

WS3 – Developing Drivers for Altera SoC

Introduction to writing linux device drivers for the
Altera SoC FPGA.



Altera SW SoC Workshop Series

- ◀ SW Workshop #1 – Altera SoC SW Development Overview
- ◀ SW Workshop #2 – Introduction to Linux on Altera SoC
- ◀ SW Workshop #3 – Developing Drivers for Altera SoC Linux

Agenda

- ◀ Essential Information Resources
- ◀ SoC Device Overview
- ◀ Developing Drivers
 - Detailed agenda later...
- ◀ Take Home Lab

Welcome. Here's What You Can Expect Today

Experienced Linux Developers

- ◀ Find a familiar embedded Linux development flow
- ◀ Standard linux device driver model
- ◀ Guide to SoC FPGA resources

New Linux Developers

- ◀ An exposure to the linux driver model
- ◀ The fundamental APIs that are leveraged by most device drivers
- ◀ Essential Linux learning and documentation resources

Hardware Developers

- ◀ HW handoff to Linux driver development flows
- ◀ SW driver implications of HW architecture

Everyone

- ◀ SoC FPGA architecture-specific information
- ◀ SoC FPGA recommendations and best practices

Essential Information Resources

Where to learn more...
...a non-exhaustive list



ALTERA®

Linux Foundation Training

Linux Developer classes are designed to help participants:

- Learn how to develop an **embedded Linux** product
- Become familiar with and learn to write **device drivers**
- Get practical experience with the Linux kernel
- Learn how to work with the Linux developer community

Developer Courses

- *LFD331 – Developing Linux Device Drivers*
- *LFD405 – Building Embedded Linux with the Yocto Project*
- *LFD411 – Embedded Linux Development*
- *LFD414 – Introduction to Embedded Android Development*
- *LFD205 – How to Participate with the Linux Community*
- *LFD211 – Introduction to Linux for Developers*
- *LFD262 – Developing with Git*
- *LFD312 – Developing Applications for Linux*
- *LFD320 – Linux Kernel Internals & Debugging*
- *LFD415 – Inside Android: An Intro to Android Internals*
- *LFD432 – Optimizing Linux Device Drivers for Power Efficiency*

<http://training.linuxfoundation.org/linux-courses/development-training>

Linux Documentation Resources

Git

- Distributed revision control system to enable distributed collaboration
- On-line documentation & training:
 - ↳ <http://git-scm.com/doc>
 - ↳ <https://training.github.com>

Denx U-Boot Manual

- Complete documentation from the folks who wrote Das U-Boot
 - ↳ <http://www.denx.de/wiki/U-Boot/Documentation>

Free-Electrons:

- Complete training materials posted free
 - ↳ <http://free-electrons.com/docs/>


Device Tree for Dummies

- <http://events.linuxfoundation.org/sites/events/files/slides/petazzoni-device-tree-dummies.pdf>

The Two Best Sources for Linux Development Information


Linux Kernel Documentation

- The most complete and most essential Linux kernel documentation
- Included with the Linux kernel source code
 - ↳ <local GIT repo>/Documentation



Documentation : bash - Konsole

```
/data/linux_repos/linux-socfpga/Documentation>>grep -ir tickless *
timers/highres.txt:The implementation leaves room for further development like full tickless
timers/NO_HZ.txt:tickless". The remainder of this document will use "dyntick-idle mode".
/data/linux_repos/linux-socfpga/Documentation>
```



File Edit View Bookmarks Tools Settings Help

NO_HZ: Reducing Scheduling-Clock Ticks

This document describes Kconfig options and boot parameters that can reduce the number of scheduling-clock interrupts, thereby improving energy efficiency and reducing OS jitter. Reducing OS jitter is important for some types of computationally intensive high-performance computing (HPC) applications and for real-time applications.

There are three main ways of managing scheduling-clock interrupts (also known as "scheduling-clock ticks" or simply "ticks"):

1. Never omit scheduling-clock ticks (CONFIG_HZ_PERIODIC=y or CONFIG_NO_HZ=n for older kernels). You normally will *not* want to choose this option.

The Two Best Sources for Linux Development Information

↳ An open source OS breeds open source information

A screenshot of a Google search results page. The search query is "boot time reduction linux". The results show two main entries:

- Embedded Linux boot time reduction - Free Electrons**
free-electrons.com/services/[boot-time/](#) ▾
Making your embedded **Linux** systems boot faster. Investigating **boot time** issues and applying optimization techniques that don't require a redesign.
- Free Electrons: Embedded Linux Experts**
free-electrons.com/ ▾
Embedded **Linux**, kernel and Android: development, training and consulting services, ...
Buildroot commercial support · Embedded **Linux boot time reduction** ... Offering our broad embedded **Linux** development experience through our ... 42,295 views; How to boot an uncompressed **Linux** kernel on ARM · 38,199 views ...
You've visited this page 4 times. Last visit: 11/1/14

At the bottom of the search results, there is a highlighted link:
[PDF] [Update on boot time reduction techniques, with figures - T...](#)
events.linuxfoundation.org/.../opdenacker-boot-time.p... ▾ Linux Foundation ▾
http://free-electrons.com/doc/training/boot-time. ▶ That's where you will find extensive

RocketBoards.org – Altera SoC Linux Community Portal

- ⬧ The source for SoC FPGA Linux info
 - Golden System Reference Design (GSRD)
 - Updates on latest releases
 - Step-by-step getting started guides
- ⬧ SoC FPGA Mailing List - RFI
 - Active community participation in answering SoC FPGA and Linux questions
- ⬧ Example Projects, Applications, and Designs
 - From Altera and the SoC community
- ⬧ Public GIT server
- ⬧ Enables the SoC community to support Linux



RocketBoards.org

RocketBoards.org Resources

The screenshot shows the RocketBoards.org homepage with several key features highlighted:

- Documentation:** The "Documentation" menu item is highlighted with a blue box and an arrow pointing to it from the left.
- MAILS:** The "MAILS" link in the top right navigation bar is highlighted with a blue box and an arrow pointing to it from the top.
- RELEASES:** The "RELEASES" link in the top right navigation bar is highlighted with a blue box and an arrow pointing to it from the top.
- GSRD and other Reference Designs:** A large blue callout box on the right side contains the text "GSRD and other Reference Designs" and features an image of a printed circuit board (PCB).
- RFI Registration & Archives:** A blue callout box on the right side contains the text "RFI Registration & Archives".
- Starting point for documentation:** A blue callout box on the left side contains the text "Starting point for documentation".
- News and Updates:** A blue callout box at the bottom center contains the text "News and Updates".
- What's new:** A section titled "What's new" is highlighted with a blue box and an arrow pointing to it from the bottom.
- Projects:** The "Projects" section is highlighted with a blue box and an arrow pointing to it from the bottom.
- Community:** The "Community" section is highlighted with a blue box and an arrow pointing to it from the bottom.

RocketBoards.org – Useful Links

◀ GSRD User Manual

- The best starting point for Linux development
- <http://www.rocketboards.org/foswiki/Documentation/GSRD>

Getting Started Guides

- [Booting Linux Using Prebuilt SD Card Image](#)
- [Connecting to Board Web Server](#)
- [Connecting to Board Using SSH](#)
- [Running Sample Linux Applications](#)
- [Compiling the Hardware Design](#)
- [Generating and Compiling the Preloader](#)
- [Generating the Device Tree](#)
- [Compiling Angstrom Linux Distribution](#)
- [Creating and Updating SD Card](#)
- [GSRD FPGA Programming](#)
- [FPGA Programming with Quartus II Programmer](#)
- [Using System Console](#)
- [Angstrom Getting Started](#)

◀ Device Tree Generator User Guide

- <http://www.rocketboards.org/foswiki/Documentation/GSRD141DeviceTreeGenerator>

◀ Programming FPGA from HPS

- <http://www.rocketboards.org/foswiki/Documentation/GSRD131ProgrammingFPGA>

SoC FPGA Linux Code Repositories on GitHub

Public git repos for SoC FPGA

<https://github.com/altera-opensource>

Migrated from RocketBoards

- Rocketboards.org git repos no longer updated

The screenshot shows the GitHub interface for the 'Altera OpenSource' organization. At the top, there's a navigation bar with links for 'Explore', 'Features', and 'Enterprise'. Below the header, the organization's logo (a blue diamond with a white rocket ship) and name are displayed, along with a link to their website (<http://rocketboards.org>). A search bar and a 'Filters' dropdown are also present. The main content area lists three repositories:

- linux-socfpga**: A repository for Linux development, described as "Linux development repository for socfpga". It was last updated 6 hours ago. The repository has 4 stars and 1 pull request.
- meta-altera**: A Yocto Layer for SoC FPGA, specifically for Angstrom. It is described as "Yocto Layer for SoC FPGA. Angstrom specific branches." It was last updated 3 days ago. The repository has 0 stars and 0 pull requests.
- angstrom-socfpga**: An Angstrom repository with updated layers file. It is described as "Angstrom repository with updated layers file". The repository has 0 stars and 0 pull requests.

Several Ways to Learn!

Instructor-led training

- Face to face with an Altera expert Training Engineer
- 20+ courses to choose from (8 hour classes)



Virtual classes (taught via WebEX)

- Can ask questions to Altera expert Training Engineer
- Course content same as instructor-led classes
(1/2 day sessions)



Online training (free and always available)

- 200+ topics available (~30 minutes in length)



Videos (free and always available)

- YouTube videos (~4 minutes each)



SoC Classes Available



Instructor-led or virtual classes

- [Designing with an ARM-based SoC](#)
- [Developing Software for an ARM-based SoC](#)

Online classes

- [Hardware Design Flow for an ARM-based SoC](#)
- [Software Design Flow for an ARM-based SoC](#)
- [SoC Hardware Overview: the Microprocessor Unit](#)
- [SoC Hardware Overview: Interconnect and Memory](#)
- [SoC Hardware Overview: System Management, Debug, and General Purpose Peripherals](#)
- [SoC Hardware Overview: Flash Controllers and Interface Protocols](#)
- [SoC Bare-metal Programming and Hardware Libraries](#)
- [Getting Started with Linux for Altera SoCs](#)

Essential SoC Hardware Documentation Resources

Hard Processor System Technical Reference Manuals

- Available in Device Handbooks:
 - ↳ <https://www.altera.com/products/soc/portfolio/cyclone-v-soc/support.html>
 - ↳ <https://www.altera.com/products/soc/portfolio/arria-v-soc/support.html>
 - ↳ <https://www.altera.com/products/soc/portfolio/arria-10-soc/support.html>
- Contain Functional Descriptions Peripheral
- Contain Control Register Address Map and Definitions
 - ↳ These are also available online at the links above in HTML and PDF formats

HPS SoC Boot Guide

- Cyclone V SoC & Arria V SoC: [AN709 - HPS SoC Boot Guide](#)
- Arria 10 SoC: included in HPS TRM in Arria 10 Device Handbook

ARM Documentation Site

- Documentation available for all ARM IP
 - ↳ Cortex-A9 & A53 MP Cores, FPU, NEON, GIC, ARM Peripherals, etc.
- Requires free registration
- Refer to HPS TRM for IP core names and revision information
- <http://infocenter.arm.com/help/index.jsp>

Essential SoC Software Documentation Resources

Altera SoC Embedded Design Software (SoC EDS) Tools

- User Guide:
https://www.altera.com/content/dam/altera-www/global/en_US/pdfs/literature/ug/ug_soc_eds.pdf
 - ↳ Linux & Baremetal Software Development Tools Overview
 - ↳ HPS Preloader User Guide
 - ↳ HPS Flash Programmer User Guide
 - ↳ SD Card Boot Utility
- Getting Started Guides: Preloader, Linux, Bare Metal, Debug, HW Library
<http://www.alterawiki.com/wiki/SoCEDSGettingStarted>
- SoC HPS Release Notes
- SoC Abstraction Layer (SoCAL) API Reference
 - ↳ <SoC EDS install dir>/ip/altera/hps/altera_hps/doc/socal/html/index.html
- Hardware Manager API Reference
 - ↳ <SoC EDS install dir>/ip/altera/hps/altera_hps/doc/hwmgr/html/index.html
- GCC Documentation
 - ↳ <SoC EDS install dir>/ds-5/documents/gcc/getting_started.html
- Bare Metal Compiler
 - ↳ <SoC EDS installation directory>/host_tools/mentor.gnu/arm/baremetal/share/doc/sourceryg++-arm-altera-eabi

Essential SoC Software Tools Online Videos

ARM DS-5 Altera Edition Toolchain

- <https://youtu.be/HV6NHr6gLx0>

DS-5 Altera Edition: Bare-metal Debug and Trace

- https://youtu.be/u_xKybPhcHI

DS-5 Altera Edition: FPGA-adaptive Linux Kernel Debug and Trace

- <https://youtu.be/lrR-SfVZd18>

Debugging Linux applications on the Altera SoC with ARM DS-5

- <https://youtu.be/ZcGQEjkYWQc>

FPGA-adaptive debug on the Altera SoC using ARM DS-5

- <https://youtu.be/2NBcUv2Txbl>

Streamline Profiling on Altera SoC FPGA. Part 1 - Setup

- <https://youtu.be/X-k9ImXQTio>

Streamline Profiling on Altera SoC FPGA. Part 2 - Running Streamline

- <https://youtu.be/Tzbd7qldKqY>



SoC Device Overview




ALTERA®

Altera Investment in Embedded Technologies

- Altera established Austin Technology Center (ATC) in 2011
- Altera's primary embedded engineering center
- Austin provides access to one of the richest embedded processing talent bases in the world




Altera SoC Product Portfolio



DEVICE AVAILABILITY

Availability of SoC devices across product portfolio ...

ARM Public Processor Offering



28nm SoC System Architecture

Processor

- Dual-core ARM® Cortex™-A9 MPCore™ processor
- Up to 5,250 MIPS (1050 MHz per core maximum)
- NEON coprocessor with double-precision FPU
- 32-KB/32-KB L1 caches per core
- 512-KB shared L2 cache

Multiport SDRAM controller


- DDR3, DDR3L, DDR2, LPDDR2
- Integrated ECC support

High-bandwidth on-chip interfaces

- > 125-Gbps HPS-to-FPGA interface
- > 125-Gbps FPGA-to-SDRAM interface

Cost- and power-optimized FPGA fabric

- Lowest power transceivers
- Up to 1,600 GMACS, 300 GFLOPS
- Up to 25Mb on-chip RAM
- More hard intellectual property (IP): PCIe® and memory controllers




Notes:


(1) Integrated direct memory access (DMA)

(2) Integrated ECC

Arria 10 HPS Block Diagram



High-Level Block Diagram



A Comparison: Cyclone V SoC, Arria V SoC, Arria 10 SoC

Metric	Cyclone V SoC	Arria V SoC	Arria 10 SoC
Technology	28nm	28nm	20nm
Processor Performance	925 MHz	1.05 GHz	1.5 GHz
Total Power Dissipation	100%	100%	60% (40% Lower)
Max PCI Express Hard IP	Gen 2 x4	Gen 2 x8	Gen 3 x8
Memory Devices Supported	DDR2, DDR3, DDR3L, LPDDR2	DDR2, DDR3, DDR3L, LPDDR2	DDR4/3, LPDDR2/3, QDRIV, RLDRAM III, Hybrid Memory Cube
Max. HPS DDR Data-Width	40-bit (32-bit + ECC)	40-bit (32-bit + ECC)	72-bit (64-bit + ECC)
EMAC Cores	EMAC x 2	EMAC x 2	EMAC x 3
NAND Device Supported	8-bit	8-bit	8-bit and 16-bit
SD/MMC devices supported	SD/SDIO/MMC	SD/SDIO/MMC	SD/SDIO/MMC 4.5 with eMMC
FPGA Logic Density Range (LEs)	25 - 110K	370 - 450K	160 - 660K
FPGA Core Performance	260 MHz	307 MHz	500 MHz

Developing Drivers

The Altera logo, featuring the word "ALTERA" in a bold, blue, sans-serif font. A registered trademark symbol (®) is located at the top right corner of the letter "A".

ALTERA®

Developing Drivers Agenda (1 of 4)

- ◀ Driver Demonstration Hardware Overview
- ◀ Hardware architecture to software development communication options
 - `sopc-create-header-files`, `sopc2dts`
- ◀ Accessing FPGA hardware from user space
 - `/dev/mem`
 - `mmap()`, `munmap()`
 - runtime device tree inspection
 - application build example
- ◀ Driver development environment setup
 - `ARCH`, `CROSS_COMPILE`, `OUT_DIR`
 - Kbuild file, Makefile
- ◀ Demonstration Module 1 - the most basic module example
 - Source code licensing
 - `module_init()`, `module_exit()`, `MODULE_LICENSE()`, `MODULE_*`()
- ◀ Demonstration Module 1t - tainted module example
 - kernel tainting, modinfo, insmod, lsmod, rmmod
 - procfs / sysfs module statistics

Developing Drivers Agenda (2 of 4)

↳ Demonstration Module 2 - introduce module parameters

- `module_param()`, `module_param_array()`, `MODULE_PARAM_DESC()`
- sysfs module parameters/* entries

↳ Demonstration Module 3 - introduce platform driver

- `of_device_id` structure
 - ↳ .compatible strings
- `MODULE_DEVICE_TABLE()`
- `platform_driver` structure
 - ↳ .probe, .remove, .driver
- `platform_driver_register()`, `platform_driver_unregister()`
- manually binding/unbinding from user space

↳ Demonstration Module 4 - extract hardware details from DT

- `platform_get_resource()`, `platform_get_irq()`, `clk_get()`, `clk_get_rate()`
- manually validating device tree information from procfs

↳ Demonstration Module 5 - reserve memory and enable IRQ

- `down_interruptible()`, `up()`, `request_mem_region()`, `release_mem_region()`, `ioremap()`, `iounmap()`,
`request_irq()`, `free_irq()`, `resource_size()`, `ioread32()`, `iowrite32()`
- basic interrupt handler

Developing Drivers Agenda (3 of 4)

↳ Demonstration Module 5t - test memory reservation and IRQ

- resource contention demonstration

↳ Demonstration Module 6 - introduce sysfs entries

- driver_create_file(), driver_remove_file(), DRIVER_ATTR(), spin_lock_irqsave(), spin_lock_irqrestore(), spin_lock(), spin_unlock(), scnprintf(), kstrtoul()
- interrupt handler
- sysfs show(), store() functions
- manual interaction with sysfs entries from user space

↳ Demonstration Module 7 - introduce misc device

- sema_init(), misc_register(), misc_deregister(), memcpys_fromio(), memcpys_toio(), copy_from_user(), copy_to_user(), get_user(), put_user(), wait_event_interruptible(), wake_up_interruptible(), pgprot_noncached(), remap_pfn_range()
- file_operations structure
 - ↳ .open, .release, .read, .write, .llseek, .unlocked_ioctl, .mmap
- miscdevice structure
- struct file.private_data, struct file.f_pos
- SEEK_SET, SEEK_CUR, SEEK_END
- ioctl() from user space


Developing Drivers Agenda (4 of 4)

- ◀ Demonstration Module 8 - introduce uio device
 - uio_info structure
 - uio_register_device(), uio_unregister_device(), down_trylock()
- ◀ Demonstration Module 9 - introduce dma
 - kmalloc(), dma_alloc_coherent(), dma_map_single(), dma_unmap_single(), dma_mapping_error()
- ◀ Demonstration Module 10 - multiple device instances
 - INIT_LIST_HEAD(), kzalloc(), list_add(), platform_set_drvdata(), iminor(), list_for_each(), list_entry()
 - struct file.f_flags
- ◀ Demonstration Module 11 - introduce custom API
 - EXPORT_SYMBOL()
- ◀ Demonstration Module 11t - test custom API

Our demonstration hardware environment for the workshop




Driver Demonstration Hardware




This is a very simplified block diagram of the Altera HPS and FPGA system that is used to demonstrate the concepts described in this workshop. Please see the Qsys system for more detail.

Driver Demonstration Hardware




The **demo_driver_subsys_0** peripheral is a custom Qsys component that presents one opaque slave interface to the system. Internally it contains a RAM, a ROM, and an Interval Timer that can generate an interrupt.

Driver Demonstration Hardware



There are three FIFO components that have a slave interface for writing data into the FIFO and another slave interface for reading data out of the FIFO. In this block diagram the FIFO is represented with only one slave interface for simplicity.

Driver Demonstration Hardware



The **memcpy_msgdma** component can be used to DMA data thru the F2H AXI bridge, which allows us to DMA data to/from anywhere in the HPS core or peripherals attached to the LWH2F and H2F AXI bridges in the FPGA.

How do we get our hardware system definition into the software domain?



Hardware details into software development flow

- One of the first challenges that we encounter if we want to write software to interact with our FPGA hardware is communicating the hardware architecture details into our software development environment.
- The Altera Quartus and Qsys tools provide various pieces of output that can help us accomplish this with automated output from the Qsys system generation flow and the Quartus compilation flow.
- The following slides describe what the hardware tools output and some utilities that can be applied to that output to create something that can be easily digested by the software development flow.

Hardware Handoff Files to Software Environment

```
01 []$ ls -r *.sopcinfo hps_isw_handoff/*
02
03 soc_system.sopcinfo
04
05 hps_isw_handoff/soc_system_hps_0:
06 tclrpt.h           sequencerDefines.h    sdram_io.h
07 tclrpt.c          sequencer.c          id
08 system.h          sequencerAutoInstInit.c hps.xml
09 soc_system_hps_0.hiof sequencerAuto.h    emif.xml
10 sequencer.h        sequencerAutoAcInit.c alt_types.h
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
```

These files are generated in the Quartus project directory as a result of Qsys system generation and Quartus compilation.

Hardware Handoff Files to Software Environment

```
01 []$ ls -r *.sopcinfo hps_isw_handoff/*
02
03 soc_system.sopcinfo ←
04
05 hps_isw_handoff/soc_system_hps_0:
06 tclrpt.h           sequencerDefines.h    sdram_io.h
07 tclrpt.c          sequencer.c          id
08 system.h          sequencerAutoInstInit.c hps.xml
09 soc_system_hps_0.hif  sequencerAuto.h      emif.xml
10 sequencer.h        sequencerAutoAcInit.c alt_types.h
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
```

The SOPCINFO file is created by Qsys during system generation. This database can be used to create system header files and device trees for the software environment.

The “sopc-create-header-files” utility can generate a variety of header file and macro file formats, and the “sopc2dts” utility can create device trees representing the Qsys system.


Hardware Handoff Files to Software Environment

```
01 []$ ls -r *.sopcinfo hps_isw_handoff/*
02
03 soc_system.sopcinfo
04
05 hps_isw_handoff/soc_system_hps_0: ←
06 tclrpt.h           sequencerDefines.h    sdram_io.h
07 tclrpt.c          sequencer.c          id
08 system.h          sequencer_auto_inst_init.c hps.xml
09 soc_system_hps_0.hiof sequencer_auto.h   emif.xml
10 sequencer.h        sequencer_auto_ac_init.c alt_types.h
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
```

The `hps_isw_handoff` directory is created by Quartus during compilation and contains information about the HPS configuration from Qsys and Quartus. This collection of files is used by the BSP generator to create a customized preloader for the HPS core to run at boot time to configure the HPS core as defined by Qsys and Quartus.

