




Linux Kernel Introduction

Free Electrons

© Copyright 2004-2014, Free Electrons.
Creative Commons BY-SA 3.0 license.
Corrections, suggestions, contributions and translations are welcome!





Linux features



History

- ▶ The Linux kernel is one component of a system, which also requires libraries and applications to provide features to end users.
- ▶ The Linux kernel was created as a hobby in 1991 by a Finnish student, Linus Torvalds.
 - ▶ Linux quickly started to be used as the kernel for free software operating systems
- ▶ Linus Torvalds has been able to create a large and dynamic developer and user community around Linux.
- ▶ Nowadays, more than one thousand people contribute to each kernel release, individuals or companies big and small.

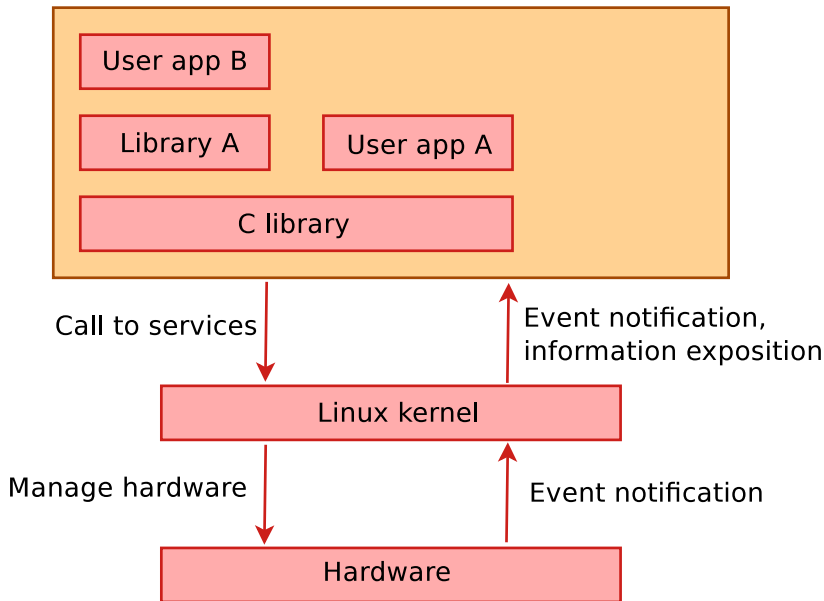


Linux kernel key features

- ▶ Portability and hardware support. Runs on most architectures.
- ▶ Scalability. Can run on super computers as well as on tiny devices (4 MB of RAM is enough).
- ▶ Compliance to standards and interoperability.
- ▶ Exhaustive networking support.
- ▶ Security. It can't hide its flaws. Its code is reviewed by many experts.
- ▶ Stability and reliability.
- ▶ Modularity. Can include only what a system needs even at run time.
- ▶ Easy to program. You can learn from existing code. Many useful resources on the net.



Linux kernel in the system





Linux kernel main roles

- ▶ **Manage all the hardware resources:** CPU, memory, I/O.
- ▶ Provide a **set of portable, architecture and hardware independent APIs** to allow user space applications and libraries to use the hardware resources.
- ▶ **Handle concurrent accesses and usage** of hardware resources from different applications.
 - ▶ Example: a single network interface is used by multiple user space applications through various network connections. The kernel is responsible to “multiplex” the hardware resource.



System calls

- ▶ The main interface between the kernel and user space is the set of system calls
- ▶ About 300 system calls that provide the main kernel services
 - ▶ File and device operations, networking operations, inter-process communication, process management, memory mapping, timers, threads, synchronization primitives, etc.
- ▶ This interface is stable over time: only new system calls can be added by the kernel developers
- ▶ This system call interface is wrapped by the C library, and user space applications usually never make a system call directly but rather use the corresponding C library function



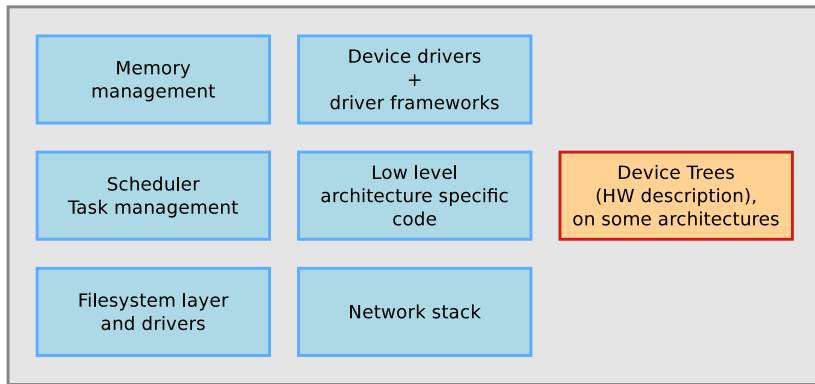
Pseudo filesystems

- ▶ Linux makes system and kernel information available in user space through **pseudo filesystems**, sometimes also called **virtual filesystems**
- ▶ Pseudo filesystems allow applications to see directories and files that do not exist on any real storage: they are created and updated on the fly by the kernel
- ▶ The two most important pseudo filesystems are
 - ▶ `proc`, usually mounted on `/proc`:
Operating system related information (processes, memory management parameters...)
 - ▶ `sysfs`, usually mounted on `/sys`:
Representation of the system as a set of devices and buses.
Information about these devices.



