Min=0, Max=SDVO dependent	Connector	TBD
		"hpos"	RANGE	Min=0, Max=SDVO dependent	Connector	TBD
		"vpos"	RANGE	Min=0, Max=SDVO dependent	Connector	TBD
		"contrast"	RANGE	Min=0, Max=SDVO dependent	Connector	TBD

		"saturation"	RANGE	Min=0, Max=SDVO dependent	Connector	TBD
		"hue"	RANGE	Min=0, Max=SDVO dependent	Connector	TBD
		"sharpness"	RANGE	Min=0, Max=SDVO dependent	Connector	TBD
		"flicker_filter"	RANGE	Min=0, Max=SDVO dependent	Connector	TBD
		"flicker_filter_adaptive"	RANGE	Min=0, Max=SDVO dependent	Connector	TBD
		"flicker_filter_2d"	RANGE	Min=0, Max=SDVO dependent	Connector	TBD
		"tv_chroma_filter"	RANGE	Min=0, Max=SDVO dependent	Connector	TBD
		"tv_luma_filter"	RANGE	Min=0, Max=SDVO dependent	Connector	TBD
		"dot_crawl"	RANGE	Min=0, Max=1	Connector	TBD
	SDVO-TV/LVDS	"brightness"	RANGE	Min=0, Max=SDVO dependent	Connector	TBD
armada	CRTC	"CSC_YUV"	ENUM	{ "Auto" , "CCIR601", "CCIR709" }	CRTC	TBD
		"CSC_RGB"	ENUM	{ "Auto", "Computer system", "Studio" }	CRTC	TBD
	Overlay	"colorkey"	RANGE	Min=0, Max=0xffffffff	Plane	TBD
		"colorkey_min"	RANGE	Min=0, Max=0xffffffff	Plane	TBD
		"colorkey_max"	RANGE	Min=0, Max=0xffffffff	Plane	TBD
		"colorkey_val"	RANGE	Min=0, Max=0xffffffff	Plane	TBD
		"colorkey_alpha"	RANGE	Min=0, Max=0xffffffff	Plane	TBD

		"col- orkey_mode"	ENUM	{ "disabled", "Y compo- nent", "U component" , "V compo- nent", "RGB", "R component", "G compo- nent", "B component" }	Plane	TBD
		"brightness"	RANGE	Min=0, Max=256 + 255	Plane	TBD
		"contrast"	RANGE	Min=0, Max=0x7fff	Plane	TBD
		"saturation"	RANGE	Min=0, Max=0x7fff	Plane	TBD
exynos	CRTC	"mode"	ENUM	{ "normal", "blank" }	CRTC	TBD
	Overlay	"zpos"	RANGE	Min=0, Max=MAX_PLANE-1	Plane	TBD
i2c/ ch7006_drv	Generic	"scale"	RANGE	Min=0, Max=2	Connector	TBD
	TV	"mode"	ENUM	{ "PAL", "PAL- M", "PAL- N"}, "PAL- Nc" , "PAL-60", "NTSC-M", "NTSC-J" }	Connector	TBD
nouveau	NV10 Over- lay	"colorkey"	RANGE	Min=0, Max=0x01ffffff	Plane	TBD
		"contrast"	RANGE	Min=0, Max=8192-1	Plane	TBD
		"brightness"	RANGE	Min=0, Max=1024	Plane	TBD
		"hue"	RANGE	Min=0, Max=359	Plane	TBD
		"saturation"	RANGE	Min=0, Max=8192-1	Plane	TBD
		"iturbt_709"	RANGE	Min=0, Max=1	Plane	TBD
	Nv04 Over- lay	"colorkey"	RANGE	Min=0, Max=0x01ffffff	Plane	TBD
		"brightness"	RANGE	Min=0, Max=1024	Plane	TBD

	Display	“dithering mode”	ENUM	{ "auto", "off", "on" }	Connector	TBD
		“dithering depth”	ENUM	{ "auto", "off", "on", "static 2x2", "dynamic 2x2", "temporal" }	Connector	TBD
		“underscan”	ENUM	{ "auto", "6 bpc", "8 bpc" }	Connector	TBD
		“underscan hborder”	RANGE	Min=0, Max=128	Connector	TBD
		“underscan vborder”	RANGE	Min=0, Max=128	Connector	TBD
		“vibrant hue”	RANGE	Min=0, Max=180	Connector	TBD
		“color vibrance”	RANGE	Min=0, Max=200	Connector	TBD
omap	Generic	“zorder”	RANGE	Min=0, Max=3	CRTC, Plane	TBD
qxl	Generic	“hot-plug_mode_update”	RANGE	Min=0, Max=1	Connector	TBD
radeon	DVI-I	“coherent”	RANGE	Min=0, Max=1	Connector	TBD
	DAC enable load detect	“load detection”	RANGE	Min=0, Max=1	Connector	TBD
	TV Standard	"tv standard"	ENUM	{ "ntsc", "pal", "pal-m", "pal-60", "ntsc-j", "scart-pal", "pal-cn", "secam" }	Connector	TBD
	legacy TMDS PLL detect	"tmids_pll"	ENUM	{ "driver", "bios" }	-	TBD
	Underscan	"underscan"	ENUM	{ "off", "on", "auto" }	Connector	TBD
		"underscan hborder"	RANGE	Min=0, Max=128	Connector	TBD
		"underscan vborder"	RANGE	Min=0, Max=128	Connector	TBD
	Audio	“audio”	ENUM	{ "off", "on", "auto" }	Connector	TBD

	FMT Dithering	"dither"	ENUM	{ "off", "on" }	Connector	TBD
rcar-du	Generic	"alpha"	RANGE	Min=0, Max=255	Plane	TBD
		"colorkey"	RANGE	Min=0, Max=0x01ffffff	Plane	TBD
		"zpos"	RANGE	Min=1, Max=7	Plane	TBD

Vertical Blanking

Vertical blanking plays a major role in graphics rendering. To achieve tear-free display, users must synchronize page flips and/or rendering to vertical blanking. The DRM API offers ioctls to perform page flips synchronized to vertical blanking and wait for vertical blanking.

The DRM core handles most of the vertical blanking management logic, which involves filtering out spurious interrupts, keeping race-free blanking counters, coping with counter wrap-around and resets and keeping use counts. It relies on the driver to generate vertical blanking interrupts and optionally provide a hardware vertical blanking counter. Drivers must implement the following operations.

- `int (*enable_vblank) (struct drm_device *dev, int crtc);`
`void (*disable_vblank) (struct drm_device *dev, int crtc);`

Enable or disable vertical blanking interrupts for the given CRTC.

- `u32 (*get_vblank_counter) (struct drm_device *dev, int crtc);`

Retrieve the value of the vertical blanking counter for the given CRTC. If the hardware maintains a vertical blanking counter its value should be returned. Otherwise drivers can use the `drm_vblank_count` helper function to handle this operation.

Drivers must initialize the vertical blanking handling core with a call to `drm_vblank_init` in their load operation. The function will set the struct `drm_device` `vblank_disable_allowed` field to 0. This will keep vertical blanking interrupts enabled permanently until the first mode set operation, where `vblank_disable_allowed` is set to 1. The reason behind this is not clear. Drivers can set the field to 1 after calling `drm_vblank_init` to make vertical blanking interrupts dynamically managed from the beginning.

Vertical blanking interrupts can be enabled by the DRM core or by drivers themselves (for instance to handle page flipping operations). The DRM core maintains a vertical blanking use count to ensure that the interrupts are not disabled while a user still needs them. To increment the use count, drivers call `drm_vblank_get`. Upon return vertical blanking interrupts are guaranteed to be enabled.

To decrement the use count drivers call `drm_vblank_put`. Only when the use count drops to zero will the DRM core disable the vertical blanking interrupts after a delay by scheduling a timer. The delay is accessible through the `vblankoffdelay` module parameter or the `drm_vblank_offdelay` global variable and expressed in milliseconds. Its default value is 5000 ms. Zero means never disable, and a negative value means disable immediately. Drivers may override the behaviour by setting the `drm_device` `vblank_disable_immediate` flag, which when set causes vblank interrupts to be disabled immediately regardless of the `drm_vblank_offdelay` value. The flag should only be set if there's a properly working hardware vblank counter present.

When a vertical blanking interrupt occurs drivers only need to call the `drm_handle_vblank` function to account for the interrupt.

Resources allocated by `drm_vblank_init` must be freed with a call to `drm_vblank_cleanup` in the driver unload operation handler.

Vertical Blanking and Interrupt Handling Functions Reference

Name

`drm_vblank_cleanup` — cleanup vblank support

Synopsis

```
void drm_vblank_cleanup (struct drm_device * dev);
```

Arguments

dev DRM device

Description

This function cleans up any resources allocated in `drm_vblank_init`.

Name

`drm_vblank_init` — initialize vblank support

Synopsis

```
int drm_vblank_init (struct drm_device * dev, unsigned int num_crtcs);
```

Arguments

dev DRM device

num_crtcs number of CRTC's supported by *dev*

Description

This function initializes vblank support for *num_crtcs* display pipelines.

Returns

Zero on success or a negative error code on failure.

Name

`drm_irq_install` — install IRQ handler

Synopsis

```
int drm_irq_install (struct drm_device * dev, int irq);
```

Arguments

dev DRM device

irq IRQ number to install the handler for

Description

Initializes the IRQ related data. Installs the handler, calling the driver `irq_preinstall` and `irq_postinstall` functions before and after the installation.

This is the simplified helper interface provided for drivers with no special needs. Drivers which need to install interrupt handlers for multiple interrupts must instead set `drm_device->irq_enabled` to signal the DRM core that vblank interrupts are available.

Returns

Zero on success or a negative error code on failure.

Name

`drm_irq_uninstall` — uninstall the IRQ handler

Synopsis

```
int drm_irq_uninstall (struct drm_device * dev);
```

Arguments

dev DRM device

Description

Calls the driver's `irq_uninstall` function and unregisters the IRQ handler. This should only be called by drivers which used `drm_irq_install` to set up their interrupt handler. Other drivers must only reset `drm_device->irq_enabled` to false.

Note that for kernel modesetting drivers it is a bug if this function fails. The sanity checks are only to catch buggy user modesetting drivers which call the same function through an `ioctl`.

Returns

Zero on success or a negative error code on failure.

Name

`drm_calc_timestamping_constants` — calculate vblank timestamp constants

Synopsis

```
void drm_calc_timestamping_constants (struct drm_crtc * crtc, const
struct drm_display_mode * mode);
```

Arguments

crtc `drm_crtc` whose timestamp constants should be updated.

mode display mode containing the scanout timings

Description

Calculate and store various constants which are later needed by vblank and swap-completion timestamping, e.g, by `drm_calc_vbltimestamp_from_scanoutpos`. They are derived from CRTC's true scanout timing, so they take things like panel scaling or other adjustments into account.

Name

`drm_calc_vbltimestamp_from_scanoutpos` — precise vblank timestamp helper

Synopsis

```
int drm_calc_vbltimestamp_from_scanoutpos (struct drm_device * dev, unsigned int pipe, int * max_error, struct timeval * vblank_time, unsigned flags, const struct drm_display_mode * mode);
```

Arguments

<i>dev</i>	DRM device
<i>pipe</i>	index of CRTC whose vblank timestamp to retrieve
<i>max_error</i>	Desired maximum allowable error in timestamps (nanosecs) On return contains true maximum error of timestamp
<i>vblank_time</i>	Pointer to struct timeval which should receive the timestamp
<i>flags</i>	Flags to pass to driver: 0 = Default, <code>DRM_CALLED_FROM_VBLIRQ</code> = If function is called from vbl IRQ handler
<i>mode</i>	mode which defines the scanout timings

Description

Implements calculation of exact vblank timestamps from given `drm_display_mode` timings and current video scanout position of a CRTC. This can be called from within `get_vblank_timestamp` implementation of a kms driver to implement the actual timestamping.

Should return timestamps conforming to the `OML_sync_control` OpenML extension specification. The timestamp corresponds to the end of the vblank interval, aka start of scanout of topmost-leftmost display pixel in the following video frame.

Requires support for optional `dev->driver->get_scanout_position` in kms driver, plus a bit of setup code to provide a `drm_display_mode` that corresponds to the true scanout timing.

The current implementation only handles standard video modes. It returns as no operation if a doublescan or interlaced video mode is active. Higher level code is expected to handle this.

Returns

Negative value on error, failure or if not supported in current

video mode

-EINVAL - Invalid CRTC. -EAGAIN - Temporary unavailable, e.g., called before initial modeset. -ENOTSUPP - Function not supported in current display mode. -EIO - Failed, e.g., due to failed scanout position query.

Returns or'ed positive status flags on success:

`DRM_VBLANKTIME_SCANOUTPOS_METHOD` - Signal this method used for timestamping. `DRM_VBLANKTIME_INVBL` - Timestamp taken while scanout was in vblank interval.

Name

`drm_vblank_count` — retrieve “cooked” vblank counter value

Synopsis

```
u32 drm_vblank_count (struct drm_device * dev, unsigned int pipe);
```

Arguments

dev DRM device

pipe index of CRTC for which to retrieve the counter

Description

Fetches the “cooked” vblank count value that represents the number of vblank events since the system was booted, including lost events due to modesetting activity.

This is the legacy version of `drm_crtc_vblank_count`.

Returns

The software vblank counter.

Name

`drm_crtc_vblank_count` — retrieve “cooked” vblank counter value

Synopsis

```
u32 drm_crtc_vblank_count (struct drm_crtc * crtc);
```

Arguments

crtc which counter to retrieve

Description

Fetches the “cooked” vblank count value that represents the number of vblank events since the system was booted, including lost events due to modesetting activity.

This is the native KMS version of `drm_vblank_count`.

Returns

The software vblank counter.

Name

`drm_vblank_count_and_time` — retrieve “cooked” vblank counter value and the system timestamp corresponding to that vblank counter value.

Synopsis

```
u32 drm_vblank_count_and_time (struct drm_device * dev, unsigned int
pipe, struct timeval * vblanktime);
```

Arguments

<i>dev</i>	DRM device
<i>pipe</i>	index of CRTC whose counter to retrieve
<i>vblanktime</i>	Pointer to struct timeval to receive the vblank timestamp.

Description

Fetches the “cooked” vblank count value that represents the number of vblank events since the system was booted, including lost events due to modesetting activity. Returns corresponding system timestamp of the time of the vblank interval that corresponds to the current vblank counter value.

This is the legacy version of `drm_crtc_vblank_count_and_time`.

Name

`drm_crtc_vblank_count_and_time` — retrieve “cooked” vblank counter value and the system timestamp corresponding to that vblank counter value

Synopsis

```
u32 drm_crtc_vblank_count_and_time (struct drm_crtc * crtc, struct
timeval * vblanktime);
```

Arguments

crtc which counter to retrieve

vblanktime Pointer to struct timeval to receive the vblank timestamp.

Description

Fetches the “cooked” vblank count value that represents the number of vblank events since the system was booted, including lost events due to modesetting activity. Returns corresponding system timestamp of the time of the vblank interval that corresponds to the current vblank counter value.

This is the native KMS version of `drm_vblank_count_and_time`.

Name

`drm_arm_vblank_event` — arm vblank event after pageflip

Synopsis

```
void drm_arm_vblank_event (struct drm_device * dev, unsigned int pipe,  
struct drm_pending_vblank_event * e);
```

Arguments

dev DRM device

pipe CRTC index

e the event to prepare to send

Description

A lot of drivers need to generate vblank events for the very next vblank interrupt. For example when the page flip interrupt happens when the page flip gets armed, but not when it actually executes within the next vblank period. This helper function implements exactly the required vblank arming behaviour.

Caller must hold event lock. Caller must also hold a vblank reference for the event *e*, which will be dropped when the next vblank arrives.

This is the legacy version of `drm_crtc_arm_vblank_event`.

Name

`drm_crtc_arm_vblank_event` — arm vblank event after pageflip

Synopsis

```
void drm_crtc_arm_vblank_event (struct drm_crtc * crtc, struct drm_pending_vblank_event * e);
```

Arguments

crtc the source CRTC of the vblank event

e the event to send

Description

A lot of drivers need to generate vblank events for the very next vblank interrupt. For example when the page flip interrupt happens when the page flip gets armed, but not when it actually executes within the next vblank period. This helper function implements exactly the required vblank arming behaviour.

Caller must hold event lock. Caller must also hold a vblank reference for the event *e*, which will be dropped when the next vblank arrives.

This is the native KMS version of `drm_arm_vblank_event`.

Name

`drm_send_vblank_event` — helper to send vblank event after pageflip

Synopsis

```
void drm_send_vblank_event (struct drm_device * dev, unsigned int pipe,  
struct drm_pending_vblank_event * e);
```

Arguments

dev DRM device

pipe CRTC index

e the event to send

Description

Updates sequence # and timestamp on event, and sends it to userspace. Caller must hold event lock.

This is the legacy version of `drm_crtc_send_vblank_event`.

Name

`drm_crtc_send_vblank_event` — helper to send vblank event after pageflip

Synopsis

```
void drm_crtc_send_vblank_event (struct drm_crtc * crtc, struct dr-  
m_pending_vblank_event * e);
```

Arguments

crtc the source CRTC of the vblank event

e the event to send

Description

Updates sequence # and timestamp on event, and sends it to userspace. Caller must hold event lock.

This is the native KMS version of `drm_send_vblank_event`.

Name

`drm_vblank_get` — get a reference count on vblank events

Synopsis

```
int drm_vblank_get (struct drm_device * dev, unsigned int pipe);
```

Arguments

dev DRM device

pipe index of CRTC to own

Description

Acquire a reference count on vblank events to avoid having them disabled while in use.

This is the legacy version of `drm_crtc_vblank_get`.

Returns

Zero on success or a negative error code on failure.

Name

`drm_crtc_vblank_get` — get a reference count on vblank events

Synopsis

```
int drm_crtc_vblank_get (struct drm_crtc * crtc);
```

Arguments

crtc which CRTC to own

Description

Acquire a reference count on vblank events to avoid having them disabled while in use.

This is the native kms version of `drm_vblank_get`.

Returns

Zero on success or a negative error code on failure.

Name

`drm_vblank_put` — release ownership of vblank events

Synopsis

```
void drm_vblank_put (struct drm_device * dev, unsigned int pipe);
```

Arguments

dev DRM device

pipe index of CRTC to release

Description

Release ownership of a given vblank counter, turning off interrupts if possible. Disable interrupts after `drm_vblank_offdelay` milliseconds.