Creating System Header Files

```
01 []$ sopc-create-header-files --help
02
03 Usage: sopc-create-header-files [<sopc>] [OPTION]...
04
05 This utility creates header files from your SOPC Builder system description.
06 By default, the header files are in cpp format and have a .h suffix.
07 Other formats may be selected with the appropriate command-line option.
08
09 Options:
10
11 ...cut...
12
13 Supported header file formats:
14 type suffix      uses          example
15 ---- -----  -----
16 h     .h      C/C++ header file for cpp #define FOO 12
17 m4    .m4      macro file for m4   m4_define("FOO", 12)
18 sh    .sh      shell scripts      FOO=12
19 mk    .mk      makefiles        FOO := 12
20 pm    .pm      Perl scripts      $macros{FOO} = 12;
21
22 By default, multiple header files are created. There is one header file for
23 the entire system and one header file for each master group in each module.
24 A master group is a set of masters in a module in the same address space.
25 In general, a module may have multiple master groups.
26 Addresses and available devices are a function of the master group.
27
28 ...cut...
29
30
```



output
formats

Creating System Header Files

```
01 []$ mkdir qsys_headers
02 []$ sopc-create-header-files soc_system.sopcinfo --output-dir qsys_headers
03 []$ ls qsys_headers
04 axi_bridge_for_acp_128_0.h    fft_sub_sgdma_to_fft.h    lw_mm_bridge.h
05 fft_ddr_bridge.h              fpga_only_master.h      memcpy_msgdma_mm_read.h
06 fft_sub_DDR.h                hps_0_arm_a9_0.h      memcpy_msgdma_mm_write.h
07 fft_sub.h                    hps_0_arm_a9_1.h      soc_system.h
08 fft_sub_mm_bridge_0.h        hps_0_bridges.h
09 fft_sub_sgdma_from_fft.h    hps_0.h
10
11 The sopc-create-header-files utility can create a memory map view from the
12 perspective of any of the masters in the Qsys system. In the default mode
13 demonstrated above, all masters in the Qsys system have a header file created for
14 their view of the memory mapped system. The macros in these headers can be useful
15 when writing applications and drivers that interact with the FPGA based peripherals.
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
```

Creating System Header Files

```
01 []$ grep "\s[^H].*_BASE" qsys_headers/hps_0_arm_a9_0.h
02 #define ONCHIP_MEMORY2_0_BASE 0xc0000000
03 #define SYSID_QSYS_BASE 0xff201000
04 #define INTR_CAPTURER_0_BASE 0xff206000
05 #define VALIDATOR_SUBSYS_0_BASE 0xff210000
06 #define MEMCPY_MSGDMA_CSR_BASE 0xff220000
07 #define MEMCPY_MSGDMA_DESCRIPTOR_SLAVE_BASE 0xff220020
08 #define DEMO_DRIVER_SUBSYS_0_BASE 0xff230000
09 #define FIFO_0_IN_CSR_BASE 0xff240000
10 #define FIFO_0_IN_BASE 0xff240020
11 #define FIFO_0_OUT_BASE 0xff240030
12 #define FIFO_1_IN_CSR_BASE 0xff244000
13 #define FIFO_1_IN_BASE 0xff244020
14 #define FIFO_1_OUT_BASE 0xff244030
15 #define FIFO_2_IN_CSR_BASE 0xff248000
16 #define FIFO_2_IN_BASE 0xff248020
17 #define FIFO_2_OUT_BASE 0xff248030
18 #define FFT_SUB_SGDMA_TO_FFT_CSR_BASE 0xff280000
19 #define FFT_SUB_SGDMA_TO_FFT_DESCRIPTOR_SLAVE_BASE 0xff290000
20 #define FFT_SUB_SGDMA_FROM_FFT_CSR_BASE 0xff2a0000
21 #define FFT_SUB_SGDMA_FROM_FFT_DESCRIPTOR_SLAVE_BASE 0xff2b0000
22 #define FFT_SUB_DATA_BASE 0xff2c0000
23 #define FFT_SUB_FFT_STADAPTER_0_BASE 0xff2d0000
24
25 This is an example of the base addresses of the FPGA based Qsys peripherals from
26 the perspective of the Cortex A9 processor. NOTE: this is a very small portion of the
27 total contents of the header file macros.
28
29
30
```

Creating System Header Files

```
01 []$ grep "\s[^H].*_BASE" qsys_headers/hps_0.h
02 #define ONCHIP_MEMORY2_0_BASE 0x0
03 #define SYSID_QSYS_BASE 0x1000
04 #define INTR_CAPTURER_0_BASE 0x6000
05 #define VALIDATOR_SUBSYS_0_BASE 0x10000
06 #define MEMCPY_MSGDMA_CSR_BASE 0x20000
07 #define MEMCPY_MSGDMA_DESCRIPTOR_SLAVE_BASE 0x20020
08 #define DEMO_DRIVER_SUBSYS_0_BASE 0x30000
09 #define FIFO_0_IN_CSR_BASE 0x40000
10 #define FIFO_0_IN_BASE 0x40020
11 #define FIFO_0_OUT_BASE 0x40030
12 #define FIFO_1_IN_CSR_BASE 0x44000
13 #define FIFO_1_IN_BASE 0x44020
14 #define FIFO_1_OUT_BASE 0x44030
15 #define FIFO_2_IN_CSR_BASE 0x48000
16 #define FIFO_2_IN_BASE 0x48020
17 #define FIFO_2_OUT_BASE 0x48030
18 #define FFT_SUB_SGDMA_TO_FFT_CSR_BASE 0x80000
19 #define FFT_SUB_SGDMA_TO_FFT_DESCRIPTOR_SLAVE_BASE 0x90000
20 #define FFT_SUB_SGDMA_FROM_FFT_CSR_BASE 0xa0000
21 #define FFT_SUB_SGDMA_FROM_FFT_DESCRIPTOR_SLAVE_BASE 0xb0000
22 #define FFT_SUB_DATA_BASE 0xc0000
23 #define FFT_SUB_FFT_STADAPTER_0_BASE 0xd0000
24
25 []$ grep "_IRQ [^0]" qsys_headers/hps_0.h
26 #define MEMCPY_MSGDMA_CSR_IRQ 7
27 #define DEMO_DRIVER_SUBSYS_0_IRQ 8
28 #define FFT_SUB_SGDMA_TO_FFT_CSR_IRQ 4
29 #define FFT_SUB_SGDMA_FROM_FFT_CSR_IRQ 3
30
```

This is an example of the base addresses of the FPGA based Qsys peripherals from the perspective of the HPS AXI bridges.

Below is an example of the IRQ values from the same perspective.

NOTE: this is a very small portion of the total contents of the header file macros.

Creating Device Trees

```
01 []$ sopc2dts --help
02 sopc2dts - 14.1 [f3b984376d464370083edffbe85dd18a2bb42a80]
03 Usage: sopc2dts <arguments>
04 Required Arguments:
05   --input <sopcinfo file>      The sopcinfo file (optional in gui mode) (Required)
06
07   -i <sopcinfo file>      Short for --input
08
09 Optional Arguments:
10   --board <boardinfo file>      board description file (can be used multiple times) (Optional)
11   --bridge-ranges <{none,bridge,child}> what to describe in bridges address translations (Optional)
12   --bridge-removal <{all,balanced,none}>      Bridge removal strategy (Optional)
13   --help                  Show this usage info and exit (Optional)
14   --verbose                Show Lots of debugging info (Optional)
15   --gui                   Run in gui mode (Optional)
16   --clocks                 Show clocks in Device Tree Source / graph (Optional)
17   --conduits                Show conduit interfaces in graph (Optional)
18   --version                 Show version information and exit (Optional)
19   --mimic-sopc-create-header-files      Try to (mis)behave like sopc-create-header-f
20   --no-timestamp            Don't add a timestamp to generated files (Optional)
21   --output <filename>      The output filename (Optional)
22   --pov <component name>      The point of view to generate from. Defaults to the first cp
23   --pov-type <{cpu,pci}>      The point of view device type (Optional)
24   --reset                  Show reset interfaces in graph (Optional)
25   --sort <{none,address,name,label}>      Sort components by (Optional)
26   --streaming                Show streaming interfaces in graph (Optional)
27   --type <{dtb,dtb-hex8,dtb-hex32,dtb-char-arr,dts,uboot,kernel}>      The type of output t
28   --bootargs <kernel-args>      Default kernel arguments for the "chosen" section of the DTS
29   --sopc-parameters <{node,cmacro,all}>      What sopc-parameters to include in DTS (Optional)
30
```

Creating Device Trees

```
01 []$ sopc2dts --input soc_system.sopcinfo --output soc_system.dts \
02     --board hps_common_board_info.xml --bridge-removal all --clocks \
03     --board board_info_ALTERA_CV_SOC.xml
04
05 []$ cat soc_system.dts
06
07 ...cut...
08
09 fifo_0: fifo@0x100040020 {
10     compatible = "ALTR,fifo-14.1", "ALTR,fifo-1.0";
11     reg = <0x00000001 0x00040020 0x00000004>,
12         <0x00000001 0x00040030 0x00000004>,
13         <0x00000001 0x00040000 0x00000020>;
14     reg-names = "in", "out", "in_csr";
15     clocks = <&c1k_0>;
16 }; //end fifo@0x100040020 (fifo_0)
17
18 ...cut...
19
20 demo_driver_subsys_0: driver@0x100030000 {
21     compatible = "demo,driver-1.0", "demo,driver-1.0";
22     reg = <0x00000001 0x00030000 0x00001000>;
23     interrupt-parent = <&hps_0_arm_gic_0>;
24     interrupts = <0 48 4>;
25     clocks = <&c1k_0>;
26 }; //end driver@0x100030000 (demo_driver_subsys_0)
27
28 ...cut...
29
30
```

The **sopc2dts** utility can create device tree entries for the FPGA based Qsys components represented in the SOPCINFO file.

These device tree entries can be compiled by the “**dtc**” compiler into a binary device tree blob that is used by the linux kernel to define the peripheral space available on the target.

NOTE: newer u-boot versions are driven by device tree definitions as well as the kernel, however u-boot and the kernel require their own unique device tree definitions, generated separately.

Device Tree Documentation

```
01 Device Tree for Dummies:  
02 http://events.linuxfoundation.org/sites/events/files/slides/  
03 petazzoni-device-tree-dummies.pdf  
04  
05  
06  
07  
08  
09 [socfpga-3.10-ltsi]$ ls -R Documentation/devicetree/  
10 Documentation/devicetree/:  
11 00-INDEX bindings booting-without-of.txt usage-model.txt  
12  
13 Documentation/devicetree/bindings:  
14 arc i2c net sound  
15 arm iio nvec spi  
16 ata input open-pic.txt staging  
17 bus interrupt-controller pci thermal  
18 c6x iommu phy timer  
19 clock leds pinctrl tty  
20 cpufreq lpddr2 power usb  
21 crypto mailbox powerpc vendor-prefixes.txt  
22 dma marvell.txt power_supply video  
23 drm media pwm virtio  
24 eeprom.txt memory-controllers regulator w1  
25 fb metag reset watchdog  
26 fpga mfd resource-names.txt x86  
27 gpio mips rng xilinx.txt  
28 gpu misc rtc  
29 hwmon mmc serial  
30 hwrng mtd serio
```

Some publicly available information on device tree.

Device Tree Documentation

```
01 [socfpga-3.10-ltsi]$ cat Documentation/devicetree/bindings/clock/altr_socfpga.txt
02 Device Tree Clock bindings for Altera's SoC FPGA platform
03
04 This binding uses the common clock binding[1].
05
06 [1] Documentation/devicetree/bindings/clock/clock-bindings.txt
07
08 Required properties:
09 - compatible : shall be one of the following:
10     "altr,socfpga-pll-clock" - for a PLL clock
11     "altr,socfpga-perip-clock" - The peripheral clock divided from the
12         PLL clock.
13     "altr,socfpga-gate-clk" - clocks that directly feed peripherals and
14         can get gated.
15
16 - reg : shall be the control register offset from CLOCK_MANAGER's base for the clock.
17 - clocks : shall be the input parent clock phandle for the clock. This is
18     either an oscillator or a pll output.
19
20 Optional properties:
21 - fixed-divider : If clocks have a fixed divider value, use this property.
22 - #clock-cells : from common clock binding; shall be set to 0.
23
24 Optional properties:
25 - fixed-divider : If clocks have a fixed divider value, use this property.
26 - clk-gate : For "socfpga-gate-clk", clk-gate contains the gating register
27     and the bit index.
28 - div-reg : For "socfpga-gate-clk", div-reg contains the divider register, bit shift,
29     and width.
```

An example of some bindings documentation.

Accessing FPGA hardware, starting off simple with a user space application.

`mmap(/dev/mem)`




Trivial access to FPGA hardware from user space.


- ⬚ Often when we begin developing a custom piece of FPGA hardware we know that we will eventually want to have a nice pretty driver that presents a convenient API into user space for applications to interact with, but in the early days of developing the hardware, we don't even know if it's working properly, and we may still have numerous iterations that we apply to the hardware before it stabilizes into the ultimate thing for which we want to create a driver.
- ⬚ At this early point in time, it can be much less overhead to simply peek and poke at the hardware with some basic stimulus thru a user space application that `mmap()`'s the hardware using `/dev/mem`.
 - In some cases this `/dev/mem` approach may be all you ever need.

demo_devmem

User Space



Kernel Space



IO Space

Color coded indicator of
user/kernel space example.

User
Space

Color Key for Code Examples

01 The general context of the code example is not highlighted at all, just black text.
02
03 The **focus** of a code example is highlighted in red text. If you ignore everything else on
04 the page in black text, make sure you take note of the red text that is highlighted as it
05 is relevant to the example.
06
07 when the **focus** of a code example refers to functions and objects provided by the kernel or
08 standard system calls there is a yellow highlight behind the red text.
09
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30

demo_devmem.c example

```
01 int main(int argc, char **argv) {  
02  
03     int devmem_fd;  
04     void *demo_driver_map;  
05     void *memcpy_msgdma_map;  
06     int result;  
07  
08     //  
09     // validate the system features  
10     //  
11     validate_system_features();  
12  
13     //  
14     // verify that DEMO_DRIVER_PHYS_BASE is page aligned  
15     //  
16     if(DEMO_DRIVER_PHYS_BASE & (sysconf(_SC_PAGE_SIZE) - 1)) {  
17     ...cut...  
18     }  
19  
20     //  
21     // verify that MEMCPY_MSGDMA_CSR_PHYS_BASE is page aligned  
22     //  
23     if(MEMCPY_MSGDMA_CSR_PHYS_BASE & (sysconf(_SC_PAGE_SIZE) - 1)) {  
24     ...cut...  
25     }  
26  
27     //  
28     // parse the command line arguments  
29     //  
30     parse_cmdline(argc, argv);
```

In the demo_devmem.c example that is provided with this workshop we can see how our user space application starts off. The first few things that it does is attempt to validate its environment.

demo_devmem.c example

```
01 //  
02 // open() the /dev/mem device  
03 //  
04 devmem_fd = open("/dev/mem", O_RDWR | O_SYNC);  
05 if(devmem_fd < 0) {  
06 ...cut...  
07 }  
08 //  
09 // mmap() the base of our demo_driver hardware  
10 //  
11 demo_driver_map = mmap(NULL, sysconf(_SC_PAGE_SIZE), PROT_READ|PROT_WRITE,  
12 // MAP_SHARED, devmem_fd, DEMO_DRIVER_PHYS_BASE);  
13 if(demo_driver_map == MAP_FAILED) {  
14 ...cut...  
15 }  
16 //  
17 // mmap() the base of our memcpy_msgdma hardware  
18 //  
19 memcpy_msgdma_map = mmap(NULL, sysconf(_SC_PAGE_SIZE), PROT_READ|PROT_WRITE,  
20 // MAP_SHARED, devmem_fd, MEMCPY_MSGDMA_CSR_PHYS_BASE);  
21 if(memcpy_msgdma_map == MAP_FAILED) {  
22 ...cut...  
23 }  
24  
25  
26  
27  
28  
29  
30
```

Then we open /dev/mem and we mmap() the base addresses to our demo_driver hardware and our memcpy_msgdma hardware.

demo_devmem.c example

```
01 //  
02 // perform the operation selected by the command line arguments  
03 //  
04 if(g_print_timer != NULL) do_print_timer(demo_driver_map);  
05 if(g_dump_rom != NULL) do_dump_rom(demo_driver_map);  
06 if(g_dump_ram != NULL) do_dump_ram(demo_driver_map);  
07 if(g_fill_ram != NULL) do_fill_ram(demo_driver_map);  
08 if(g_dma_rom_ram != NULL) do_dma_rom_ram(memcpy_msgdma_map);  
09 if(g_help != NULL) do_help();  
10 //  
11 // munmap everything and close the /dev/mem file descriptor  
12 //  
13 result = munmap(demo_driver_map, sysconf(_SC_PAGE_SIZE));  
14 if(result < 0) {  
15 ...cut...  
16 }  
17  
18 result = munmap(memcpy_msgdma_map, sysconf(_SC_PAGE_SIZE));  
19 if(result < 0) {  
20 ...cut...  
21 }  
22  
23  
24 close(devmem_fd);  
25 exit(EXIT_SUCCESS);  
26  
27  
28  
29  
30
```

Then we can use those pointers that we mmap()'ed and then munmap() them when we're finished and close /dev/mem.

demo_devmem.c example

```
01 void do_print_timer(void *demo_driver_map) {
02     volatile unsigned long *timer_base = demo_driver_map + TIMER_OFST;
03     double f_timeout_period = (double)(0.010) / ((double)(1) / (double)(DEMO_DRIVER_FR
04     unsigned long timeout_period = f_timeout_period;
05     unsigned long timer_snaps_l[4];
06     unsigned long timer_snaps_h[4];
07     int i;
08
09     //
10     // if timer is not running, start it
11     //
12     if((timer_base[ALTERA_AVALON_TIMER_STATUS_REG] & ALTERA_AVALON_TIMER_STATUS_RUN_MS
13         printf("Timer not currently running, initializing and starting timer.\n");
14         timeout_period--;
15         timer_base[ALTERA_AVALON_TIMER_PERIODL_REG] = timeout_period;
16         timer_base[ALTERA_AVALON_TIMER_PERIODH_REG] = timeout_period >> 16;
17         timer_base[ALTERA_AVALON_TIMER_CONTROL_REG] = ALTERA_AVALON_TIMER_CONTROL_START;
18     }
19
20     //
21     // Dump the timer registers
22     //
23     printf("  status = 0x%08lx\n", timer_base[ALTERA_AVALON_TIMER_STATUS_REG]);
24     printf("  control = 0x%08lx\n", timer_base[ALTERA_AVALON_TIMER_CONTROL_REG]);
25     printf("  period_l = 0x%08lx\n", timer_base[ALTERA_AVALON_TIMER_PERIODL_REG]);
26     printf("  period_h = 0x%08lx\n", timer_base[ALTERA_AVALON_TIMER_PERIODH_REG]);
27     printf("  snap_l = 0x%08lx\n", timer_base[ALTERA_AVALON_TIMER_SNAPL_REG]);
28     printf("  snap_h = 0x%08lx\n", timer_base[ALTERA_AVALON_TIMER_SNAPH_REG]);
```

We use the mmap()'ed pointer
in our C code just like any other
pointer.

demo_devmem.c example

```
01 void do_dump_rom(void *demo_driver_map) {
02     write(STDOUT_FILENO, (void *) (demo_driver_map + ROM_OFST), ROM_SPAN);
03 }
04
05 void do_dump_ram(void *demo_driver_map) {
06     write(STDOUT_FILENO, (void *) (demo_driver_map + RAM_OFST), RAM_SPAN);
07 }
08
09 void do_fill_ram(void *demo_driver_map) {
10     read(STDIN_FILENO, (void *) (demo_driver_map + RAM_OFST), RAM_SPAN);
11 }
12
13 void do_dma_rom_ram(void *memcpy_msgdma_map) {
14     volatile unsigned long *memcpy_msgdma_csr_base = memcpy_msgdma_map;
15     volatile unsigned long *memcpy_msgdma_desc_base = memcpy_msgdma_map +
16                                         MEMCPY_MSGDMA_DESC_PHYS_OFST;
17     unsigned long temp;
18
19     //
20     // make sure the DMA is not busy and the descriptor buffer is empty
21     //
22     temp = memcpy_msgdma_csr_base[ALTERA_MSGDMA_CSR_STATUS_REG];
23     if((temp & ALTERA_MSGDMA_CSR_BUSY_MASK) != 0x00) {
24         error(0, 0, "dma is busy before first use");
25         exit(EXIT_FAILURE);
26     }
27     if((temp & ALTERA_MSGDMA_CSR_DESCRIPTOR_BUFFER_EMPTY_MASK) == 0x00) {
28         error(0, 0, "dma descriptor buffer is not empty before first use");
29         exit(EXIT_FAILURE);
30     }
31 }
```

A few more examples.

demo_devmem.c example

```
01 void validate_system_features(void) {
02 ...cut...
03     //
04     // test to see that the demo_driver device entry exists in the sysfs
05     //
06     dirname = DEMO_DRIVER_SYSFS_ENTRY_DIR;
07     dp = opendir(dirname);
08 ...cut...
09     if(closedir(dp)) {
10 ...cut...
11         //
12         // test to see that the demo_driver device entry exists in the procfs
13         //
14         dirname = DEMO_DRIVER_PROCFS_ENTRY_DIR;
15         dp = opendir(dirname);
16 ...cut...
17         if(closedir(dp)) {
18 ...cut...
19             //
20             // fetch the clocks value out of our device entry
21             //
22             filename = DEMO_DRIVER_CLOCKS_ENTRY;
23             fd = open(filename, O_RDONLY);
24 ...cut...
25             result = read(fd, clocks_array, 4);
26 ...cut...
27             if(close(fd)) {
28 ...cut...
29
30
```

We use the run time device tree models to validate the system that we're running in before we start touching any hardware.

demo_devmem.h example

```

01 //  

02 // expected values for the physical base address and clock frequency of the  

03 // demo driver hardware  

04 //  

05 #define DEMO_DRIVER_PHYS_BASE (0xFF230000)  

06 #define DEMO_DRIVER_SYSFS_ENTRY_DIR "/sys/bus/platform/devices/ff230000.driver"  

07 #define DEMO_DRIVER_PROCFS_ENTRY_DIR "/proc/device-tree/sopc@0/bridge@0xc0000000/driver@0x1000000000000000"  

08 #define DEMO_DRIVER_FREQ      (50000000)  

09 #define DEMO_DRIVER_CLOCKS_ENTRY "/proc/device-tree/sopc@0/bridge@0xc0000000/..cut../clocks"  

10 #define H2F_USER1_CLOCK_PHANDLE_ENTRY "/proc/device-tree/clocks/clk_0/linux,phandle"  

11 //  

12 // expected values for the physical base addresses of the memcpy_mgdma core  

13 //  

14 #define MEMCPY_MSGDMA_CSR_PHYS_BASE (0xFF220000)  

15 #define MEMCPY_MSGDMA_DESC_PHYS_BASE (0xFF220020)  

16 #define MEMCPY_MSGDMA_DESC_PHYS_OFST (MEMCPY_MSGDMA_DESC_PHYS_BASE - MEMCPY_MSGDMA_CSR_PHYS)  

17 #define MEMCPY_MSGDMA_REG_NAMES_ENTRY "/proc/device-tree/sopc@0/bridge@0xc0000000/msgdma@0x1000000000000000"  

18 #define MEMCPY_MSGDMA_REG_NAMES_VALUE { 0x63, 0x73, 0x72, 0x00, 0x64, 0x65, 0x73, 0x63, 0x72,  

19 0x6f, 0x72, 0x5f, 0x73, 0x6c, 0x61, 0x76, 0x65, 0x00 }  

20 #define MEMCPY_MSGDMA_REG_ENTRY "/proc/device-tree/sopc@0/bridge@0xc0000000/msgdma@0x1000200000000000"  

21 #define MEMCPY_MSGDMA_REG_VALUE { 0x00, 0x00, 0x00, 0x01, 0x00, 0x02, 0x00, 0x00, 0x00, 0x00,  

22 0x00, 0x00, 0x01, 0x00, 0x02, 0x00, 0x20, 0x00, 0x00, 0x00, 0x10 }  

23 //  

24 // demo driver hardware map  

25 //  

26 #define ROM_OFST    (0)  

27 #define ROM_SPAN   (1024)  

28 #define RAM_OFST    (ROM_OFST + ROM_SPAN)  

29 #define RAM_SPAN   (1024)  

30 #define TIMER_OFST  (RAM_OFST + RAM_SPAN)

```

Various system parameters that define our hardware, and device tree values we can validate against at run time from user space.