Inside the Linux kernel

Linux Kernel



Implemented mainly in C,
a little bit of assembly.



Written in a Device Tree
specific language.



Supported hardware architectures


- ▶ See the `arch/` directory in the kernel sources
- ▶ Minimum: 32 bit processors, with or without MMU, and `gcc` support
- ▶ 32 bit architectures (`arch/` subdirectories)
Examples: `arm`, `avr32`, `blackfin`, `c6x`, `m68k`, `microblaze`, `mips`, `score`, `sparc`, `um`
- ▶ 64 bit architectures:
Examples: `alpha`, `arm64`, `ia64`, `tile`
- ▶ 32/64 bit architectures
Examples: `powerpc`, `x86`, `sh`, `sparc`
- ▶ Find details in kernel sources: `arch/<arch>/Kconfig`, `arch/<arch>/README`, or `Documentation/<arch>/`



Embedded Linux Kernel Usage

Free Electrons

© Copyright 2004-2014, Free Electrons.
Creative Commons BY-SA 3.0 license.
Corrections, suggestions, contributions and translations are welcome!





Linux kernel sources



Location of kernel sources

- ▶ The official versions of the Linux kernel, as released by Linus Torvalds, are available at <http://www.kernel.org>
 - ▶ These versions follow the development model of the kernel
 - ▶ However, they may not contain the latest development from a specific area yet. Some features in development might not be ready for mainline inclusion yet.
- ▶ Many chip vendors supply their own kernel sources
 - ▶ Focusing on hardware support first
 - ▶ Can have a very important delta with mainline Linux
 - ▶ Useful only when mainline hasn't caught up yet.
- ▶ Many kernel sub-communities maintain their own kernel, with usually newer but less stable features
 - ▶ Architecture communities (ARM, MIPS, PowerPC, etc.), device drivers communities (I2C, SPI, USB, PCI, network, etc.), other communities (real-time, etc.)
 - ▶ No official releases, only development trees are available.



Getting Linux sources

- ▶ The kernel sources are available from <http://kernel.org/pub/linux/kernel> as **full tarballs** (complete kernel sources) and **patches** (differences between two kernel versions).
- ▶ However, more and more people use the `git` version control system. Absolutely needed for kernel development!
 - ▶ Fetch the entire kernel sources and history

```
git clone git://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git
```
 - ▶ Create a branch that starts at a specific stable version

```
git checkout -b <name-of-branch> v3.11
```
 - ▶ Web interface available at <http://git.kernel.org/cgit/linux/kernel/git/torvalds/linux.git/tree/>.
 - ▶ Read more about Git at <http://git-scm.com/>



Linux kernel size (1)

- ▶ Linux 3.10 sources:
 - Raw size: 573 MB (43,000 files, approx 15,800,000 lines)
 - gzip compressed tar archive: 105 MB
 - bzip2 compressed tar archive: 83 MB (better)
 - xz compressed tar archive: 69 MB (best)
- ▶ Minimum Linux 3.17 compiled kernel size, booting on the ARM Versatile board (hard drive on PCI, ext2 filesystem, ELF executable support, framebuffer console and input devices): 876 KB (compressed), 2.3 MB (raw)
- ▶ Why are these sources so big?
 - Because they include thousands of device drivers, many network protocols, support many architectures and filesystems...
- ▶ The Linux core (scheduler, memory management...) is pretty small!



Linux kernel size (2)

As of kernel version 3.10.

- ▶ `drivers/`: 49.4%
- ▶ `arch/`: 21.9%
- ▶ `fs/`: 6.0%
- ▶ `include/`: 4.7%
- ▶ `sound/`: 4.4%
- ▶ `Documentation/`: 4.0%
- ▶ `net/`: 3.9%
- ▶ `firmware/`: 1.0%
- ▶ `kernel/`: 1.0%
- ▶ `tools/`: 0.9%
- ▶ `scripts/`: 0.5%
- ▶ `mm/`: 0.5%
- ▶ `crypto/`: 0.4%
- ▶ `security/`: 0.4%
- ▶ `lib/`: 0.4%
- ▶ `block/`: 0.2%
- ▶ ...