This is the legacy version of `drm_crtc_vblank_put`.

Name

`drm_crtc_vblank_put` — give up ownership of vblank events

Synopsis

```
void drm_crtc_vblank_put (struct drm_crtc * crtc);
```

Arguments

crtc which counter to give up

Description

Release ownership of a given vblank counter, turning off interrupts if possible. Disable interrupts after `drm_vblank_offdelay` milliseconds.

This is the native kms version of `drm_vblank_put`.

Name

`drm_wait_one_vblank` — wait for one vblank

Synopsis

```
void drm_wait_one_vblank (struct drm_device * dev, unsigned int pipe);
```

Arguments

dev DRM device

pipe CRTC index

Description

This waits for one vblank to pass on *pipe*, using the irq driver interfaces. It is a failure to call this when the vblank irq for *pipe* is disabled, e.g. due to lack of driver support or because the crtc is off.

Name

`drm_crtc_wait_one_vblank` — wait for one vblank

Synopsis

```
void drm_crtc_wait_one_vblank (struct drm_crtc * crtc);
```

Arguments

crtc DRM crtc

Description

This waits for one vblank to pass on *crtc*, using the irq driver interfaces. It is a failure to call this when the vblank irq for *crtc* is disabled, e.g. due to lack of driver support or because the crtc is off.

Name

`drm_vblank_off` — disable vblank events on a CRTC

Synopsis

```
void drm_vblank_off (struct drm_device * dev, unsigned int pipe);
```

Arguments

dev DRM device

pipe CRTC index

Description

Drivers can use this function to shut down the vblank interrupt handling when disabling a crtc. This function ensures that the latest vblank frame count is stored so that `drm_vblank_on` can restore it again.

Drivers must use this function when the hardware vblank counter can get reset, e.g. when suspending.

This is the legacy version of `drm_crtc_vblank_off`.

Name

`drm_crtc_vblank_off` — disable vblank events on a CRTC

Synopsis

```
void drm_crtc_vblank_off (struct drm_crtc * crtc);
```

Arguments

crtc CRTC in question

Description

Drivers can use this function to shut down the vblank interrupt handling when disabling a *crtc*. This function ensures that the latest vblank frame count is stored so that `drm_vblank_on` can restore it again.

Drivers must use this function when the hardware vblank counter can get reset, e.g. when suspending.

This is the native kms version of `drm_vblank_off`.

Name

`drm_crtc_vblank_reset` — reset vblank state to off on a CRTC

Synopsis

```
void drm_crtc_vblank_reset (struct drm_crtc * crtc);
```

Arguments

crtc CRTC in question

Description

Drivers can use this function to reset the vblank state to off at load time. Drivers should use this together with the `drm_crtc_vblank_off` and `drm_crtc_vblank_on` functions. The difference compared to `drm_crtc_vblank_off` is that this function doesn't save the vblank counter and hence doesn't need to call any driver hooks.

Name

`drm_vblank_on` — enable vblank events on a CRTC

Synopsis

```
void drm_vblank_on (struct drm_device * dev, unsigned int pipe);
```

Arguments

dev DRM device

pipe CRTC index

Description

This functions restores the vblank interrupt state captured with `drm_vblank_off` again. Note that calls to `drm_vblank_on` and `drm_vblank_off` can be unbalanced and so can also be unconditionally called in driver load code to reflect the current hardware state of the crtc.

This is the legacy version of `drm_crtc_vblank_on`.

Name

`drm_crtc_vblank_on` — enable vblank events on a CRTC

Synopsis

```
void drm_crtc_vblank_on (struct drm_crtc * crtc);
```

Arguments

crtc CRTC in question

Description

This functions restores the vblank interrupt state captured with `drm_vblank_off` again. Note that calls to `drm_vblank_on` and `drm_vblank_off` can be unbalanced and so can also be unconditionally called in driver load code to reflect the current hardware state of the crtc.

This is the native kms version of `drm_vblank_on`.

Name

`drm_vblank_pre_modeset` — account for vblanks across mode sets

Synopsis

```
void drm_vblank_pre_modeset (struct drm_device * dev, unsigned int
pipe);
```

Arguments

dev DRM device

pipe CRTC index

Description

Account for vblank events across mode setting events, which will likely reset the hardware frame counter.

This is done by grabbing a temporary vblank reference to ensure that the vblank interrupt keeps running across the modeset sequence. With this the software-side vblank frame counting will ensure that there are no jumps or discontinuities.

Unfortunately this approach is racy and also doesn't work when the vblank interrupt stops running, e.g. across system suspend resume. It is therefore highly recommended that drivers use the newer `drm_vblank_off` and `drm_vblank_on` instead. `drm_vblank_pre_modeset` only works correctly when using “cooked” software vblank frame counters and not relying on any hardware counters.

Drivers must call `drm_vblank_post_modeset` when re-enabling the same crtc again.

Name

`drm_vblank_post_modeset` — undo `drm_vblank_pre_modeset` changes

Synopsis

```
void drm_vblank_post_modeset (struct drm_device * dev, unsigned int  
pipe);
```

Arguments

dev DRM device

pipe CRTC index

Description

This function again drops the temporary vblank reference acquired in `drm_vblank_pre_modeset`.

Name

`drm_handle_vblank` — handle a vblank event

Synopsis

```
bool drm_handle_vblank (struct drm_device * dev, unsigned int pipe);
```

Arguments

dev DRM device

pipe index of CRTC where this event occurred

Description

Drivers should call this routine in their vblank interrupt handlers to update the vblank counter and send any signals that may be pending.

This is the legacy version of `drm_crtc_handle_vblank`.

Name

`drm_crtc_handle_vblank` — handle a vblank event

Synopsis

```
bool drm_crtc_handle_vblank (struct drm_crtc * crtc);
```

Arguments

crtc where this event occurred

Description

Drivers should call this routine in their vblank interrupt handlers to update the vblank counter and send any signals that may be pending.

This is the native KMS version of `drm_handle_vblank`.

Returns

True if the event was successfully handled, false on failure.

Name

`drm_vblank_no_hw_counter` — "No hw counter" implementation of `.get_vblank_counter`

Synopsis

```
u32 drm_vblank_no_hw_counter (struct drm_device * dev, unsigned int
    pipe);
```

Arguments

dev DRM device

pipe CRTC for which to read the counter

Description

Drivers can plug this into the `.get_vblank_counter` function if there is no useable hardware frame counter available.

Returns

0

Name

`drm_crtc_vblank_waitqueue` — get vblank waitqueue for the CRTC

Synopsis

```
wait_queue_head_t * drm_crtc_vblank_waitqueue (struct drm_crtc * crtc);
```

Arguments

crtc which CRTC's vblank waitqueue to retrieve

Description

This function returns a pointer to the vblank waitqueue for the CRTC. Drivers can use this to implement vblank waits using `wait_event` & co.

Open/Close, File Operations and IOCTLs

Open and Close

```
int (*firstopen) (struct drm_device *);  
void (*lastclose) (struct drm_device *);  
int (*open) (struct drm_device *, struct drm_file *);  
void (*preclose) (struct drm_device *, struct drm_file *);  
void (*postclose) (struct drm_device *, struct drm_file *);
```

Open and close handlers. None of those methods are mandatory.

The `firstopen` method is called by the DRM core for legacy UMS (User Mode Setting) drivers only when an application opens a device that has no other opened file handle. UMS drivers can implement it to acquire device resources. KMS drivers can't use the method and must acquire resources in the `load` method instead.

Similarly the `lastclose` method is called when the last application holding a file handle opened on the device closes it, for both UMS and KMS drivers. Additionally, the method is also called at module unload time or, for hot-pluggable devices, when the device is unplugged. The `firstopen` and `lastclose` calls can thus be unbalanced.

The `open` method is called every time the device is opened by an application. Drivers can allocate per-file private data in this method and store them in the struct `drm_file` `driver_priv` field. Note that the `open` method is called before `firstopen`.

The close operation is split into `preclose` and `postclose` methods. Drivers must stop and cleanup all per-file operations in the `preclose` method. For instance pending vertical blanking and page flip events must be cancelled. No per-file operation is allowed on the file handle after returning from the `preclose` method.

Finally the `postclose` method is called as the last step of the close operation, right before calling the `lastclose` method if no other open file handle exists for the device. Drivers that have allocated per-file private data in the `open` method should free it here.

The `lastclose` method should restore CRTC and plane properties to default value, so that a subsequent open of the device will not inherit state from the previous user. It can also be used to execute delayed power switching state changes, e.g. in conjunction with the `vga_switcheroo` infrastructure (see Part III,

“vga_switcheroo”). Beyond that KMS drivers should not do any further cleanup. Only legacy UMS drivers might need to clean up device state so that the vga console or an independent fbdev driver could take over.

File Operations

```
const struct file_operations *fops
```

File operations for the DRM device node.

Drivers must define the file operations structure that forms the DRM userspace API entry point, even though most of those operations are implemented in the DRM core. The `open`, `release` and `ioctl` operations are handled by

```
.owner = THIS_MODULE,
.open = drm_open,
.release = drm_release,
.unlocked_ioctl = drm_ioctl,
#ifdef CONFIG_COMPAT
.compat_ioctl = drm_compat_ioctl,
#endif
```

Drivers that implement private `ioctls` that requires 32/64bit compatibility support must provide their own `compat_ioctl` handler that processes private `ioctls` and calls `drm_compat_ioctl` for core `ioctls`.

The `read` and `poll` operations provide support for reading DRM events and polling them. They are implemented by

```
.poll = drm_poll,
.read = drm_read,
.llseek = no_llseek,
```

The memory mapping implementation varies depending on how the driver manages memory. Pre-GEM drivers will use `drm_mmap`, while GEM-aware drivers will use `drm_gem_mmap`. See the section called “The Graphics Execution Manager (GEM)”.

```
.mmap = drm_gem_mmap,
```

No other file operation is supported by the DRM API.

IOCTLs

```
struct drm_ioctl_desc *ioctls;
int num_ioctls;
```

Driver-specific `ioctls` descriptors table.

Driver-specific `ioctls` numbers start at `DRM_COMMAND_BASE`. The `ioctls` descriptors table is indexed by the `ioctl` number offset from the base value. Drivers can use the `DRM_IOCTL_DEF_DRV()` macro to initialize the table entries.

```
DRM_IOCTL_DEF_DRV(ioctl, func, flags)
```

ioctl is the ioctl name. Drivers must define the `DRM_##ioctl` and `DRM_IOCTL_##ioctl` macros to the ioctl number offset from `DRM_COMMAND_BASE` and the ioctl number respectively. The first macro is private to the device while the second must be exposed to userspace in a public header.

func is a pointer to the ioctl handler function compatible with the `drm_ioctl_t` type.

```
typedef int drm_ioctl_t(struct drm_device *dev, void *data,
    struct drm_file *file_priv);
```

flags is a bitmask combination of the following values. It restricts how the ioctl is allowed to be called.

- `DRM_AUTH` - Only authenticated callers allowed
- `DRM_MASTER` - The ioctl can only be called on the master file handle
- `DRM_ROOT_ONLY` - Only callers with the `SYSADMIN` capability allowed
- `DRM_CONTROL_ALLOW` - The ioctl can only be called on a control device
- `DRM_UNLOCKED` - The ioctl handler will be called without locking the DRM global mutex. This is the enforced default for kms drivers (i.e. using the `DRIVER_MODESET` flag) and hence shouldn't be used any more for new drivers.

Name

`drm_noop` — DRM no-op ioctl implementation

Synopsis

```
int drm_noop (struct drm_device * dev, void * data, struct drm_file  
* file_priv);
```

Arguments

<i>dev</i>	DRM device for the ioctl
<i>data</i>	data pointer for the ioctl
<i>file_priv</i>	DRM file for the ioctl call

Description

This no-op implementation for drm ioctls is useful for deprecated functionality where we can't return a failure code because existing userspace checks the result of the ioctl, but doesn't care about the action.

Always returns successfully with 0.

Name

`drm_invalid_op` — DRM invalid ioctl implementation

Synopsis

```
int drm_invalid_op (struct drm_device * dev, void * data, struct drm_file  
* file_priv);
```

Arguments

<i>dev</i>	DRM device for the ioctl
<i>data</i>	data pointer for the ioctl
<i>file_priv</i>	DRM file for the ioctl call

Description

This no-op implementation for drm ioctls is useful for deprecated functionality where we really don't want to allow userspace to call the ioctl any more. This is the case for old ums interfaces for drivers that transitioned to kms gradually and so kept the old legacy tables around. This only applies to radeon and i915 kms drivers, other drivers shouldn't need to use this function.

Always fails with a return value of `-EINVAL`.

Name

`drm_ioctl` — ioctl callback implementation for DRM drivers

Synopsis

```
long drm_ioctl (struct file * filp, unsigned int cmd, unsigned long arg);
```

Arguments

filp file this ioctl is called on

cmd ioctl cmd number

arg user argument

Looks up the ioctl function in the

`:ioctls` table, checking for root privileges if so required, and dispatches to the respective function.

Returns

Zero on success, negative error code on failure.

Name

`drm_ioctl_flags` — Check for core ioctl and return ioctl permission flags

Synopsis

```
bool drm_ioctl_flags (unsigned int nr, unsigned int * flags);
```

Arguments

nr ioctl number

flags where to return the ioctl permission flags

Description

This ioctl is only used by the `vmwgfx` driver to augment the access checks done by the drm core and insofar a pretty decent layering violation. This shouldn't be used by any drivers.

Returns

True if the *nr* corresponds to a DRM core ioctl numer, false otherwise.

Legacy Support Code

The section very briefly covers some of the old legacy support code which is only used by old DRM drivers which have done a so-called shadow-attach to the underlying device instead of registering as a real driver. This also includes some of the old generic buffer management and command submission code. Do not use any of this in new and modern drivers.

Legacy Suspend/Resume

The DRM core provides some suspend/resume code, but drivers wanting full suspend/resume support should provide `save()` and `restore()` functions. These are called at suspend, hibernate, or resume time, and should perform any state save or restore required by your device across suspend or hibernate states.

```
int (*suspend) (struct drm_device *, pm_message_t state);
int (*resume) (struct drm_device *);
```

Those are legacy suspend and resume methods which *only* work with the legacy shadow-attach driver registration functions. New driver should use the power management interface provided by their bus type (usually through the struct `device_driver` `dev_pm_ops`) and set these methods to `NULL`.

Legacy DMA Services

This should cover how DMA mapping etc. is supported by the core. These functions are deprecated and should not be used.

Chapter 3. Userland interfaces

The DRM core exports several interfaces to applications, generally intended to be used through corresponding libdrm wrapper functions. In addition, drivers export device-specific interfaces for use by user-space drivers & device-aware applications through ioctls and sysfs files.

External interfaces include: memory mapping, context management, DMA operations, AGP management, vblank control, fence management, memory management, and output management.

Cover generic ioctls and sysfs layout here. We only need high-level info, since man pages should cover the rest.

Render nodes

DRM core provides multiple character-devices for user-space to use. Depending on which device is opened, user-space can perform a different set of operations (mainly ioctls). The primary node is always created and called `card<num>`. Additionally, a currently unused control node, called `controlID<num>` is also created. The primary node provides all legacy operations and historically was the only interface used by userspace. With KMS, the control node was introduced. However, the planned KMS control interface has never been written and so the control node stays unused to date.

With the increased use of offscreen renderers and GPGPU applications, clients no longer require running compositors or graphics servers to make use of a GPU. But the DRM API required unprivileged clients to authenticate to a DRM-Master prior to getting GPU access. To avoid this step and to grant clients GPU access without authenticating, render nodes were introduced. Render nodes solely serve render clients, that is, no modesetting or privileged ioctls can be issued on render nodes. Only non-global rendering commands are allowed. If a driver supports render nodes, it must advertise it via the `DRIVER_RENDER` DRM driver capability. If not supported, the primary node must be used for render clients together with the legacy `drmAuth` authentication procedure.

If a driver advertises render node support, DRM core will create a separate render node called `renderD<num>`. There will be one render node per device. No ioctls except `PRIME`-related ioctls will be allowed on this node. Especially `GEM_OPEN` will be explicitly prohibited. Render nodes are designed to avoid the buffer-leaks, which occur if clients guess the flink names or `mmap` offsets on the legacy interface. Additionally to this basic interface, drivers must mark their driver-dependent render-only ioctls as `DRM_RENDER_ALLOW` so render clients can use them. Driver authors must be careful not to allow any privileged ioctls on render nodes.

With render nodes, user-space can now control access to the render node via basic file-system access-modes. A running graphics server which authenticates clients on the privileged primary/legacy node is no longer required. Instead, a client can open the render node and is immediately granted GPU access. Communication between clients (or servers) is done via `PRIME`. `FLINK` from render node to legacy node is not supported. New clients must not use the insecure `FLINK` interface.

Besides dropping all modeset/global ioctls, render nodes also drop the DRM-Master concept. There is no reason to associate render clients with a DRM-Master as they are independent of any graphics server. Besides, they must work without any running master, anyway. Drivers must be able to run without a master object if they support render nodes. If, on the other hand, a driver requires shared state between clients which is visible to user-space and accessible beyond open-file boundaries, they cannot support render nodes.

VBlank event handling

The DRM core exposes two vertical blank related ioctls:

DRM_IOCTL_WAIT_VBLANK	This takes a struct <code>drm_wait_vblank</code> structure as its argument, and it is used to block or request a signal when a specified vblank event occurs.
DRM_IOCTL_MODESET_CTL	This was only used for user-mode-setting drivers around modesetting changes to allow the kernel to update the vblank interrupt after mode setting, since on many devices the vertical blank counter is reset to 0 at some point during modeset. Modern drivers should not call this any more since with kernel mode setting it is a no-op.

Part II. DRM Drivers

This second part of the GPU Driver Developer's Guide documents driver code, implementation details and also all the driver-specific userspace interfaces. Especially since all hardware-acceleration interfaces to userspace are driver specific for efficiency and other reasons these interfaces can be rather substantial. Hence every driver has its own chapter.