Various system parameters that define our hardware, and device tree values we can validate against at run time from user space.

demo_devmem build script

```
01 #!/bin/sh
02
03 type -t ${CROSS_COMPILE:?}gcc > /dev/null 2>&1 || {
04     echo ""
05     echo "ERROR: cross compiler tools are not visible in the environment."
06     echo ""
07     exit 1
08 }
09
10 ${CROSS_COMPILE:?}gcc \
11     -march=armv7-a \
12     -mfloat-abi=hard \
13     -mfpu=vfp3 \
14     -mthumb-interwork \
15     -mthumb \
16     -O2 \
17     -g \
18     -feliminate-unused-debug-types \
19     -std=gnu99 \
20     -w \
21     -wall \
22     -werror \
23     -Wc++-compat \
24     -Wwrite-strings \
25     -Wstrict-prototypes \
26     -pedantic \
27     -o demo_devmem \
28     demo_devmem.c
29
30
```

A typical gcc invocation is all it takes to build the application.

demo_devmem target demonstration


```
01 alias ls='ls --color=never'          # if you'd like to disabled colored output for 'ls'
02 cd /examples/drivers                # this is where the example programs are installed
03
04 ./demo_devmem -h
05 ./demo_devmem -t
06 ./demo_devmem -o | hexdump -Cv
07 ./demo_devmem -a | hexdump -Cv
08 dd if=/dev/zero | ./demo_devmem -f
09 ./demo_devmem -a | hexdump -Cv
10 ./demo_devmem -d
11 ./demo_devmem -a | hexdump -Cv
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
```

demo_devmem target demonstration

```
01 ./demo_devmem -t  
02  
03 CPU reads from timer registers.
```

```
04  
05  
06  
07  
08  
09
```

```
10
```




demo_devmem target demonstration

```
01 ./demo_devmem -o  
02  
03 CPU reads from ROM.
```

```
04  
05  
06  
07  
08  
09
```

```
10
```




demo_devmem target demonstration

```
01 ./demo_devmem -a  
02  
03 CPU reads from RAM.
```

```
04  
05  
06  
07  
08  
09
```

```
10
```




demo_devmem target demonstration

```
01 ./demo_devmem -d  
02  
03 MSGDMA reads from ROM.
```

```
04  
05  
06  
07  
08  
09
```

```
10
```




demo_devmem target demonstration

```
01 ./demo_devmem -d  
02  
03 MSGDMA writes to RAM.
```

```
04  
05  
06  
07  
08  
09
```

```
10
```



/dev/mem summary

- Using /dev/mem to mmap() IO space into your user space application can be quick and convenient with little overhead.
- mmap()'ing /dev/mem allows you to have peek and poke access into the IO space of your FPGA hardware.
- mmap()'ing /dev/mem will NOT allow you have access to kernel space memory, or any other user space memory outside your own application's memory.
- mmap()'ing /dev/mem will NOT allow you to register an interrupt handler or interact with hardware interrupts in any way.

Beginning device driver development

The Altera logo, featuring the word "ALTERA" in a bold, blue, sans-serif font. A registered trademark symbol (®) is located at the top right corner of the letter "A".

ALTERA®

Some helpful pointers and documentation

- ☛ The Altera linux socfpga repository is located on rocketboards.org.
 - `git clone git://git.rocketboards.org/linux-socfpga.git`
- ☛ Module writers should be familiar with the Kbuild environment used by the kernel build flow.
 - `linux-socfpga/Documentation/kbuild`
 - `linux-socfpga/Documentation/kbuild/modules.txt`
- ☛ Module writers should be familiar with the kernel coding style, especially if you plan to upstream your work.
 - `linux-socfpga/Documentation/CodingStyle`
 - `${OUT_DIR}/scripts/Lindent <source file>`
 - `${OUT_DIR}/scripts/checkpatch.pl --file --strict <source file>`
- ☛ Kernel tainting.
 - `linux-socfpga/Documentation/oops-tracing.txt`
- ☛ GIT resources and documentation.
 - <http://git-scm.com/>
 - <http://gitref.org/>

Setup the development environment

```
01 #
02 # ARCH should be arm for the Cortex A9
03 #
04 export ARCH=arm
05
06 #
07 # CROSS_COMPILE should point into the target cross compiler tools directory and
08 # contain the prefix used by all the target tools in that directory
09 #
10 export CROSS_COMPILE=...cut.../soc_workshop/toolchain/bin/arm-linux-gnueabihf-
11
12 #
13 # OUT_DIR should point into the linux kernel build directory, that's where the
14 # .config file resides for the kernel that we're building against.
15 #
16 export OUT_DIR=...cut.../soc_workshop/socfpga-3.10-ltsi
17
18
19 When we're going to be compiling drivers against the kernel that we built for our target
20 platform, it's rather convenient to define these three variables in our environment so the
21 kernel makefiles have access to their values.
22
23
24 To compile a kernel module, our architecture is always going to be "arm", we need to have
25 access to the cross compiler tools that the kernel was built with, and we need to have
26 access to the kernel output directory, so our module can inherit the kernel configuration that
27 we're running on the target.
28
29
30
```

Setup the development environment

```
01 []$ echo ${CROSS_COMPILE}  
02 ...cut.../soc_workshop/toolchain/bin/arm-linux-gnueabihf-  
03 []$ ls ...cut.../soc_workshop/toolchain/bin/  
04 arm-linux-gnueabihf-addr2line      arm-linux-gnueabihf-gfortran  
05 arm-linux-gnueabihf-ar            arm-linux-gnueabihf-gprof  
06 arm-linux-gnueabihf-as            arm-linux-gnueabihf-ld  
07 arm-linux-gnueabihf-c++          arm-linux-gnueabihf-ld.bfd  
08 arm-linux-gnueabihf-c++filt       arm-linux-gnueabihf-ldd  
09 arm-linux-gnueabihf-cpp          arm-linux-gnueabihf-ld.gold  
10 arm-linux-gnueabihf-ct-ng.config  arm-linux-gnueabihf-nm  
11 arm-linux-gnueabihf-dwp           arm-linux-gnueabihf-objcopy  
12 arm-linux-gnueabihf-elfedit        arm-linux-gnueabihf-objdump  
13 arm-linux-gnueabihf-g++          arm-linux-gnueabihf-pkg-config  
14 arm-linux-gnueabihf-gcc          arm-linux-gnueabihf-pkg-config-real  
15 arm-linux-gnueabihf-gcc-4.9.2     arm-linux-gnueabihf-ranlib  
16 arm-linux-gnueabihf-gcc-ar        arm-linux-gnueabihf-readelf  
17 arm-linux-gnueabihf-gcc-nm        arm-linux-gnueabihf-size  
18 arm-linux-gnueabihf-gcc-ranlib    arm-linux-gnueabihf-strings  
19 arm-linux-gnueabihf-gcov         arm-linux-gnueabihf-strip  
20 arm-linux-gnueabihf-gdb  
21 []$ ls ${OUT_DIR}  
22 arch          firmware  kernel          Module.symvers  sound  
23 block         fs         lib              net             System.map  
24 COPYING        include   MAINTAINERS      README          tools  
25 CREDITS       init       Makefile        REPORTING-BUGS  usr  
26 crypto         ipc       mm              samples         virt  
27 Documentation Kbuild    modules.builtin  scripts        vmlinux  
28 drivers        Kconfig   modules.order    security       vmlinux.o  
29  
30
```

The typical Linaro
GCC tools for ARM
linux development.

The standard kernel
source tree/output
directory.

Kbuild file

```
01 []$ cat Kbuild
02 obj-m := demo_module_01.o
03 obj-m += demo_module_01t.o
04 obj-m += demo_module_02.o
05 obj-m += demo_module_03.o
06 obj-m += demo_module_04.o
07 obj-m += demo_module_05.o
08 obj-m += demo_module_05t.o
09 obj-m += demo_module_06.o
10 obj-m += demo_module_07.o
11 obj-m += demo_module_08.o
12 obj-m += demo_module_09.o
13 obj-m += demo_module_10.o
14 obj-m += demo_module_11.o
15 obj-m += demo_module_11t.o
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
```

We start with a standard Kbuild file to define what modules we wish to build in our development directory.

Makefile

```
01 []$ cat Makefile
02 ifndef OUT_DIR
03     $(error OUT_DIR is undefined, bad environment,
04 you point OUT_DIR to the linux kernel build output
05 directory)
06 endif
07
08 KDIR ?= $(OUT_DIR)
09
10 default:
11     $(MAKE) -C $(KDIR) M=$$PWD
12
13 clean:
14     $(MAKE) -C $(KDIR) M=$$PWD clean
15
16 help:
17     $(MAKE) -C $(KDIR) M=$$PWD help
18
19 modules:
20     $(MAKE) -C $(KDIR) M=$$PWD modules
21
22 modules_install:
23     $(MAKE) -C $(KDIR) M=$$PWD modules_install
24
25
26
27
28
29
30
```

We use a standard module Makefile in our development directory.

Ready to write some code



demo_module_01

User Space

demo_module_01

Kernel Space

IO Space

Licensing

```
01 /*
02 * Copyright (C) 2015 Altera Corporation
03 *
04 * This program is free software; you can redistribute it and/or modify it
05 * under the terms of the GNU General Public License as published by the Free
06 * Software Foundation; either version 2 of the License, or (at your option)
07 * any later version.
08 *
09 * This program is distributed in the hope that it will be useful, but WITHOUT
10 * ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or
11 * FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for
12 * more details.
13 *
14 * You should have received a copy of the GNU General Public License along with
15 * this program; if not, write to the Free Software Foundation, Inc., 59 Temple
16 * Place, Suite 330, Boston, MA 02111-1307 USA
17 */
```

The last thing that most folks care to think about when they start writing a piece of software is the legal licensing requirements of the software they are about to write. But if you ever plan to share your software with anyone else this is one of the most important pieces of the source file, as it should clearly specify what rights anyone else has to use this code and what obligations they assume if they do.

The demonstration modules provided in this workshop all carry the GPLv2 license shown above. For more information on open source software licenses please see:
<http://opensource.org/licenses>

demo_module_01.c – the most basic module

```
01 #include <linux/module.h>
02
03 static int demo_init(void)
04 {
05     pr_info("demo_init enter\n");
06     pr_info("demo_init exit\n");
07     return 0;
08 }
09
10 static void demo_exit(void)
11 {
12     pr_info("demo_exit enter\n");
13     pr_info("demo_exit exit\n");
14 }
15
16 module_init(demo_init);
17 module_exit(demo_exit);
18
19 MODULE_LICENSE("GPL");
20 MODULE_AUTHOR("Driver Student One <dso@company.com>");
21 MODULE_AUTHOR("Driver Student Two <dst@organization.org>");
22 MODULE_DESCRIPTION("Demonstration Module 1 - the most basic module example");
23 MODULE_VERSION("1.0");
24
25
26
27
28
29
30
```

Every module has an init and an exit function. There are some macros that you could invoke if you just want some rather standard boilerplate init and exit functionality, but every module will have an init and exit entry point.

demo_module_01.c – the most basic module

```
01 #include <linux/module.h>
02
03 static int demo_init(void)
04 {
05     pr_info("demo_init enter\n");
06     pr_info("demo_init exit\n");
07     return 0;
08 }
09
10 static void demo_exit(void)
11 {
12     pr_info("demo_exit enter\n");
13     pr_info("demo_exit exit\n");
14 }
15
16 module_init(demo_init);
17 module_exit(demo_exit);
18
19 MODULE_LICENSE("GPL");
20 MODULE_AUTHOR("Driver Student One <dso@company.com>");
21 MODULE_AUTHOR("Driver Student Two <dst@organization.org>");
22 MODULE_DESCRIPTION("Demonstration Module 1 - the most basic module example");
23 MODULE_VERSION("1.0");
24
25
26
27
28
29
30
```

These macros export the symbols for the init and exit functions such that the kernel code that loads your module can identify these entry points.

demo_module_01.c – the most basic module

```
01 #include <linux/module.h>
02
03 static int demo_init(void)
04 {
05     pr_info("demo_init enter\n");
06     pr_info("demo_init exit\n");
07     return 0;
08 }
09
10 static void demo_exit(void)
11 {
12     pr_info("demo_exit enter\n");
13     pr_info("demo_exit exit\n");
14 }
15
16 module_init(demo_init);
17 module_exit(demo_exit);
18
19 MODULE_LICENSE("GPL");
20 MODULE_AUTHOR("Driver Student One <dso@company.com>");
21 MODULE_AUTHOR("Driver Student Two <dst@organization.org>");
22 MODULE_DESCRIPTION("Demonstration Module 1 - the most basic module example");
23 MODULE_VERSION("1.0");
24
25
26
27
28
29
30
```

There are a collection of macros used to identify various attributes about a module. These strings get packaged into the module and can be accessed by various tools.

The most important module description macro is the **MODULE_LICENSE** macro. If this macro is not set to some sort of GPL license tag, then the kernel will become “tainted” when you load your module.

demo_module_01t.c – a tainted module

```
01 []$ diff -u demo_module_01.c demo_module_01t.c
02 --- demo_module_01.c
03 +++ demo_module_01t.c
04 @@ -34,8 +34,10 @@
05 module_init(demo_init);
06 module_exit(demo_exit);
07
08 +/*
09 MODULE_LICENSE("GPL");
10 +*/
11 MODULE_AUTHOR("Driver Student One <dso@company.com>");  

12 MODULE_AUTHOR("Driver Student Two <dst@organization.org>");  

13 -MODULE_DESCRIPTION("Demonstration Module 1 - the most basic module example");
14 +MODULE_DESCRIPTION("Demonstration Module 1t - tainted module example");
15 MODULE_VERSION("1.0");
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
```

demo_module_01t.c demonstrates a tainted module, simply by commenting out the MODULE_LICENCE macro.

Building the module

- Once we have the environment set up with pointers to the cross compiler tools and the kernel output directory, and we've created a Kbuild file and a Makefile for our project, and we've written our module source C file, then all we need to do to build it is run "make".
- The output of a successful module build will be a *.ko file.
- This is the file that we can copy over to our target and load as a module to run it on the target.

modinfo

```
01 []$ modinfo demo_module_01.ko
02 filename: ...cut.../demo_module_01.ko
03 version: 1.0
04 description: Demonstration Module 1 - the most basic module example
05 author: Driver Student Two <dst@organization.org>
06 author: Driver Student One <dso@company.com>
07 license: GPL
08 srcversion: 53D0236960C81C1164C0F6D
09 depends:
10 vermagic: 3.10.31-ltsi SMP mod_unload ARMv7 p2v8
11 []$ modinfo demo_module_01t.ko
12 filename: ...cut.../demo_module_01t.ko
13 version: 1.0
14 description: Demonstration Module 1t - tainted module example
15 author: Driver Student Two <dst@organization.org>
16 author: Driver Student One <dso@company.com>
17 srcversion: 53C8A51D7CA7ED5FB6C7217
18 depends:
19 vermagic: 3.10.31-ltsi SMP mod_unload ARMv7 p2v8
20
21
22
23
```

Once our module is built, we can use modinfo to see what the MODULE_* macros defined about our module. Notice the lack of a license property for the demo_module_01t.ko module.

demo_module_01.ko target demonstration

```
01 cat /proc/sys/kernel/tainted          # should be 0 = untainted
02 lsmod                                # should be "Not tainted"
03 cat /proc/modules                      # this is what lsmod formats for you
04 insmod demo_module_01.ko              # insert our module
05 cat /proc/sys/kernel/tainted          # should be 4096 = GPL, out of tree
06 lsmod                                # should be "Tainted: G"
07 insmod demo_module_01t.ko             # insert our tainted module
08 cat /proc/sys/kernel/tainted          # should be 4097 = proprietary, oot
09 lsmod                                # should be "Tainted: P"
10 cat /proc/modules                      # our modules appear (P) and (O)
11 find /sys -name "*demo*"            # find our modules in sysfs
12 ls /sys/module/demo_module_01        # see what the directory contains
13 cat /sys/module/demo_module_01/taint # should be "0" = GPL, out of tree
14 cat /sys/module/demo_module_01t/taint # should be "P0" = proprietary, oot
15 rmmod demo_module_01t                # remove our modules
16 rmmod demo_module_01
17 lsmod                                # still tainted
18 cat /proc/sys/kernel/tainted          # still tainted
19 ls /sys/module/                      # these are all the modules, in tree
20 tail /var/log/messages               # the printk messages from console
21
22
23
24 If you're target has not been freshly booted, your kernel may already be tainted. If so,
25 consider rebooting the target to start with a fresh untainted environment.
26
27
28
29
30
```


Module parameters

For more information:
include/linux/moduleparam.h



demo_module_02

User Space



Kernel Space

IO Space

demo_module_02.c – module parameters

```
01 /* define module parameters */  
02 static unsigned char param_byte = 0xFF;  
03 static short param_short = 0xFFFF;  
04 static unsigned short param_ushort = 0xFFFF;  
05 static int param_int = 0xFFFFFFFF;  
06 static unsigned int param_uint = 0xFFFFFFFF;  
07 static long param_long = 0xFFFFFFFF;  
08 static unsigned long param_ulong = 0xFFFFFFFF;  
09 static bool param_bool = 1;  
10 static char *param_charp;  
11  
12 module_param(param_byte, byte, S_IRUGO);  
13 module_param(param_short, short, S_IRUGO);  
14 module_param(param_ushort, ushort, S_IRUGO);  
15 module_param(param_int, int, S_IRUGO);  
16 module_param(param_uint, uint, S_IRUGO);  
17 module_param(param_long, long, S_IRUGO);  
18 module_param(param_ulong, ulong, S_IRUGO);  
19 module_param(param_bool, bool, S_IRUGO);  
20 module_param(param_charp, charp, S_IRUGO);  
21  
22  
23 static int param_int_array[] = { 1, 2, 3, 4 };  
24  
25 static unsigned int param_int_array_count;  
26  
27 module_param_array(param_int_array, int, &param_int_array_count, S_IRUGO);  
28  
29  
30
```

We can define parameters for our module which allows us to pass in command line arguments to override these default values at runtime when we insmod or modprobe the module.

We start by defining initialized global variables, these are the default values unless overridden when loaded.

Then this `module_param()` macro informs the kernel that these variables are parameters to the module.

This is an array parameter example here.

demo_module_02.c – module parameters

```
01 MODULE_PARM_DESC(param_byte, "a byte parameter");
02 MODULE_PARM_DESC(param_short, "a short parameter");
03 MODULE_PARM_DESC(param_ushort, "a ushort parameter");
04 MODULE_PARM_DESC(param_int, "a int parameter");
05 MODULE_PARM_DESC(param_uint, "a uint parameter");
06 MODULE_PARM_DESC(param_long, "a long parameter");
07 MODULE_PARM_DESC(param_ulong, "a ulong parameter");
08 MODULE_PARM_DESC(param_bool, "a bool parameter");
09 MODULE_PARM_DESC(param_charp, "a charp parameter");
10 MODULE_PARM_DESC(param_int_array, "an array of int parameters");
11
12
13
14
15
16
17
18
19
20
21
22
23
24 Remember those MODULE_* macros that we saw before. MODULE_PARAM_DESC will
25 put these strings in the module, which makes them visible to modinfo dumps.
26
27
28
29
30
```

demo_module_02.c – module parameters

```
01 static int demo_init(void)
02 {
03     pr_info("demo_init enter\n");
04
05     pr_info("\n");
06     pr_info("param_byte    = 0x%02X : %u\n", param_byte, param_byte);
07     pr_info("param_short   = 0x%04X : %d\n", param_short, param_short);
08     pr_info("param_ushort  = 0x%04X : %u\n", param_ushort, param_ushort);
09     pr_info("param_int     = 0x%08X : %d\n", param_int, param_int);
10    pr_info("param_uint    = 0x%08X : %u\n", param_uint, param_uint);
11    pr_info("param_long    = 0x%08lx : %ld\n", param_long, param_long);
12    pr_info("param_ulong   = 0x%08lx : %lu\n", param_ulong, param_ulong);
13    pr_info("param_bool    = %d\n", param_bool);
14    pr_info("param_charp   = \\'%s\\'\n", param_charp);
15    pr_info("\n");
16    pr_info("param_int_array_count = %d\n", param_int_array_count);
17    pr_info("param_int_array[0] = %d\n", param_int_array[0]);
18    pr_info("param_int_array[1] = %d\n", param_int_array[1]);
19    pr_info("param_int_array[2] = %d\n", param_int_array[2]);
20    pr_info("param_int_array[3] = %d\n", param_int_array[3]);
21    pr_info("\n");
22
23    pr_info("demo_init exit\n");
24    return 0;
25 }
```

26
27 This example simply dumps all the parameter values from it's init routine.
28
29
30

demo_module_02.ko target demonstration

```
01 insmod demo_module_02.ko          # insert the module without parameters
02 rmmod demo_module_02             # remove the module
03
04 insmod demo_module_02.ko \
05 param_byte=0x12 \
06 param_short=0x3456 \
07 param_ushort=0x789A \
08 param_int=0x3CDEF012 \
09 param_uint=0x3456789A \
10 param_long=0x3CDEF012 \
11 param_ulong=0x3456789A \
12 param_bool=n \
13 param_charp=hello-tim \
14 param_int_array=5000,-6000,7000      # insert the module with parameters
15
16 ls /sys/module/demo_module_02/        # notice the parameters directory
17 ls /sys/module/demo_module_02/parameters/    # notice the parameter files
18 cat /sys/module/demo_module_02/parameters/param_bool
19 cat /sys/module/demo_module_02/parameters/param_byte | hexdump -C
20 cat /sys/module/demo_module_02/parameters/param_charp
21 cat /sys/module/demo_module_02/parameters/param_int
22 cat /sys/module/demo_module_02/parameters/param_int_array
23 cat /sys/module/demo_module_02/parameters/param_long
24 printf "0x%08X\n" $(cat /sys/module/demo_module_02/parameters/param_long)
25
26 rmmod demo_module_02             # remove the module
27
28
29 Observe the output on the printk console.
30
```

Platform drivers

For more information:

Documentation/driver-model/platform.txt

Documentation/devicetree/



demo_module_03

User Space

demo_module_03

demo_driver_3

Kernel Space

demo_driver_subsys_0
demo_driver-1.0

IO Space

demo_module_03.c – platform driver

```
01 #include <linux/module.h>
02 #include <linux/platform_device.h>
03
04 static struct of_device_id demo_driver_dt_ids[] = {
05     {
06         .compatible = "demo,driver-1.0"
07     },
08     { /* end of table */ }
09 };
10
11 MODULE_DEVICE_TABLE(of, demo_driver_dt_ids);
12
13 static struct platform_driver the_platform_driver = {
14     .probe = platform_probe,
15     .remove = platform_remove,
16     .driver = {
17         .name = "demo_driver_3",
18         .owner = THIS_MODULE,
19         .of_match_table = demo_driver_dt_ids,
20     },
21     /*
22         .shutdown = unused,
23         .suspend = unused,
24         .resume = unused,
25         .id_table = unused,
26     */
27 };
28
29
30
```

We create an array of of_device_id's where we specify ".compatible" strings that the kernel will use to bind our driver to any devices represented in the device tree when our module is inserted. This will automatically trigger our driver probe if the device tree contains a compatible device entry.