Practical lab - Get Linux Kernel Source Code




- ▶ Clone the mainline Linux source tree with git



Kernel Source Code

Free Electrons

© Copyright 2004-2014, Free Electrons.
Creative Commons BY-SA 3.0 license.
Corrections, suggestions, contributions and translations are welcome!





Linux Code and Device Drivers



Programming language

- ▶ Implemented in C like all Unix systems. (C was created to implement the first Unix systems)
- ▶ A little Assembly is used too:
 - ▶ CPU and machine initialization, exceptions
 - ▶ Critical library routines.
- ▶ No C++ used, see <http://www.tux.org/lkml/#s15-3>
- ▶ All the code compiled with gcc
 - ▶ Many gcc specific extensions used in the kernel code, any ANSI C compiler will not compile the kernel
 - ▶ A few alternate compilers are supported (Intel and Marvell)
 - ▶ See <http://gcc.gnu.org/onlinedocs/gcc-4.9.0/gcc/C-Extensions.html>



No C library

- ▶ The kernel has to be standalone and can't use user space code.
- ▶ User space is implemented on top of kernel services, not the opposite.
- ▶ Kernel code has to supply its own library implementations (string utilities, cryptography, uncompression ...)
- ▶ So, you can't use standard C library functions in kernel code. (`printf()`, `memset()`, `malloc()`, ...).
- ▶ Fortunately, the kernel provides similar C functions for your convenience, like `printk()`, `memset()`, `kmalloc()`, ...



Portability

- ▶ The Linux kernel code is designed to be portable
- ▶ All code outside `arch/` should be portable
- ▶ To this aim, the kernel provides macros and functions to abstract the architecture specific details
 - ▶ Endianness
 - ▶ `cpu_to_be32()`
 - ▶ `cpu_to_le32()`
 - ▶ `be32_to_cpu()`
 - ▶ `le32_to_cpu()`
 - ▶ I/O memory access
 - ▶ Memory barriers to provide ordering guarantees if needed
 - ▶ DMA API to flush and invalidate caches if needed



No floating point computation

- ▶ Never use floating point numbers in kernel code. Your code may be run on a processor without a floating point unit (like on certain ARM CPUs).
- ▶ Don't be confused with floating point related configuration options
 - ▶ They are related to the emulation of floating point operation performed by the user space applications, triggering an exception into the kernel.
 - ▶ Using soft-float, i.e. emulation in user space, is however recommended for performance reasons.



No stable Linux internal API

- ▶ The internal kernel API to implement kernel code can undergo changes between two releases.
- ▶ In-tree drivers are updated by the developer proposing the API change: works great for mainline code.
- ▶ An out-of-tree driver compiled for a given version may no longer compile or work on a more recent one.
- ▶ See `Documentation/stable_api_nonsense.txt` in kernel sources for reasons why.
- ▶ Of course, the kernel to userspace API does not change (system calls, `/proc`, `/sys`), as it would break existing programs.



Kernel memory constraints

- ▶ No memory protection
- ▶ Accessing illegal memory locations result in (often fatal) kernel oopses.
- ▶ Fixed size stack (8 or 4 KB). Unlike in user space, there's no way to make it grow.
- ▶ Kernel memory can't be swapped out (for the same reasons).



Linux kernel licensing constraints

- ▶ The Linux kernel is licensed under the GNU General Public License version 2
 - ▶ This license gives you the right to use, study, modify and share the software freely
- ▶ However, when the software is redistributed, either modified or unmodified, the GPL requires that you redistribute the software under the same license, with the source code
 - ▶ If modifications are made to the Linux kernel (for example to adapt it to your hardware), it is a derivative work of the kernel, and therefore must be released under GPLv2
 - ▶ The validity of the GPL on this point has already been verified in courts
- ▶ However, you're only required to do so
 - ▶ At the time the device starts to be distributed
 - ▶ To your customers, not to the entire world



Proprietary code and the kernel

- ▶ It is illegal to distribute a binary kernel that includes statically compiled proprietary drivers
- ▶ The kernel modules are a gray area: are they derived works of the kernel or not?
 - ▶ The general opinion of the kernel community is that proprietary drivers are bad: <http://j.mp/fbyuuH>
 - ▶ From a legal point of view, each driver is probably a different case
 - ▶ Is it really useful to keep your drivers secret?
- ▶ There are some examples of proprietary drivers, like the Nvidia graphics drivers
 - ▶ They use a wrapper between the driver and the kernel
 - ▶ Unclear whether it makes it legal or not



Advantages of GPL drivers

- ▶ You don't have to write your driver from scratch. You can reuse code from similar free software drivers.
- ▶ You could get free community contributions, support, code review and testing, though this generally only happens with code submitted for the mainline kernel.
- ▶ Your drivers can be freely and easily shipped by others (for example by Linux distributions or embedded Linux build systems).
- ▶ Pre-compiled drivers work with only one kernel version and one specific configuration, making life difficult for users who want to change the kernel version.
- ▶ Legal certainty, you are sure that a GPL driver is fine from a legal point of view.