Table of Contents

4. drm/i915 Intel GFX Driver	637
Core Driver Infrastructure	637
Runtime Power Management	637
Interrupt Handling	654
Intel GVT-g Guest Support(vGPU)	657
Display Hardware Handling	660
Mode Setting Infrastructure	660
Frontbuffer Tracking	660
Display FIFO Underrun Reporting	668
Plane Configuration	674
Atomic Plane Helpers	674
Output Probing	679
Hotplug	679
High Definition Audio	683
Panel Self Refresh PSR (PSR/SRD)	691
Frame Buffer Compression (FBC)	697
Display Refresh Rate Switching (DRRS)	704
DPIO	710
CSR firmware support for DMC	711
Memory Management and Command Submission	714
Batchbuffer Parsing	714
Batchbuffer Pools	720
Logical Rings, Logical Ring Contexts and Execlists	723
Global GTT views	735
GTT Fences and Swizzling	741
Object Tiling IOCTLs	750
Buffer Object Eviction	752
Buffer Object Memory Shrinking	754
GuC	757
GuC-specific firmware loader	757
GuC-based command submission	760
GuC Firmware Layout	766
Tracing	766
i915_ppgtt_create and i915_ppgtt_release	766
i915_context_create and i915_context_free	766
switch_mm	766

Chapter 4. drm/i915 Intel GFX Driver

The drm/i915 driver supports all (with the exception of some very early models) integrated GFX chipsets with both Intel display and rendering blocks. This excludes a set of SoC platforms with an SGX rendering unit, those have basic support through the gma500 drm driver.

Core Driver Infrastructure

This section covers core driver infrastructure used by both the display and the GEM parts of the driver.

Runtime Power Management

The i915 driver supports dynamic enabling and disabling of entire hardware blocks at runtime. This is especially important on the display side where software is supposed to control many power gates manually on recent hardware, since on the GT side a lot of the power management is done by the hardware. But even there some manual control at the device level is required.

Since i915 supports a diverse set of platforms with a unified codebase and hardware engineers just love to shuffle functionality around between power domains there's a sizeable amount of indirection required. This file provides generic functions to the driver for grabbing and releasing references for abstract power domains. It then maps those to the actual power wells present for a given platform.

Name

`__intel_display_power_is_enabled` — unlocked check for a power domain

Synopsis

```
bool __intel_display_power_is_enabled (struct drm_i915_private * dev_priv,
enum intel_display_power_domain domain);
```

Arguments

dev_priv i915 device instance

domain power domain to check

Description

This is the unlocked version of `intel_display_power_is_enabled` and should only be used from error capture and recovery code where deadlocks are possible.

Returns

True when the power domain is enabled, false otherwise.

Name

`intel_display_power_is_enabled` — check for a power domain

Synopsis

```
bool intel_display_power_is_enabled (struct drm_i915_private * dev_priv, enum intel_display_power_domain domain);
```

Arguments

dev_priv i915 device instance

domain power domain to check

Description

This function can be used to check the hw power domain state. It is mostly used in hardware state readout functions. Everywhere else code should rely upon explicit power domain reference counting to ensure that the hardware block is powered up before accessing it.

Callers must hold the relevant modesetting locks to ensure that concurrent threads can't disable the power well while the caller tries to read a few registers.

Returns

True when the power domain is enabled, false otherwise.

Name

`intel_display_set_init_power` — set the initial power domain state

Synopsis

```
void intel_display_set_init_power (struct drm_i915_private * dev_priv,  
bool enable);
```

Arguments

dev_priv i915 device instance

enable whether to enable or disable the initial power domain state

Description

For simplicity our driver load/unload and system suspend/resume code assumes that all power domains are always enabled. This functions controls the state of this little hack. While the initial power domain state is enabled runtime pm is effectively disabled.

Name

`intel_display_power_get` — grab a power domain reference

Synopsis

```
void intel_display_power_get (struct drm_i915_private * dev_priv, enum  
intel_display_power_domain domain);
```

Arguments

dev_priv i915 device instance

domain power domain to reference

Description

This function grabs a power domain reference for *domain* and ensures that the power domain and all its parents are powered up. Therefore users should only grab a reference to the innermost power domain they need.

Any power domain reference obtained by this function must have a symmetric call to `intel_display_power_put` to release the reference again.

Name

`intel_display_power_put` — release a power domain reference

Synopsis

```
void intel_display_power_put (struct drm_i915_private * dev_priv, enum  
intel_display_power_domain domain);
```

Arguments

dev_priv i915 device instance

domain power domain to reference

Description

This function drops the power domain reference obtained by `intel_display_power_get` and might power down the corresponding hardware block right away if this is the last reference.

Name

`intel_power_domains_init` — initializes the power domain structures

Synopsis

```
int intel_power_domains_init (struct drm_i915_private * dev_priv);
```

Arguments

dev_priv i915 device instance

Description

Initializes the power domain structures for *dev_priv* depending upon the supported platform.

Name

`intel_power_domains_fini` — finalizes the power domain structures

Synopsis

```
void intel_power_domains_fini (struct drm_i915_private * dev_priv);
```

Arguments

dev_priv i915 device instance

Description

Finalizes the power domain structures for *dev_priv* depending upon the supported platform. This function also disables runtime pm and ensures that the device stays powered up so that the driver can be re-loaded.

Name

`intel_power_domains_init_hw` — initialize hardware power domain state

Synopsis

```
void intel_power_domains_init_hw (struct drm_i915_private * dev_priv,  
bool resume);
```

Arguments

dev_priv i915 device instance

resume -- undescribed --

Description

This function initializes the hardware power domain state and enables all power domains using `intel_display_set_init_power`.

Name

`intel_power_domains_suspend` — suspend power domain state

Synopsis

```
void intel_power_domains_suspend (struct drm_i915_private * dev_priv);
```

Arguments

dev_priv i915 device instance

Description

This function prepares the hardware power domain state before entering system suspend. It must be paired with `intel_power_domains_init_hw`.

Name

`intel_runtime_pm_get` — grab a runtime pm reference

Synopsis

```
void intel_runtime_pm_get (struct drm_i915_private * dev_priv);
```

Arguments

dev_priv i915 device instance

Description

This function grabs a device-level runtime pm reference (mostly used for GEM code to ensure the GTT or GT is on) and ensures that it is powered up.

Any runtime pm reference obtained by this function must have a symmetric call to `intel_runtime_pm_put` to release the reference again.

Name

`intel_runtime_pm_get_noresume` — grab a runtime pm reference

Synopsis

```
void intel_runtime_pm_get_noresume (struct drm_i915_private * dev_priv);
```

Arguments

dev_priv i915 device instance

Description

This function grabs a device-level runtime pm reference (mostly used for GEM code to ensure the GTT or GT is on).

It will not power up the device but instead only check that it's powered on. Therefore it is only valid to call this functions from contexts where the device is known to be powered up and where trying to power it up would result in hilarity and deadlocks. That pretty much means only the system suspend/resume code where this is used to grab runtime pm references for delayed setup down in work items.

Any runtime pm reference obtained by this function must have a symmetric call to `intel_runtime_pm_put` to release the reference again.

Name

`intel_runtime_pm_put` — release a runtime pm reference

Synopsis

```
void intel_runtime_pm_put (struct drm_i915_private * dev_priv);
```

Arguments

dev_priv i915 device instance

Description

This function drops the device-level runtime pm reference obtained by `intel_runtime_pm_get` and might power down the corresponding hardware block right away if this is the last reference.

Name

`intel_runtime_pm_enable` — enable runtime pm

Synopsis

```
void intel_runtime_pm_enable (struct drm_i915_private * dev_priv);
```

Arguments

dev_priv i915 device instance

Description

This function enables runtime pm at the end of the driver load sequence.

Note that this function does currently not enable runtime pm for the subordinate display power domains. That is only done on the first modeset using `intel_display_set_init_power`.

Name

`intel_uncore_forcewake_get` — grab forcewake domain references

Synopsis

```
void intel_uncore_forcewake_get (struct drm_i915_private * dev_priv,  
enum forcewake_domains fw_domains);
```

Arguments

dev_priv i915 device instance

fw_domains forcewake domains to get reference on

Description

This function can be used get GT's forcewake domain references. Normal register access will handle the forcewake domains automatically. However if some sequence requires the GT to not power down a particular forcewake domains this function should be called at the beginning of the sequence. And subsequently the reference should be dropped by symmetric call to `intel_unforce_forcewake_put`. Usually caller wants all the domains to be kept awake so the *fw_domains* would be then `FORCEWAKE_ALL`.

Name

`intel_uncore_forcewake_get__locked` — grab forcewake domain references

Synopsis

```
void intel_uncore_forcewake_get__locked (struct drm_i915_private * dev_priv, enum forcewake_domains fw_domains);
```

Arguments

dev_priv i915 device instance

fw_domains forcewake domains to get reference on

Description

See `intel_uncore_forcewake_get`. This variant places the onus on the caller to explicitly handle the `dev_priv->uncore.lock` spinlock.

Name

`intel_uncore_forcewake_put` — release a forcewake domain reference

Synopsis

```
void intel_uncore_forcewake_put (struct drm_i915_private * dev_priv,  
enum forcewake_domains fw_domains);
```

Arguments

dev_priv i915 device instance

fw_domains forcewake domains to put references

Description

This function drops the device-level forcewakes for specified domains obtained by `intel_uncore_forcewake_get`.

Name

`intel_uncore_forcewake_put__locked` — grab forcewake domain references

Synopsis

```
void intel_uncore_forcewake_put__locked (struct drm_i915_private * dev_priv, enum forcewake_domains fw_domains);
```

Arguments

dev_priv i915 device instance

fw_domains forcewake domains to get reference on

Description

See `intel_uncore_forcewake_put`. This variant places the onus on the caller to explicitly handle the `dev_priv->uncore.lock` spinlock.

Interrupt Handling

These functions provide the basic support for enabling and disabling the interrupt handling support. There's a lot more functionality in `i915_irq.c` and related files, but that will be described in separate chapters.

Name

`intel_irq_init` — initializes irq support

Synopsis

```
void intel_irq_init (struct drm_i915_private * dev_priv);
```

Arguments

dev_priv i915 device instance

Description

This function initializes all the irq support including work items, timers and all the vtables. It does not setup the interrupt itself though.

Name

`intel_runtime_pm_disable_interrupts` — runtime interrupt disabling

Synopsis

```
void intel_runtime_pm_disable_interrupts (struct drm_i915_private * dev_priv);
```

Arguments

dev_priv i915 device instance

Description

This function is used to disable interrupts at runtime, both in the runtime pm and the system suspend/resume code.

Name

`intel_runtime_pm_enable_interrupts` — runtime interrupt enabling

Synopsis

```
void intel_runtime_pm_enable_interrupts (struct drm_i915_private * dev_priv);
```

Arguments

dev_priv i915 device instance

Description

This function is used to enable interrupts at runtime, both in the runtime pm and the system suspend/resume code.

Intel GVT-g Guest Support(vGPU)

Intel GVT-g is a graphics virtualization technology which shares the GPU among multiple virtual machines on a time-sharing basis. Each virtual machine is presented a virtual GPU (vGPU), which has equivalent features as the underlying physical GPU (pGPU), so i915 driver can run seamlessly in a virtual machine. This file provides vGPU specific optimizations when running in a virtual machine, to reduce the complexity of vGPU emulation and to improve the overall performance.

A primary function introduced here is so-called “address space ballooning” technique. Intel GVT-g partitions global graphics memory among multiple VMs, so each VM can directly access a portion of the memory without hypervisor's intervention, e.g. filling textures or queuing commands. However with the partitioning an unmodified i915 driver would assume a smaller graphics memory starting from address ZERO, then requires vGPU emulation module to translate the graphics address between 'guest view' and 'host view', for all registers and command opcodes which contain a graphics memory address. To reduce the complexity, Intel GVT-g introduces “address space ballooning”, by telling the exact partitioning knowledge to each guest i915 driver, which then reserves and prevents non-allocated portions from allocation. Thus vGPU emulation module only needs to scan and validate graphics addresses without complexity of address translation.

Name

`i915_check_vgpu` — detect virtual GPU

Synopsis

```
void i915_check_vgpu (struct drm_device * dev);
```

Arguments

dev drm device *

Description

This function is called at the initialization stage, to detect whether running on a vGPU.

Name

intel_vgt_deballoon — deballoon reserved graphics address trunks

Synopsis

```
void intel_vgt_deballoon ( void );
```

Arguments

void no arguments

Description

This function is called to deallocate the ballooned-out graphic memory, when driver is unloaded or when ballooning fails.

Name

intel_vgt_balloon — balloon out reserved graphics address trunks

Synopsis

```
int intel_vgt_balloon (struct drm_device * dev);
```

Arguments

dev drm device

Description

This function is called at the initialization stage, to balloon out the graphic address space allocated to other vGPUs, by marking these spaces as reserved. The ballooning related knowledge(starting address and size of the mappable/unmappable graphic memory) is described in the vgt_if structure in a reserved mmio range.

To give an example, the drawing below depicts one typical scenario after ballooning. Here the vGPU1 has 2 pieces of graphic address spaces ballooned out each for the mappable and the non-mappable part. From the vGPU1 point of view, the total size is the same as the physical one, with the start address of its graphic space being zero. Yet there are some portions ballooned out(the shadow part, which are marked as reserved by drm allocator). From the host point of view, the graphic address space is partitioned by multiple vGPUs in different VMs.

```
vGPU1 view Host view 0 -----> +-----+ +-----+ ^ | | | | | | | | vGPU3 | | | | | | | | +-----+
| | | | | | | | vGPU2 | | +-----+ +-----+ mappable GM | available | ==> | vGPU1 | | +-----+
+-----+ | | | | | | | | | v | | | | | | | | Host | +-----+ +-----+ +-----+ +-----+ ^ | | | | | | | |
vGPU3 | | | | | | | | +-----+ | | | | | | | | vGPU2 | | +-----+ +-----+ unmappable GM | available
| ==> | vGPU1 | | +-----+ +-----+ | | | | | | | | | | | | | | Host | v | | | | | | | | | total GM size -----
> +-----+ +-----+
```

Returns

zero on success, non-zero if configuration invalid or ballooning failed

Display Hardware Handling

This section covers everything related to the display hardware including the mode setting infrastructure, plane, sprite and cursor handling and display, output probing and related topics.

Mode Setting Infrastructure

The i915 driver is thus far the only DRM driver which doesn't use the common DRM helper code to implement mode setting sequences. Thus it has its own tailor-made infrastructure for executing a display configuration change.

Frontbuffer Tracking

Many features require us to track changes to the currently active frontbuffer, especially rendering targeted at the frontbuffer.

To be able to do so GEM tracks frontbuffers using a bitmask for all possible frontbuffer slots through `i915_gem_track_fb`. The function in this file are then called when the contents of the frontbuffer are invalidated, when frontbuffer rendering has stopped again to flush out all the changes and when the frontbuffer is exchanged with a flip. Subsystems interested in frontbuffer changes (e.g. PSR, FBC, DRRS) should directly put their callbacks into the relevant places and filter for the frontbuffer slots that they are interested in.

On a high level there are two types of powersaving features. The first one work like a special cache (FBC and PSR) and are interested when they should stop caching and when to restart caching. This is done by placing callbacks into the invalidate and the flush functions: At invalidate the caching must be stopped and at flush time it can be restarted. And maybe they need to know when the frontbuffer changes (e.g. when the hw doesn't initiate an invalidate and flush on its own) which can be achieved with placing callbacks into the flip functions.

The other type of display power saving feature only cares about busyness (e.g. DRRS). In that case all three (invalidate, flush and flip) indicate busyness. There is no direct way to detect idleness. Instead an idle timer work delayed work should be started from the flush and flip functions and cancelled as soon as busyness is detected.

Note that there's also an older frontbuffer activity tracking scheme which just tracks general activity. This is done by the various `mark_busy` and `mark_idle` functions. For display power management features using these functions is deprecated and should be avoided.

Name

`intel_fb_obj_invalidate` — invalidate framebuffer object

Synopsis

```
void intel_fb_obj_invalidate (struct drm_i915_gem_object * obj, enum  
fb_op_origin origin);
```

Arguments

obj GEM object to invalidate

origin which operation caused the invalidation

Description

This function gets called every time rendering on the given object starts and framebuffer caching (fbc, low refresh rate for DRRS, panel self refresh) must be invalidated. For `ORIGIN_CS` any subsequent invalidation will be delayed until the rendering completes or a flip on this framebuffer plane is scheduled.

Name

intel_frontbuffer_flush — flush frontbuffer

Synopsis

```
void intel_frontbuffer_flush (struct drm_device * dev, unsigned front-  
buffer_bits, enum fb_op_origin origin);
```

Arguments

<i>dev</i>	DRM device
<i>frontbuffer_bits</i>	frontbuffer plane tracking bits
<i>origin</i>	which operation caused the flush

Description

This function gets called every time rendering on the given planes has completed and frontbuffer caching can be started again. Flushes will get delayed if they're blocked by some outstanding asynchronous rendering.

Can be called without any locks held.

Name

`intel_fb_obj_flush` — flush frontbuffer object

Synopsis

```
void intel_fb_obj_flush (struct drm_i915_gem_object * obj, bool retire,  
enum fb_op_origin origin);
```

Arguments

obj GEM object to flush

retire set when retiring asynchronous rendering

origin which operation caused the flush

Description

This function gets called every time rendering on the given object has completed and frontbuffer caching can be started again. If *retire* is true then any delayed flushes will be unblocked.

Name

intel_frontbuffer_flip_prepare — prepare asynchronous frontbuffer flip

Synopsis

```
void intel_frontbuffer_flip_prepare (struct drm_device * dev, unsigned
frontbuffer_bits);
```

Arguments

<i>dev</i>	DRM device
<i>frontbuffer_bits</i>	frontbuffer plane tracking bits

Description

This function gets called after scheduling a flip on *obj*. The actual frontbuffer flushing will be delayed until completion is signalled with `intel_frontbuffer_flip_complete`. If an invalidate happens in between this flush will be cancelled.

Can be called without any locks held.

Name

`intel_frontbuffer_flip_complete` — complete asynchronous frontbuffer flip

Synopsis

```
void intel_frontbuffer_flip_complete (struct drm_device * dev, unsigned  
frontbuffer_bits);
```

Arguments

<i>dev</i>	DRM device
<i>frontbuffer_bits</i>	frontbuffer plane tracking bits

Description

This function gets called after the flip has been latched and will complete on the next vblank. It will execute the flush if it hasn't been cancelled yet.

Can be called without any locks held.