We also define an instance for our platform_driver.

demo_module_03.c – platform driver

```
01 static int platform_probe(struct platform_device *pdev)
02 {
03     pr_info("platform_probe enter\n");
04     pr_info("platform_probe exit\n");
05     return 0;
06 }
07
08 static int platform_remove(struct platform_device *pdev)
09 {
10     pr_info("platform_remove enter\n");
11     pr_info("platform_remove exit\n");
12     return 0;
13 }
```

We define our probe and remove functions.

demo_module_03.c – platform driver

```
01 static int demo_init(void)
02 {
03     int ret_val;
04     pr_info("demo_init enter\n");
05
06     ret_val = platform_driver_register(&the_platform_driver);
07     if (ret_val != 0) {
08         pr_err("platform_driver_register returned %d\n", ret_val);
09         return ret_val;
10    }
11
12    pr_info("demo_init exit\n");
13    return 0;
14 }
15
16 static void demo_exit(void)
17 {
18     pr_info("demo_exit enter\n");
19
20     platform_driver_unregister(&the_platform_driver);
21
22     pr_info("demo_exit exit\n");
23 }
```

We register our platform driver in our module init function and we unregister our platform driver in our module exit function.

demo_module_03.ko target demonstration

```
01 insmod demo_module_03.ko          # insert the module
02 find /sys -name "*demo*"        # find sysfs entries
03 ls -l /sys/module/demo_module_03/drivers/ # this is a link to the driver entry
04 ls /sys/module/demo_module_03      # this looks like the module we saw before
05 ls /sys/bus/platform/drivers/demo_driver_3  # this is a platform driver entry
06
07 # this links to the device entry
08 ls -l /sys/bus/platform/drivers/demo_driver_3/ff230000.driver
09
10 # we can unbind the device
11 echo "ff230000.driver" > /sys/bus/platform/drivers/demo_driver_3/unbind
12 ls /sys/bus/platform/drivers/demo_driver_3      # and we see the result
13
14 # we can bind the device
15 echo "ff230000.driver" > /sys/bus/platform/drivers/demo_driver_3/bind
16 ls /sys/bus/platform/drivers/demo_driver_3      # and we see the result
17 rmmod demo_module_03               # remove the device
18
19
20
21
22
23
24
25
26
27 Observe the output on the printk console.
28
29
30
```

Extracting information from the device tree

For more information:

Documentation/devicetree/

Documentation/clk.txt



Retrieving our properties from the device tree

- ⬚ There are three properties that are of significant interest to most hardware blocks and the software drivers that control them.
 - Memory region
 - ⬚ every hardware peripheral will likely consume some span of the address map, and the driver needs to know where that is so that it can map it to a pointer that can interact with that region.
 - Interrupt signal
 - ⬚ peripherals that rely on a hardware interrupt signal need to know what interrupt line that event is signaled on so that an interrupt handler may be registered to handle it.
 - Clock
 - ⬚ some peripherals that are sensitive to their clock properties will need to know what those details are.

demo_module_04

User Space

demo_module_04

demo_driver_4

Kernel Space

demo_driver_subsys_0
demo_driver-1.0

IO Space

demo_module_04.c – extracting from device tree

```
01 static int platform_probe(struct platform_device *pdev)
02 {
03     int ret_val;
04     struct resource *r;
05     int irq;
06     struct clk *clk;
07     unsigned long clk_rate;
08
09     pr_info("platform_probe enter\n");
10
11     ret_val = -EINVAL;
12
13     /* get our first memory resource */
14     r = platform_get_resource(pdev, IORESOURCE_MEM, 0);
15     if (r != NULL) {
16         pr_info("r->start = 0x%08lx\n", (long unsigned int)r->start);
17         pr_info("r->end = 0x%08lx\n", (long unsigned int)r->end);
18         pr_info("r->name = %s\n", r->name);
19         if (r->start & ~PAGE_MASK) {
20             ...cut...
21             }
22             if (((r->end - r->start) + 1) > PAGE_SIZE) {
23             ...cut...
24             }
25         } else {
26             pr_err("IORESOURCE_MEM, 0 does not exist\n");
27             goto bad_exit_return;
28         }
29
30         g_demo_driver_base_addr = r->start;
```

We retrieve the memory regions allocated by our device by calling `platform_get_resource()`.

demo_module_04.c – extracting from device tree

```
01      /* get our interrupt resource */
02      irq = platform_get_irq(pdev, 0);
03      if (irq < 0) {
04 ...cut...
05          } else {
06              pr_info("irq = %d\n", irq);
07          }
08
09      g_demo_driver_irq = irq;
10
11      /* get our clock resource */
12      clk = clk_get(&pdev->dev, NULL);
13      if (IS_ERR(clk)) {
14 ...cut...
15          } else {
16              clk_rate = clk_get_rate(clk);
17              pr_info("clk = %lu HZ\n", clk_rate);
18          }
19
20      g_demo_driver_clk_rate = clk_rate;
21
22      pr_info("platform_probe exit\n");
23      return 0;
24
25 bad_exit_return:
26      pr_info("platform_probe bad_exit\n");
27      return ret_val;
28 }
29
30 }
```

We retrieve the irq that our hardware is allocated by calling `platform_get_irq()`.

We retrieve our clock information by calling `clk_get()` and then `clk_get_rate()`.

demo_module_04.ko target demonstration (1)

```
01 insmod demo_module_03.ko          # insert the first module, notice probe
02 insmod demo_module_04.ko          # insert the second module, no probe
03 rmmod demo_module_03             # remove the first module
04 rmmod demo_module_04             # remove the module
05 insmod demo_module_04.ko          # insert the module notice probe printk
06 rmmod demo_module_04             # remove the module
07
08 ls /proc/device-tree/
09 ls /proc/device-tree/sopc@0/
10 ls /proc/device-tree/sopc@0/bridge@0xc0000000/
11 ls /proc/device-tree/sopc@0/bridge@0xc0000000/driver@0x100030000/
12 hexdump -C /proc/device-tree/sopc@0/bridge@0xc0000000/driver@0x100030000/compatible
13 hexdump -C /proc/device-tree/sopc@0/bridge@0xc0000000/driver@0x100030000/name
14
15 # reg = [0:7] bridge address, [8:11] span
16 hexdump -C /proc/device-tree/sopc@0/bridge@0xc0000000/driver@0x100030000/reg
17
18 # ranges = [0:7] bridge address, [8:11] processor address, [12:15] span
19 hexdump -C /proc/device-tree/sopc@0/bridge@0xc0000000/ranges
20
21 # interrupts = [0:3] type, [4:7] number, [8:11] flags
22 hexdump -C /proc/device-tree/sopc@0/bridge@0xc0000000/driver@0x100030000/interrupts
23 printf "%d\n" <number>           # convert the hex interrupt number to decimal
24 expr <number> + 32 # add 32 to the interrupt number for actual GIC IRQ
25
26 # this is a phandle
27 hexdump -C /proc/device-tree/sopc@0/bridge@0xc0000000/driver@0x100030000/interrupt-parent
28
29 Observe the output on the printk console.
30
```

demo_module_04.ko target demonstration (2)

```
01 # list all the phandles, sorted
02 find /proc/device-tree -name "linux,phandle" | sort
03
04 # list all the phandle values, sorted
05 # then you can locate the interrupt-parrent phandle
06 find /proc/device-tree -name "linux,phandle" | sort | xargs hexdump -C
07
08 # verify the phandle
09 ls /proc/device-tree/sopc@0/intc@0xffffed000/
10 hexdump -C /proc/device-tree/sopc@0/intc@0xffffed000/linux,phandle
11
12 # this is a phandle
13 hexdump -C /proc/device-tree/sopc@0/bridge@0xc0000000/driver@0x100030000/clocks
14
15 # verify the phandle
16 ls /proc/device-tree/clocks/c1k_0/
17 hexdump -C /proc/device-tree/clocks/c1k_0/linux,phandle
18
19 # verify the clock frequency
20 hexdump -C /proc/device-tree/clocks/c1k_0/clock-frequency
21 printf "%d\n" <frequency>
22
23
24
25
26
```

27 Observe the output on the printk console.

Reserving memory regions and registering IRQ handlers



demo_module_05

User Space

demo_module_05
demo_driver_5

Kernel Space

ioremap
irq
demo_driver_subsys_0
demo_driver-1.0

IO Space

demo_module_05.c – reserve memory, register irq

```
01 static int platform_probe(struct platform_device *pdev)
02 {
03     ...cut...
04     pr_info("platform_probe enter\n");
05
06     ret_val = -EBUSY;
07
08     /* acquire the probe lock */
09     if (down_interruptible(&g_dev_probe_sem))
10         return -ERESTARTSYS;
11
12     if (g_platform_probe_flag != 0)
13         goto bad_exit_return;
14
15     ret_val = -EINVAL;
16
17     /* get our first memory resource */
18     r = platform_get_resource(pdev, IORESOURCE_MEM, 0);
19     if (r == NULL) {
20         pr_err("IORESOURCE_MEM, 0 does not exist\n");
21         goto bad_exit_return;
22     }
23
24     g_demo_driver_base_addr = r->start;
25     g_demo_driver_size = resource_size(r);
26
27
28
29
30
```

We start by retrieving our memory resource.

Note that we acquire a semaphore for the driver resources with the `down_interruptible()` call.

demo_module_05.c – reserve memory, register irq

```
01      /* get our interrupt resource */
02      irq = platform_get_irq(pdev, 0);
03      if (irq < 0) {
04          pr_err("irq not available\n");
05          goto bad_exit_return;
06      }
07
08      g_demo_driver_irq = irq;
09
10     /* get our clock resource */
11     clk = clk_get(&pdev->dev, NULL);
12     if (IS_ERR(clk)) {
13         pr_err("clk not available\n");
14         goto bad_exit_return;
15     } else {
16         clk_rate = clk_get_rate(clk);
17     }
18
19
20
21
22
23
24
25
26
27
28
29
30
```

Then we retrieve our IRQ number and our clock information.

demo_module_05.c – reserve memory, register irq

```
01     ret_val = -EBUSY;  
02  
03     /* reserve our memory region */  
04     demo_driver_mem_region = request_mem_region(g_demo_driver_base_addr,  
05                                                 g_demo_driver_size,  
06                                                 "demo_driver_hw_region");  
07     if (demo_driver_mem_region == NULL) {  
08         pr_err("request_mem_region failed: g_demo_driver_base_addr\n");  
09         goto bad_exit_return;  
10    }  
11  
12     /* ioremap our memory region */  
13     g_ioremap_addr = ioremap(g_demo_driver_base_addr, g_demo_driver_size);  
14     if (g_ioremap_addr == NULL) {  
15         pr_err("ioremap failed: g_demo_driver_base_addr\n");  
16         goto bad_exit_release_mem_region;  
17    }  
18  
19     g_timer_base = g_ioremap_addr + TIMER_OFST;  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30
```

Now we can reserve our memory region and remap it into an IO pointer that our driver can use.

demo_module_05.c – reserve memory, register irq

```
01  /* initialize our peripheral timer hardware */
02  io_result = ioread32(IOADDR_ALTERA_AVALON_TIMER_STATUS(g_timer_base));
03  io_result &= ALTERA_AVALON_TIMER_STATUS_TO_MSK |
04      ALTERA_AVALON_TIMER_STATUS_RUN_MSK;
05  if (io_result != 0) {
06      pr_err("peripheral timer hardware, incorrect initial state");
07      goto bad_exit_iounmap;
08  }
09
10 period_100ms = (clk_rate / 10) - 1;
11 iowrite32(period_100ms,
12             IOADDR_ALTERA_AVALON_TIMER_PERIODL(g_timer_base));
13 iowrite32(period_100ms >> 16,
14             IOADDR_ALTERA_AVALON_TIMER_PERIODH(g_timer_base));
```

Now that we have an IO pointer we can initialize our timer hardware.

demo_module_05.c – reserve memory, register irq

```
01  /* register our interrupt handler */
02  ret_val = request_irq(g_demo_driver_irq,
03                      demo_driver_interrupt_handler,
04                      0,
05                      the_platform_driver.driver.name,
06                      &the_platform_driver);
07
08  if (ret_val) {
09      pr_err("request_irq failed");
10      goto bad_exit_iounmap;
11  }
12
13  ret_val = -EBUSY;
14
15  /* start our timer and enable our timer hardware interrupts */
16  iowrite32(ALTERA_AVALON_TIMER_CONTROL_I TO_MSK |
17             ALTERA_AVALON_TIMER_CONTROL_CONT_MSK |
18             ALTERA_AVALON_TIMER_CONTROL_START_MSK,
19             IOADDR_ALTERA_AVALON_TIMER_CONTROL(g_timer_base));
20
21  io_result = ioread32(IOADDR_ALTERA_AVALON_TIMER_STATUS(g_timer_base));
22  io_result &= ALTERA_AVALON_TIMER_STATUS_RUN_MSK;
23  if (io_result == 0) {
24      /* stop our timer and disable our timer hardware interrupts */
25      iowrite32(ALTERA_AVALON_TIMER_CONTROL_STOP_MSK,
26                 IOADDR_ALTERA_AVALON_TIMER_CONTROL(g_timer_base));
27
28      pr_err("peripheral timer hardware, failed to start");
29      goto bad_exit_freeirq;
30  }
```

Now we can register our IRQ handler and start our timer hardware to generate interrupts.

demo_module_05.c – reserve memory, register irq

```
01     g_platform_probe_flag = 1;
02     up(&g_dev_probe_sem);
03     pr_info("platform_probe exit\n");
04     return 0;
05
06 bad_exit_freeirq:
07     free_irq(g_demo_driver_irq, &the_platform_driver);
08 bad_exit_iounmap:
09     iounmap(g_ioremap_addr);
10 bad_exit_release_mem_region:
11     release_mem_region(g_demo_driver_base_addr, g_demo_driver_size);
12 bad_exit_return:
13     up(&g_dev_probe_sem);
14     pr_info("platform_probe bad_exit\n");
15     return ret_val;
16 }
```

And we can clean things up and return from our probe function.

Note the bad exit return labels at the bottom of the function.

This construct allows us to unwind any allocations we've made in the opposite order that we created them in.

You can also see that we release the semaphore with up().

demo_module_05.c – reserve memory, register irq

```
01 irqreturn_t demo_driver_interrupt_handler(int irq, void *dev_id)
02 {
03     /* clear the interrupt */
04     iowrite32(0, IOADDR_ALTERA_AVALON_TIMER_STATUS(g_timer_base));
05
06     return IRQ_HANDLED;
07 }
```

This is what our IRQ
handler looks like.

demo_module_05.c – reserve memory, register irq

```
01 static int platform_remove(struct platform_device *pdev)
02 {
03     uint32_t io_result;
04     pr_info("platform_remove enter\n");
05
06     /* stop our timer and disable our timer hardware interrupts */
07     iowrite32(ALTERA_AVALON_TIMER_CONTROL_STOP_MSK,
08               IOADDR_ALTERA_AVALON_TIMER_CONTROL(g_timer_base));
09     /* ensure there is no pending IRQ */
10     do {
11         io_result =
12             ioread32(IOADDR_ALTERA_AVALON_TIMER_STATUS(g_timer_base));
13         io_result &= ALTERA_AVALON_TIMER_STATUS_TO_MSK;
14     } while (io_result != 0);
15
16     free_irq(g_demo_driver_irq, &the_platform_driver);
17     iounmap(g_ioremap_addr);
18     release_mem_region(g_demo_driver_base_addr, g_demo_driver_size);
19
20     if (down_interruptible(&g_dev_probe_sem))
21         return -ERESTARTSYS;
22
23     g_platform_probe_flag = 0;
24     up(&g_dev_probe_sem);
25
26     pr_info("platform_remove exit\n");
27     return 0;
28 }
```

Our release function simply releases all the resources that we allocated in probe, and it shuts down our timer hardware.

demo_module_05.ko target demonstration

```
01 ./demo_devmem -s                                # stop the timer hardware
02 insmod demo_module_05.ko                         # insert module
03 cat /proc/interrupts | grep "demo"               # observe IRQ occurring
04 cat /proc/interrupts | grep "demo"               # observe IRQ occurring
05 cat /proc/interrupts | grep "demo"               # observe IRQ occurring
06 cat /proc/iomem | grep "demo"                   # observe reserved memory
07
08 # demo_module_05t.ko is a test module that we force to reserve the same resources that
09 # demo_module_05.ko has already reserved.  we can see which of those actions fail and which
10 # of those actions succeed.
11 insmod demo_module_05t.ko base=0xff230000 irq=80      # observe printk
12
13
14
15 rmmod demo_module_05                            # request_mem_region failed
16
17
18
19
20
21
22
23
24
25
26
27 Observe the output on the printk console.          # ioremap succeeded
28
29
30
```



Creating sysfs files

The Altera logo, featuring the word "ALTERA" in a bold, blue, sans-serif font. A registered trademark symbol (®) is located at the top right corner of the letter "A".

ALTERA®

demo_module_06

User Space



Kernel Space

IO Space

demo_module_06.c – sysfs files

```
01 static int platform_probe(struct platform_device *pdev)
02 {
03     ...cut...
04         /* acquire the probe lock */
05     ...cut...
06         /* get our first memory resource */
07     ...cut...
08         /* get our clock resource */
09     ...cut...
10         /* reserve our memory region */
11     ...cut...
12         /* ioremap our memory region */
13     ...cut...
14         /* initialize our peripheral timer hardware */
15     ...cut...
16         /* register our interrupt handler */
17     ...cut...
18         /* start our timer and enable our timer hardware interrupts */
19     ...cut...
20
21
22
23
24
25
26
27
28
29
30
```

We start our probe function in much the same way as before.

demo_module_06.c – sysfs files

```
01  /* create the sysfs entries */
02  ret_val = driver_create_file(&the_platform_driver.driver,
03                                &driver_attr_irq_delays);
04  if (ret_val != 0) {
05      pr_err("failed to create irq_delays sysfs entry");
06      goto bad_exit_stop_timer;
07  }
08
09  ret_val = driver_create_file(&the_platform_driver.driver,
10                                &driver_attr_interval);
11  if (ret_val != 0) {
12      pr_err("failed to create interval sysfs entry");
13      goto bad_exit_remove_irq_delays;
14  }
15
16  g_platform_probe_flag = 1;
17  up(&g_dev_probe_sem);
18  pr_info("platform_probe exit\n");
19  return 0;
20
21
22
23
24
25
26
27
28
29
30
```

Then we create a couple
of sysfs files with the
driver_create_file() call.

demo_module_06.c – sysfs files

```
01 bad_exit_remove_irq_delays:  
02     driver_remove_file(&the_platform_driver.driver,  
03                           &driver_attr_irq_delays);  
04 bad_exit_stop_timer:  
05     iowrite32(ALTERA_AVALON_TIMER_CONTROL_STOP_MSK,  
06                 IOADDR_ALTERA_AVALON_TIMER_CONTROL(g_timer_base));  
07     /* ensure there is no pending IRQ */  
08     do {  
09         io_result =  
10             ioread32(IOADDR_ALTERA_AVALON_TIMER_STATUS(g_timer_base));  
11         io_result &= ALTERA_AVALON_TIMER_STATUS_TO_MSK;  
12     } while (io_result != 0);  
13 bad_exit_freeirq:  
14     free_irq(g_demo_driver_irq, &the_platform_driver);  
15 bad_exit_iounmap:  
16     iounmap(g_ioremap_addr);  
17 bad_exit_release_mem_region:  
18     release_mem_region(g_demo_driver_base_addr, g_demo_driver_size);  
19 bad_exit_return:  
20     up(&g_dev_probe_sem);  
21     pr_info("platform_probe bad_exit\n");  
22     return ret_val;  
23 }  
24  
25  
26  
27  
28  
29  
30
```

If we get an error during probe then we may need to remove the sysfs files we've already created.

demo_module_06.c – sysfs files

```
01 static int platform_remove(struct platform_device *pdev)
02 {
03     ...cut...
04     driver_remove_file(&the_platform_driver.driver, &driver_attr_interval);
05     driver_remove_file(&the_platform_driver.driver,
06                         &driver_attr_irq_delays);
07
08     /* stop our timer and disable our timer hardware interrupts */
09     iowrite32(ALTERA_AVALON_TIMER_CONTROL_STOP_MSK,
10               IOADDR_ALTERA_AVALON_TIMER_CONTROL(g_timer_base));
11     /* ensure there is no pending IRQ */
12     do {
13         io_result =
14             ioread32(IOADDR_ALTERA_AVALON_TIMER_STATUS(g_timer_base));
15         io_result &= ALTERA_AVALON_TIMER_STATUS_TO_MSK;
16     } while (io_result != 0);
17
18     free_irq(g_demo_driver_irq, &the_platform_driver);
19     iounmap(g_ioremap_addr);
20     release_mem_region(g_demo_driver_base_addr, g_demo_driver_size);
21
22     if (down_interruptible(&g_dev_probe_sem))
23         return -ERESTARTSYS;
24
25     g_platform_probe_flag = 0;
26     up(&g_dev_probe_sem);
27
28     pr_info("platform_remove exit\n");
29     return 0;
30 }
```

Our release function
needs to remove the sysfs
files in addition to what it
already did.

demo_module_06.c – sysfs files

```
01 static ssize_t irq_delays_show(struct device_driver *driver, char *buf)
02 {
03     unsigned long flags;
04     uint32_t max_irq_delay;
05     uint32_t min_irq_delay;
06
07     /* acquire the irq_lock */
08     spin_lock_irqsave(&g_irq_lock, flags);
09
10    /* capture the shared data values */
11    max_irq_delay = g_max_irq_delay;
12    min_irq_delay = g_min_irq_delay;
13
14    /* release the irq_lock */
15    spin_unlock_irqrestore(&g_irq_lock, flags);
16
17    if (max_irq_delay == 0)
18        return scnprintf(buf, PAGE_SIZE, "no IRQ delays yet\n");
19
20    return scnprintf(buf, PAGE_SIZE,
21                     "max: 0x%08x = %u\n"
22                     "min: 0x%08x = %u\n",
23                     max_irq_delay, max_irq_delay,
24                     min_irq_delay, min_irq_delay);
25 }
26
27 DRIVER_ATTR(irq_delays, (S_IRUGO), irq_delays_show, NULL);
28
29
30 }
```

The `irq_delays` sysfs file that we create registers a `show()` function, so this file can be read from user space. This file allows us to read the max and min IRQ delay stats, stored by the IRQ handler.