Advantages of in-tree kernel drivers

- ▶ Once your sources are accepted in the mainline tree, they are maintained by people making changes.
- ▶ Near cost-free maintenance, security fixes and improvements.
- ▶ Easy access to your sources by users.
- ▶ Many more people reviewing your code.



User space device drivers 1/3

- ▶ In some cases, it is possible to implement device drivers in user space!
- ▶ Can be used when
 - ▶ The kernel provides a mechanism that allows userspace applications to directly access the hardware.
 - ▶ There is no need to leverage an existing kernel subsystem such as the networking stack or filesystems.
 - ▶ There is no need for the kernel to act as a “multiplexer” for the device: only one application accesses the device.



User space device drivers 2/3

- ▶ Possibilities for userspace device drivers:
 - ▶ USB with *libusb*, <http://www.libusb.org/>
 - ▶ SPI with *spidev*, [Documentation/spi/spidev](#)
 - ▶ I2C with *i2cdev*, [Documentation/i2c/dev-interface](#)
 - ▶ Memory-mapped devices with *UIO*, including interrupt handling, [Documentation/DocBook/uio-howto/](#)
- ▶ Certain classes of devices (printers, scanners, 2D/3D graphics acceleration) are typically handled partly in kernel space, partly in user space.



User space device drivers 3/3

► Advantages

- No need for kernel coding skills. Easier to reuse code between devices.
- Drivers can be written in any language, even Perl!
- Drivers can be kept proprietary.
- Driver code can be killed and debugged. Cannot crash the kernel.
- Can be swapped out (kernel code cannot be).
- Can use floating-point computation.
- Less in-kernel complexity.
- Potentially higher performance, especially for memory-mapped devices, thanks to the avoidance of system calls.

► Drawbacks

- Less straightforward to handle interrupts.
- Increased interrupt latency vs. kernel code.



Linux sources



Linux sources structure 1/5

- ▶ `arch/<ARCH>`
 - ▶ Architecture specific code
 - ▶ `arch/<ARCH>/mach-<machine>`, machine/board specific code
 - ▶ `arch/<ARCH>/include/asm`, architecture-specific headers
 - ▶ `arch/<ARCH>/boot/dts`, Device Tree source files, for some architectures
- ▶ `block/`
 - ▶ Block layer core
- ▶ `COPYING`
 - ▶ Linux copying conditions (GNU GPL)
- ▶ `CREDITS`
 - ▶ Linux main contributors
- ▶ `crypto/`
 - ▶ Cryptographic libraries



Linux sources structure 2/5

- ▶ `Documentation/`
 - ▶ Kernel documentation. Don't miss it!
- ▶ `drivers/`
 - ▶ All device drivers except sound ones (usb, pci...)
- ▶ `firmware/`
 - ▶ Legacy: firmware images extracted from old drivers
- ▶ `fs/`
 - ▶ Filesystems (`fs/ext3/`, etc.)
- ▶ `include/`
 - ▶ Kernel headers
- ▶ `include/linux/`
 - ▶ Linux kernel core headers
- ▶ `include/uapi/`
 - ▶ User space API headers
- ▶ `init/`
 - ▶ Linux initialization (including `main.c`)
- ▶ `ipc/`
 - ▶ Code used for process communication



Linux sources structure 3/5

- ▶ `Kbuild`
 - ▶ Part of the kernel build system
- ▶ `Kconfig`
 - ▶ Top level description file for configuration parameters
- ▶ `kernel/`
 - ▶ Linux kernel core (very small!)
- ▶ `lib/`
 - ▶ Misc library routines (zlib, crc32...)
- ▶ `MAINTAINERS`
 - ▶ Maintainers of each kernel part. Very useful!
- ▶ `Makefile`
 - ▶ Top Linux Makefile (sets arch and version)
- ▶ `mm/`
 - ▶ Memory management code (small too!)