Name

`intel_frontbuffer_flip` — synchronous frontbuffer flip

Synopsis

```
void intel_frontbuffer_flip (struct drm_device * dev, unsigned front-  
buffer_bits);
```

Arguments

dev DRM device

frontbuffer_bits frontbuffer plane tracking bits

Description

This function gets called after scheduling a flip on *obj*. This is for synchronous plane updates which will happen on the next vblank and which will not get delayed by pending gpu rendering.

Can be called without any locks held.

Name

i915_gem_track_fb — update frontbuffer tracking

Synopsis

```
void i915_gem_track_fb (struct drm_i915_gem_object * old, struct dr-
m_i915_gem_object * new, unsigned frontbuffer_bits);
```

Arguments

old current GEM buffer for the frontbuffer slots

new new GEM buffer for the frontbuffer slots

frontbuffer_bits bitmask of frontbuffer slots

Description

This updates the frontbuffer tracking bits *frontbuffer_bits* by clearing them from *old* and setting them in *new*. Both *old* and *new* can be NULL.

Display FIFO Underrun Reporting

The i915 driver checks for display fifo underruns using the interrupt signals provided by the hardware. This is enabled by default and fairly useful to debug display issues, especially watermark settings.

If an underrun is detected this is logged into dmesg. To avoid flooding logs and occupying the cpu underrun interrupts are disabled after the first occurrence until the next modeset on a given pipe.

Note that underrun detection on gmch platforms is a bit more ugly since there is no interrupt (despite that the signalling bit is in the PIPESTAT pipe interrupt register). Also on some other platforms underrun interrupts are shared, which means that if we detect an underrun we need to disable underrun reporting on all pipes.

The code also supports underrun detection on the PCH transcoder.

Name

`intel_set_cpu_fifo_underrun_reporting` — set cpu fifo underrun reporting state

Synopsis

```
bool intel_set_cpu_fifo_underrun_reporting (struct drm_i915_private *  
dev_priv, enum pipe pipe, bool enable);
```

Arguments

dev_priv i915 device instance

pipe (CPU) pipe to set state for

enable whether underruns should be reported or not

Description

This function sets the fifo underrun state for *pipe*. It is used in the modeset code to avoid false positives since on many platforms underruns are expected when disabling or enabling the pipe.

Notice that on some platforms disabling underrun reports for one pipe disables for all due to shared interrupts. Actual reporting is still per-pipe though.

Returns the previous state of underrun reporting.

Name

`intel_set_pch_fifo_underrun_reporting` — set PCH fifo underrun reporting state

Synopsis

```
bool intel_set_pch_fifo_underrun_reporting (struct drm_i915_private *  
dev_priv, enum transcoder pch_transcoder, bool enable);
```

Arguments

<i>dev_priv</i>	i915 device instance
<i>pch_transcoder</i>	the PCH transcoder (same as pipe on IVB and older)
<i>enable</i>	whether underruns should be reported or not

Description

This function makes us disable or enable PCH fifo underruns for a specific PCH transcoder. Notice that on some PCHs (e.g. CPT/PPT), disabling FIFO underrun reporting for one transcoder may also disable all the other PCH error interrupts for the other transcoders, due to the fact that there's just one interrupt mask/enable bit for all the transcoders.

Returns the previous state of underrun reporting.

Name

`intel_cpu_fifo_underrun_irq_handler` — handle CPU fifo underrun interrupt

Synopsis

```
void intel_cpu_fifo_underrun_irq_handler (struct drm_i915_private * dev_priv, enum pipe pipe);
```

Arguments

dev_priv i915 device instance

pipe (CPU) pipe to set state for

Description

This handles a CPU fifo underrun interrupt, generating an underrun warning into dmesg if underrun reporting is enabled and then disables the underrun interrupt to avoid an irq storm.

Name

`intel_pch_fifo_underrun_irq_handler` — handle PCH fifo underrun interrupt

Synopsis

```
void intel_pch_fifo_underrun_irq_handler (struct drm_i915_private * dev_priv, enum transcoder pch_transcoder);
```

Arguments

dev_priv i915 device instance

pch_transcoder the PCH transcoder (same as pipe on IVB and older)

Description

This handles a PCH fifo underrun interrupt, generating an underrun warning into dmesg if underrun reporting is enabled and then disables the underrun interrupt to avoid an irq storm.

Name

`intel_check_cpu_fifo_underruns` — check for CPU fifo underruns immediately

Synopsis

```
void intel_check_cpu_fifo_underruns (struct drm_i915_private * dev_priv);
```

Arguments

dev_priv i915 device instance

Description

Check for CPU fifo underruns immediately. Useful on IVB/HSW where the shared error interrupt may have been disabled, and so CPU fifo underruns won't necessarily raise an interrupt, and on GMCH platforms where underruns never raise an interrupt.

Name

`intel_check_pch_fifo_underruns` — check for PCH fifo underruns immediately

Synopsis

```
void intel_check_pch_fifo_underruns (struct drm_i915_private * dev_priv);
```

Arguments

dev_priv i915 device instance

Description

Check for PCH fifo underruns immediately. Useful on CPT/PPT where the shared error interrupt may have been disabled, and so PCH fifo underruns won't necessarily raise an interrupt.

Plane Configuration

This section covers plane configuration and composition with the primary plane, sprites, cursors and overlays. This includes the infrastructure to do atomic vsync'ed updates of all this state and also tightly coupled topics like watermark setup and computation, framebuffer compression and panel self refresh.

Atomic Plane Helpers

The functions here are used by the atomic plane helper functions to implement legacy plane updates (i.e., `drm_plane->update_plane` and `drm_plane->disable_plane`). This allows plane updates to use the atomic state infrastructure and perform plane updates as separate prepare/check/commit/cleanup steps.

Name

`intel_create_plane_state` — create plane state object

Synopsis

```
struct intel_plane_state * intel_create_plane_state (struct drm_plane
* plane);
```

Arguments

plane drm plane

Description

Allocates a fresh plane state for the given plane and sets some of the state values to sensible initial values.

Returns

A newly allocated plane state, or NULL on failure

Name

`intel_plane_duplicate_state` — duplicate plane state

Synopsis

```
struct drm_plane_state * intel_plane_duplicate_state (struct drm_plane
* plane);
```

Arguments

plane drm plane

Description

Allocates and returns a copy of the plane state (both common and Intel-specific) for the specified plane.

Returns

The newly allocated plane state, or NULL on failure.

Name

`intel_plane_destroy_state` — destroy plane state

Synopsis

```
void intel_plane_destroy_state (struct drm_plane * plane, struct dr-  
m_plane_state * state);
```

Arguments

plane drm plane

state state object to destroy

Description

Destroys the plane state (both common and Intel-specific) for the specified plane.

Name

`intel_plane_atomic_get_property` — fetch plane property value

Synopsis

```
int intel_plane_atomic_get_property (struct drm_plane * plane, const
struct drm_plane_state * state, struct drm_property * property, uint64_t
* val);
```

Arguments

<i>plane</i>	plane to fetch property for
<i>state</i>	state containing the property value
<i>property</i>	property to look up
<i>val</i>	pointer to write property value into

Description

The DRM core does not store shadow copies of properties for atomic-capable drivers. This entrypoint is used to fetch the current value of a driver-specific plane property.

Name

`intel_plane_atomic_set_property` — set plane property value

Synopsis

```
int intel_plane_atomic_set_property (struct drm_plane * plane, struct
drm_plane_state * state, struct drm_property * property, uint64_t val);
```

Arguments

<i>plane</i>	plane to set property for
<i>state</i>	state to update property value in
<i>property</i>	property to set
<i>val</i>	value to set property to

Description

Writes the specified property value for a plane into the provided atomic state object.

Returns 0 on success, -EINVAL on unrecognized properties

Output Probing

This section covers output probing and related infrastructure like the hotplug interrupt storm detection and mitigation code. Note that the i915 driver still uses most of the common DRM helper code for output probing, so those sections fully apply.

Hotplug

Simply put, hotplug occurs when a display is connected to or disconnected from the system. However, there may be adapters and docking stations and Display Port short pulses and MST devices involved, complicating matters.

Hotplug in i915 is handled in many different levels of abstraction.

The platform dependent interrupt handling code in `i915_irq.c` enables, disables, and does preliminary handling of the interrupts. The interrupt handlers gather the hotplug detect (HPD) information from relevant registers into a platform independent mask of hotplug pins that have fired.

The platform independent interrupt handler `intel_hpd_irq_handler` in `intel_hotplug.c` does hotplug irq storm detection and mitigation, and passes further processing to appropriate bottom halves (Display Port specific and regular hotplug).

The Display Port work function `i915_digport_work_func` calls into `intel_dp_hpd_pulse` via hooks, which handles DP short pulses and DP MST long pulses, with failures and non-MST long pulses triggering regular hotplug processing on the connector.

The regular hotplug work function `i915_hotplug_work_func` calls connector detect hooks, and, if connector status changes, triggers sending of hotplug uevent to userspace via `drm_kms_helper_hotplug_event`.

Finally, the userspace is responsible for triggering a modeset upon receiving the hotplug uevent, disabling or enabling the crtc as needed.

The hotplug interrupt storm detection and mitigation code keeps track of the number of interrupts per hotplug pin per a period of time, and if the number of interrupts exceeds a certain threshold, the interrupt is disabled for a while before being re-enabled. The intention is to mitigate issues raising from broken hardware triggering massive amounts of interrupts and grinding the system to a halt.

Current implementation expects that hotplug interrupt storm will not be seen when display port sink is connected, hence on platforms whose DP callback is handled by `i915_digport_work_func` reenabling of hpd is not performed (it was never expected to be disabled in the first place ;)) this is specific to DP sinks handled by this routine and any other display such as HDMI or DVI enabled on the same port will have proper logic since it will use `i915_hotplug_work_func` where this logic is handled.

Name

`intel_hpd_irq_storm_detect` — gather stats and detect HPD irq storm on a pin

Synopsis

```
bool intel_hpd_irq_storm_detect (struct drm_i915_private * dev_priv,
enum hpd_pin pin);
```

Arguments

dev_priv private driver data pointer

pin the pin to gather stats on

Description

Gather stats about HPD irqs from the specified *pin*, and detect irq storms. Only the pin specific stats and state are changed, the caller is responsible for further action.

HPD_STORM_THRESHOLD irqs are allowed within *HPD_STORM_DETECT_PERIOD* ms, otherwise it's considered an irq storm, and the irq state is set to *HPD_MARK_DISABLED*.

Return true if an irq storm was detected on *pin*.

Name

intel_hpd_irq_handler — main hotplug irq handler

Synopsis

```
void intel_hpd_irq_handler (struct drm_device * dev, u32 pin_mask, u32  
long_mask);
```

Arguments

dev drm device

pin_mask a mask of hpd pins that have triggered the irq

long_mask a mask of hpd pins that may be long hpd pulses

Description

This is the main hotplug irq handler for all platforms. The platform specific irq handlers call the platform specific hotplug irq handlers, which read and decode the appropriate registers into bitmasks about hpd pins that have triggered (*pin_mask*), and which of those pins may be long pulses (*long_mask*). The *long_mask* is ignored if the port corresponding to the pin is not a digital port.

Here, we do hotplug irq storm detection and mitigation, and pass further processing to appropriate bottom halves.

Name

`intel_hpd_init` — initializes and enables hpd support

Synopsis

```
void intel_hpd_init (struct drm_i915_private * dev_priv);
```

Arguments

dev_priv i915 device instance

Description

This function enables the hotplug support. It requires that interrupts have already been enabled with `intel_irq_init_hw`. From this point on hotplug and poll request can run concurrently to other code, so locking rules must be obeyed.

This is a separate step from interrupt enabling to simplify the locking rules in the driver load and resume code.

High Definition Audio

The graphics and audio drivers together support High Definition Audio over HDMI and Display Port. The audio programming sequences are divided into audio codec and controller enable and disable sequences. The graphics driver handles the audio codec sequences, while the audio driver handles the audio controller sequences.

The disable sequences must be performed before disabling the transcoder or port. The enable sequences may only be performed after enabling the transcoder and port, and after completed link training. Therefore the audio enable/disable sequences are part of the modeset sequence.

The codec and controller sequences could be done either parallel or serial, but generally the ELDV/PD change in the codec sequence indicates to the audio driver that the controller sequence should start. Indeed, most of the co-operation between the graphics and audio drivers is handled via audio related registers. (The notable exception is the power management, not covered here.)

The struct `i915_audio_component` is used to interact between the graphics and audio drivers. The struct `i915_audio_component_ops *ops` in it is defined in graphics driver and called in audio driver. The struct `i915_audio_component_audio_ops *audio_ops` is called from i915 driver.

Name

`intel_audio_codec_enable` — Enable the audio codec for HD audio

Synopsis

```
void intel_audio_codec_enable (struct intel_encoder * intel_encoder);
```

Arguments

intel_encoder encoder on which to enable audio

Description

The enable sequences may only be performed after enabling the transcoder and port, and after completed link training.

Name

`intel_audio_codec_disable` — Disable the audio codec for HD audio

Synopsis

```
void intel_audio_codec_disable (struct intel_encoder * intel_encoder);
```

Arguments

intel_encoder encoder on which to disable audio

Description

The disable sequences must be performed before disabling the transcoder or port.

Name

`intel_init_audio` — Set up chip specific audio functions

Synopsis

```
void intel_init_audio (struct drm_device * dev);
```

Arguments

dev drm device

Name

`i915_audio_component_init` — initialize and register the audio component

Synopsis

```
void i915_audio_component_init (struct drm_i915_private * dev_priv);
```

Arguments

dev_priv i915 device instance

Description

This will register with the component framework a child component which will bind dynamically to the `snd_hda_intel` driver's corresponding master component when the latter is registered. During binding the child initializes an instance of `struct i915_audio_component` which it receives from the master. The master can then start to use the interface defined by this struct. Each side can break the binding at any point by deregistering its own component after which each side's component unbind callback is called.

We ignore any error during registration and continue with reduced functionality (i.e. without HDMI audio).

Name

`i915_audio_component_cleanup` — deregister the audio component

Synopsis

```
void i915_audio_component_cleanup (struct drm_i915_private * dev_priv);
```

Arguments

dev_priv i915 device instance

Description

Deregisters the audio component, breaking any existing binding to the corresponding `snd_hda_intel` driver's master component.

Name

struct i915_audio_component_ops — Ops implemented by i915 driver, called by hda driver

Synopsis

```
struct i915_audio_component_ops {
    struct module * owner;
    void (* get_power) (struct device *);
    void (* put_power) (struct device *);
    void (* codec_wake_override) (struct device *, bool enable);
    int (* get_cdclk_freq) (struct device *);
    int (* sync_audio_rate) (struct device *, int port, int rate);
    int (* get_eld) (struct device *, int port, bool *enabled, unsigned char *buf, int
};
```

Members

owner	i915 module
get_power	get the POWER_DOMAIN_AUDIO power well Request the power well to be turned on.
put_power	put the POWER_DOMAIN_AUDIO power well Allow the power well to be turned off.
codec_wake_override	Enable/disable codec wake signal
get_cdclk_freq	Get the Core Display Clock in kHz
sync_audio_rate	set n/cts based on the sample rate Called from audio driver. After audio driver sets the sample rate, it will call this function to set n/cts
get_eld	fill the audio state and ELD bytes for the given port Called from audio driver to get the HDMI/DP audio state of the given digital port, and also fetch ELD bytes to the given pointer. It returns the byte size of the original ELD (not the actually copied size), zero for an invalid ELD, or a negative error code. Note that the returned size may be over <i>max_bytes</i> . Then it implies that only a part of ELD has been copied to the buffer.

Name

struct i915_audio_component_audio_ops — Ops implemented by hda driver, called by i915 driver

Synopsis

```
struct i915_audio_component_audio_ops {  
    void * audio_ptr;  
    void (* pin_eld_notify) (void *audio_ptr, int port);  
};
```

Members

audio_ptr	Pointer to be used in call to pin_eld_notify
pin_eld_notify	Notify the HDA driver that pin sense and/or ELD information has changed Called when the i915 driver has set up audio pipeline or has just begun to tear it down. This allows the HDA driver to update its status accordingly (even when the HDA controller is in power save mode).

Name

struct i915_audio_component — Used for direct communication between i915 and hda drivers

Synopsis

```
struct i915_audio_component {
    struct device * dev;
    int aud_sample_rate[MAX_PORTS];
    const struct i915_audio_component_ops * ops;
    const struct i915_audio_component_audio_ops * audio_ops;
};
```

Members

dev	i915 device, used as parameter for ops
aud_sample_rate[MAX_PORTS]	the array of audio sample rate per port
ops	Ops implemented by i915 driver, called by hda driver
audio_ops	Ops implemented by hda driver, called by i915 driver

Panel Self Refresh PSR (PSR/SRD)

Since Haswell Display controller supports Panel Self-Refresh on display panels which have a remote frame buffer (RFB) implemented according to PSR spec in eDP1.3. PSR feature allows the display to go to lower standby states when system is idle but display is on as it eliminates display refresh request to DDR memory completely as long as the frame buffer for that display is unchanged.

Panel Self Refresh must be supported by both Hardware (source) and Panel (sink).

PSR saves power by caching the framebuffer in the panel RFB, which allows us to power down the link and memory controller. For DSI panels the same idea is called “manual mode”.

The implementation uses the hardware-based PSR support which automatically enters/exits self-refresh mode. The hardware takes care of sending the required DP aux message and could even retrain the link (that part isn't enabled yet though). The hardware also keeps track of any framebuffer changes to know when to exit self-refresh mode again. Unfortunately that part doesn't work too well, hence why the i915 PSR support uses the software framebuffer tracking to make sure it doesn't miss a screen update. For this integration `intel_psr_invalidate` and `intel_psr_flush` get called by the framebuffer tracking code. Note that because of locking issues the self-refresh re-enable code is done from a work queue, which must be correctly synchronized/cancelled when shutting down the pipe."

Name

`intel_psr_enable` — Enable PSR

Synopsis

```
void intel_psr_enable (struct intel_dp * intel_dp);
```

Arguments

intel_dp Intel DP

Description

This function can only be called after the pipe is fully trained and enabled.

Name

`intel_psr_disable` — Disable PSR

Synopsis

```
void intel_psr_disable (struct intel_dp * intel_dp);
```

Arguments

intel_dp Intel DP

Description

This function needs to be called before disabling pipe.