Note the use of spinlocks.

Note the `DRIVER_ATTR` macro, registering show function only.

demo_module_06.c – sysfs files

```
01 static ssize_t interval_show(struct device_driver *driver, char *buf)
02 {
03     ...cut...
04     /* acquire the irq_lock */
05     spin_lock_irqsave(&g_irq_lock, flags);
06
07     /* capture the relevant hardware registers */
08     raw_status = ioread32(IOADDR_ALTERA_AVALON_TIMER_STATUS(g_timer_base));
09     raw_periodl =
10         ioread32(IOADDR_ALTERA_AVALON_TIMER_PERIODL(g_timer_base));
11     raw_periodh =
12         ioread32(IOADDR_ALTERA_AVALON_TIMER_PERIODH(g_timer_base));
13
14     /* release the irq_lock */
15     spin_unlock_irqrestore(&g_irq_lock, flags);
16
17     /* calculate the current timer interval */
18     raw_status &= ALTERA_AVALON_TIMER_STATUS_RUN_MSK;
19     if (raw_status == 0) {
20         interval = 0;
21     } else {
22         period = (raw_periodl & 0x0000FFFF) |
23             ((raw_periodh << 16) & 0xFFFF0000);
24         period += 1;
25         interval = g_demo_driver_clk_rate / period;
26     }
27
28     return scnprintf(buf, PAGE_SIZE,
29                         "irq interval: %u per second\n", interval);
30 }
```

The interval sysfs file that we create registers a show() function and a store() function, so this file can be read and written from user space. This show function allows us to read the current timer IRQ interval.

Note the use of spinlocks.

demo_module_06.c – sysfs files

```
01 static ssize_t interval_store(struct device_driver *driver, const char *buf,
02                               size_t count)
03 {
04     ...cut...
05     /* convert the input string to the requested new interval value */
06     result = kstrtoul(buf, 0, &new_interval);
07     if (result != 0)
08         return -EINVAL;
09
10    /* range check the requested new interval value */
11    if (new_interval > 100)
12        return -EINVAL;
13
14    /* calculate the new period value */
15    if (new_interval > 0)
16        new_period = (g_demo_driver_clk_rate / new_interval) - 1;
17    else
18        new_period = g_demo_driver_clk_rate;
19
20    /* acquire the irq_lock */
21    spin_lock_irqsave(&g_irq_lock, flags);
22
23    /* stop the interval timer */
24    iowrite32(ALTERA_AVALON_TIMER_CONTROL_STOP_MSK,
25              IOADDR_ALTERA_AVALON_TIMER_CONTROL(g_timer_base));
```

This is the interval store() function. It allows us to change the IRQ interval of the timer hardware.

Note the use of spinlocks.

demo_module_06.c – sysfs files

```
01  /* ensure there is no pending IRQ that we are blocking */
02  the_status = ioread32(IOADDR_ALTERA_AVALON_TIMER_STATUS(g_timer_base));
03  the_status &= ALTERA_AVALON_TIMER_STATUS_TO_MSK;
04  if (the_status != 0) {
05      do {
06          /*
07              if we are blocking, release the lock to allow IRQ
08              handler to execute, acquire the lock and check again
09          */
10         spin_unlock_irqrestore(&g_irq_lock, flags);
11         spin_lock_irqsave(&g_irq_lock, flags);
12         the_status =
13             ioread32(IOADDR_ALTERA_AVALON_TIMER_STATUS
14                     (g_timer_base));
15         the_status &= ALTERA_AVALON_TIMER_STATUS_TO_MSK;
16     } while (the_status != 0);
17 }
18
19 /* write the new period value */
20 iowrite32(new_period, IOADDR_ALTERA_AVALON_TIMER_PERIODL(g_timer_base));
21 iowrite32(new_period >> 16,
22           IOADDR_ALTERA_AVALON_TIMER_PERIODH(g_timer_base));
23
24 /* initialize the MAX/MIN variables */
25 g_max_irq_delay = 0;
26 g_min_irq_delay = 0xFFFFFFFF;
```

Note the use of spinlocks.

demo_module_06.c – sysfs files

```
01     /* start the timer */
02     if (new_interval > 0)
03         iowrite32(ALTERA_AVALON_TIMER_CONTROL_ITO_MSK |
04                     ALTERA_AVALON_TIMER_CONTROL_CONT_MSK |
05                     ALTERA_AVALON_TIMER_CONTROL_START_MSK,
06                     IOADDR_ALTERA_AVALON_TIMER_CONTROL(g_timer_base));
07
08     /* release the irq_lock */
09     spin_unlock_irqrestore(&g_irq_lock, flags);
10     return count;
11 }
12
13 DRIVER_ATTR(interval, (S_IWUGO | S_IRUGO), interval_show, interval_store);
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
```

The end of our store()
function.

Note the use of spinlocks.

Note the DRIVER_ATTR
macro registering a show
and store function.

demo_module_06.c – sysfs files

```
01 irqreturn_t demo_driver_interrupt_handler(int irq, void *dev_id)
02 {
03     ...cut...
04     spin_lock(&g_irq_lock);
05
06     /* get the current timer value */
07     iowrite32(0, IOADDR_ALTERA_AVALON_TIMER_SNAPL(g_timer_base));
08     raw_snapl = ioread32(IOADDR_ALTERA_AVALON_TIMER_SNAPL(g_timer_base));
09     raw_snaph = ioread32(IOADDR_ALTERA_AVALON_TIMER_SNAPH(g_timer_base));
10     snap = (raw_snapl & 0x0000FFFF) | ((raw_snaph << 16) & 0xFFFF0000);
11
12     /* get the current period value */
13     raw_periodl = ioread32(IOADDR_ALTERA_AVALON_TIMER_PERIODL(g_timer_base));
14     raw_periodh = ioread32(IOADDR_ALTERA_AVALON_TIMER_PERIODH(g_timer_base));
15     period = (raw_periodl & 0x0000FFFF) | ((raw_periodh << 16) & 0xFFFF0000);
16
17     /* calculate response delay and update MAX/MIN variables */
18     elapsed_ticks = period - snap;
19     if (elapsed_ticks > g_max_irq_delay)
20         g_max_irq_delay = elapsed_ticks;
21     if (elapsed_ticks < g_min_irq_delay)
22         g_min_irq_delay = elapsed_ticks;
23
24     /* clear the interrupt */
25     iowrite32(0, IOADDR_ALTERA_AVALON_TIMER_STATUS(g_timer_base));
26
27     spin_unlock(&g_irq_lock);
28     return IRQ_HANDLED;
29 }
```

Our IRQ handler gets a little more complicated in this example.

Note the use of spinlocks.

This handler saves off some performance statistics that are read out with the `irq_delays` file.

demo_module_06.ko target demonstration

```
01 insmod demo_module_06.ko                      # insert the module
02 find /sys -name "*demo*"                     # find all sysfs entries
03 ls /sys/bus/platform/drivers/demo_driver_6    # view the sysfs files
04
05 # dump the sysfs files
06 cat /sys/bus/platform/drivers/demo_driver_6/interval
07 cat /sys/bus/platform/drivers/demo_driver_6/irq_delays
08
09 # run this about 5 times, once per second to verify the 10 IRQ per second speed
10 cat /proc/interrupts | grep "demo"
11
12 # change the IRQ speed to 100 per second
13 echo "100" > /sys/bus/platform/drivers/demo_driver_6/interval
14
15 # run this about 5 times, once per second to verify the 100 IRQ per second speed
16 cat /proc/interrupts | grep "demo"
17
18 # dump the sysfs files
19 cat /sys/bus/platform/drivers/demo_driver_6/irq_delays
20 cat /sys/bus/platform/drivers/demo_driver_6/interval
21
22 # change the IRQ speed back to 10 per second
23 echo "10" > /sys/bus/platform/drivers/demo_driver_6/interval
24
25 # run this about 5 times, once per second to verify the 10 IRQ per second speed
26 cat /proc/interrupts | grep "demo"
27
28 cat /sys/bus/platform/drivers/demo_driver_6/interval
29 cat /sys/bus/platform/drivers/demo_driver_6/irq_delays
30 rmmod demo_module_06
```

Use a calculator to compute the number of CPU clocks the IRQ response takes.

Misc device

For more information:


[Documentation/ioctl/ioctl-decoding.txt](#)

[Documentation/ioctl/ioctl-number.txt](#)



demo_module_07

User Space



Kernel Space

IO Space

demo_module_07.c – misc device

```
01 static int platform_probe(struct platform_device *pdev)
02 {
03 ...cut...
04     /* register misc device dev_rom */
05     sema_init(&the_demo_rom_dev.sem, 1);
06     ret_val = misc_register(&demo_rom_dev_device);
07     if (ret_val != 0) {
08 ...cut...
09     }
10     /* register misc device dev_ram */
11     sema_init(&the_demo_ram_dev.sem, 1);
12     ret_val = misc_register(&demo_ram_dev_device);
13     if (ret_val != 0) {
14 ...cut...
15     }
16     /* register misc device dev_tmr */
17     sema_init(&the_demo_tmr_dev.sem, 1);
18     ret_val = misc_register(&demo_tmr_dev_device);
19     if (ret_val != 0) {
20 ...cut...
21     }
22     /* register misc device dev_map */
23     sema_init(&the_demo_map_dev.sem, 1);
24     ret_val = misc_register(&demo_map_dev_device);
25     if (ret_val != 0) {
26 ...cut...
27     }
```

We start our probe function in much the same way as before. We use `misc_register()` to register 4 different devices with the kernel.

demo_module_07.c – misc device

```
01 ...cut...
02 bad_exit_deregister_demo_tmr:
03     misc_deregister(&demo_tmr_dev_device);
04 bad_exit_deregister_demo_ram:
05     misc_deregister(&demo_ram_dev_device);
06 bad_exit_deregister_demo_rom:
07     misc_deregister(&demo_rom_dev_device);
08 ...cut...
09 }
10
11 static int platform_remove(struct platform_device *pdev)
12 {
13 ...cut...
14     misc_deregister(&demo_map_dev_device);
15     misc_deregister(&demo_tmr_dev_device);
16     misc_deregister(&demo_ram_dev_device);
17     misc_deregister(&demo_rom_dev_device);
18 ...cut...
19 }
```

Our probe function ends in much the same way as before. We may have to deregister our misc devices on error conditions.

Our remove function simply removes our misc devices with `misc_deregister()`.

demo_module_07.c – misc device

```
01 static const struct file_operations demo_rom_dev_fops = {  
02     .owner = THIS_MODULE,  
03     .open = demo_rom_dev_open,  
04     .release = demo_rom_dev_release,  
05     .read = demo_rom_dev_read,  
06     .llseek = demo_rom_dev_llseek,  
07 };  
08  
09 static struct miscdevice demo_rom_dev_device = {  
10     .minor = MISC_DYNAMIC_MINOR,  
11     .name = "demo_rom",  
12     .fops = &demo_rom_dev_fops,  
13 };  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30
```

Our `demo_rom` device registers an open, release, read and llseek function.

demo_module_07.c – misc device

```
01 static int demo_rom_dev_open(struct inode *ip, struct file *fp)
02 {
03     struct demo_rom_dev *dev = &the_demo_rom_dev;
04 ...cut...
05     if (down_interruptible(&dev->sem))
06         return -ERESTARTSYS;
07
08     fp->private_data = dev;
09     dev->open_count++;
10
11     up(&dev->sem);
12     pr_info("demo_rom_dev_open exit\n");
13     return 0;
14 }
15
16 static int demo_rom_dev_release(struct inode *ip, struct file *fp)
17 {
18     struct demo_rom_dev *dev = fp->private_data;
19 ...cut...
20     if (down_interruptible(&dev->sem))
21         return -ERESTARTSYS;
22
23     dev->release_count++;
24
25     up(&dev->sem);
26     pr_info("demo_rom_dev_release exit\n");
27     return 0;
28 }
```

In this example our open and release functions don't really do much, other than count some statistics for demonstration.

demo_module_07.c – misc device

```
01 static ssize_t
02 demo_rom_dev_read(struct file *fp, char __user *user_buffer,
03                     size_t count, loff_t *offset)
04 {
05     ...cut...
06     if (down_interruptible(&dev->sem)) {
07         pr_info("demo_rom_dev_read sem interrupted exit\n");
08         return -ERESTARTSYS;
09     }
10
11     dev->read_count++;
12
13     if (*offset > max_offset) {
14         up(&dev->sem);
15         pr_info("demo_rom_dev_read offset > max_offset exit\n");
16         return -EINVAL;
17     }
18
19     if (*offset == max_offset) {
20         up(&dev->sem);
21         pr_info("demo_rom_dev_read offset == max_offset exit\n");
22         return 0;
23     }
24
25     if (next_offset > max_offset)
26         count -= next_offset - max_offset;
27
28     temp_count = count;
29     rom_ptr = g_ioremap_addr + ROM_OFST;
30     rom_ptr += *offset;
```

Our `demo_rom` read routine allows user space to read from our ROM memory in our hardware.

demo_module_07.c – misc device

```
01     while (temp_count > 0) {
02         int this_loop_count = IO_BUF_SIZE;
03         if (temp_count < IO_BUF_SIZE)
04             this_loop_count = temp_count;
05
06         memcpy_fromio(&dev->io_buf, rom_ptr, this_loop_count);
07         if (copy_to_user(user_buffer, &dev->io_buf, this_loop_count)) {
08             up(&dev->sem);
09             pr_info("demo_rom_dev_read copy_to_user exit\n");
10             return -EFAULT;
11         }
12         temp_count -= this_loop_count;
13         user_buffer += this_loop_count;
14         rom_ptr += this_loop_count;
15     }
16
17     dev->read_byte_count += count;
18     *offset += count;
19
20     up(&dev->sem);
21     pr_info("demo_rom_dev_read exit\n");
22     return count;
23 }
```

In this example we demonstrate `memcpy_fromio()` along with `copy_to_user()`, but in fact on the Cortex A9 it would be safe to simply use `copy_to_user()` alone. That may not be the case on all CPU architectures however.

demo_module_07.c – misc device

```
01 static loff_t demo_rom_dev_llseek(struct file *fp, loff_t offset, int mode)
02 {
03     struct demo_rom_dev *dev = fp->private_data;
04     loff_t max_offset = ROM_SPAN;
05     loff_t next_offset;
06 ...cut...
07     if (down_interruptible(&dev->sem)) {
08         pr_info("demo_rom_dev_llseek sem interrupted exit\n");
09         return -ERESTARTSYS;
10     }
11
12     dev->llseek_count++;
13
14     switch (mode) {
15     case SEEK_SET:
16         next_offset = offset;
17         break;
18     case SEEK_CUR:
19         next_offset = fp->f_pos + offset;
20         break;
21     case SEEK_END:
22         next_offset = max_offset;
23         break;
24     default:
25         up(&dev->sem);
26         pr_info("demo_rom_dev_llseek bad mode exit\n");
27         return -EINVAL;
28     }
29
30 }
```

Our demo_rom llseek routine allows user space to seek within our ROM space.

demo_module_07.c – misc device

```
01     if (next_offset < 0) {
02         up(&dev->sem);
03         pr_info("demo_rom_dev_llseek negative offset exit\n");
04         return -EINVAL;
05     }
06
07     if (next_offset > max_offset)
08         next_offset = max_offset;
09
10     fp->f_pos = next_offset;
11
12     up(&dev->sem);
13     pr_info("demo_rom_dev_llseek exit\n");
14     return next_offset;
15 }
```

demo_module_07.c – misc device

```
01 static const struct file_operations demo_ram_dev_fops = {  
02     .owner = THIS_MODULE,  
03     .open = demo_ram_dev_open,  
04     .release = demo_ram_dev_release,  
05     .read = demo_ram_dev_read,  
06     .write = demo_ram_dev_write,  
07     .llseek = demo_ram_dev_llseek,  
08 };  
09  
10 static struct miscdevice demo_ram_dev_device = {  
11     .minor = MISC_DYNAMIC_MINOR,  
12     .name = "demo_ram",  
13     .fops = &demo_ram_dev_fops,  
14 };  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30
```

Our `demo_ram` device registers a write function in addition to the functions that our `demo_rom` device registered. This allows user space to write and read our RAM hardware.

demo_module_07.c – misc device

```
01 static ssize_t
02 demo_ram_dev_write(struct file *fp,
03                      const char __user *user_buffer, size_t count,
04                      loff_t *offset)
05 {
06     struct demo_ram_dev *dev = fp->private_data;
07     loff_t max_offset = RAM_SPAN;
08     loff_t next_offset = *offset + count;
09     size_t temp_count;
10     void *ram_ptr;
11 ...cut...
12     if (down_interruptible(&dev->sem)) {
13         pr_info("demo_ram_dev_write sem interrupted exit\n");
14         return -ERESTARTSYS;
15     }
16
17     dev->write_count++;
18
19     if (*offset > max_offset) {
20         up(&dev->sem);
21         pr_info("demo_ram_dev_write offset > max_offset exit\n");
22         return -EINVAL;
23     }
24
25     if (*offset == max_offset) {
26         up(&dev->sem);
27         pr_info("demo_ram_dev_write offset == max_offset exit\n");
28         return -ENOSPC;
29     }
30 }
```

Our `demo_ram` write function is similar to the read function, just implementing the opposite data flow.

demo_module_07.c – misc device

```
01     if (next_offset > max_offset)
02         count -= next_offset - max_offset;
03
04     temp_count = count;
05     ram_ptr = g_ioremap_addr + RAM_OFST;
06     ram_ptr += *offset;
07
08     while (temp_count > 0) {
09         int this_loop_count = IO_BUF_SIZE;
10         if (temp_count < IO_BUF_SIZE)
11             this_loop_count = temp_count;
12
13         if (copy_from_user
14             (&dev->io_buf, user_buffer, this_loop_count)) {
15             up(&dev->sem);
16             pr_info("demo_ram_dev_write copy_from_user exit\n");
17             return -EFAULT;
18         }
19         memcpy_toio(ram_ptr, &dev->io_buf, this_loop_count);
20         temp_count -= this_loop_count;
21         user_buffer += this_loop_count;
22         ram_ptr += this_loop_count;
23     }
24
25     dev->write_byte_count += count;
26     *offset += count;
27
28     up(&dev->sem);
29     pr_info("demo_ram_dev_write exit\n");
30     return count;
```

In our write we
copy_from_user() and
memcpy_toio().

demo_module_07.c – misc device

```
01 static const struct file_operations demo_tmr_dev_fops = {  
02     .owner = THIS_MODULE,  
03     .open = demo_tmr_dev_open,  
04     .release = demo_tmr_dev_release,  
05     .unlocked_ioctl = demo_tmr_dev_ioctl,  
06 };  
07  
08 static struct miscdevice demo_tmr_dev_device = {  
09     .minor = MISC_DYNAMIC_MINOR,  
10     .name = "demo_tmr",  
11     .fops = &demo_tmr_dev_fops,  
12 };  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30
```

Our `demo_tmr` device registers an open, release and ioctl function. This allows user space to interact with our timer hardware.

demo_module_07.c – misc device

```
01 static long
02 demo_tmr_dev_ioctl(struct file *fp, unsigned int cmd, unsigned long arg)
03 {
04     ...cut...
05     if (down_interruptible(&dev->sem)) {
06         pr_info("demo_tmr_dev_ioctl sem interrupted exit\n");
07         return -ERESTARTSYS;
08     }
09     dev->ioctl_count++;
10
11     switch (cmd) {
12     case IOC_SET_INTERVAL:
13     ...cut...
14     case IOC_GET_INTERVAL:
15     ...cut...
16     case IOC_GET_MAX_DELAY:
17     ...cut...
18     case IOC_GET_MIN_DELAY:
19     ...cut...
20     default:
21         up(&dev->sem);
22         pr_info("demo_tmr_dev_ioctl bad cmd exit\n");
23         return -EINVAL;
24     }
25
26     up(&dev->sem);
27     pr_info("demo_tmr_dev_ioctl exit\n");
28     return 0;
29 }
```

Our ioctl boils down to a switch statement.

demo_module_07.c – misc device

```
01 case IOC_SET_INTERVAL:
02     if (get_user(new_interval, (uint32_t *)arg) < 0) {
03         up(&dev->sem);
04         pr_info("demo_tmr_dev_ioctl get_user exit\n");
05         return -EFAULT;
06     }
07     /* range check the requested new interval value */
08     if (new_interval > 100) {
09         up(&dev->sem);
10         pr_info("demo_tmr_dev_ioctl new_interval > 100 exit\n");
11         return -EINVAL;
12     }
13     /* calculate the new period value */
14     if (new_interval > 0)
15         new_period =
16             (g_demo_driver_clk_rate / new_interval) - 1;
17     else
18         new_period = g_demo_driver_clk_rate;
19
20     /* acquire the irq_lock */
21     spin_lock_irqsave(&g_irq_lock, flags);
22
23     /* stop the interval timer */
24     iowrite32(ALTERA_AVALON_TIMER_CONTROL_STOP_MSK,
25               IOADDR_ALTERA_AVALON_TIMER_CONTROL(g_timer_base));
```

Details of switch case SET_INTERVAL. Uses get_user() instead of copy_from_user().