Linux sources structure 4/5

- ▶ `net/`
 - ▶ Network support code (not drivers)
- ▶ `README`
 - ▶ Overview and building instructions
- ▶ `REPORTING-BUGS`
 - ▶ Bug report instructions
- ▶ `samples/`
 - ▶ Sample code (markers, kprobes, kobjects...)
- ▶ `scripts/`
 - ▶ Scripts for internal or external use
- ▶ `security/`
 - ▶ Security model implementations (SELinux...)
- ▶ `sound/`
 - ▶ Sound support code and drivers
- ▶ `tools/`
 - ▶ Code for various user space tools (mostly C)



Linux sources structure 5/5

- ▶ `usr/`
 - ▶ Code to generate an initramfs cpio archive
- ▶ `virt/`
 - ▶ Virtualization support (KVM)



Kernel source management tools



Cscope

- ▶ Tool to browse source code (mainly C, but also C++ or Java)
- ▶ Supports huge projects like the Linux kernel. Typically takes less than 1 min. to index the whole Linux sources.
- ▶ In Linux kernel sources, two ways of running it:
 - ▶ `cscope -Rk`
All files for all architectures at once
 - ▶ `make cscope`
`cscope -d cscope.out`
Only files for your current architecture
- ▶ Allows searching for a symbol, a definition, functions, strings, files, etc.
- ▶ Integration with editors like `vim` and `emacs`.
- ▶ Dedicated graphical front-end: `KScope`
- ▶ <http://cscope.sourceforge.net/>



Cscope screenshot

```
xterm
C symbol: request_irq

File          Function      Line
0 omap_udc.c   omap_udc_probe 2821 status = request_irq(pdev->resource[1].start, omap_udc_irq,
1 omap_udc.c   omap_udc_probe 2830 status = request_irq(pdev->resource[2].start, omap_udc_pio_irq,
2 omap_udc.c   omap_udc_probe 2838 status = request_irq(pdev->resource[3].start, omap_udc_iso_irq,
3 pxa2xx_udc.c pxa2xx_udc_probe 2517 retval = request_irq(IRQ_USB, pxa2xx_udc_irq,
4 pxa2xx_udc.c pxa2xx_udc_probe 2528 retval = request_irq(LUBBOCK_USB_DISC_IRQ,
5 pxa2xx_udc.c pxa2xx_udc_probe 2539 retval = request_irq(LUBBOCK_USB_IRQ,
6 hc_crisv10.c etrax_usb_hc_init 4423 if (request_irq(ETRAX_USB_HC_IRQ, etrax_usb_hc_interrupt_top_half,
    0,
7 hc_crisv10.c etrax_usb_hc_init 4431 if (request_irq(ETRAX_USB_RX_IRQ, etrax_usb_rx_interrupt, 0,
8 hc_crisv10.c etrax_usb_hc_init 4439 if (request_irq(ETRAX_USB_TX_IRQ, etrax_usb_tx_interrupt, 0,
9 amifb.c       amifb_init      2431 if (request_irq(IRQ_AMIGA_COPPER, amifb_interrupt, 0,
a arcfb.c       arcfb_probe      564 if (request_irq(par->irq, &arcfb_interrupt, SA_SHIRQ,
b atafb.c       atafb_init       2720 request_irq(IRQ_AUTO_4, falcon_vbl_switcher, IRQ_TYPE_PRIO,
c atyfb_base.c  aty_enable_irq   1562 if (request_irq(par->irq, aty_irq, SA_SHIRQ, "atyfb", par)) {

* 155 more lines - press the space bar to display more *
Find this C symbol:
Find this global definition:
Find functions called by this function:
Find functions calling this function:
Find this text string:
Change this text string:
Find this egrep pattern:
Find this file:
Find files #including this file:
```

[Tab]: move the cursor between search results and commands

[Ctrl] [D]: exit cscope



LXR: Linux Cross Reference

- ▶ Generic source indexing tool and code browser
- ▶ Web server based, very easy and fast to use
- ▶ Very easy to find the declaration, implementation or usage of symbols
- ▶ Supports C and C++
- ▶ Supports huge code projects such as the Linux kernel (431 MB of source code in version 3.0).
- ▶ Takes a little time and patience to setup (configuration, indexing, web server configuration)
- ▶ You don't need to set up LXR by yourself. Use our <http://lxr.free-electrons.com> server!
- ▶ <http://sourceforge.net/projects/lxr>



LXR screenshot



Linux Cross Reference

Free Electrons

Embedded Linux Experts

• [Source Navigation](#) • [Diff Markup](#) • [Identifier Search](#) • [Freetext Search](#) •

Version: 2.6.32 2.6.33 2.6.34 2.6.35 2.6.36 2.6.37 2.6.38 2.6.39 3.0 3.1 3.2 3.3 3.4 3.5 3.6 3.7 3.8 3.9 3.10 3.11 3.12