Name

`intel_psr_single_frame_update` — Single Frame Update

Synopsis

```
void intel_psr_single_frame_update (struct drm_device * dev, unsigned  
frontbuffer_bits);
```

Arguments

dev DRM device

frontbuffer_bits frontbuffer plane tracking bits

Description

Some platforms support a single frame update feature that is used to send and update only one frame on Remote Frame Buffer. So far it is only implemented for Valleyview and Cherryview because hardware requires this to be done before a page flip.

Name

intel_psr_invalidate — Invalidate PSR

Synopsis

```
void intel_psr_invalidate (struct drm_device * dev, unsigned front-  
buffer_bits);
```

Arguments

<i>dev</i>	DRM device
<i>frontbuffer_bits</i>	frontbuffer plane tracking bits

Description

Since the hardware frontbuffer tracking has gaps we need to integrate with the software frontbuffer tracking. This function gets called every time frontbuffer rendering starts and a buffer gets dirtied. PSR must be disabled if the frontbuffer mask contains a buffer relevant to PSR.

Dirty frontbuffers relevant to PSR are tracked in `busy_frontbuffer_bits`."

Name

intel_psr_flush — Flush PSR

Synopsis

```
void intel_psr_flush (struct drm_device * dev, unsigned front-  
buffer_bits, enum fb_op_origin origin);
```

Arguments

<i>dev</i>	DRM device
<i>frontbuffer_bits</i>	frontbuffer plane tracking bits
<i>origin</i>	which operation caused the flush

Description

Since the hardware frontbuffer tracking has gaps we need to integrate with the software frontbuffer tracking. This function gets called every time frontbuffer rendering has completed and flushed out to memory. PSR can be enabled again if no other frontbuffer relevant to PSR is dirty.

Dirty frontbuffers relevant to PSR are tracked in `busy_frontbuffer_bits`.

Name

`intel_psr_init` — Init basic PSR work and mutex.

Synopsis

```
void intel_psr_init (struct drm_device * dev);
```

Arguments

dev DRM device

Description

This function is called only once at driver load to initialize basic PSR stuff.

Frame Buffer Compression (FBC)

FBC tries to save memory bandwidth (and so power consumption) by compressing the amount of memory used by the display. It is total transparent to user space and completely handled in the kernel.

The benefits of FBC are mostly visible with solid backgrounds and variation-less patterns. It comes from keeping the memory footprint small and having fewer memory pages opened and accessed for refreshing the display.

i915 is responsible to reserve stolen memory for FBC and configure its offset on proper registers. The hardware takes care of all compress/decompress. However there are many known cases where we have to forcibly disable it to allow proper screen updates.

Name

intel_fbc_is_active — Is FBC active?

Synopsis

```
bool intel_fbc_is_active (struct drm_i915_private * dev_priv);
```

Arguments

dev_priv i915 device instance

Description

This function is used to verify the current state of FBC.

FIXME

This should be tracked in the plane config eventually instead of queried at runtime for most callers.

Name

`__intel_fbc_update` — activate/deactivate FBC as needed, unlocked

Synopsis

```
void __intel_fbc_update (struct intel_crtc * crtc);
```

Arguments

crtc the CRTC that triggered the update

Description

This function completely reevaluates the status of FBC, then activates, deactivates or maintains it on the same state.

Name

`intel_fbc_enable` —

Synopsis

```
void intel_fbc_enable (struct intel_crtc * crtc);
```

Arguments

crtc the CRTC

Description

This function checks if it's possible to enable FBC on the following CRTC, then enables it. Notice that it doesn't activate FBC.

Name

`__intel_fbc_disable` — disable FBC

Synopsis

```
void __intel_fbc_disable (struct drm_i915_private * dev_priv);
```

Arguments

dev_priv i915 device instance

Description

This is the low level function that actually disables FBC. Callers should grab the FBC lock.

Name

`intel_fbc_disable_crtc` — disable FBC if it's associated with `crtc`

Synopsis

```
void intel_fbc_disable_crtc (struct intel_crtc * crtc);
```

Arguments

crtc the CRTC

Description

This function disables FBC if it's associated with the provided CRTC.

Name

`intel_fbc_disable` — globally disable FBC

Synopsis

```
void intel_fbc_disable (struct drm_i915_private * dev_priv);
```

Arguments

dev_priv i915 device instance

Description

This function disables FBC regardless of which CRTC is associated with it.

Name

`intel_fbc_init` — Initialize FBC

Synopsis

```
void intel_fbc_init (struct drm_i915_private * dev_priv);
```

Arguments

dev_priv the i915 device

Description

This function might be called during PM init process.

Display Refresh Rate Switching (DRRS)

Display Refresh Rate Switching (DRRS) is a power conservation feature which enables switching between low and high refresh rates, dynamically, based on the usage scenario. This feature is applicable for internal panels.

Indication that the panel supports DRRS is given by the panel EDID, which would list multiple refresh rates for one resolution.

DRRS is of 2 types - static and seamless. Static DRRS involves changing refresh rate (RR) by doing a full modeset (may appear as a blink on screen) and is used in dock-undock scenario. Seamless DRRS involves changing RR without any visual effect to the user and can be used during normal system usage. This is done by programming certain registers.

Support for static/seamless DRRS may be indicated in the VBT based on inputs from the panel spec.

DRRS saves power by switching to low RR based on usage scenarios.

eDP DRRS:- The implementation is based on frontbuffer tracking implementation. When there is a disturbance on the screen triggered by user activity or a periodic system activity, DRRS is disabled (RR is changed to high RR). When there is no movement on screen, after a timeout of 1 second, a switch to low RR is made. For integration with frontbuffer tracking code, `intel_edp_drrs_invalidate` and `intel_edp_drrs_flush` are called.

DRRS can be further extended to support other internal panels and also the scenario of video playback wherein RR is set based on the rate requested by userspace.

Name

`intel_dp_set_drrs_state` — program registers for RR switch to take effect

Synopsis

```
void intel_dp_set_drrs_state (struct drm_device * dev, int re-  
fresh_rate);
```

Arguments

dev DRM device

refresh_rate RR to be programmed

Description

This function gets called when refresh rate (RR) has to be changed from one frequency to another. Switches can be between high and low RR supported by the panel or to any other RR based on media playback (in this case, RR value needs to be passed from user space).

The caller of this function needs to take a lock on `dev_priv->drrs`.

Name

`intel_edp_drrs_enable` — init drrs struct if supported

Synopsis

```
void intel_edp_drrs_enable (struct intel_dp * intel_dp);
```

Arguments

intel_dp DP struct

Description

Initializes `frontbuffer_bits` and `drrs.dp`

Name

`intel_edp_drrs_disable` — Disable DRRS

Synopsis

```
void intel_edp_drrs_disable (struct intel_dp * intel_dp);
```

Arguments

intel_dp DP struct

Name

`intel_edp_drrs_invalidate` — Disable Idleness DRRS

Synopsis

```
void intel_edp_drrs_invalidate (struct drm_device * dev, unsigned front-  
buffer_bits);
```

Arguments

dev DRM device

frontbuffer_bits frontbuffer plane tracking bits

Description

This function gets called everytime rendering on the given planes start. Hence DRRS needs to be Up-clocked, i.e. (LOW_RR -> HIGH_RR).

Dirty frontbuffers relevant to DRRS are tracked in `busy_frontbuffer_bits`.

Name

intel_edp_drrs_flush — Restart Idleness DRRS

Synopsis

```
void intel_edp_drrs_flush (struct drm_device * dev, unsigned front-  
buffer_bits);
```

Arguments

dev DRM device

frontbuffer_bits frontbuffer plane tracking bits

Description

This function gets called every time rendering on the given planes has completed or flip on a crtc is completed. So DRRS should be upclocked (LOW_RR -> HIGH_RR). And also Idleness detection should be started again, if no other planes are dirty.

Dirty frontbuffers relevant to DRRS are tracked in *busy_frontbuffer_bits*.

Name

`intel_dp_drrs_init` — Init basic DRRS work and mutex.

Synopsis

```
struct drm_display_mode * intel_dp_drrs_init (struct intel_connector *  
intel_connector, struct drm_display_mode * fixed_mode);
```

Arguments

intel_connector eDP connector

fixed_mode preferred mode of panel

Description

This function is called only once at driver load to initialize basic DRRS stuff.

Returns

Downclock mode if panel supports it, else return NULL. DRRS support is determined by the presence of downclock mode (apart from VBT setting).

DPIO

VLV, CHV and BXT have slightly peculiar display PHYs for driving DP/HDMI ports. DPIO is the name given to such a display PHY. These PHYs don't follow the standard programming model using direct MMIO registers, and instead their registers must be accessed through IOSF sideband. VLV has one such PHY for driving ports B and C, and CHV adds another PHY for driving port D. Each PHY responds to specific IOSF-SB port.

Each display PHY is made up of one or two channels. Each channel houses a common lane part which contains the PLL and other common logic. CH0 common lane also contains the IOSF-SB logic for the Common Register Interface (CRI) ie. the DPIO registers. CRI clock must be running when any DPIO registers are accessed.

In addition to having their own registers, the PHYs are also controlled through some dedicated signals from the display controller. These include PLL reference clock enable, PLL enable, and CRI clock selection, for example.

Each channel also has two splines (also called data lanes), and each spline is made up of one Physical Access Coding Sub-Layer (PCS) block and two TX lanes. So each channel has two PCS blocks and four TX lanes. The TX lanes are used as DP lanes or TMDS data/clock pairs depending on the output type.

Additionally the PHY also contains an AUX lane with AUX blocks for each channel. This is used for DP AUX communication, but this fact isn't really relevant for the driver since AUX is controlled from the display controller side. No DPIO registers need to be accessed during AUX communication,

Generally on VLV/CHV the common lane corresponds to the pipe and the spline (PCS/TX) corresponds to the port.

For dual channel PHY (VLV/CHV):

pipe A == CMN/PLL/REF CH0

pipe B == CMN/PLL/REF CH1

port B == PCS/TX CH0

port C == PCS/TX CH1

This is especially important when we cross the streams ie. drive port B with pipe B, or port C with pipe A.

For single channel PHY (CHV):

pipe C == CMN/PLL/REF CH0

port D == PCS/TX CH0

On BXT the entire PHY channel corresponds to the port. That means the PLL is also now associated with the port rather than the pipe, and so the clock needs to be routed to the appropriate transcoder. Port A PLL is directly connected to transcoder EDP and port B/C PLLs can be routed to any transcoder A/B/C.

Note: DDI0 is digital port B, DDI1 is digital port C, and DDI2 is digital port D (CHV) or port A (BXT).

Dual channel PHY (VLV/CHV/BXT) ----- | CH0 | CH1 | | CMN/PLL/REF
 | CMN/PLL/REF | |-----|-----| Display PHY | PCS01 | PCS23 | PCS01 | PCS23 |
 |-----|-----|-----|-----| TX0|TX1|TX2|TX3|TX0|TX1|TX2|TX3| ----- | DDI0 |
 DDI1 | DP/HDMI ports -----

Single channel PHY (CHV/BXT) ----- | CH0 | | CMN/PLL/REF | |-----| Display PHY |
 PCS01 | PCS23 | |-----|-----| TX0|TX1|TX2|TX3| ----- | DDI2 | DP/HDMI port -----

CSR firmware support for DMC

Display Context Save and Restore (CSR) firmware support added from gen9 onwards to drive newly added DMC (Display microcontroller) in display engine to save and restore the state of display engine when it enter into low-power state and comes back to normal.

Firmware loading status will be one of the below states: FW_UNINITIALIZED, FW_LOADED, FW_FAILED.

Once the firmware is written into the registers status will be moved from FW_UNINITIALIZED to FW_LOADED and for any erroneous condition status will be moved to FW_FAILED.

Name

`intel_csr_load_program` — write the firmware from memory to register.

Synopsis

```
void intel_csr_load_program (struct drm_i915_private * dev_priv);
```

Arguments

dev_priv i915 drm device.

Description

CSR firmware is read from a .bin file and kept in internal memory one time. Everytime display comes back from low power state this function is called to copy the firmware from internal memory to registers.

Name

`intel_csr_ucode_init` — initialize the firmware loading.

Synopsis

```
void intel_csr_ucode_init (struct drm_i915_private * dev_priv);
```

Arguments

dev_priv i915 drm device.

Description

This function is called at the time of loading the display driver to read firmware from a .bin file and copied into a internal memory.

Name

`intel_csr_ucode_fini` — unload the CSR firmware.

Synopsis

```
void intel_csr_ucode_fini (struct drm_i915_private * dev_priv);
```

Arguments

dev_priv i915 drm device.

Description

Firmware unloading includes freeing the internal memory and reset the firmware loading status.

Memory Management and Command Submission

This section covers all things related to the GEM implementation in the i915 driver.

Batchbuffer Parsing

Motivation: Certain OpenGL features (e.g. transform feedback, performance monitoring) require userspace code to submit batches containing commands such as `MI_LOAD_REGISTER_IMM` to access various registers. Unfortunately, some generations of the hardware will noop these commands in “unsecure” batches (which includes all userspace batches submitted via i915) even though the commands may be safe and represent the intended programming model of the device.

The software command parser is similar in operation to the command parsing done in hardware for unsecure batches. However, the software parser allows some operations that would be noop'd by hardware, if the parser determines the operation is safe, and submits the batch as “secure” to prevent hardware parsing.

Threats: At a high level, the hardware (and software) checks attempt to prevent granting userspace undue privileges. There are three categories of privilege.

First, commands which are explicitly defined as privileged or which should only be used by the kernel driver. The parser generally rejects such commands, though it may allow some from the drm master process.

Second, commands which access registers. To support correct/enhanced userspace functionality, particularly certain OpenGL extensions, the parser provides a whitelist of registers which userspace may safely access (for both normal and drm master processes).

Third, commands which access privileged memory (i.e. GGTT, HWS page, etc). The parser always rejects such commands.

The majority of the problematic commands fall in the `MI_*` range, with only a few specific commands on each ring (e.g. `PIPE_CONTROL` and `MI_FLUSH_DW`).

Implementation: Each ring maintains tables of commands and registers which the parser uses in scanning batch buffers submitted to that ring.

Since the set of commands that the parser must check for is significantly smaller than the number of commands supported, the parser tables contain only those commands required by the parser. This generally

works because command opcode ranges have standard command length encodings. So for commands that the parser does not need to check, it can easily skip them. This is implemented via a per-ring length decoding vfunc.

Unfortunately, there are a number of commands that do not follow the standard length encoding for their opcode range, primarily amongst the MI_* commands. To handle this, the parser provides a way to define explicit “skip” entries in the per-ring command tables.

Other command table entries map fairly directly to high level categories mentioned above: rejected, master-only, register whitelist. The parser implements a number of checks, including the privileged memory checks, via a general bitmasking mechanism.

Name

`i915_cmd_parser_init_ring` — set cmd parser related fields for a ringbuffer

Synopsis

```
int i915_cmd_parser_init_ring (struct intel_engine_cs * ring);
```

Arguments

ring the ringbuffer to initialize

Description

Optionally initializes fields related to batch buffer command parsing in the struct `intel_engine_cs` based on whether the platform requires software command parsing.

Return

non-zero if initialization fails

Name

`i915_cmd_parser_fini_ring` — clean up cmd parser related fields

Synopsis

```
void i915_cmd_parser_fini_ring (struct intel_engine_cs * ring);
```

Arguments

ring the ringbuffer to clean up

Description

Releases any resources related to command parsing that may have been initialized for the specified ring.

Name

`i915_needs_cmd_parser` — should a given ring use software command parsing?

Synopsis

```
bool i915_needs_cmd_parser (struct intel_engine_cs * ring);
```

Arguments

ring the ring in question

Description

Only certain platforms require software batch buffer command parsing, and only when enabled via module parameter.

Return

true if the ring requires software command parsing

Name

i915_parse_cmds — parse a submitted batch buffer for privilege violations

Synopsis

```
int i915_parse_cmds (struct intel_engine_cs * ring, struct dr-
m_i915_gem_object * batch_obj, struct drm_i915_gem_object * shad-
ow_batch_obj, u32 batch_start_offset, u32 batch_len, bool is_master);
```

Arguments

<i>ring</i>	the ring on which the batch is to execute
<i>batch_obj</i>	the batch buffer in question
<i>shadow_batch_obj</i>	copy of the batch buffer in question
<i>batch_start_offset</i>	byte offset in the batch at which execution starts
<i>batch_len</i>	length of the commands in batch_obj
<i>is_master</i>	is the submitting process the drm master?

Description

Parses the specified batch buffer looking for privilege violations as described in the overview.

Return

non-zero if the parser finds violations or otherwise fails; -EACCES if the batch appears legal but should use hardware parsing

Name

`i915_cmd_parser_get_version` — get the cmd parser version number

Synopsis

```
int i915_cmd_parser_get_version ( void );
```

Arguments

`void` no arguments

Description

The cmd parser maintains a simple increasing integer version number suitable for passing to userspace clients to determine what operations are permitted.

Return

the current version number of the cmd parser

Batchbuffer Pools

In order to submit batch buffers as 'secure', the software command parser must ensure that a batch buffer cannot be modified after parsing. It does this by copying the user provided batch buffer contents to a kernel owned buffer from which the hardware will actually execute, and by carefully managing the address space bindings for such buffers.

The batch pool framework provides a mechanism for the driver to manage a set of scratch buffers to use for this purpose. The framework can be extended to support other uses cases should they arise.

Name

`i915_gem_batch_pool_init` — initialize a batch buffer pool

Synopsis

```
void i915_gem_batch_pool_init (struct drm_device * dev, struct  
i915_gem_batch_pool * pool);
```

Arguments

dev the drm device

pool the batch buffer pool

Name

`i915_gem_batch_pool_fini` — clean up a batch buffer pool

Synopsis

```
void i915_gem_batch_pool_fini (struct i915_gem_batch_pool * pool);
```

Arguments

pool the pool to clean up

Note

Callers must hold the `struct_mutex`.

Name

`i915_gem_batch_pool_get` — allocate a buffer from the pool

Synopsis

```
struct    drm_i915_gem_object    *    i915_gem_batch_pool_get    (struct
i915_gem_batch_pool * pool, size_t size);
```

Arguments

pool the batch buffer pool

size the minimum desired size of the returned buffer

Description

Returns an inactive buffer from *pool* with at least *size* bytes, with the pages pinned. The caller must `i915_gem_object_unpin_pages` on the returned object.