demo_module_07.c – misc device

```
01  /* ensure there is no pending IRQ that we are blocking */
02  the_status =
03      ioread32(IOADDR_ALTERA_AVALON_TIMER_STATUS(g_timer_base));
04  the_status &= ALTERA_AVALON_TIMER_STATUS_TO_MSK;
05  if (the_status != 0) {
06      do {
07          /*
08              if we are blocking, release the lock to allow
09              IRQ handler to execute, acquire the lock and
10              check again
11          */
12          spin_unlock_irqrestore(&g_irq_lock, flags);
13          spin_lock_irqsave(&g_irq_lock, flags);
14          the_status =
15              ioread32(IOADDR_ALTERA_AVALON_TIMER_STATUS
16                      (g_timer_base));
17          the_status &= ALTERA_AVALON_TIMER_STATUS_TO_MSK;
18      } while (the_status != 0);
19  }
20  /* write the new period value */
21  iowrite32(new_period,
22            IOADDR_ALTERA_AVALON_TIMER_PERIODL(g_timer_base));
23  iowrite32(new_period >> 16,
24            IOADDR_ALTERA_AVALON_TIMER_PERIODH(g_timer_base));

25  /* initialize the MAX/MIN variables */
26  g_max_irq_delay = 0;
27  g_min_irq_delay = 0xFFFFFFFF;
```

manage potential race for spin lock
with irq handler

demo_module_07.c – misc device

```
01      /* start the timer */
02      if (new_interval > 0)
03          iowrite32(ALTERA_AVALON_TIMER_CONTROL_ITO_MSK |
04                      ALTERA_AVALON_TIMER_CONTROL_CONT_MSK |
05                      ALTERA_AVALON_TIMER_CONTROL_START_MSK,
06                      IOADDR_ALTERA_AVALON_TIMER_CONTROL
07                      (g_timer_base));
08
09      /* release the irq_lock */
10      spin_unlock_irqrestore(&g_irq_lock, flags);
11      break;
12
13  case IOC_GET_MAX_DELAY:
14      /* acquire the irq_lock */
15      spin_lock_irqsave(&g_irq_lock, flags);
16
17      /* capture the shared data values */
18      max_irq_delay = g_max_irq_delay;
19
20      /* release the irq_lock */
21      spin_unlock_irqrestore(&g_irq_lock, flags);
22
23      if (put_user(max_irq_delay, (uint32_t *)arg) < 0) {
24          up(&dev->sem);
25          pr_info("demo_tmr_dev_ioctl put_user exit\n");
26          return -EFAULT;
27      }
28      break;
```

Details of switch case GET_MAX_DELAY. Uses put_user() instead of copy_to_user().

demo_module_07.c – misc device

```
01 static const struct file_operations demo_map_dev_fops = {  
02     .owner = THIS_MODULE,  
03     .open = demo_map_dev_open,  
04     .release = demo_map_dev_release,  
05     .read = demo_map_dev_read,  
06     .mmap = demo_map_dev_mmap,  
07 };  
08  
09 static struct miscdevice demo_map_dev_device = {  
10     .minor = MISC_DYNAMIC_MINOR,  
11     .name = "demo_map",  
12     .fops = &demo_map_dev_fops,  
13 };  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30
```

Our `demo_map` device registers functions for `open`, `release`, `read` and `mmap`. This allows user space to `mmap()` our device, similar to what we did in the `/dev/mem` example. The `read` function allows us to emulate IRQ polling in much the same way as a UIO device does, we'll demonstrate a UIO device later.

demo_module_07.c – misc device

```
01 static ssize_t
02 demo_map_dev_read(struct file *fp, char __user *user_buffer,
03                     size_t count, loff_t *offset)
04 {
05     struct demo_map_dev *dev = fp->private_data;
06     uint32_t cur_irq_cnt;
07 ...cut...
08     if (down_interruptible(&dev->sem)) {
09         pr_info("demo_map_dev_read sem interrupted exit\n");
10         return -ERESTARTSYS;
11     }
12
13     dev->read_count++;
14
15     if (*offset != 0) {
16         up(&dev->sem);
17         pr_info("demo_map_dev_read offset != 0 exit\n");
18         return -EINVAL;
19     }
20
21     if (count != 4) {
22         up(&dev->sem);
23         pr_info("demo_map_dev_read count != 4 exit\n");
24         return -EINVAL;
25     }
26
27 The demo_map read function starts by validating the input arguments for the exact count
28 requirements that we expect to fulfill, anything else is considered an error.
29
30
```

demo_module_07.c – misc device

```
01     cur_irq_cnt = get_current_irq_count();
02     while (cur_irq_cnt == get_current_irq_count()) {
03         up(&dev->sem);
04         if (wait_event_interruptible(g_irq_wait_queue,
05                                     (cur_irq_cnt !=
06                                         get_current_irq_count())))
07             pr_info("demo_map_dev_read wait interrupted exit\n");
08             return -ERESTARTSYS;
09     }
10     if (down_interruptible(&dev->sem)) {
11         pr_info("demo_map_dev_read sem interrupted exit\n");
12         return -ERESTARTSYS;
13     }
14 }
15
16 cur_irq_cnt = get_current_irq_count();
17 if (copy_to_user(user_buffer, &cur_irq_cnt, count)) {
18     up(&dev->sem);
19     pr_info("demo_map_dev_read copy_to_user exit\n");
20     return -EFAULT;
21 }
22
23 dev->read_byte_count += count;
24
25 up(&dev->sem);
26 pr_info("demo_map_dev_read exit\n");
27 return count;
28 }
```

The heart of the demo_map read function is waiting for the next IRQ event to occur.

demo_module_07.c – misc device

```
01 static inline uint32_t get_current_irq_count(void)
02 {
03     uint32_t current_count;
04     unsigned long flags;
05
06     spin_lock_irqsave(&g_irq_lock, flags);
07     current_count = g_irq_count;
08     spin_unlock_irqrestore(&g_irq_lock, flags);
09     return current_count;
10 }
11
12 irqreturn_t demo_driver_interrupt_handler(int irq, void *dev_id)
13 {
14     ...cut...
15     spin_lock(&g_irq_lock);
16     ...cut...
17     spin_unlock(&g_irq_lock);
18     wake_up_interruptible(&g_irq_wait_queue);
19     return IRQ_HANDLED;
20 }
21
22
23
24
25
26
27
28
29
30
```

The read function leveraged this helper function to query the IRQ count in a safe fashion to avoid race conditions.

The IRQ handler for this example is similar to before with the addition of the wake up event signal, which is what the read function will ultimately pend on.

demo_module_07.c – misc device

```
01 static int demo_map_dev_mmap(struct file *fp, struct vm_area_struct *vma)
02 {
03     struct demo_map_dev *dev = fp->private_data;
04 ...cut...
05     if (down_interruptible(&dev->sem))
06         return -ERESTARTSYS;
07
08     dev->mmap_count++;
09
10    if (vma->vm_end - vma->vm_start != PAGE_SIZE) {
11        up(&dev->sem);
12        return -EINVAL;
13    }
14
15    vma->vm_page_prot = pgprot_noncached(vma->vm_page_prot);
16
17    vma->vm_pgoff = g_demo_driver_base_addr >> PAGE_SHIFT;
18
19    if (remap_pfn_range(vma, vma->vm_start, vma->vm_pgoff, PAGE_SIZE,
20                        vma->vm_page_prot)) {
21        up(&dev->sem);
22        return -EAGAIN;
23    }
24
25    up(&dev->sem);
26    pr_info("demo_map_dev_mmap exit\n");
27    return 0;
28 }
```

The `demo_map mmap` function maps our IO region into a user space safe pointer as noncached memory.

demo_map_test.c

```
01 int main(int argc, char **argv) {  
02     ...cut...  
03     //  
04     // parse the command line arguments  
05     //  
06     parse_cmdline(argc, argv);  
07     //  
08     // open() the /dev/demo_map device  
09     //  
10     dev_demo_map_fd = open("/dev/demo_map", O_RDWR | O_SYNC);  
11     if(dev_demo_map_fd < 0) {  
12         perror("dev_demo_map open");  
13         exit(EXIT_FAILURE);  
14     }  
15     //  
16     // mmap() the base of our demo_driver hardware  
17     //  
18     demo_driver_map = mmap(NULL, sysconf(_SC_PAGE_SIZE), PROT_READ|PROT_WRITE,  
19                             MAP_SHARED, dev_demo_map_fd, 0);  
20     if(demo_driver_map == MAP_FAILED) {  
21         perror("dev_demo_map mmap");  
22         close(dev_demo_map_fd);  
23         exit(EXIT_FAILURE);  
24     }  
25     //  
26     //  
27     //  
28     //  
29     //  
30 }
```

We provide a test application to drive the mmap functionality of our demo_map driver. This looks basically the same as our /dev/mem example, but this example uses /dev/demo_map instead.

ioctl_test.c

```
01 int main(int argc, char **argv) {  
02     int dev_demo_tmr_fd;  
03     //  
04     // parse the command line arguments  
05     //  
06     parse_cmdline(argc, argv);  
07     //  
08     // open() the /dev/demo_tmr device  
09     //  
10    dev_demo_tmr_fd = open("/dev/demo_tmr", O_RDWR | O_SYNC);  
11    if(dev_demo_tmr_fd < 0) {  
12        perror("dev_demo_tmr open");  
13        exit(EXIT_FAILURE);  
14    }  
15    //  
16    // perform the operation selected by the command line arguments  
17    //  
18    if(g_get_interval != NULL) do_get_interval(dev_demo_tmr_fd);  
19    if(g_set_interval != NULL) do_set_interval(dev_demo_tmr_fd);  
20    if(g_get_max_delay != NULL) do_get_max_delay(dev_demo_tmr_fd);  
21    if(g_get_min_delay != NULL) do_get_min_delay(dev_demo_tmr_fd);  
22    if(g_help != NULL) do_help();  
23    //  
24    close(dev_demo_tmr_fd);  
25    exit(EXIT_SUCCESS);  
26 }  
27  
28  
29  
30 }
```

We provide a test application to drive the ioctl functionality of our demo_tmr driver.

ioctl_test.c

```
01 void do_get_interval(int dev_demo_tmr_fd) {  
02  
03     int result;  
04     unsigned long interval;  
05  
06     result = ioctl(dev_demo_tmr_fd, IOC_GET_INTERVAL, &interval);  
07     if(result != 0) {  
08         error(1, errno, "%s:%d ioctl failed", __FILE__, __LINE__);  
09     }  
10  
11     printf("Current IRQ interval is %lu interrupts per second.\n", interval);  
12 }  
13  
14 void do_set_interval(int dev_demo_tmr_fd) {  
15  
16     int result;  
17  
18     result = ioctl(dev_demo_tmr_fd, IOC_SET_INTERVAL, &g_new_interval);  
19     if(result != 0) {  
20         error(1, errno, "%s:%d ioctl failed", __FILE__, __LINE__);  
21     }  
22  
23     printf("IRQ interval set to %lu interrupts per second.\n", g_new_interval);  
24 }
```

Here are the get and set interval functions.

ioctl_test.c

```
01 void do_get_max_delay(int dev_demo_tmr_fd) {  
02  
03     int result;  
04     unsigned long delay;  
05  
06     result = ioctl(dev_demo_tmr_fd, IOC_GET_MAX_DELAY, &delay);  
07     if(result != 0) {  
08         error(1, errno, "%s:%d ioctl failed", __FILE__, __LINE__);  
09     }  
10  
11     printf("Maximum IRQ service delay is %lu ticks.\n", delay);  
12 }  
13  
14 void do_get_min_delay(int dev_demo_tmr_fd) {  
15  
16     int result;  
17     unsigned long delay;  
18  
19     result = ioctl(dev_demo_tmr_fd, IOC_GET_MIN_DELAY, &delay);  
20     if(result != 0) {  
21         error(1, errno, "%s:%d ioctl failed", __FILE__, __LINE__);  
22     }  
23  
24     printf("Minimum IRQ service delay is %lu ticks.\n", delay);  
25 }
```

Here are the get max and min delay functions.

demo_module_07.ko target demonstration

```
01 insmod demo_module_07.ko
02 ls /dev | grep "demo"
03 dd if=/dev/demo_rom | hexdump -C
04 dd if=/dev/zero of=/dev/demo_rom
05 dd if=/dev/demo_ram | hexdump -C
06 dd if=/dev/zero of=/dev/demo_ram
07 dd if=/dev/demo_ram | hexdump -C
08 dd if=/dev/demo_ram bs=64 count=1 | hexdump -Cv
09 dd if=/dev/demo_rom bs=64 count=1 | hexdump -C
10
11 # demonstrate lseek by copying 4 bytes from offset 56 in rom to offset 16 in ram
12 dd if=/dev/demo_rom of=/dev/demo_ram bs=1 count=4 skip=56 seek=16
13 dd if=/dev/demo_ram bs=64 count=1 | hexdump -Cv
14 ./demo_map_test -h
15 ./demo_map_test -o | hexdump -C
16 ./demo_map_test -a | hexdump -C
17 dd if=/dev/urandom | ./demo_map_test -f
18 ./demo_map_test -a | hexdump -C
19 ./demo_map_test -t
20 for I in 1 2 3 4 ; do
21 dd if=/dev/demo_map bs=4 count=1 | hexdump -C
22 done
23 cat /sys/bus/platform/drivers/demo_driver_7/interval
24 ./ioctl_test -g
25 echo "100" > /sys/bus/platform/drivers/demo_driver_7/interval
26 ./ioctl_test -g
27 cat /sys/bus/platform/drivers/demo_driver_7/irq_delays
28 ./ioctl_test -x
29 ./ioctl_test -n
30 rmmod demo_module_07
```

insert module
observe devfs entries
read from rom
write to rom
read from ram
write to ram
verify write
dump first 64 ram bytes
dump first 64 rom bytes

verify the operation
run demo_map_test
dump rom
dump ram
fill ram
verify ram
dump the timer registers
demonstrate demo_map read
the result is the irq count
returned at the next IRQ event
demonstrate ioctl

Observe the output on the printk console.

UIO device

User Space I/O



UIO templates in the kernel source tree


```
01 [socfpga-3.10-ltsi]$ find drivers/uio -name "*.c"
02 drivers/uio/uio.c
03 drivers/uio/uio_netx.c
04 drivers/uio/uio_sercos3.c
05 drivers/uio/uio_cif.c
06 drivers/uio/uio_pdrv.c
07 drivers/uio/uio_dmem_genirq.c
08 drivers/uio/uio_pruss.c
09 drivers/uio/uio_pdrv_genirq.c
10 drivers/uio/uio_aec.c
11 drivers/uio/uio_pci_generic.c
12
13
14
```

It's worth noting that there are some useful templates in the kernel source tree for simple platform driver implementations driven by a device tree entry, these can be easily modified to provide basic functionality by simply adding a compatibility string to them. The fundamental concept of a UIO driver is that it simply exposes the raw hardware registers to user space and does nothing within kernel space to really control any of the hardware details. Our UIO example does not use the UIO framework in that basic way, instead it builds on the framework that we started with in the previous demonstration modules such that our kernel space code does interact more directly with the hardware in certain situations. We just plug that functionality into the framework that is defined by UIO.

An example of the basic UIO template usage is demonstrated in the lab work that accompanies the WS3 session.

demo_module_08

User Space



Kernel Space

IO Space

demo_module_08.c – uio device

```
01 static int platform_probe(struct platform_device *pdev)
02 {
03     ...cut...
04         /* acquire the probe lock */
05     ...cut...
06         /* get our first memory resource */
07     ...cut...
08         /* get our interrupt resource */
09     ...cut...
10         /* get our clock resource */
11     ...cut...
12         /* reserve our memory region */
13     ...cut...
14         /* ioremap our memory region */
15     ...cut...
16         /* initialize uio_info struct uio_mem array */
17         the_uio_info.mem[0].memtype = UIO_MEM_PHYS;
18         the_uio_info.mem[0].addr = r->start;
19         the_uio_info.mem[0].size = resource_size(r);
20         the_uio_info.mem[0].name = "demo_uio_driver_hw_region";
21         the_uio_info.mem[0].internal_addr = g_ioremap_addr;
22
23         for (i = 1; i < MAX_UIO_MAPS; i++)
24             the_uio_info.mem[i].size = 0;
25
26         /* initialize uio_info irq */
27         the_uio_info.irq = g_demo_driver_irq;
28
29
30
```

We start our probe function in much the same way as before. We reserve our memory region and then configure the `uio_info` struct.

A basic UIO driver would not be concerned about reserving the memory region or ioremap'ing the IO space since the kernel space UIO driver would not be interacting with the hardware directly.

demo_module_08.c – uio device

```
01     /* register the uio device */
02     sema_init(&g_demo_uio_dev_sem, 1);
03     ret_val = uio_register_device(&pdev->dev, &the_uio_info);
04     if (ret_val != 0) {
05         pr_warn("Could not register device \\\"demo_uio\\\"...");
06         goto bad_exit_iounmap;
07     }
08
09     g_platform_probe_flag = 1;
10     up(&g_dev_probe_sem);
11     pr_info("platform_probe exit\n");
12     return 0;
13
14 bad_exit_iounmap:
15     iounmap(g_ioremap_addr);
16 bad_exit_release_mem_region:
17     release_mem_region(g_demo_driver_base_addr, g_demo_driver_size);
18 bad_exit_return:
19     up(&g_dev_probe_sem);
20     pr_info("platform_probe bad_exit\n");
21     return ret_val;
22 }
```

Then we register our uio device.

demo_module_08.c – uio device

```
01 static int platform_remove(struct platform_device *pdev)
02 {
03     pr_info("platform_remove enter\n");
04
05     uio_unregister_device(&the_uio_info);
06
07     iounmap(g_ioremap_addr);
08     release_mem_region(g_demo_driver_base_addr, g_demo_driver_size);
09
10    if (down_interruptible(&g_dev_probe_sem))
11        return -ERESTARTSYS;
12
13    g_platform_probe_flag = 0;
14    up(&g_dev_probe_sem);
15
16    pr_info("platform_remove exit\n");
17    return 0;
18 }
```

Our remove function
unregisters the
uio_device.

demo_module_08.c – uio device

```
01 static struct uio_info the_uio_info = {  
02     .name = "demo_uio",  
03     .version = "1.0",  
04     .irq_flags = 0,  
05     .handler = demo_uio_interrupt_handler,  
06     .open = demo_uio_open,  
07     .release = demo_uio_release,  
08     .irqcontrol = demo_uio_irqcontrol,  
09 };
```

Our uio_info struct.

demo_module_08.c – uio device

```
01 static int demo_uio_open(struct uio_info *info, struct inode *inode)
02 {
03     ...cut...
04     if (down_trylock(&g_demo_uio_dev_sem) != 0)
05         return -EAGAIN;
06
07     /* initialize our peripheral timer hardware */
08     io_result = ioread32(IOADDR_ALTERA_AVALON_TIMER_STATUS(g_timer_base));
09     io_result &= ALTERA_AVALON_TIMER_STATUS_TO_MSK |
10             ALTERA_AVALON_TIMER_STATUS_RUN_MSK;
11     if (io_result != 0) {
12         pr_err("peripheral timer hardware, incorrect initial state");
13         return -EIO;
14     }
15
16     period_1s = (g_demo_driver_clk_rate) - 1;
17     iowrite32(period_1s, IOADDR_ALTERA_AVALON_TIMER_PERIODL(g_timer_base));
18     iowrite32(period_1s >> 16,
19               IOADDR_ALTERA_AVALON_TIMER_PERIODH(g_timer_base));
20
21     /* start our timer and enable our timer hardware interrupts */
22     iowrite32(ALTERA_AVALON_TIMER_CONTROL_ITO_MSK |
23             ALTERA_AVALON_TIMER_CONTROL_CONT_MSK |
24             ALTERA_AVALON_TIMER_CONTROL_START_MSK,
25             IOADDR_ALTERA_AVALON_TIMER_CONTROL(g_timer_base));
26
27     return 0;
28 }
```

Our open routine.

A basic UIO driver would rely on user space code to perform this once the device was mmap()'ed.

demo_module_08.c – uio device

```
01 static int demo_uio_release(struct uio_info *info, struct inode *inode)
02 {
03     uint32_t io_result;
04
05     /* stop our timer and disable our timer hardware interrupts */
06     iowrite32(ALTERA_AVALON_TIMER_CONTROL_STOP_MSK,
07               IOADDR_ALTERA_AVALON_TIMER_CONTROL(g_timer_base));
08     /* ensure there is no pending IRQ */
09     do {
10         io_result =
11             ioread32(IOADDR_ALTERA_AVALON_TIMER_STATUS(g_timer_base));
12         io_result &= ALTERA_AVALON_TIMER_STATUS_TO_MSK;
13     } while (io_result != 0);
14
15     up(&g_demo_uio_dev_sem);
16     return 0;
17 }
```

Our release routine.

A basic UIO driver would rely on user space code to perform this prior the device being munmap()'ed.

demo_module_08.c – uio device

```
01 irqreturn_t demo_uio_interrupt_handler(int irq, struct uio_info *dev_info)
02 {
03     /* snapshot the current timer value */
04     iowrite32(0, IOADDR_ALTERA_AVALON_TIMER_SNAPL(g_timer_base));
05
06     /* clear the interrupt */
07     iowrite32(0, IOADDR_ALTERA_AVALON_TIMER_STATUS(g_timer_base));
08
09     return IRQ_HANDLED;
10 }
```

Our irq handler.

A basic UIO driver would rely on user space code to perform this following a read() from the driver.

demo_module_08.c – uio device

```
01 static int demo_uio_irqcontrol(struct uio_info *info, s32 irq_on)
02 {
03     uint32_t io_result;
04
05     if (irq_on) {
06         /* start our timer and enable our timer hardware interrupts */
07         iowrite32(ALTERA_AVALON_TIMER_CONTROL_ITO_MSK |
08                    ALTERA_AVALON_TIMER_CONTROL_CONT_MSK |
09                    ALTERA_AVALON_TIMER_CONTROL_START_MSK,
10                    IOADDR_ALTERA_AVALON_TIMER_CONTROL(g_timer_base));
11    } else {
12        /* stop our timer and disable our timer hardware interrupts */
13        iowrite32(ALTERA_AVALON_TIMER_CONTROL_STOP_MSK,
14                  IOADDR_ALTERA_AVALON_TIMER_CONTROL(g_timer_base));
15        /* ensure there is no pending IRQ */
16        do {
17            io_result =
18                ioread32(IOADDR_ALTERA_AVALON_TIMER_STATUS
19                          (g_timer_base));
20            io_result &= ALTERA_AVALON_TIMER_STATUS_TO_MSK;
21        } while (io_result != 0);
22    }
23    return 0;
24 }
```

Our irqcontrol routine.