Architecture: x86 arm avr32 blackfin m68k m68knommu microblaze mips powerpc sh

Linux/kernel/user.c

```
1 /*
2  * The "user cache".
3  *
4  * (C) Copyright 1991-2000 Linus Torvalds
5  *
6  * We have a per-user structure to keep track of how many
7  * processes, files etc the user has claimed, in order to be
8  * able to have per-user limits for system resources.
9  */
10
11 #include <linux/init.h>
12 #include <linux/sched.h>
13 #include <linux/slab.h>
14 #include <linux/bitops.h>
15 #include <linux/key.h>
16 #include <linux/interrupt.h>
17 #include <linux/export.h>
18 #include <linux/user_namespace.h>
19 #include <linux/proc_ns.h>
20
21 /*
22  * users count is 1 for root user, 1 for init_uts_ns,
23  * and 1 for... ?
24  */
25 struct user_namespace init_user_ns = {
26     .uid_map = {
27         .nr_extents = 1,
28         .extent[0] = {
29             .first = 0,
30             .lower_first = 0,
31             .count = 4294967295U,
```



Practical lab - Kernel Source Code - Exploring



- ▶ Explore kernel sources manually
- ▶ Use automated tools to explore the source code



Kernel configuration



Kernel configuration and build system

- ▶ The kernel configuration and build system is based on multiple Makefiles
- ▶ One only interacts with the main `Makefile`, present at the **top directory** of the kernel source tree
- ▶ Interaction takes place
 - ▶ using the `make` tool, which parses the Makefile
 - ▶ through various **targets**, defining which action should be done (configuration, compilation, installation, etc.). Run `make help` to see all available targets.
- ▶ Example
 - ▶ `cd linux-3.6.x/`
 - ▶ `make <target>`



Kernel configuration (1)

- ▶ The kernel contains thousands of device drivers, filesystem drivers, network protocols and other configurable items
- ▶ Thousands of options are available, that are used to selectively compile parts of the kernel source code
- ▶ The kernel configuration is the process of defining the set of options with which you want your kernel to be compiled
- ▶ The set of options depends
 - ▶ On your hardware (for device drivers, etc.)
 - ▶ On the capabilities you would like to give to your kernel (network capabilities, filesystems, real-time, etc.)



Kernel configuration (2)

- ▶ The configuration is stored in the `.config` file at the root of kernel sources
 - ▶ Simple text file, `key=value` style
- ▶ As options have dependencies, typically never edited by hand, but through graphical or text interfaces:
 - ▶ `make xconfig`, `make gconfig` (graphical)
 - ▶ `make menuconfig`, `make nconfig` (text)
 - ▶ You can switch from one to another, they all load/save the same `.config` file, and show the same set of options
- ▶ To modify a kernel in a GNU/Linux distribution: the configuration files are usually released in `/boot/`, together with kernel images: `/boot/config-3.2.0-31-generic`



Kernel or module?

- ▶ The **kernel image** is a **single file**, resulting from the linking of all object files that correspond to features enabled in the configuration
 - ▶ This is the file that gets loaded in memory by the bootloader
 - ▶ All included features are therefore available as soon as the kernel starts, at a time where no filesystem exists
- ▶ Some features (device drivers, filesystems, etc.) can however be compiled as **modules**
 - ▶ These are *plugins* that can be loaded/unloaded dynamically to add/remove features to the kernel
 - ▶ Each **module is stored as a separate file in the filesystem**, and therefore access to a filesystem is mandatory to use modules
 - ▶ This is not possible in the early boot procedure of the kernel, because no filesystem is available



Kernel option types

- ▶ There are different types of options
 - ▶ `bool` options, they are either
 - ▶ *true* (to include the feature in the kernel) or
 - ▶ *false* (to exclude the feature from the kernel)
 - ▶ `tristate` options, they are either
 - ▶ *true* (to include the feature in the kernel image) or
 - ▶ *module* (to include the feature as a kernel module) or
 - ▶ *false* (to exclude the feature)
 - ▶ `int` options, to specify integer values
 - ▶ `hex` options, to specify hexadecimal values
 - ▶ `string` options, to specify string values



Kernel option dependencies

- ▶ There are dependencies between kernel options
- ▶ For example, enabling a network driver requires the network stack to be enabled
- ▶ Two types of dependencies
 - ▶ `depends on` dependencies. In this case, option A that depends on option B is not visible until option B is enabled
 - ▶ `select` dependencies. In this case, with option A depending on option B, when option A is enabled, option B is automatically enabled
 - ▶ `make xconfig` allows to see all options, even the ones that cannot be selected because of missing dependencies. In this case, they are displayed in gray




make xconfig

`make xconfig`

- ▶ The most common graphical interface to configure the kernel.
- ▶ Make sure you read
`help -> introduction: useful options!`
- ▶ File browser: easier to load configuration files
- ▶ Search interface to look for parameters
- ▶ Required Debian / Ubuntu packages: `libqt4-dev g++`
(`libqt3-mt-dev` for older kernel releases)