Note

Callers must hold the `struct_mutex`

Return

the buffer object or an error pointer

Logical Rings, Logical Ring Contexts and Execlists

Motivation: GEN8 brings an expansion of the HW contexts: “Logical Ring Contexts”. These expanded contexts enable a number of new abilities, especially “Execlists” (also implemented in this file).

One of the main differences with the legacy HW contexts is that logical ring contexts incorporate many more things to the context's state, like PDPs or ringbuffer control registers:

The reason why PDPs are included in the context is straightforward: as PPGTTs (per-process GTTs) are actually per-context, having the PDPs contained there mean you don't need to do a `ppgtt->switch_mm` yourself, instead, the GPU will do it for you on the context switch.

But, what about the ringbuffer control registers (head, tail, etc..)? shouldn't we just need a set of those per engine command streamer? This is where the name “Logical Rings” starts to make sense: by virtualizing the rings, the engine cs shifts to a new “ring buffer” with every context switch. When you want to submit a workload to the GPU you: A) choose your context, B) find its appropriate virtualized ring, C) write commands to it and then, finally, D) tell the GPU to switch to that context.

Instead of the legacy `MI_SET_CONTEXT`, the way you tell the GPU to switch to a contexts is via a context execution list, ergo “Execlists”.

LRC implementation: Regarding the creation of contexts, we have:

- One global default context. - One local default context for each opened fd. - One local extra context for each context create ioctl call.

Now that ringbuffers belong per-context (and not per-engine, like before) and that contexts are uniquely tied to a given engine (and not reusable, like before) we need:

- One ringbuffer per-engine inside each context. - One backing object per-engine inside each context.

The global default context starts its life with these new objects fully allocated and populated. The local default context for each opened fd is more complex, because we don't know at creation time which engine is going to use them. To handle this, we have implemented a deferred creation of LR contexts:

The local context starts its life as a hollow or blank holder, that only gets populated for a given engine once we receive an execbuffer. If later on we receive another execbuffer ioctl for the same context but a different engine, we allocate/populate a new ringbuffer and context backing object and so on.

Finally, regarding local contexts created using the ioctl call: as they are only allowed with the render ring, we can allocate & populate them right away (no need to defer anything, at least for now).

Execlists implementation: Execlists are the new method by which, on gen8+ hardware, workloads are submitted for execution (as opposed to the legacy, ringbuffer-based, method). This method works as follows:

When a request is committed, its commands (the BB start and any leading or trailing commands, like the seqno breadcrumbs) are placed in the ringbuffer for the appropriate context. The tail pointer in the hardware context is not updated at this time, but instead, kept by the driver in the ringbuffer structure. A structure representing this request is added to a request queue for the appropriate engine: this structure contains a copy of the context's tail after the request was written to the ring buffer and a pointer to the context itself.

If the engine's request queue was empty before the request was added, the queue is processed immediately. Otherwise the queue will be processed during a context switch interrupt. In any case, elements on the queue will get sent (in pairs) to the GPU's ExecLists Submit Port (ELSP, for short) with a globally unique 20-bits submission ID.

When execution of a request completes, the GPU updates the context status buffer with a context complete event and generates a context switch interrupt. During the interrupt handling, the driver examines the events in the buffer: for each context complete event, if the announced ID matches that on the head of the request queue, then that request is retired and removed from the queue.

After processing, if any requests were retired and the queue is not empty then a new execution list can be submitted. The two requests at the front of the queue are next to be submitted but since a context may not occur twice in an execution list, if subsequent requests have the same ID as the first then the two requests must be combined. This is done simply by discarding requests at the head of the queue until either only one requests is left (in which case we use a NULL second context) or the first two requests have unique IDs.

By always executing the first two requests in the queue the driver ensures that the GPU is kept as busy as possible. In the case where a single context completes but a second context is still executing, the request for this second context will be at the head of the queue when we remove the first one. This request will then be resubmitted along with a new request for a different context, which will cause the hardware to continue executing the second request and queue the new request (the GPU detects the condition of a context getting preempted with the same context and optimizes the context switch flow by not doing preemption, but just sampling the new tail pointer).

Name

`intel_sanitise_enable_execlists` — sanitize `i915.enable_execlists`

Synopsis

```
int intel_sanitise_enable_execlists (struct drm_device * dev, int enable_execlists);
```

Arguments

dev DRM device.

enable_execlists value of `i915.enable_execlists` module parameter.

Description

Only certain platforms support Execlists (the prerequisites being support for Logical Ring Contexts and Aliasing PPGTT or better).

Return

1 if Execlists is supported and has to be enabled.

Name

`intel_execlists_ctx_id` — get the Execlists Context ID

Synopsis

```
u32 intel_execlists_ctx_id (struct drm_i915_gem_object * ctx_obj);
```

Arguments

ctx_obj Logical Ring Context backing object.

Description

Do not confuse with `ctx->id`! Unfortunately we have a name overload

here

the old context ID we pass to userspace as a handler so that they can refer to a context, and the new context ID we pass to the ELSP so that the GPU can inform us of the context status via interrupts.

Return

20-bits globally unique context ID.

Name

`intel_lrc_irq_handler` — handle Context Switch interrupts

Synopsis

```
void intel_lrc_irq_handler (struct intel_engine_cs * ring);
```

Arguments

ring Engine Command Streamer to handle.

Description

Check the unread Context Status Buffers and manage the submission of new contexts to the ELSP accordingly.

Name

`intel_logical_ring_begin` — prepare the logical ringbuffer to accept some commands

Synopsis

```
int intel_logical_ring_begin (struct drm_i915_gem_request * req, int
num_dwords);
```

Arguments

req The request to start some new work for

num_dwords number of DWORDs that we plan to write to the ringbuffer.

Description

The ringbuffer might not be ready to accept the commands right away (maybe it needs to be wrapped, or wait a bit for the tail to be updated). This function takes care of that and also preallocates a request (every workload submission is still mediated through requests, same as it did with legacy ringbuffer submission).

Return

non-zero if the ringbuffer is not ready to be written to.

Name

`intel_execlists_submission` — submit a batchbuffer for execution, Execlists style

Synopsis

```
int intel_execlists_submission (struct i915_execbuffer_params * params,
                                struct drm_i915_gem_execbuffer2 * args, struct list_head * vmas);
```

Arguments

params -- undescribed --

args execbuffer call arguments.

*vm*as list of vmas.

Description

This is the evil twin version of `i915_gem_ringbuffer_submission`. It abstracts away the submission details of the execbuffer ioctl call.

Return

non-zero if the submission fails.

Name

`gen8_init_indirectctx_bb` — initialize indirect ctx batch with WA

Synopsis

```
int gen8_init_indirectctx_bb (struct intel_engine_cs * ring, struct
i915_wa_ctx_bb * wa_ctx, uint32_t *const batch, uint32_t * offset);
```

Arguments

ring only applicable for RCS

wa_ctx structure representing *wa_ctx*

batch page in which WA are loaded

offset This field specifies the start of the batch, it should be cache-aligned otherwise it is adjusted accordingly. Typically we only have one *indirect_ctx* and *per_ctx* batch buffer which are initialized at the beginning and shared across all contexts but this field helps us to have multiple batches at different offsets and select them based on a criteria. At the moment this batch always start at the beginning of the page and at this point we don't have multiple *wa_ctx* batch buffers.

offset

specifies start of the batch, should be cache-aligned. This is updated with the offset value received as input.

size

size of the batch in DWORDS but HW expects in terms of cachelines

Description

The number of WA applied are not known at the beginning; we use this field to return the no of DWORDS written.

It is to be noted that this batch does not contain `MI_BATCH_BUFFER_END` so it adds NOOPs as padding to make it cacheline aligned. `MI_BATCH_BUFFER_END` will be added to *perctx* batch and both of them together makes a complete batch buffer.

Return

non-zero if we exceed the `PAGE_SIZE` limit.

Name

`gen8_init_perctx_bb` — initialize per ctx batch with WA

Synopsis

```
int gen8_init_perctx_bb (struct intel_engine_cs * ring, struct
i915_wa_ctx_bb * wa_ctx, uint32_t *const batch, uint32_t * offset);
```

Arguments

ring only applicable for RCS

wa_ctx structure representing wa_ctx

batch page in which WA are loaded

offset This field specifies the start of this batch. This batch is started immediately after indirect_ctx batch. Since we ensure that indirect_ctx ends on a cacheline this batch is aligned automatically.

offset

specifies start of the batch, should be cache-aligned.

size

size of the batch in DWORDS but HW expects in terms of cachelines

Description

The number of DWORDS written are returned using this field.

This batch is terminated with `MI_BATCH_BUFFER_END` and so we need not add padding to align it with cacheline as padding after `MI_BATCH_BUFFER_END` is redundant.

Name

`intel_logical_ring_cleanup` — deallocate the Engine Command Streamer

Synopsis

```
void intel_logical_ring_cleanup (struct intel_engine_cs * ring);
```

Arguments

ring Engine Command Streamer.

Name

`intel_logical_rings_init` — allocate, populate and init the Engine Command Streamers

Synopsis

```
int intel_logical_rings_init (struct drm_device * dev);
```

Arguments

dev DRM device.

Description

This function inits the engines for an Execlists submission style (the equivalent in the legacy ringbuffer submission world would be `i915_gem_init_rings`). It does it only for those engines that are present in the hardware.

Return

non-zero if the initialization failed.

Name

`intel_lr_context_free` — free the LRC specific bits of a context

Synopsis

```
void intel_lr_context_free (struct intel_context * ctx);
```

Arguments

ctx the LR context to free.

The real context freeing is done in `i915_gem_context_free`

this only

takes care of the bits that are LRC related

the per-engine backing objects and the logical ringbuffer.

Name

`intel_lr_context_deferred_alloc` — create the LRC specific bits of a context

Synopsis

```
int intel_lr_context_deferred_alloc (struct intel_context * ctx, struct
intel_engine_cs * ring);
```

Arguments

ctx LR context to create.

ring engine to be used with the context.

Description

This function can be called more than once, with different engines, if we plan to use the context with them. The context backing objects and the ringbuffers (specially the ringbuffer backing objects) suck a lot of memory up, and that's why

the creation is a deferred call

it's better to make sure first that we need to use a given ring with the context.

Return

non-zero on error.

Global GTT views

Background and previous state

Historically objects could exist (be bound) in global GTT space only as singular instances with a view representing all of the object's backing pages in a linear fashion. This view will be called a normal view.

To support multiple views of the same object, where the number of mapped pages is not equal to the backing store, or where the layout of the pages is not linear, concept of a GGTT view was added.

One example of an alternative view is a stereo display driven by a single image. In this case we would have a framebuffer looking like this (2x2 pages):

12 34

Above would represent a normal GGTT view as normally mapped for GPU or CPU rendering. In contrast, fed to the display engine would be an alternative view which could look something like this:

1212 3434

In this example both the size and layout of pages in the alternative view is different from the normal view.

Implementation and usage

GGTT views are implemented using VMAs and are distinguished via enum `i915_gggtt_view_type` and struct `i915_gggtt_view`.

A new flavour of core GEM functions which work with GGTT bound objects were added with the `_ggtt_` infix, and sometimes with `_view` postfix to avoid renaming in large amounts of code. They take the struct `i915_ggtt_view` parameter encapsulating all metadata required to implement a view.

As a helper for callers which are only interested in the normal view, globally const `i915_ggtt_view_normal` singleton instance exists. All old core GEM API functions, the ones not taking the view parameter, are operating on, or with the normal GGTT view.

Code wanting to add or use a new GGTT view needs to:

1. Add a new enum with a suitable name.
2. Extend the metadata in the `i915_ggtt_view` structure if required.
3. Add support to `i915_get_vma_pages`.

New views are required to build a scatter-gather table from within the `i915_get_vma_pages` function. This table is stored in the `vma.ggtt_view` and exists for the lifetime of an VMA.

Core API is designed to have copy semantics which means that passed in struct `i915_ggtt_view` does not need to be persistent (left around after calling the core API functions).

Name

`gen8_ppgtt_alloc_pagetabs` — Allocate page tables for VA range.

Synopsis

```
int gen8_ppgtt_alloc_pagetabs (struct i915_address_space * vm, struct
i915_page_directory * pd, uint64_t start, uint64_t length, unsigned long
* new_pts);
```

Arguments

<i>vm</i>	Master vm structure.
<i>pd</i>	Page directory for this address range.
<i>start</i>	Starting virtual address to begin allocations.
<i>length</i>	Size of the allocations.
<i>new_pts</i>	Bitmap set by function with new allocations. Likely used by the caller to free on error.

Description

Allocate the required number of page tables. Extremely similar to `gen8_ppgtt_alloc_page_directories`. The main difference is here we are limited by the page directory boundary (instead of the page directory pointer). That boundary is 1GB virtual. Therefore, unlike `gen8_ppgtt_alloc_page_directories`, it is possible, and likely that the caller will need to use multiple calls of this function to achieve the appropriate allocation.

Return

0 if success; negative error code otherwise.

Name

`gen8_ppgtt_alloc_page_directories` — Allocate page directories for VA range.

Synopsis

```
int gen8_ppgtt_alloc_page_directories (struct i915_address_space * vm,
    struct i915_page_directory_pointer * pdp, uint64_t start, uint64_t
    length, unsigned long * new_pds);
```

Arguments

<i>vm</i>	Master vm structure.
<i>pdp</i>	Page directory pointer for this address range.
<i>start</i>	Starting virtual address to begin allocations.
<i>length</i>	Size of the allocations.
<i>new_pds</i>	Bitmap set by function with new allocations. Likely used by the caller to free on error.

Description

Allocate the required number of page directories starting at the pde index of *start*, and ending at the pde index *start + length*. This function will skip over already allocated page directories within the range, and only allocate new ones, setting the appropriate pointer within the pdp as well as the correct position in the bitmap *new_pds*.

The function will only allocate the pages within the range for a give page directory pointer. In other words, if *start + length* straddles a virtually addressed PDP boundary (512GB for 4k pages), there will be more allocations required by the caller, This is not currently possible, and the BUG in the code will prevent it.

Return

0 if success; negative error code otherwise.

Name

`gen8_ppgtt_alloc_page_dirpointers` — Allocate pdps for VA range.

Synopsis

```
int gen8_ppgtt_alloc_page_dirpointers (struct i915_address_space * vm,
struct i915_pml4 * pml4, uint64_t start, uint64_t length, unsigned long
* new_pdps);
```

Arguments

<i>vm</i>	Master vm structure.
<i>pml4</i>	Page map level 4 for this address range.
<i>start</i>	Starting virtual address to begin allocations.
<i>length</i>	Size of the allocations.
<i>new_pdps</i>	Bitmap set by function with new allocations. Likely used by the caller to free on error.

Description

Allocate the required number of page directory pointers. Extremely similar to `gen8_ppgtt_alloc_page_directories` and `gen8_ppgtt_alloc_pagetabs`. The main difference is here we are limited by the pml4 boundary (instead of the page directory pointer).

Return

0 if success; negative error code otherwise.

Name

`i915_vma_bind` — Sets up PTEs for an VMA in it's corresponding address space.

Synopsis

```
int i915_vma_bind (struct i915_vma * vma, enum i915_cache_level
cache_level, u32 flags);
```

Arguments

<i>vma</i>	VMA to map
<i>cache_level</i>	mapping cache level
<i>flags</i>	flags like global or local mapping

Description

DMA addresses are taken from the scatter-gather table of this object (or of this VMA in case of non-default GGTT views) and PTE entries set up. Note that DMA addresses are also the only part of the SG table we care about.

Name

`i915_ggtt_view_size` — Get the size of a GGTT view.

Synopsis

```
size_t i915_ggtt_view_size (struct drm_i915_gem_object * obj, const  
struct i915_ggtt_view * view);
```

Arguments

obj Object the view is of.

view The view in question.

Description

return The size of the GGTT view in bytes.

GTT Fences and Swizzling

Name

`i915_gem_object_put_fence` — force-remove fence for an object

Synopsis

```
int i915_gem_object_put_fence (struct drm_i915_gem_object * obj);
```

Arguments

obj object to map through a fence reg

Description

This function force-removes any fence from the given object, which is useful if the kernel wants to do untiled GTT access.

Returns

0 on success, negative error code on failure.

Name

`i915_gem_object_get_fence` — set up fencing for an object

Synopsis

```
int i915_gem_object_get_fence (struct drm_i915_gem_object * obj);
```

Arguments

obj object to map through a fence reg

Description

When mapping objects through the GTT, userspace wants to be able to write to them without having to worry about swizzling if the object is tiled. This function walks the fence regs looking for a free one for *obj*, stealing one if it can't find any.

It then sets up the reg based on the object's properties: address, pitch and tiling format.

For an untiled surface, this removes any existing fence.

Returns

0 on success, negative error code on failure.

Name

`i915_gem_object_pin_fence` — pin fencing state

Synopsis

```
bool i915_gem_object_pin_fence (struct drm_i915_gem_object * obj);
```

Arguments

obj object to pin fencing for

Description

This pins the fencing state (whether tiled or untiled) to make sure the object is ready to be used as a scanout target. Fencing status must be synchronize first by calling `i915_gem_object_get_fence`:

The resulting fence pin reference must be released again with `i915_gem_object_unpin_fence`.

Returns

True if the object has a fence, false otherwise.

Name

`i915_gem_object_unpin_fence` — unpin fencing state

Synopsis

```
void i915_gem_object_unpin_fence (struct drm_i915_gem_object * obj);
```

Arguments

obj object to unpin fencing for

Description

This releases the fence pin reference acquired through `i915_gem_object_pin_fence`. It will handle both objects with and without an attached fence correctly, callers do not need to distinguish this.

Name

`i915_gem_restore_fences` — restore fence state

Synopsis

```
void i915_gem_restore_fences (struct drm_device * dev);
```

Arguments

dev DRM device

Description

Restore the hw fence state to match the software tracking again, to be called after a gpu reset and on resume.

Name

`i915_gem_detect_bit_6_swizzle` — detect bit 6 swizzling pattern

Synopsis

```
void i915_gem_detect_bit_6_swizzle (struct drm_device * dev);
```

Arguments

dev DRM device

Description

Detects bit 6 swizzling of address lookup between IGD access and CPU access through main memory.