A basic UIO driver would rely on user space code to perform this coordinated with a write to the driver.

demo_module_08.ko target demonstration (1)

```
01 insmod demo_module_08.ko          # insert module
02 find /dev -name "uio*"           # find the device node
03 find /sys/ -name "uio*"          # find sysfs entries
04 ls /sys/class/uio/uio0           # display uio contents
05 ls /sys/class/uio/uio0/maps/    # display maps entry
06 ls /sys/class/uio/uio0/maps/map0/ # display map0 entry
07 cat /sys/class/uio/uio0/maps/map0/* # dump the map0 entries
08
09 # run this block as one copy/paste operation
10 exec 7</>/dev/uio0              # open /dev/uio0 for rd/wr
11 for I in 1 2 3 4 5
12 do
13 dd bs=4 count=1 0<&7 | hexdump -Cv          # read an IRQ event
14 done
15 dd if=/dev/zero bs=4 count=1 >&7            # disable IRQs
16 dd bs=4 count=1 0<&7 | hexdump -Cv          # read an IRQ event
17
18 # the last command of the previous block should stall forever, ^C and run this
19 # block next
20 dd if=/dev/zero bs=4 count=1 | tr '\000' '\377' >&7      # enable IRQs
21 for I in 1 2 3 4 5
22 do
23 dd bs=4 count=1 0<&7 | hexdump -Cv          # read an IRQ event
24 done
25
26 # finally run this command
27 exec 7</>/dev/null                # close /dev/uio0
28
29
30
```

demo_module_08.ko target demonstration (2)

```
01 ./demo_uio_test -h                                # run demo_uio_test
02 ./demo_uio_test -t                                # exercise timer regs
03 ./demo_uio_test -o | hexdump -C                  # dump rom
04 ./demo_uio_test -a | hexdump -C                  # dump ram
05 dd if=/dev/urandom | ./demo_uio_test -f          # fill ram
06 ./demo_uio_test -a | hexdump -C                  # dump ram to verify
07 ls /sys/class/uio/uio0/                           # observe the event entry
08 cat /sys/class/uio/uio0/event                     # dump the event entry
09 dd if=/dev/uio0 bs=4 count=1 | hexdump -Cv      # read another IRQ
10 printf "%d\n" <hex number>                      # convert the hex number to dec
11 cat /sys/class/uio/uio0/event                     # verify event count matches
12 rmmod demo_module_08                             # remove module
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
```


DMA

The Altera logo consists of the word "ALTERA" in a bold, blue, sans-serif font. A registered trademark symbol (®) is located at the top right corner of the letter "A".

ALTERA®

demo_module_09

User Space



Kernel Space

IO Space

demo_module_09.c – DMA

```
01 static struct of_device_id demo_driver_dt_ids[] = {  
02     {  
03         .compatible = "demo,memcpy_msgdma"},  
04     { /* end of table */ }  
05 };  
06  
07  
08  
09  
10  
11  
12  
13  
14 static int platform_probe(struct platform_device *pdev)  
15 {  
16     ...cut...  
17     /* acquire the probe lock */  
18     ...cut...  
19     /* get our csr memory resource */  
20     ...cut...  
21     /* reserve our csr memory region */  
22     ...cut...  
23     /* ioremap our csr memory region */  
24     ...cut...  
25     /* initialize the DMA controller */  
26     ...cut...  
27     /* get our desc memory resource */  
28     ...cut...  
29     /* ioremap our desc memory region */  
30     ...cut...
```

For this example we bind with the memcpy_msgdma device.

Our probe function starts off pretty much the same.

demo_module_09.c – DMA

```
01  /* allocate some memory buffers */
02  g_kmalloc_ptr_4k = kmalloc(PAGE_SIZE, GFP_KERNEL);
03  if (g_kmalloc_ptr_4k == NULL) {
04      pr_err("kmalloc failed\n");
05      goto bad_exit_iounmap_desc;
06  }
07
08  the_demo_dma_xx(pdev).pdev_dev = &pdev->dev;
09  g_coherent_ptr_4k =
10     dma_alloc_coherent(the_demo_dma_xx(pdev).pdev_dev, PAGE_SIZE,
11                         &g_dma_handle_4k, GFP_KERNEL);
12  if (g_coherent_ptr_4k == NULL) {
13      pr_err("dma_alloc_coherent failed 4KB\n");
14      goto bad_exit_kfree;
15  }
16
17  g_coherent_ptr_1m = dma_alloc_coherent(the_demo_dma_xx(pdev).pdev_dev,
18                                         DMA_DEVICE_BUFFER_SIZE,
19                                         &g_dma_handle_1m, GFP_KERNEL);
20  if (g_coherent_ptr_1m == NULL) {
21      pr_err("dma_alloc_coherent failed 1MB\n");
22      goto bad_exit_dma_free_coherent_4k;
23  }
24
25
26
27
28
29
30
```

We kmalloc a buffer for our streaming driver and we allocate a coherent buffer for our coherent driver.

coherent = non-cacheable.

Then we allocate a coherent buffer for our 1M file emulation buffer.

demo_module_09.c – DMA

```
01      /* register our interrupt handler */
02      init_waitqueue_head(&g_irq_wait_queue);
03 ...cut...
04      /* enable the DMA global IRQ mask */
05 ...cut...
06      /* register misc device demo_dma_co */
07      sema_init(&the_demo_dma_xx.sem, 1);
08      ret_val = misc_register(&demo_dma_co_device);
09      if (ret_val != 0) {
10          pr_warn("Could not register device \"demo_dma_co\"...");
11          goto bad_exit_freeirq;
12      }
13      /* register misc device demo_dma_st */
14      ret_val = misc_register(&demo_dma_st_device);
15      if (ret_val != 0) {
16          pr_warn("Could not register device \"demo_dma_st\"...");
17          goto bad_exit_deregister_demo_dma_co;
18      }
19 ...cut...
20 }
21
22
23
24
25
26
27
28
29
30
```

Then we register a couple misc devices, a coherent and streaming driver.

demo_module_09.c – DMA

```
01 irqreturn_t demo_driver_interrupt_handler(int irq, void *dev_id)
02 {
03     spin_lock(&g_irq_lock);                                This is our IRQ handler.
04
05     /* clear the IRQ state */
06     iowrite32(ALTERA_MSGDMA_CSR_IRQ_SET_MASK,
07                g_ioremap_csr_addr + CSR_STATUS_REG);
08
09     spin_unlock(&g_irq_lock);
10    wake_up_interruptible(&g_irq_wait_queue);
11    return IRQ_HANDLED;
12 }
```

13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30

demo_module_09.c – DMA

```
01 static const struct file_operations demo_dma_st_fops = {  
02     .owner = THIS_MODULE,  
03     .read = demo_dma_st_read,  
04     .write = demo_dma_st_write,  
05     .llseek = demo_dma_xx_llseek,  
06 };  
07  
08 static struct miscdevice demo_dma_st_device = {  
09     .minor = MISC_DYNAMIC_MINOR,  
10     .name = "demo_dma_st",  
11     .fops = &demo_dma_st_fops,  
12 };  
13  
14 static const struct file_operations demo_dma_co_fops = {  
15     .owner = THIS_MODULE,  
16     .read = demo_dma_co_read,  
17     .write = demo_dma_co_write,  
18     .llseek = demo_dma_xx_llseek,  
19 };  
20  
21 static struct miscdevice demo_dma_co_device = {  
22     .minor = MISC_DYNAMIC_MINOR,  
23     .name = "demo_dma_co",  
24     .fops = &demo_dma_co_fops,  
25 };  
26  
27  
28  
29  
30
```

Our streaming and coherent drivers register the same functions, the llseek is actually the same function for both implementations.

demo_module_09.c – DMA

```
01 static ssize_t demo_dma_co_read(struct file *fp, char __user *user_buffer,
02                                     size_t count, loff_t *offset)
03 {
04     ...cut...
05     ram_ptr = g_dma_handle_1m;
06     ram_ptr += *offset;
07     ...cut...
08     iowrite32(ram_ptr, g_ioremap_desc_addr + DESC_READ_ADDRESS_REG);
09     iowrite32(g_dma_handle_4k + next_dma_io_buf_ofst,
10             g_ioremap_desc_addr + DESC_WRITE_ADDRESS_REG);
11     iowrite32(this_loop_count, g_ioremap_desc_addr + DESC_LENGTH_REG);
12     iowrite32(START_DMA_MASK, g_ioremap_desc_addr + DESC_CONTROL_REG);
13
14     temp_dma_count -= this_loop_count;
15     ram_ptr += this_loop_count;
16     next_dma_io_buf_ofst += this_loop_count;
17     if (next_dma_io_buf_ofst >= PAGE_SIZE)
18         next_dma_io_buf_ofst = 0;
19
20
21
22
23
24
25
26
27
28
29
30
```

descriptor configuration

Our coherent read starts to directly interact with the `g_dma_handle_4k`.

demo_module_09.c – DMA

```
01     while (temp_user_count > 0) {
02         if (temp_dma_count > 0) {
03             iowrite32(ram_ptr,
04                         g_ioremap_desc_addr + DESC_READ_ADDRESS_REG);
05             iowrite32(g_dma_handle_4k + next_dma_io_buf_ofst,
06                         g_ioremap_desc_addr + DESC_WRITE_ADDRESS_REG);
07             iowrite32(this_loop_count,
08                         g_ioremap_desc_addr + DESC_LENGTH_REG);
09             iowrite32(START_DMA_MASK,
10                         g_ioremap_desc_addr + DESC_CONTROL_REG);
11
12             temp_dma_count -= this_loop_count;
13             ram_ptr += this_loop_count;
14             next_dma_io_buf_ofst += this_loop_count;
15             if (next_dma_io_buf_ofst >= PAGE_SIZE)
16                 next_dma_io_buf_ofst = 0;
17         } else {
18             while (get_dma_busy() != 0) {
19                 if (wait_event_interruptible(g_irq_wait_queue,
20                                         (get_dma_busy() ==
21                                         0))) {
22 ...cut...
23
24         }
25     }
26
27
28
29
30 }
```

descriptor configuration

Continues to directly interact with the g_dma_handle_4k.

demo_module_09.c – DMA

```
01     while (get_dma_fill_level() > 0) {
02         if (wait_event_interruptible(g_irq_wait_queue,
03                                     (get_dma_fill_level() ==
04                                     0))) {
05             ...cut...
06             }
07         }
08
09         if (copy_to_user(user_buffer, g_coherent_ptr_4k +
10             next_user_io_buf_ofst, this_loop_count)) {
11             up(&dev->sem);
12             ...cut...
13             }
14             temp_user_count -= this_loop_count;
15             user_buffer += this_loop_count;
16             next_user_io_buf_ofst += this_loop_count;
17             if (next_user_io_buf_ofst >= PAGE_SIZE)
18                 next_user_io_buf_ofst = 0;
19             }
20             ...cut...
21     }
22
23
24
25
26
27
28
29
30
```

Continues to directly interact
with the **g_dma_handle_4k**.

demo_module_09.c – DMA

```
01 static ssize_t demo_dma_co_write(struct file *fp,
02                                     const char __user *user_buffer, size_t count,
03                                     loff_t *offset)
04 {
05     ...cut...
06     while (temp_count > 0) {
07         while (get_dma_fill_level() > 2) {
08             if (wait_event_interruptible(g_irq_wait_queue,
09                                         (get_dma_fill_level() <=
10                                         2))) {
11     ...cut...
12         }
13
14         if (copy_from_user(g_coherent_ptr_4k +
15                             next_io_buf_ofst,
16                             user_buffer, this_loop_count)) {
17     ...cut...
18         }
19
20         iowrite32(g_dma_handle_4k + next_io_buf_ofst,
21                   g_ioremap_desc_addr + DESC_READ_ADDRESS_REG);
22         iowrite32(ram_ptr,
23                   g_ioremap_desc_addr + DESC_WRITE_ADDRESS_REG);
24         iowrite32(this_loop_count,
25                   g_ioremap_desc_addr + DESC_LENGTH_REG);
26         iowrite32(START_DMA_MASK,
27                   g_ioremap_desc_addr + DESC_CONTROL_REG);
28     ...cut...
29 }
30 }
```

descriptor configuration

Our coherent write directly interacts with the `g_dma_handle_4k` as well.

demo_module_09.c – DMA

```
01 static ssize_t demo_dma_st_read(struct file *fp, char __user *user_buffer,
02                                     size_t count, loff_t *offset)
03 {
04     ...cut...
05     ram_ptr = g_dma_handle_1m;
06     ram_ptr += *offset;
07     ...cut...
08     dma_handle = dma_map_single(dev->pdev_dev,
09                                 g_kmalloc_ptr_4k + next_dma_io_buf_ofst,
10                                 this_loop_count, DMA_FROM_DEVICE);
11
12     if (dma_mapping_error(dev->pdev_dev, dma_handle)) {
13         up(&dev->sem);
14     ...cut...
15     }
16
17     last_dma_handle_1 = last_dma_handle_0;
18     last_dma_handle_0 = dma_handle;
19
20     iowrite32(ram_ptr, g_ioremap_desc_addr + DESC_READ_ADDRESS_REG);
21     iowrite32(dma_handle, g_ioremap_desc_addr + DESC_WRITE_ADDRESS_REG);
22     iowrite32(this_loop_count, g_ioremap_desc_addr + DESC_LENGTH_REG);
23     iowrite32(START_DMA_MASK, g_ioremap_desc_addr + DESC_CONTROL_REG);
24
25     temp_dma_count -= this_loop_count;
26     ram_ptr += this_loop_count;
27     next_dma_io_buf_ofst += this_loop_count;
28     if (next_dma_io_buf_ofst >= PAGE_SIZE)
29         next_dma_io_buf_ofst = 0;
30 }
```

descriptor configuration

Our streaming read first dma_maps the g_malloc_ptr_4k to a dma_handle before the DMA. This performs the cache maintenance for us.

demo_module_09.c – DMA

```
01     while (temp_user_count > 0) {
02         if (temp_dma_count > 0) {
03             dma_handle = dma_map_single(dev->pdev_dev,
04                                         g_kmalloc_ptr_4k +
05                                         next_dma_io_buf_ofst,
06                                         this_loop_count,
07                                         DMA_FROM_DEVICE);
08
09             if (dma_mapping_error(dev->pdev_dev, dma_handle)) {
10                 up(&dev->sem);
11             }
12
13
14             last_dma_handle_1 = last_dma_handle_0;
15             last_dma_handle_0 = dma_handle;
16
17             iowrite32(ram_ptr,
18                         g_ioremap_desc_addr + DESC_READ_ADDRESS_REG);
19             iowrite32(dma_handle,
20                         g_ioremap_desc_addr + DESC_WRITE_ADDRESS_REG);
21             iowrite32(this_loop_count,
22                         g_ioremap_desc_addr + DESC_LENGTH_REG);
23             iowrite32(START_DMA_MASK,
24                         g_ioremap_desc_addr + DESC_CONTROL_REG);
25
26
27
28
29
30 }
```

descriptor configuration

Then it continues to **dma_map** the **g_malloc_ptr_4k** again.

demo_module_09.c – DMA

```
01 } else {
02     while (get_dma_busy() != 0) {
03         if (wait_event_interruptible(g_irq_wait_queue,
04                                     (get_dma_busy() ==
05                                         0))) {
06     ...cut...
07 }
08
09     while (get_dma_fill_level() > 0) {
10         if (wait_event_interruptible(g_irq_wait_queue,
11                                     (get_dma_fill_level() ==
12                                         0))) {
13     ...cut...
14 }
15
16     if (last_dma_handle_1 != 0) {
17         dma_unmap_single(dev->pdev_dev, last_dma_handle_1,
18                           this_loop_count, DMA_FROM_DEVICE);
19         last_dma_handle_1 = 0;
20     } else if (last_dma_handle_0 != 0) {
21         dma_unmap_single(dev->pdev_dev, last_dma_handle_0,
22                           this_loop_count, DMA_FROM_DEVICE);
23         last_dma_handle_0 = 0;
24     }
25
26     if (copy_to_user(user_buffer, g_kmalloc_ptr_4k +
27                      next_user_io_buf_ofst, this_loop_count)) {
28     ...cut...
29 }
```

Then it dma_unmaps the dma_handle before using the g_malloc_ptr_4k.

demo_module_09.c – DMA

```
01 static ssize_t demo_dma_st_write(struct file *fp,
02                                     const char __user *user_buffer, size_t count,
03                                     loff_t *offset)
04 {
05     ...cut...
06     ram_ptr = g_dma_handle_1m;
07     ram_ptr += *offset;
08     ...cut...
09     while (temp_count > 0) {
10         while (get_dma_fill_level() > 2) {
11             if (wait_event_interruptible(g_irq_wait_queue,
12                                         (get_dma_fill_level() <=
13                                         2))) {
14     ...cut...
15         }
16
17         if (last_dma_handle_3 != 0) {
18             dma_unmap_single(pdev->pdev_dev, last_dma_handle_3,
19                               this_loop_count, DMA_TO_DEVICE);
20             last_dma_handle_3 = 0;
21         }
22
23         if (copy_from_user(g_kmalloc_ptr_4k +
24                             next_io_buf_ofst,
25                             user_buffer, this_loop_count)) {
26             up(&dev->sem);
27             pr_info("demo_dma_st_write copy_from_user exit\n");
28             return -EFAULT;
29         }
30     }
```

The streaming write uses the `g_kmalloc_ptr_4k` first.

demo_module_09.c – DMA

```
01     dma_handle = dma_map_single(pdev->pdev, g_kmalloc_ptr_4k + next_io_buf_ofst, this_loop_count, DMA_TO_DEVICE);  
02  
03     if (dma_mapping_error(pdev, dma_handle)) {  
04         ...cut...  
05     }  
06  
07     last_dma_handle_3 = last_dma_handle_2;  
08     last_dma_handle_2 = last_dma_handle_1;  
09     last_dma_handle_1 = last_dma_handle_0;  
10     last_dma_handle_0 = dma_handle;  
11  
12     ...cut...  
13  
14     iowrite32(dma_handle, g_ioremap_desc_addr + DESC_READ_ADDRESS_REG);  
15     iowrite32(ram_ptr, g_ioremap_desc_addr + DESC_WRITE_ADDRESS_REG);  
16     iowrite32(this_loop_count, g_ioremap_desc_addr + DESC_LENGTH_REG);  
17     iowrite32(START_DMA_MASK, g_ioremap_desc_addr + DESC_CONTROL_REG);  
18  
19     ...cut...  
20     }  
21  
22     while (get_dma_busy() != 0) {  
23         if (wait_event_interruptible(g_irq_wait_queue, (get_dma_busy() == 0))) {  
24             ...cut...  
25         }  
26     }
```

Then it dma_maps it to a
dma_handle before the
DMA.

descriptor configuration

demo_module_09.c – DMA

```
01     if (last_dma_handle_3 != 0)
02         dma_unmap_single(dev->pdev_dev, last_dma_handle_3,
03                             this_loop_count, DMA_TO_DEVICE);
04
05     if (last_dma_handle_2 != 0)
06         dma_unmap_single(dev->pdev_dev, last_dma_handle_2,
07                             this_loop_count, DMA_TO_DEVICE);
08
09     if (last_dma_handle_1 != 0)
10        dma_unmap_single(dev->pdev_dev, last_dma_handle_1,
11                           this_loop_count, DMA_TO_DEVICE);
12
13     if (last_dma_handle_0 != 0)
14        dma_unmap_single(dev->pdev_dev, last_dma_handle_0,
15                           this_loop_count, DMA_TO_DEVICE);
16 ...cut...
17 }
```

Then it dma_unmaps the
dma_handle after the DMA.

demo_module_09.ko target demonstration

```
01 insmod demo_module_09.ko
02
03 # create a couple random files
04 dd if=/dev/urandom of=random_co.bin bs=1M count=1
05 dd if=/dev/urandom of=random_st.bin bs=1M count=1
06
07 # write to coherent DMA device and read from coherent DMA device
08 dd if=random_co.bin of=/dev/demo_dma_co bs=1024
09 dd if=/dev/demo_dma_co of=random_co_data.bin bs=1024
10
11 # write to streaming DMA device and read from streaming DMA device
12 dd if=random_st.bin of=/dev/demo_dma_st bs=1024
13 dd if=/dev/demo_dma_st of=random_st_data.bin bs=1024
14
15 # write to coherent DMA device and read from streaming DMA device
16 dd if=random_co.bin of=/dev/demo_dma_co bs=1024
17 dd if=/dev/demo_dma_st of=random_co2st_data.bin bs=1024
18
19 # write to streaming DMA device and read from coherent DMA device
20 dd if=random_st.bin of=/dev/demo_dma_st bs=1024
21 dd if=/dev/demo_dma_co of=random_st2co_data.bin bs=1024
22
23 # compare all the files
24 md5sum *.bin
25 rmmod demo_module_09
26
27
28
29
30
```


Multiple Device Instances

The Altera logo, featuring the word "ALTERA" in a bold, blue, sans-serif font. A registered trademark symbol (®) is located at the top right corner of the letter "A".

ALTERA®

demo_module_10

User Space



Kernel Space

IO Space

demo_module_10.c – Multiple Device Instances

```
01 static struct of_device_id demo_driver_dt_ids[] = {  
02     {  
03         .compatible = "ALTR,fifo-14.1",  
04     } /* end of table */ }  
05 };  
06  
07  
08  
09  
10  
11  
12  
13  
14 static int demo_init(void)  
15 {  
16     int ret_val;  
17  
18     INIT_LIST_HEAD(&g_dev_list);  
19     sema_init(&g_dev_list_sem, 1);  
20     g_dev_index = 0;  
21  
22     ret_val = platform_driver_register(&the_platform_driver);  
23     if (ret_val != 0) {  
24         pr_err("platform_driver_register returned %d\n", ret_val);  
25         return ret_val;  
26     }  
27  
28     return 0;  
29 }  
30 }
```

We bind this example with devices matching this compatible string.