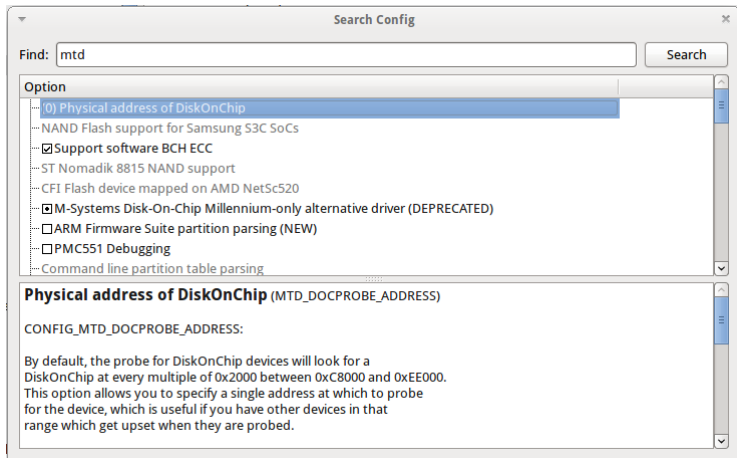
make xconfig screenshot





make xconfig search interface

Looks for a keyword in the parameter name. Allows to select or unselect found parameters.





Kernel configuration options

Compiled as a module (separate file)

`CONFIG_ISO9660_FS=m`

Driver options

`CONFIG_JOLIET=y`

`CONFIG_ZISOFS=y`

Compiled statically into the kernel

`CONFIG_UDF_FS=y`

- ☐ ISO 9660 CDROM file system support
 - ☒ Microsoft Joliet CDROM extensions
 - ☒ Transparent decompression extension
- ☒ UDF file system support



Corresponding .config file excerpt

Options are grouped by sections and are prefixed with `CONFIG_`.


```
#  
# CD-ROM/DVD Filesystems  
#  
CONFIG_ISO9660_FS=m  
CONFIG_JOLIET=y  
CONFIG_ZISOFS=y  
CONFIG_UDF_FS=y  
CONFIG_UDF_NLS=y  
  
#  
# DOS/FAT/NT Filesystems  
#  
# CONFIG_MSDOS_FS is not set  
# CONFIG_VFAT_FS is not set  
CONFIG_NTFS_FS=m  
# CONFIG_NTFS_DEBUG is not set  
CONFIG_NTFS_RW=y
```




make gconfig

make gconfig

- ▶ **GTK** based graphical configuration interface. Functionality similar to that of `make xconfig`.
- ▶ Just lacking a search functionality.
- ▶ Required Debian packages:
`libglade2-dev`





make menuconfig

make menuconfig

- ▶ Useful when no graphics are available. Pretty convenient too!
- ▶ Same interface found in other tools: BusyBox, Buildroot...
- ▶ Required Debian packages:
`libncurses-dev`

```
.config - Linux/arm 3.0.6 Kernel Configuration

System type
Arrow keys navigate the menu. <Enter> selects submenus --->. Highlighted
letters are hotkeys. Pressing <Y> includes, <N> excludes, <M> modularizes
features. Press <Esc><Esc> to exit, <F> for Help, <F> for Search. Legend:
[*] built-in [ ] excluded <M> module < > module capable

[*] MMU-based Paged Memory Management Support
ARM System type (TI OMAP) --->
TI OMAP Common Features --->
TI OMAP2/3/4 Specific Features --->
*** System MMU ***
*** Processor Type ***
Marvell Sheeva CPU Architecture
*** Processor Features ***
[*] Support Thumb user binaries (NEW)
[ ] Inable ThumbEE CPU extension (NEW)
[ ] Run BES kernel on a little endian machine (NEW)
[ ] Disable I-Cache (1-bit) (NEW)
[ ] Disable D-Cache (C-bit) (NEW)
[ ] Disable branch prediction (NEW)
[*] Inable lazy flush for v6 smp (NEW)
[*] stop_machine function can livelock (NEW)
[ ] Spinlocks using LDREX and STREX instructions can livelock (NEW)
[ ] Inable S/W handling for Unaligned Access (NEW)
[*] Inable the L2x0 outer cache controller (NEW)
[ ] ARM errata: Invalidation of the Instruction Cache operation can fail

<Select> <Exit> <Help>
```