Name

`i915_gem_object_do_bit_17_swizzle` — fixup bit 17 swizzling

Synopsis

```
void i915_gem_object_do_bit_17_swizzle (struct drm_i915_gem_object *  
obj);
```

Arguments

obj i915 GEM buffer object

Description

This function fixes up the swizzling in case any page frame number for this object has changed in bit 17 since that state has been saved with `i915_gem_object_save_bit_17_swizzle`.

This is called when pinning backing storage again, since the kernel is free to move unpinned backing storage around (either by directly moving pages or by swapping them out and back in again).

Name

`i915_gem_object_save_bit_17_swizzle` — save bit 17 swizzling

Synopsis

```
void i915_gem_object_save_bit_17_swizzle (struct drm_i915_gem_object *  
obj);
```

Arguments

obj i915 GEM buffer object

Description

This function saves the bit 17 of each page frame number so that swizzling can be fixed up later on with `i915_gem_object_do_bit_17_swizzle`. This must be called before the backing storage can be unpinned.

Global GTT Fence Handling

Important to avoid confusions: “fences” in the i915 driver are not execution fences used to track command completion but hardware detiler objects which wrap a given range of the global GTT. Each platform has only a fairly limited set of these objects.

Fences are used to detile GTT memory mappings. They're also connected to the hardware frontbuffer render tracking and hence interact with frontbuffer compression. Furthermore on older platforms fences are required for tiled objects used by the display engine. They can also be used by the render engine - they're required for blitter commands and are optional for render commands. But on gen4+ both display (with the exception of fbc) and rendering have their own tiling state bits and don't need fences.

Also note that fences only support X and Y tiling and hence can't be used for the fancier new tiling formats like W, Ys and Yf.

Finally note that because fences are such a restricted resource they're dynamically associated with objects. Furthermore fence state is committed to the hardware lazily to avoid unnecessary stalls on gen2/3. Therefore code must explicitly call `i915_gem_object_get_fence` to synchronize fencing status for cpu access. Also note that some code wants an unfenced view, for those cases the fence can be removed forcefully with `i915_gem_object_put_fence`.

Internally these functions will synchronize with userspace access by removing CPU ptes into GTT mmaps (not the GTT ptes themselves) as needed.

Hardware Tiling and Swizzling Details

The idea behind tiling is to increase cache hit rates by rearranging pixel data so that a group of pixel accesses are in the same cacheline. Performance improvement from doing this on the back/depth buffer are on the order of 30%.

Intel architectures make this somewhat more complicated, though, by adjustments made to addressing of data when the memory is in interleaved mode (matched pairs of DIMMS) to improve memory bandwidth. For interleaved memory, the CPU sends every sequential 64 bytes to an alternate memory channel so it can get the bandwidth from both.

The GPU also rearranges its accesses for increased bandwidth to interleaved memory, and it matches what the CPU does for non-tiled. However, when tiled it does it a little differently, since one walks addresses not just in the X direction but also Y. So, along with alternating channels when bit 6 of the address flips, it also alternates when other bits flip -- Bits 9 (every 512 bytes, an X tile scanline) and 10 (every two X tile scanlines) are common to both the 915 and 965-class hardware.

The CPU also sometimes XORs in higher bits as well, to improve bandwidth doing strided access like we do so frequently in graphics. This is called “Channel XOR Randomization” in the MCH documentation. The result is that the CPU is XORing in either bit 11 or bit 17 to bit 6 of its address decode.

All of this bit 6 XORing has an effect on our memory management, as we need to make sure that the 3d driver can correctly address object contents.

If we don't have interleaved memory, all tiling is safe and no swizzling is required.

When bit 17 is XORed in, we simply refuse to tile at all. Bit 17 is not just a page offset, so as we page an object out and back in, individual pages in it will have different bit 17 addresses, resulting in each 64 bytes being swapped with its neighbor!

Otherwise, if interleaved, we have to tell the 3d driver what the address swizzling it needs to do is, since it's writing with the CPU to the pages (bit 6 and potentially bit 11 XORed in), and the GPU is reading from the pages (bit 6, 9, and 10 XORed in), resulting in a cumulative bit swizzling required by the CPU of XORing in bit 6, 9, 10, and potentially 11, in order to match what the GPU expects.

Object Tiling IOCTLs

Name

`i915_gem_set_tiling` — IOCTL handler to set tiling mode

Synopsis

```
int i915_gem_set_tiling (struct drm_device * dev, void * data, struct
drm_file * file);
```

Arguments

dev DRM device

data data pointer for the ioctl

file DRM file for the ioctl call

Description

Sets the tiling mode of an object, returning the required swizzling of bit 6 of addresses in the object.

Called by the user via ioctl.

Returns

Zero on success, negative errno on failure.

Name

`i915_gem_get_tiling` — IOCTL handler to get tiling mode

Synopsis

```
int i915_gem_get_tiling (struct drm_device * dev, void * data, struct
drm_file * file);
```

Arguments

dev DRM device

data data pointer for the ioctl

file DRM file for the ioctl call

Description

Returns the current tiling mode and required bit 6 swizzling for the object.

Called by the user via ioctl.

Returns

Zero on success, negative errno on failure.

`i915_gem_set_tiling` and `i915_gem_get_tiling` is the userspace interface to declare fence register requirements.

In principle GEM doesn't care at all about the internal data layout of an object, and hence it also doesn't care about tiling or swizzling. There's two exceptions:

- For X and Y tiling the hardware provides detilers for CPU access, so called fences. Since there's only a limited amount of them the kernel must manage these, and therefore userspace must tell the kernel the object tiling if it wants to use fences for detiling. - On gen3 and gen4 platforms have a swizzling pattern for tiled objects which depends upon the physical page frame number. When swapping such objects the page frame number might change and the kernel must be able to fix this up and hence now the tiling. Note that on a subset of platforms with asymmetric memory channel population the swizzling pattern changes in an unknown way, and for those the kernel simply forbids swapping completely.

Since neither of this applies for new tiling layouts on modern platforms like W, Ys and Yf tiling GEM only allows object tiling to be set to X or Y tiled. Anything else can be handled in userspace entirely without the kernel's involvement.

Buffer Object Eviction

This section documents the interface functions for evicting buffer objects to make space available in the virtual gpu address spaces. Note that this is mostly orthogonal to shrinking buffer objects caches, which has the goal to make main memory (shared with the gpu through the unified memory architecture) available.

Name

i915_gem_evict_something — Evict vmas to make room for binding a new one

Synopsis

```
int i915_gem_evict_something (struct drm_device * dev, struct i915_ad-
dress_space * vm, int min_size, unsigned alignment, unsigned cache_lev-
el, unsigned long start, unsigned long end, unsigned flags);
```

Arguments

<i>dev</i>	drm_device
<i>vm</i>	address space to evict from
<i>min_size</i>	size of the desired free space
<i>alignment</i>	alignment constraint of the desired free space
<i>cache_level</i>	cache_level for the desired space
<i>start</i>	start (inclusive) of the range from which to evict objects
<i>end</i>	end (exclusive) of the range from which to evict objects
<i>flags</i>	additional flags to control the eviction algorithm

Description

This function will try to evict vmas until a free space satisfying the requirements is found. Callers must check first whether any such hole exists already before calling this function.

This function is used by the object/vma binding code.

Since this function is only used to free up virtual address space it only ignores pinned vmas, and not object where the backing storage itself is pinned. Hence obj->pages_pin_count does not protect against eviction.

To clarify

This is for freeing up virtual address space, not for freeing memory in e.g. the shrinker.

Name

`i915_gem_evict_vm` — Evict all idle vmas from a vm

Synopsis

```
int i915_gem_evict_vm (struct i915_address_space * vm, bool do_idle);
```

Arguments

`vm` Address space to cleanse

`do_idle` Boolean directing whether to idle first.

Description

This function evicts all idles vmas from a vm. If all unpinned vmas should be evicted the `do_idle` needs to be set to true.

This is used by the execbuf code as a last-ditch effort to defragment the address space.

To clarify

This is for freeing up virtual address space, not for freeing memory in e.g. the shrinker.

Buffer Object Memory Shrinking

This section documents the interface function for shrinking memory usage of buffer object caches. Shrinking is used to make main memory available. Note that this is mostly orthogonal to evicting buffer objects, which has the goal to make space in gpu virtual address spaces.

Name

`i915_gem_shrink` — Shrink buffer object caches

Synopsis

```
unsigned long i915_gem_shrink (struct drm_i915_private * dev_priv, un-  
signed long target, unsigned flags);
```

Arguments

dev_priv i915 device

target amount of memory to make available, in pages

flags control flags for selecting cache types

Description

This function is the main interface to the shrinker. It will try to release up to *target* pages of main memory backing storage from buffer objects. Selection of the specific caches can be done with *flags*. This is e.g. useful when purgeable objects should be removed from caches preferentially.

Note that it's not guaranteed that released amount is actually available as free system memory - the pages might still be in-used to due to other reasons (like cpu mmaps) or the mm core has reused them before we could grab them. Therefore code that needs to explicitly shrink buffer objects caches (e.g. to avoid deadlocks in memory reclaim) must fall back to `i915_gem_shrink_all`.

Also note that any kind of pinning (both per-vma address space pins and backing storage pins at the buffer object level) result in the shrinker code having to skip the object.

Returns

The number of pages of backing storage actually released.

Name

`i915_gem_shrink_all` — Shrink buffer object caches completely

Synopsis

```
unsigned long i915_gem_shrink_all (struct drm_i915_private * dev_priv);
```

Arguments

dev_priv i915 device

Description

This is a simple wrapper around `i915_gem_shrink` to aggressively shrink all caches completely. It also first waits for and retires all outstanding requests to also be able to release backing storage for active objects.

This should only be used in code to intentionally quiescent the gpu or as a last-ditch effort when memory seems to have run out.

Returns

The number of pages of backing storage actually released.

Name

`i915_gem_shrinker_init` — Initialize i915 shrinker

Synopsis

```
void i915_gem_shrinker_init (struct drm_i915_private * dev_priv);
```

Arguments

dev_priv i915 device

Description

This function registers and sets up the i915 shrinker and OOM handler.

GuC

GuC-specific firmware loader

`intel_guc`: Top level structure of guc. It handles firmware loading and manages client pool and doorbells. `intel_guc` owns a `i915_guc_client` to replace the legacy ExecList submission.

Firmware versioning: The firmware build process will generate a version header file with major and minor version defined. The versions are built into CSS header of firmware. i915 kernel driver set the minimal firmware version required per platform. The firmware installation package will install (symbolic link) proper version of firmware.

GuC address space: GuC does not allow any gfx GGTT address that falls into range [0, WOPCM_TOP), which is reserved for Boot ROM, SRAM and WOPCM. Currently this top address is 512K. In order to exclude 0-512K address space from GGTT, all gfx objects used by GuC is pinned with PIN_OFFSET_BIAS along with size of WOPCM.

Firmware log: Firmware log is enabled by setting `i915.guc_log_level` to non-negative level. Log data is printed out via reading debugfs `i915_guc_log_dump`. Reading from `i915_guc_load_status` will print out firmware loading status and scratch registers value.

Name

`intel_guc_ucose_load` — load GuC uCode into the device

Synopsis

```
int intel_guc_ucose_load (struct drm_device * dev);
```

Arguments

dev drm device

Description

Called from `gem_init_hw` during driver loading and also after a GPU reset.

The firmware image should have already been fetched into memory by the earlier call to `intel_guc_ucose_init`, so here we need only check that is succeeded, and then transfer the image to the h/w.

Return

non-zero code on error

Name

`intel_guc_ucose_init` — define parameters and fetch firmware

Synopsis

```
void intel_guc_ucose_init (struct drm_device * dev);
```

Arguments

dev drm device

Description

Called early during driver load, but after GEM is initialised.

The firmware will be transferred to the GuC's memory later, when `intel_guc_ucose_load` is called.

Name

`intel_guc_ucose_fini` — clean up all allocated resources

Synopsis

```
void intel_guc_ucose_fini (struct drm_device * dev);
```

Arguments

dev drm device

GuC-based command submission

`i915_guc_client`: We use the term client to avoid confusion with contexts. A `i915_guc_client` is equivalent to GuC object `guc_context_desc`. This context descriptor is allocated from a pool of 1024 entries. Kernel driver will allocate doorbell and workqueue for it. Also the process descriptor (`guc_process_desc`), which is mapped to client space. So the client can write Work Item then ring the doorbell.

To simplify the implementation, we allocate one gem object that contains all pages for doorbell, process descriptor and workqueue.

The Scratch registers: There are 16 MMIO-based registers start from 0xC180. The kernel driver writes a value to the action register (`SOFT_SCRATCH_0`) along with any data. It then triggers an interrupt on the GuC via another register write (0xC4C8). Firmware writes a success/fail code back to the action register after processes the request. The kernel driver polls waiting for this update and then proceeds. See `host2guc_action`

Doorbells: Doorbells are interrupts to uKernel. A doorbell is a single cache line (QW) mapped into process space.

Work Items: There are several types of work items that the host may place into a workqueue, each with its own requirements and limitations. Currently only `WQ_TYPE_INORDER` is needed to support legacy submission via GuC, which represents in-order queue. The kernel driver packs ring tail pointer and an ELSP context descriptor dword into Work Item. See `guc_add_workqueue_item`

Name

`i915_guc_submit` — Submit commands through GuC

Synopsis

```
int i915_guc_submit (struct i915_guc_client * client, struct dr-  
m_i915_gem_request * rq);
```

Arguments

client the guc client where commands will go through

rq request associated with the commands

Return

0 if succeed

Name

`gem_allocate_guc_obj` — Allocate gem object for GuC usage

Synopsis

```
struct drm_i915_gem_object * gem_allocate_guc_obj (struct drm_device *  
dev, u32 size);
```

Arguments

dev drm device

size size of object

Description

This is a wrapper to create a gem obj. In order to use it inside GuC, the object needs to be pinned lifetime. Also we must pin it to gtt space other than [0, GUC_WOPCM_TOP) because this range is reserved inside GuC.

Return

A `drm_i915_gem_object` if successful, otherwise NULL.

Name

`gem_release_guc_obj` — Release gem object allocated for GuC usage

Synopsis

```
void gem_release_guc_obj (struct drm_i915_gem_object * obj);
```

Arguments

obj gem obj to be released

Name

`guc_client_alloc` — Allocate an `i915_guc_client`

Synopsis

```
struct i915_guc_client * guc_client_alloc (struct drm_device * dev,
uint32_t priority, struct intel_context * ctx);
```

Arguments

<i>dev</i>	drm device
<i>priority</i>	four levels priority <code>_CRITICAL</code> , <code>_HIGH</code> , <code>_NORMAL</code> and <code>_LOW</code> The kernel client to replace ExecList submission is created with <code>NORMAL</code> priority. Priority of a client for scheduler can be <code>HIGH</code> , while a preemption context can use <code>CRITICAL</code> .
<i>ctx</i>	the context that owns the client (we use the default render context)

Return

An `i915_guc_client` object if success.

Name

`intel_guc_suspend` — notify GuC entering suspend state

Synopsis

```
int intel_guc_suspend (struct drm_device * dev);
```

Arguments

dev drm device

Name

intel_guc_resume — notify GuC resuming from suspend state

Synopsis

```
int intel_guc_resume (struct drm_device * dev);
```

Arguments

dev drm device

GuC Firmware Layout

The GuC firmware layout looks like this:

```
+-----+ |   guc_css_header   | |   contains   major/minor   version   |
+-----+ |   uCode   | +-----+ |   RSA   signature   |
+-----+ | modulus key | +-----+ | exponent val |
+-----+
```

The firmware may or may not have modulus key and exponent data. The header, uCode and RSA signature are must-have components that will be used by driver. Length of each components, which is all in dwords, can be found in header. In the case that modulus and exponent are not present in fw, a.k.a truncated image, the length value still appears in header.

Driver will do some basic fw size validation based on the following rules:

1. Header, uCode and RSA are must-have components. 2. All firmware components, if they present, are in the sequence illustrated in the layout table above. 3. Length info of each component can be found in header, in dwords. 4. Modulus and exponent key are not required by driver. They may not appear in fw. So driver will load a truncated firmware in this case.

Tracing

This sections covers all things related to the tracepoints implemented in the i915 driver.

i915_ppgtt_create and i915_ppgtt_release

With full ppgtt enabled each process using drm will allocate at least one translation table. With these traces it is possible to keep track of the allocation and of the lifetime of the tables; this can be used during testing/debug to verify that we are not leaking ppgtts. These traces identify the ppgtt through the vm pointer, which is also printed by the i915_vma_bind and i915_vma_unbind tracepoints.

i915_context_create and i915_context_free

These tracepoints are used to track creation and deletion of contexts. If full ppgtt is enabled, they also print the address of the vm assigned to the context.

switch_mm

This tracepoint allows tracking of the mm switch, which is an important point in the lifetime of the vm in the legacy submission path. This tracepoint is called only if full ppgtt is enabled.

Part III. vga_switcheroo

vga_switcheroo is the Linux subsystem for laptop hybrid graphics. These come in two flavors:

* muxed: Dual GPUs with a multiplexer chip to switch outputs between GPUs. * muxless: Dual GPUs but only one of them is connected to outputs. The other one is merely used to offload rendering, its results are copied over PCIe into the framebuffer. On Linux this is supported with DRI PRIME.

Hybrid graphics started to appear in the late Naughties and were initially all muxed. Newer laptops moved to a muxless architecture for cost reasons. A notable exception is the MacBook Pro which continues to use a mux. Muxes come with varying capabilities: Some switch only the panel, others can also switch external displays. Some switch all display pins at once while others can switch just the DDC lines. (To allow EDID probing for the inactive GPU.) Also, muxes are often used to cut power to the discrete GPU while it is not used.

DRM drivers register GPUs with vga_switcheroo, these are henceforth called clients. The mux is called the handler. Muxless machines also register a handler to control the power state of the discrete GPU, its ->switchto callback is a no-op for obvious reasons. The discrete GPU is often equipped with an HDA controller for the HDMI/DP audio signal, this will also register as a client so that vga_switcheroo can take care of the correct suspend/resume order when changing the discrete GPU's power state. In total there can thus be up to three clients: Two vga clients (GPUs) and one audio client (on the discrete GPU). The code is mostly prepared to support machines with more than two GPUs should they become available. The GPU to which the outputs are currently switched is called the active client in vga_switcheroo parlance. The GPU not in use is the inactive client.