Our init routine clears a global index that we track each instance of our driver that be bind to each device.

demo_module_10.c – Multiple Device Instances

```
01 static int platform_probe(struct platform_device *pdev)
02 {
03     int ret_val;
04     struct resource *r0 = NULL;
05     struct resource *r1 = NULL;
06     struct resource *r2 = NULL;
07     struct resource *temp_res = NULL;
08     struct demo_fifo_dev *the_demo_fifo_dev;
09     uint32_t fifo_level;
10
11     if (down_interruptible(&g_dev_list_sem))
12         return -ERESTARTSYS;
13
14     ret_val = -ENOMEM;
15
16     /* allocate a demo_fifo_dev structure */
17     the_demo_fifo_dev = kzalloc(sizeof(struct demo_fifo_dev), GFP_KERNEL);
18     if (the_demo_fifo_dev == NULL) {
19         pr_err("kzalloc failed\n");
20         goto bad_exit_return;
21     }
22
23
24
25
26
27
28
29
30
```

Our probe routine allocates the storage for our device structure rather than using a statically allocated structure like the previous examples.

demo_module_10.c – Multiple Device Instances

```
01  /* initialize the demo_fifo_dev structure */
02  scnprintf(the_demo_fifo_dev->name, NAME_BUF_SIZE, "demo_fifo_%d",
03  g_dev_index);
04
05  INIT_LIST_HEAD(&the_demo_fifo_dev->dev_list);
06  sema_init(&the_demo_fifo_dev->dev_sem, 1);
07  init_waitqueue_head(&the_demo_fifo_dev->wait_queue);
08
09  the_demo_fifo_dev->in_res = NULL;
10  the_demo_fifo_dev->out_res = NULL;
11  the_demo_fifo_dev->in_csr_res = NULL;
12
13  the_demo_fifo_dev->ioremap_in_addr = NULL;
14  the_demo_fifo_dev->ioremap_out_addr = NULL;
15  the_demo_fifo_dev->ioremap_in_csr_addr = NULL;
16
17  the_demo_fifo_dev->open_for_read = 0;
18  the_demo_fifo_dev->open_for_write = 0;
19
20  the_demo_fifo_dev->miscdev.minor = MISC_DYNAMIC_MINOR;
21  the_demo_fifo_dev->miscdev.name = the_demo_fifo_dev->name;
22  the_demo_fifo_dev->miscdev.fops = &demo_fifo_fops;
23
24  ret_val = -EINVAL;
25
26
27
28
29
30
```

Then we initialize the newly allocated device structure. The name being the most apparent item that we will see in user space.

demo_module_10.c – Multiple Device Instances

```
01     /* get our three expected memory resources */
02     r0 = platform_get_resource(pdev, IORESOURCE_MEM, 0);
03     if (r0 == NULL) {
04         pr_err("IORESOURCE_MEM, 0 does not exist\n");
05         goto bad_exit_kfree_the_demo_fifo_dev;      We expect the FIFO devices
06     }                                              to present three memory
07
08     r1 = platform_get_resource(pdev, IORESOURCE_MEM, 1); resources.
09     if (r1 == NULL) {
10         pr_err("IORESOURCE_MEM, 1 does not exist\n");
11         goto bad_exit_kfree_the_demo_fifo_dev;
12     }
13
14     r2 = platform_get_resource(pdev, IORESOURCE_MEM, 2);
15     if (r2 == NULL) {
16         pr_err("IORESOURCE_MEM, 2 does not exist\n");
17         goto bad_exit_kfree_the_demo_fifo_dev;
18     }
19
20
21
22
23
24
25
26
27
28
29
30
```

demo_module_10.c – Multiple Device Instances

```
01      /* associate the resources for in, out and in_csr */  
02  
03      if (!strcmp(r0->name, "in"))  
04          the_demo_fifo_dev->in_res = r0;  
05      else if (!strcmp(r1->name, "in"))  
06          the_demo_fifo_dev->in_res = r1;  
07      else if (!strcmp(r2->name, "in"))  
08          the_demo_fifo_dev->in_res = r2;  
09  
10      if (!strcmp(r0->name, "out"))  
11          the_demo_fifo_dev->out_res = r0;  
12      else if (!strcmp(r1->name, "out"))  
13          the_demo_fifo_dev->out_res = r1;  
14      else if (!strcmp(r2->name, "out"))  
15          the_demo_fifo_dev->out_res = r2;  
16  
17      if (!strcmp(r0->name, "in_csr"))  
18          the_demo_fifo_dev->in_csr_res = r0;  
19      else if (!strcmp(r1->name, "in_csr"))  
20          the_demo_fifo_dev->in_csr_res = r1;  
21      else if (!strcmp(r2->name, "in_csr"))  
22          the_demo_fifo_dev->in_csr_res = r2;  
23  
24  
25  
26  
27  
28  
29  
30
```

We don't assume the order in which the resources are presented, our software needs to know which resource belongs to which register set. So we figure that out by the name associated with each resource.

demo_module_10.c – Multiple Device Instances

```
01  /* verify that we found all three resources */
02  if (the_demo_fifo_dev->in_res == NULL) {
03      pr_err("no resource found for \"in\"\n");
04      goto bad_exit_kfree_the_demo_fifo_dev;
05  }
06
07  if (the_demo_fifo_dev->out_res == NULL) {
08      pr_err("no resource found for \"out\"\n");
09      goto bad_exit_kfree_the_demo_fifo_dev;
10  }
11
12  if (the_demo_fifo_dev->in_csr_res == NULL) {
13      pr_err("no resource found for \"in_csr\"\n");
14      goto bad_exit_kfree_the_demo_fifo_dev;
15  }
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
```

And after all that we just verify that we actually found each named resource that we were looking for.

demo_module_10.c – Multiple Device Instances

```
01             /* reserve our memory regions */
02 ...cut...
03             /* ioremap our memory regions */
04 ...cut...
05             /* initialize the FIFO hardware */
06     fifo_level = ioread32(the_demo_fifo_dev->ioremap_in_csr_addr +
07                           FIFO_LEVEL_REG);
08     while (fifo_level > 0) {
09         ioread32(the_demo_fifo_dev->ioremap_out_addr + FIFO_DATA_REG);
10         fifo_level--;
11     }
12
13     fifo_level = ioread32(the_demo_fifo_dev->ioremap_in_csr_addr +
14                           FIFO_LEVEL_REG);
15
16     if (fifo_level != 0) {
17         pr_err("fifo initialization failed");
18         goto bad_exit_iounmap_in_csr;
19     }
20
21     ret_val = -EINVAL;
22
23     /* register misc device demo_fifo */
24 ...cut...
25     /* clean up and exit */
26     list_add(&the_demo_fifo_dev->dev_list, &g_dev_list);
27     g_dev_index++;
28     platform_set_drvdata(pdev, the_demo_fifo_dev);
29     up(&g_dev_list_sem);
30     return 0;
```

We initialize our FIFO hardware.

And finally we add our device to a global list that we track our device instances with, increment the g_dev_index and store our device structure into the platform driver data.

demo_module_10.c – Multiple Device Instances

```
01 static const struct file_operations demo_fifo_fops = {  
02     .owner = THIS_MODULE,  
03     .open = demo_fifo_open,  
04     .release = demo_fifo_release,  
05     .read = demo_fifo_read,  
06     .write = demo_fifo_write,  
07 };
```

We can open, release, read and write our FIFO device.

demo_module_10.c – Multiple Device Instances

```
01 static int demo_fifo_open(struct inode *ip, struct file *fp)
02 {
03     struct demo_fifo_dev *the_demo_fifo_dev = NULL;
04     uint32_t this_minor;
05     uint32_t access_mode;
06     struct list_head *next_list_entry;
07     int found_it = 0;
08
09     if (down_interruptible(&g_dev_list_sem))
10         return -ERESTARTSYS;
11
12     this_minor = iminor(ip);
13
14     list_for_each(next_list_entry, &g_dev_list) {
15         the_demo_fifo_dev = list_entry(next_list_entry,
16                                         struct demo_fifo_dev, dev_list);
17         if (the_demo_fifo_dev->miscdev.minor == this_minor) {
18             found_it = 1;
19             break;
20         }
21     }
22
23     up(&g_dev_list_sem);
24
25     if (found_it == 0)
26         return -ENXIO;
27
28
29
30 }
```

Our open routine needs to search thru our device list to locate our device instance that matches the device that's being opened.

The device minor number is what identifies this for us.

demo_module_10.c – Multiple Device Instances

```
01 access_mode = fp->f_flags & O_ACCMODE;
02 switch (access_mode) {
03     case (O_RDONLY):
04         if (the_demo_fifo_dev->open_for_read != 0)
05             return -EBUSY;
06         the_demo_fifo_dev->open_for_read = 1;
07         break;
08     case (O_WRONLY):
09         if (the_demo_fifo_dev->open_for_write != 0)
10             return -EBUSY;
11         the_demo_fifo_dev->open_for_write = 1;
12         break;
13     case (O_RDWR):
14         if (the_demo_fifo_dev->open_for_read != 0)
15             return -EBUSY;
16         if (the_demo_fifo_dev->open_for_write != 0)
17             return -EBUSY;
18         the_demo_fifo_dev->open_for_read = 1;
19         the_demo_fifo_dev->open_for_write = 1;
20         break;
21     default:
22         return -EINVAL;
23 }
24
25 fp->private_data = the_demo_fifo_dev;
26
27 return 0;
28 }
```

We ensure that we only allow one file descriptor be associated with reading or writing the device at a time.

We determine this from the access mode of the file.

demo_module_10.c – Multiple Device Instances

```
01 static ssize_t demo_fifo_read(struct file *fp, char __user *user_buffer,
02                               size_t count, loff_t *offset)
03 {
04     struct demo_fifo_dev *the_demo_fifo_dev = fp->private_data;
05 ...cut...
06     if (down_interruptible(&the_demo_fifo_dev->dev_sem)) {
07 ...cut...
08         }
09
10     if (count & (4 - 1)) {  

11 ...cut...
12         }
13
14     while (ioread32(the_demo_fifo_dev->ioremap_in_csr_addr + FIFO_LEVEL_REG)
15           == 0) {
16         up(&the_demo_fifo_dev->dev_sem);
17         if (wait_event_interruptible(the_demo_fifo_dev->wait_queue,
18                                     ioread32
19                                     (the_demo_fifo_dev->
20                                      ioremap_in_csr_addr
21                                      + FIFO_LEVEL_REG) != 0))) {
22             pr_info("demo_fifo_read wait interrupted exit\n");
23             return -ERESTARTSYS;
24         }
25         if (down_interruptible(&the_demo_fifo_dev->dev_sem)) {
26             pr_info("demo_fifo_read sem interrupted exit\n");
27             return -ERESTARTSYS;
28         }
29     }
30 }
```

We only accept a modulo 4 count.

Our read blocks on an empty FIFO and waits for the write routine to clear the wait queue.

demo_module_10.c – Multiple Device Instances

```
01     fifo_level = ioread32(the_demo_fifo_dev->ioremap_in_csr_addr +
02                             FIFO_LEVEL_REG);
03
04     while (fifo_level > 0) {
05         temp_data = ioread32(the_demo_fifo_dev->ioremap_out_addr +
06                               FIFO_DATA_REG);
07         if (copy_to_user(user_buffer, &temp_data, 4)) {
08             up(&the_demo_fifo_dev->dev_sem);
09             pr_info("demo_fifo_read copy_to_user exit\n");
10             return -EFAULT;
11         }
12         user_buffer += 4;
13         fifo_level--;
14         this_count += 4;
15         count -= 4;
16         if (count == 0)
17             break;
18     }
19
20     up(&the_demo_fifo_dev->dev_sem);
21     wake_up_interruptible(&the_demo_fifo_dev->wait_queue);
22
23     return this_count;
24 }
```

Then we read whatever we can out of the FIFO, up to count.

And we wake the wait queue when we're finished.

demo_module_10.c – Multiple Device Instances

```
01 static ssize_t demo_fifo_write(struct file *fp,
02                                 const char __user *user_buffer, size_t count,
03                                 loff_t *offset)
04 {
05     struct demo_fifo_dev *the_demo_fifo_dev = fp->private_data;
06 ...cut...
07     if (down_interruptible(&the_demo_fifo_dev->dev_sem)) {
08 ...cut...
09     }
10
11     if (count & (4 - 1)) {  

12 ...cut...
13     }
14
15     while (ioread32(the_demo_fifo_dev->ioremap_in_csr_addr + FIFO_LEVEL_REG)
16             >= FIFO_MAX_FILL_LEVEL) {
17         up(&the_demo_fifo_dev->dev_sem);
18         if (wait_event_interruptible(the_demo_fifo_dev->wait_queue,
19                                     ioread32
20                                     (the_demo_fifo_dev->
21                                      ioremap_in_csr_addr
22                                      + FIFO_LEVEL_REG) <
23                                      FIFO_MAX_FILL_LEVEL))) {
24             pr_info("demo_fifo_write wait interrupted exit\n");
25             return -ERESTARTSYS;
26         }
27         if (down_interruptible(&the_demo_fifo_dev->dev_sem)) {
28 ...cut...
29     }
30 }
```

We only accept a modulo 4 count.

Our write blocks on a full FIFO and waits for the read routine to clear the wait queue.

demo_module_10.c – Multiple Device Instances

```
01     fifo_level = ioread32(the_demo_fifo_dev->ioremap_in_csr_addr +
02                             FIFO_LEVEL_REG);
03
04     while (fifo_level < FIFO_MAX_FILL_LEVEL) {
05         if (copy_from_user(&temp_data, user_buffer, 4)) {
06             up(&the_demo_fifo_dev->dev_sem);
07             pr_info("demo_fifo_write copy_to_user exit\n");
08             return -EFAULT;
09         }
10         iowrite32(temp_data, the_demo_fifo_dev->ioremap_in_addr +
11                   FIFO_DATA_REG);
12
13         user_buffer += 4;
14         fifo_level++;
15         this_count += 4;
16         count -= 4;
17         if (count == 0)
18             break;
19     }
20
21     up(&the_demo_fifo_dev->dev_sem);
22     wake_up_interruptible(&the_demo_fifo_dev->wait_queue);
23
24     return this_count;
25 }
```

Then we write whatever we can into the FIFO, up to count.

And we wake the wait queue when we're finished.

demo_module_10.ko target demonstration

```
01 insmod demo_module_10.ko                      # insert the module
02 find /dev -name "*demo*"                     # observe the device nodes
03 find /sys -name "*demo*"                     # observe the sysfs files
04 ls /sys/bus/platform/drivers/demo_driver_10   # notice the 3 devices bound
05
06 # create 3 random data files
07 dd if=/dev/urandom of=random0.bin bs=256K count=1
08 dd if=/dev/urandom of=random1.bin bs=256K count=1
09 dd if=/dev/urandom of=random2.bin bs=256K count=1
10
11 # start 3 separate processes to write the random files to a different fifo
12 dd if=random0.bin of=/dev/demo_fifo_0 bs=1024 &
13 dd if=random1.bin of=/dev/demo_fifo_1 bs=1024 &
14 dd if=random2.bin of=/dev/demo_fifo_2 bs=1024 &
15
16 # read back the first 1K of data from each fifo
17 dd if=/dev/demo_fifo_0 of=random0_out.bin bs=1024 count=1
18 dd if=/dev/demo_fifo_1 of=random1_out.bin bs=1024 count=1
19 dd if=/dev/demo_fifo_2 of=random2_out.bin bs=1024 count=1
20
21 # read back the remaining data from all the fifos
22 dd if=/dev/demo_fifo_0 of=random0_out.bin bs=1024 count=255 conv=notrunc seek=1 &
23 dd if=/dev/demo_fifo_1 of=random1_out.bin bs=1024 count=255 conv=notrunc seek=1 &
24 dd if=/dev/demo_fifo_2 of=random2_out.bin bs=1024 count=255 conv=notrunc seek=1 &
25
26 wait                                         # wait for background processes
27 md5sum *.bin                                 # validate the in/out data files
28 rmmod demo_module_10                         # remove the module
29
30
```


Custom API



ALTERA®

demo_module_11

User Space




Kernel Space



IO Space

demo_module_11t

User Space



Kernel Space



IO Space

demo_module_11.c – Custom API

```
01 void demo_increment(uint32_t *target)
02 {
03     *target += 1;
04 }
05 EXPORT_SYMBOL(demo_increment);
06
07 void demo_decrement(uint32_t *target)
08 {
09     *target -= 1;
10 }
11 EXPORT_SYMBOL(demo_decrement);
12
13 void demo_complement(uint32_t *target)
14 {
15     *target = ~(*target);
16 }
17 EXPORT_SYMBOL(demo_complement);
18
19 static int demo_init(void)
20 {
21     pr_info("demo_init enter\n");
22     pr_info("demo_init exit\n");
23     return 0;
24 }
25
26 static void demo_exit(void)
27 {
28     pr_info("demo_exit enter\n");
29     pr_info("demo_exit exit\n");
30 }
```

We simply export some entry points that become available in kernel space once our module is loaded.

demo_module_11t.c – Custom API

```
01 static int demo_init(void)
02 {
03     ...cut...
04         ret_val = platform_driver_register(&the_platform_driver);
05     ...cut...
06         /* create the sysfs entries */
07         ret_val = driver_create_file(&the_platform_driver.driver,
08                                     &driver_attr_target);
09     ...cut...
10         ret_val = driver_create_file(&the_platform_driver.driver,
11                                     &driver_attr_increment);
12     ...cut...
13         ret_val = driver_create_file(&the_platform_driver.driver,
14                                     &driver_attr_decrement);
15     ...cut...
16         ret_val = driver_create_file(&the_platform_driver.driver,
17                                     &driver_attr_complement);
18     ...cut...
19         return 0;
20
21 bad_exit_remove_decrement_file:
22     driver_remove_file(&the_platform_driver.driver, &driver_attr_decrement);
23 bad_exit_remove_increment_file:
24     driver_remove_file(&the_platform_driver.driver, &driver_attr_increment);
25 bad_exit_remove_target_file:
26     driver_remove_file(&the_platform_driver.driver, &driver_attr_target);
27 bad_exit_platform_driver_unregister:
28     platform_driver_unregister(&the_platform_driver);
29 bad_exit_return:
30     return ret_val;
```

This test module defines a number of sysfs files that we can interact with our custom API through.

demo_module_11t.c – Custom API

```
01 static uint32_t g_target;
02
03 static ssize_t target_show(struct device_driver *driver, char *buf)
04 {
05     return scnprintf(buf, PAGE_SIZE, "0x%08X\n", g_target);
06 }
07
08 static ssize_t target_store(struct device_driver *driver, const char *buf,
09                             size_t count)
10 {
11     int result;
12     unsigned long new_target;
13
14     /* convert the input string to the requested new target value */
15     result = kstrtoul(buf, 0, &new_target);
16     if (result != 0)
17         return -EINVAL;
18
19     g_target = new_target;
20
21     return count;
22 }
23
24 DRIVER_ATTR(target, (S_IWUGO | S_IRUGO), target_show, target_store);
25
26
27
28
29
30
```

These are our “target”
variable access
routines.

demo_module_11t.c – Custom API

```
01 static ssize_t increment_show(struct device_driver *driver, char *buf)
02 {
03     demo_increment(&g_target);
04
05     return scnprintf(buf, PAGE_SIZE, "0x%08X\n", g_target);
06 }
07
08 DRIVER_ATTR(increment, (S_IRUGO), increment_show, NULL);
09
10 static ssize_t decrement_show(struct device_driver *driver, char *buf)
11 {
12     demo_decrement(&g_target);
13
14     return scnprintf(buf, PAGE_SIZE, "0x%08X\n", g_target);
15 }
16
17 DRIVER_ATTR(decrement, (S_IRUGO), decrement_show, NULL);
18
19 static ssize_t complement_show(struct device_driver *driver, char *buf)
20 {
21     demo_complement(&g_target);
22
23     return scnprintf(buf, PAGE_SIZE, "0x%08X\n", g_target);
24 }
25
26 DRIVER_ATTR(complement, (S_IRUGO), complement_show, NULL);
27
```

These are our target “increment”, “decrement” and “complement” access routines.

demo_module_11.ko target demonstration

```
01 insmod demo_module_11t.ko          # observe module insert fails
02 insmod demo_module_11.ko          # insert module
03 insmod demo_module_11t.ko          # insert module
04 find /sys -name "*demo*"          # observe entries created
05 ls /sys/bus/platform/drivers/demo_driver_11t    # observe sysfs files
06
07 # exercise the sysfs entries
08 cat /sys/bus/platform/drivers/demo_driver_11t/target
09 cat /sys/bus/platform/drivers/demo_driver_11t/increment
10 cat /sys/bus/platform/drivers/demo_driver_11t/increment
11 cat /sys/bus/platform/drivers/demo_driver_11t/increment
12 cat /sys/bus/platform/drivers/demo_driver_11t/increment
13 cat /sys/bus/platform/drivers/demo_driver_11t/decrement
14 cat /sys/bus/platform/drivers/demo_driver_11t/decrement
15 cat /sys/bus/platform/drivers/demo_driver_11t/complement
16 cat /sys/bus/platform/drivers/demo_driver_11t/increment
17 cat /sys/bus/platform/drivers/demo_driver_11t/increment
18 cat /sys/bus/platform/drivers/demo_driver_11t/increment
19 cat /sys/bus/platform/drivers/demo_driver_11t/decrement
20 cat /sys/bus/platform/drivers/demo_driver_11t/complement
21 rmmod demo_module_11          # observe remove module fails
22 rmmod demo_module_11t         # remove module
23 rmmod demo_module_11          # remove module
24
25
26
27
28
29
30
```

Take Home Lab



ALTERA®

The Altera logo consists of the word "ALTERA" in a bold, blue, sans-serif font. A registered trademark symbol (®) is located at the top right corner of the letter "A".

What You'll Need

- ◀ A supported development kit
 - ALTERA_AV_SOC - Altera
 - ALTERA_CV_SOC - Altera
 - ARROW_SOCKIT - Arrow
 - CRITICALLINK_MITYSOM_DEVKIT - Critical Link
 - DE0_NANO_SOC - Terasic
 - MACNICA_HELIOTROPE_12, 13 or 14 - Macnica
- ◀ A blank SD Card for your development kit
 - To install the workshop target environment that boots the above boards.
- ◀ Communication cables
 - Provides connectivity between host and target.
 - ◀ Console UART
 - ◀ USB gadget mass storage and networking
- ◀ Linux machine, native or VM
 - 4GByte RAM minimum
- ◀ Serial terminal application
 - Putty, minicom, or similar, for console UART communication with target

Obtaining Lab Files and Instructions

↳ SoC SW workshop series getting started page

<http://www.rocketboards.org/foswiki/Projects/SoCSWWorkshopSeriesGettingStarted>

↳ WS3 lab instructions are posted to Rocketboards.org

<http://www.rocketboards.org/foswiki/Projects/WS3DevelopingDriversForAlteraSoCLinux>

The screenshot shows a foswiki page for the "WS3 Developing Drivers for Altera SoC Linux" project. At the top, there's a navigation bar with links for Home, Documentation, Community, Projects, and Partners. A search bar and user account links are also present. The main content area has a header "WS3 Developing Drivers for Altera SoC Linux" and a sub-header "This is the third of three SoC workshops providing an overview of the SoC Linux driver development concepts." Below this, a message states "Last modified by ProFromDover on 01 Jun 2015 - 00:59 - r26". The "Details" section contains a "State: running" status update. The "Full description" section lists several items: Introduction, Presentation Material, Setup for Labs, What you will accomplish, What you will need, Obtain Development Board, Obtain SD card Image, Locate WS3 Lab README, and Target Orientation.

What You Will Accomplish

- ☛ Configure your host development environment.
- ☛ Build an example misc driver.
- ☛ Build a test application for the misc driver.
- ☛ Modify the uio_pdrv_genirq template and build it.
- ☛ Build a test application for the UIO driver.
- ☛ Load all of the above onto your dev kit target.
- ☛ Run the test applications to verify proper driver functionality.
- ☛ Run a validation program.
 - Verifies the lab was completed.
- ☛ Submit results & feedback.

Thank You

© 2015 Altera Corporation—Public

All rights reserved. ALTERA, ARRIA, CYCLONE, ENPIRION, MAX, MEGACORE, NIOS, QUARTUS and STRATIX words and logos are trademarks of Altera Corporation and registered in the U.S. Patent and Trademark Office and in other countries. All other words and logos identified as trademarks or service marks are the property of their respective holders as described at www.altera.com/legal.

The Altera logo consists of the word "ALTERA" in a bold, blue, sans-serif font. The letters are outlined in blue and filled with a lighter shade of blue. A registered trademark symbol (®) is located at the top right corner of the letter "A".