Table of Contents

5. Modes of Use	770
Manual switching and manual power control	770
Driver power control	770
6. API	771
Public functions	771
Public structures	782
Public constants	784
Private structures	786
7. Handlers	789
apple-gmux Handler	789
Graphics mux	789
Power control	790
Backlight control	790

Chapter 5. Modes of Use

Manual switching and manual power control

In this mode of use, the file `/sys/kernel/debug/vgaswitcheroo/switch` can be read to retrieve the current `vga_switcheroo` state and commands can be written to it to change the state. The file appears as soon as two GPU drivers and one handler have registered with `vga_switcheroo`. The following commands are understood:

* OFF: Power off the device not in use. * ON: Power on the device not in use. * IGD: Switch to the integrated graphics device. Power on the integrated GPU if necessary, power off the discrete GPU. Pre-requisite is that no user space processes (e.g. Xorg, alsactl) have opened device files of the GPUs or the audio client. If the switch fails, the user may invoke `lsof(8)` or `fuser(1)` on `/dev/dri/` and `/dev/snd/controlC1` to identify processes blocking the switch. * DIS: Switch to the discrete graphics device. * DIGD: Delayed switch to the integrated graphics device. This will perform the switch once the last user space process has closed the device files of the GPUs and the audio client. * DDIS: Delayed switch to the discrete graphics device. * MIGD: Mux-only switch to the integrated graphics device. Does not remap console or change the power state of either gpu. If the integrated GPU is currently off, the screen will turn black. If it is on, the screen will show whatever happens to be in VRAM. Either way, the user has to blindly enter the command to switch back. * MDIS: Mux-only switch to the discrete graphics device.

For GPUs whose power state is controlled by the driver's runtime pm, the ON and OFF commands are a no-op (see next section).

For muxless machines, the IGD/DIS, DIGD/DDIS and MIGD/MDIS commands should not be used.

Driver power control

In this mode of use, the discrete GPU automatically powers up and down at the discretion of the driver's runtime pm. On muxed machines, the user may still influence the muxer state by way of the debugfs interface, however the ON and OFF commands become a no-op for the discrete GPU.

This mode is the default on Nvidia HybridPower/Optimus and ATI PowerXpress. Specifying `nouveau.runpm=0`, `radeon.runpm=0` or `amdgpu.runpm=0` on the kernel command line disables it.

When the driver decides to power up or down, it notifies `vga_switcheroo` thereof so that it can (a) power the audio device on the GPU up or down, and (b) update its internal power state representation for the device. This is achieved by `vga_switcheroo_set_dynamic_switch`.

After the GPU has been suspended, the handler needs to be called to cut power to the GPU. Likewise it needs to reinstate power before the GPU can resume. This is achieved by `vga_switcheroo_init_domain_pm_ops`, which augments the GPU's suspend/resume functions by the requisite calls to the handler.

When the audio device resumes, the GPU needs to be woken. This is achieved by `vga_switcheroo_init_domain_pm_optimus_hdmi_audio`, which augments the audio device's resume function.

On muxed machines, if the mux is initially switched to the discrete GPU, the user ends up with a black screen when the GPU powers down after boot. As a workaround, the mux is forced to the integrated GPU on runtime suspend, cf. https://bugs.freedesktop.org/show_bug.cgi?id=75917

Chapter 6. API

Public functions

Name

`vga_switcheroo_register_handler` — register handler

Synopsis

```
int vga_switcheroo_register_handler (const struct vga_switcheroo_handler * handler);
```

Arguments

handler handler callbacks

Description

Register handler. Enable `vga_switcheroo` if two vga clients have already registered.

Return

0 on success, `-EINVAL` if a handler was already registered.

Name

`vga_switcheroo_unregister_handler` — unregister handler

Synopsis

```
void vga_switcheroo_unregister_handler ( void );
```

Arguments

void no arguments

Description

Unregister handler. Disable `vga_switcheroo`.

Name

`vga_switcheroo_register_client` — register vga client

Synopsis

```
int vga_switcheroo_register_client (struct pci_dev * pdev, const struct  
vga_switcheroo_client_ops * ops, bool driver_power_control);
```

Arguments

<i>pdev</i>	client pci device
<i>ops</i>	client callbacks
<i>driver_power_control</i>	whether power state is controlled by the driver's runtime pm

Description

Register vga client (GPU). Enable `vga_switcheroo` if another GPU and a handler have already registered. The power state of the client is assumed to be ON.

Return

0 on success, `-ENOMEM` on memory allocation error.

Name

`vga_switcheroo_register_audio_client` — register audio client

Synopsis

```
int vga_switcheroo_register_audio_client (struct pci_dev * pdev, const
struct vga_switcheroo_client_ops * ops, enum vga_switcheroo_client_id
id);
```

Arguments

pdev client pci device

ops client callbacks

id client identifier

Description

Register audio client (audio device on a GPU). The power state of the client is assumed to be ON.

Return

0 on success, -ENOMEM on memory allocation error.

Name

`vga_switcheroo_get_client_state` — obtain power state of a given client

Synopsis

```
enum    vga_switcheroo_state    vga_switcheroo_get_client_state    (struct  
pci_dev * pdev);
```

Arguments

pdev client pci device

Description

Obtain power state of a given client as seen from `vga_switcheroo`. The function is only called from `hda_intel.c`.

Return

Power state.

Name

`vga_switcheroo_unregister_client` — unregister client

Synopsis

```
void vga_switcheroo_unregister_client (struct pci_dev * pdev);
```

Arguments

pdev client pci device

Description

Unregister client. Disable `vga_switcheroo` if this is a vga client (GPU).

Name

`vga_switcheroo_client_fb_set` — set framebuffer of a given client

Synopsis

```
void vga_switcheroo_client_fb_set (struct pci_dev * pdev, struct fb_info  
* info);
```

Arguments

pdev client pci device

info framebuffer

Description

Set framebuffer of a given client. The console will be remapped to this on switching.

Name

vga_switcheroo_process_delayed_switch — helper for delayed switching

Synopsis

```
int vga_switcheroo_process_delayed_switch ( void );
```

Arguments

void no arguments

Manual switching and manual power control

Process a delayed switch if one is pending. DRM drivers should call this from their `->lastclose` callback.

Return

0 on success. `-EINVAL` if no delayed switch is pending, if the client has unregistered in the meantime or if there are other clients blocking the switch. If the actual switch fails, an error is reported and 0 is returned.

Name

`vga_switcheroo_set_dynamic_switch` — helper for driver power control

Synopsis

```
void vga_switcheroo_set_dynamic_switch (struct pci_dev * pdev, enum  
vga_switcheroo_state dynamic);
```

Arguments

pdev client pci device

dynamic new power state

Description

Helper for GPUs whose power state is controlled by the driver's runtime pm. When the driver decides to power up or down, it notifies `vga_switcheroo` thereof using this helper so that it can (a) power the audio device on the GPU up or down, and (b) update its internal power state representation for the device.

Name

`vga_switcheroo_init_domain_pm_ops` — helper for driver power control

Synopsis

```
int vga_switcheroo_init_domain_pm_ops (struct device * dev, struct de-  
v_pm_domain * domain);
```

Arguments

dev vga client device

domain power domain

Description

Helper for GPUs whose power state is controlled by the driver's runtime pm. After the GPU has been suspended, the handler needs to be called to cut power to the GPU. Likewise it needs to reinstate power before the GPU can resume. To this end, this helper augments the suspend/resume functions by the requisite calls to the handler. It needs only be called on platforms where the power switch is separate to the device being powered down.

Name

`vga_switcheroo_init_domain_pm_optimus_hdmi_audio` — helper for driver power control

Synopsis

```
int vga_switcheroo_init_domain_pm_optimus_hdmi_audio (struct device *  
dev, struct dev_pm_domain * domain);
```

Arguments

dev audio client device

domain power domain

Description

Helper for GPUs whose power state is controlled by the driver's runtime pm. When the audio device resumes, the GPU needs to be woken. This helper augments the audio device's resume function to do that.

Return

0 on success, -EINVAL if no power management operations are defined for this device.

Public structures

Name

struct vga_switcheroo_handler — handler callbacks

Synopsis

```
struct vga_switcheroo_handler {  
    int (* init) (void);  
    int (* switchto) (enum vga_switcheroo_client_id id);  
    int (* power_state) (enum vga_switcheroo_client_id id, enum vga_switcheroo_state  
        enum vga_switcheroo_client_id (* get_client_id) (struct pci_dev *pdev);  
};
```

Members

init	initialize handler. Optional. This gets called when vga_switcheroo is enabled, i.e. when two vga clients have registered. It allows the handler to perform some delayed initialization that depends on the existence of the vga clients. Currently only the radeon and amdgpu drivers use this. The return value is ignored
switchto	switch outputs to given client. Mandatory. For muxless machines this should be a no-op. Returning 0 denotes success, anything else failure (in which case the switch is aborted)
power_state	cut or reinstate power of given client. Optional. The return value is ignored
get_client_id	determine if given pci device is integrated or discrete GPU. Mandatory

Description

Handler callbacks. The multiplexer itself. The *switchto* and *get_client_id* methods are mandatory, all others may be set to NULL.

Name

struct vga_switcheroo_client_ops — client callbacks

Synopsis

```
struct vga_switcheroo_client_ops {  
    void (* set_gpu_state) (struct pci_dev *dev, enum vga_switcheroo_state);  
    void (* reprobe) (struct pci_dev *dev);  
    bool (* can_switch) (struct pci_dev *dev);  
};
```

Members

set_gpu_state	do the equivalent of suspend/resume for the card. Mandatory. This should not cut power to the discrete GPU, which is the job of the handler
reprobe	poll outputs. Optional. This gets called after waking the GPU and switching the outputs to it
can_switch	check if the device is in a position to switch now. Mandatory. The client should return false if a user space process has one of its device files open

Description

Client callbacks. A client can be either a GPU or an audio device on a GPU. The *set_gpu_state* and *can_switch* methods are mandatory, *reprobe* may be set to NULL. For audio clients, the *reprobe* member is bogus.

Public constants

Name

enum vga_switcheroo_client_id — client identifier

Synopsis

```
enum vga_switcheroo_client_id {  
    VGA_SWITCHEROO_UNKNOWN_ID,  
    VGA_SWITCHEROO_IGD,  
    VGA_SWITCHEROO_DIS,  
    VGA_SWITCHEROO_MAX_CLIENTS  
};
```

Constants

VGA_SWITCHEROO_UNKNOWN_ID	initial identifier assigned to vga clients. Determining the id requires the handler, so GPUs are given their true id in a delayed fashion in <code>vga_switcheroo_enable</code>
VGA_SWITCHEROO_IGD	integrated graphics device
VGA_SWITCHEROO_DIS	discrete graphics device
VGA_SWITCHEROO_MAX_CLIENTS	currently no more than two GPUs are supported

Description

Client identifier. Audio clients use the same identifier & 0x100.

Name

enum vga_switcheroo_state — client power state

Synopsis

```
enum vga_switcheroo_state {
    VGA_SWITCHEROO_OFF,
    VGA_SWITCHEROO_ON,
    VGA_SWITCHEROO_NOT_FOUND
};
```

Constants

VGA_SWITCHEROO_OFF off

VGA_SWITCHEROO_ON on

VGA_SWITCHEROO_NOT_FOUND Client has not registered with vga_switcheroo. Only used in vga_switcheroo_get_client_state which in turn is only called from hda_intel.c

Description

Client power state.

Private structures

Name

struct vgasr_priv — vga_switcheroo private data

Synopsis

```
struct vgasr_priv {
    bool active;
    bool delayed_switch_active;
    enum vga_switcheroo_client_id delayed_client_id;
    struct dentry * debugfs_root;
    struct dentry * switch_file;
    int registered_clients;
    struct list_head clients;
    const struct vga_switcheroo_handler * handler;
};
```

Members

active	whether vga_switcheroo is enabled. Prerequisite is the registration of two GPUs and a handler
delayed_switch_active	whether a delayed switch is pending
delayed_client_id	client to which a delayed switch is pending
debugfs_root	directory for vga_switcheroo debugfs interface
switch_file	file for vga_switcheroo debugfs interface
registered_clients	number of registered GPUs (counting only vga clients, not audio clients)
clients	list of registered clients
handler	registered handler

Description

vga_switcheroo private data. Currently only one vga_switcheroo instance per system is supported.

Name

struct vga_switcheroo_client — registered client

Synopsis

```
struct vga_switcheroo_client {
    struct pci_dev * pdev;
    struct fb_info * fb_info;
    enum vga_switcheroo_state pwr_state;
    const struct vga_switcheroo_client_ops * ops;
    enum vga_switcheroo_client_id id;
    bool active;
    bool driver_power_control;
    struct list_head list;
};
```

Members

pdev	client pci device
fb_info	framebuffer to which console is remapped on switching
pwr_state	current power state
ops	client callbacks
id	client identifier. Determining the id requires the handler, so gpus are initially assigned VGA_SWITCHEROO_UNKNOWN_ID and later given their true id in vga_switcheroo_enable
active	whether the outputs are currently switched to this client
driver_power_control	whether power state is controlled by the driver's runtime pm. If true, writing ON and OFF to the vga_switcheroo debugfs interface is a no-op so as not to interfere with runtime pm
list	client list

Description

Registered client. A client can be either a GPU or an audio device on a GPU. For audio clients, the *fb_info*, *active* and *driver_power_control* members are bogus.

Chapter 7. Handlers

apple-gmux Handler

:1: <http://www.latticesemi.com/en/Products/FPGAandCPLD/LatticeXP2.aspx> :2: <http://www.renesas.com/products/mpumcu/h8s/h8s2100/h8s2113/index.jsp>

gmux is a microcontroller built into the MacBook Pro to support dual GPUs: A {1}[Lattice XP2] on pre-retinas, a {2}[Renesas R4F2113] on retinas.

(The MacPro6,1 2013 also has a gmux, however it is unclear why since it has dual GPUs but no built-in display.)

gmux is connected to the LPC bus of the southbridge. Its I/O ports are accessed differently depending on the microcontroller: Driver functions to access a pre-retina gmux are infixed ``_pio_``, those for a retina gmux are infixed ``_index_``.

gmux is also connected to a GPIO pin of the southbridge and thereby is able to trigger an ACPI GPE. On the MBP5 2008/09 it's GPIO pin 22 of the Nvidia MCP79, on all following generations it's GPIO pin 6 of the Intel PCH. The GPE merely signals that an interrupt occurred, the actual type of event is identified by reading a gmux register.

Graphics mux

:5: <http://pimg-fpiw.uspto.gov/fdd/07/870/086/0.pdf> :6: http://www.nxp.com/documents/data_sheet/CBTL06141.pdf :7: <http://www.ti.com/lit/ds/symlink/hd3ss212.pdf> :8: <https://www.pericom.com/assets/Datasheets/PI3VDP12412.pdf> :9: <http://www.ti.com/lit/ds/symlink/sn74lv4066a.pdf> :10: <http://pdf.datasheetarchive.com/indexerfiles/Datasheets-SW16/DSASW00308511.pdf> :11: <http://www.ti.com/lit/ds/symlink/ts3ds10224.pdf>

On pre-retinas, the LVDS outputs of both GPUs feed into gmux which muxes either of them to the panel. One of the tricks gmux has up its sleeve is to lengthen the blanking interval of its output during a switch to synchronize it with the GPU switched to. This allows for a flicker-free switch that is imperceptible by the user ({5}[US 8,687,007 B2]).

On retinas, muxing is no longer done by gmux itself, but by a separate chip which is controlled by gmux. The chip is triple sourced, it is either an {6}[NXP CBTL06142], {7}[TI HD3SS212] or {8}[Pericom PI3VDP12412]. The panel is driven with eDP instead of LVDS since the pixel clock required for retina resolution exceeds LVDS' limits.

Pre-retinas are able to switch the panel's DDC pins separately. This is handled by a {9}[TI SN74LV4066A] which is controlled by gmux. The inactive GPU can thus probe the panel's EDID without switching over the entire panel. Retinas lack this functionality as the chips used for eDP muxing are incapable of switching the AUX channel separately (see the linked data sheets, Pericom would be capable but this is unused). However the retina panel has the `NO_AUX_HANDSHAKE_LINK_TRAINING` bit set in its DPCD, allowing the inactive GPU to skip the AUX handshake and set up the output with link parameters pre-calibrated by the active GPU.

The external DP port is only fully switchable on the first two unibody MacBook Pro generations, MBP5 2008/09 and MBP6 2010. This is done by an {6}[NXP CBTL06141] which is controlled by gmux. It's the predecessor of the eDP mux on retinas, the difference being support for 2.7 versus 5.4 Gbit/s.

The following MacBook Pro generations replaced the external DP port with a combined DP/Thunderbolt port and lost the ability to switch it between GPUs, connecting it either to the discrete GPU or the Thunderbolt controller. Oddly enough, while the full port is no longer switchable, AUX and HPD are still switchable by way of an {10}[NXP CBTL03062] (on pre-retinas MBP8 2011 and MBP9 2012) or two {11}[TI TS3DS10224] (on retinas) under the control of gmux. Since the integrated GPU is missing the main link, external displays appear to it as phantoms which fail to link-train.

gmux receives the HPD signal of all display connectors and sends an interrupt on hotplug. On generations which cannot switch external ports, the discrete GPU can then be woken to drive the newly connected display. The ability to switch AUX on these generations could be used to improve reliability of hotplug detection by having the integrated GPU poll the ports while the discrete GPU is asleep, but currently we do not make use of this feature.

gmux' initial switch state on bootup is user configurable via the EFI variable `gpu-power-prefs-fa4ce28d-b62f-4c99-9cc3-6815686e30f9` (5th byte, 1 = IGD, 0 = DIS). Based on this setting, the EFI firmware tells gmux to switch the panel and the external DP connector and allocates a framebuffer for the selected GPU.

Power control

gmux is able to cut power to the discrete GPU. It automatically takes care of the correct sequence to tear down and bring up the power rails for core voltage, VRAM and PCIe.

Backlight control

:3: <http://www.ti.com/lit/ds/symlink/lp8543.pdf> :4: <http://www.ti.com/lit/ds/symlink/lp8545.pdf>

On single GPU MacBooks, the PWM signal for the backlight is generated by the GPU. On dual GPU MacBook Pros by contrast, either GPU may be suspended to conserve energy. Hence the PWM signal needs to be generated by a separate backlight driver which is controlled by gmux. The earliest generation MBP5 2008/09 uses a {3}[TI LP8543] backlight driver. All newer models use a {4}[TI LP8545].