make nconfig

make nconfig

- ▶ A newer, similar text interface
- ▶ More user friendly (for example, easier to access help information).
- ▶ Required Debian packages:
`libncurses-dev`

```
.config - Linux/x86_64 3.0.0 Kernel Configuration
Linux/x86_64 3.0.0 Kernel Configuration

General setup --->
[ ] Enable loadable module support --->
[*] Enable the block layer --->
Processor type and features --->
Power management and ACPI options --->
Bus options (PCI etc.) --->
Executable file formats / Emulations --->
[ ] Networking support --->
Device Drivers --->
Firmware Drivers --->
File systems --->
Kernel hacking --->
Security options --->
[ ] Cryptographic API --->
[ ] Virtualization --->
Library routines --->

F1 Help F2 Sys Info F3 Insts F4 Config F5 Back F6 Save F7 Load F8 Sys Search F9 Exit
```



make oldconfig

`make oldconfig`

- ▶ Needed very often!
- ▶ Useful to upgrade a `.config` file from an earlier kernel release
- ▶ Issues warnings for configuration parameters that no longer exist in the new kernel.
- ▶ Asks for values for new parameters

If you edit a `.config` file by hand, it's strongly recommended to run `make oldconfig` afterwards!



Undoing configuration changes

A frequent problem:

- ▶ After changing several kernel configuration settings, your kernel no longer works.
- ▶ If you don't remember all the changes you made, you can get back to your previous configuration:

```
$ cp .config.old .config
```

- ▶ All the configuration interfaces of the kernel (`xconfig`, `menuconfig`, `oldconfig`...) keep this `.config.old` backup copy.



Configuration per architecture

- ▶ The set of configuration options is architecture dependent
 - ▶ Some configuration options are very architecture-specific
 - ▶ Most of the configuration options (global kernel options, network subsystem, filesystems, most of the device drivers) are visible in all architectures.
- ▶ By default, the kernel build system assumes that the kernel is being built for the host architecture, i.e. native compilation
- ▶ The architecture is not defined inside the configuration, but at a higher level
- ▶ We will see later how to override this behaviour, to allow the configuration of kernels for a different architecture



Compiling and installing the kernel for the host system



Kernel compilation

- ▶ `make`
 - ▶ in the main kernel source directory
 - ▶ Remember to run multiple jobs in parallel if you have multiple CPU cores. Example: `make -j 4`
 - ▶ No need to run as root!
- ▶ Generates
 - ▶ `vmlinux`, the raw uncompressed kernel image, at the ELF format, useful for debugging purposes, but cannot be booted
 - ▶ `arch/<arch>/boot/*Image`, the final, usually compressed, kernel image that can be booted
 - ▶ `bzImage` for x86, `zImage` for ARM, `vmImage.gz` for Blackfin, etc.
 - ▶ `arch/<arch>/boot/dts/*.dtb`, compiled Device Tree files (on some architectures)
 - ▶ All kernel modules, spread over the kernel source tree, as `.ko` files.



Kernel installation

- ▶ `make install`

- ▶ Does the installation for the host system by default, so needs to be run as root. Generally not used when compiling for an embedded system, and it installs files on the development workstation.

- ▶ Installs

- ▶ `/boot/vmlinuz-<version>`

Compressed kernel image. Same as the one in `arch/<arch>/boot`

- ▶ `/boot/System.map-<version>`

Stores kernel symbol addresses

- ▶ `/boot/config-<version>`

Kernel configuration for this version

- ▶ Typically re-runs the bootloader configuration utility to take the new kernel into account.



Module installation

- ▶ `make modules_install`
 - ▶ Does the installation for the host system by default, so needs to be run as root
- ▶ Installs all modules in `/lib/modules/<version>/`
 - ▶ `kernel/`
Module `.ko` (Kernel Object) files, in the same directory structure as in the sources.
 - ▶ `modules.alias`
Module aliases for module loading utilities. Example line:
`alias sound-service-?-0 snd_mixer_oss`
 - ▶ `modules.dep`
Module dependencies
 - ▶ `modules.symbols`
Tells which module a given symbol belongs to.



Kernel cleanup targets

- ▶ Clean-up generated files (to force re-compilation):
`make clean`
- ▶ Remove all generated files. Needed when switching from one architecture to another. Caution: it also removes your `.config` file!
`make mrproper`
- ▶ Also remove editor backup and patch reject files (mainly to generate patches):
`make distclean`





Cross-compiling the kernel



Cross-compiling the kernel

When you compile a Linux kernel for another CPU architecture

- ▶ Much faster than compiling natively, when the target system is much slower than your GNU/Linux workstation.
- ▶ Much easier as development tools for your GNU/Linux workstation are much easier to find.
- ▶ To make the difference with a native compiler, cross-compiler executables are prefixed by the name of the target system, architecture and sometimes library. Examples:

`mips-linux-gcc`, the prefix is `mips-linux-`

`arm-linux-gnueabi-gcc`, the prefix is `arm-linux-gnueabi-`



Specifying cross-compilation (1)

The CPU architecture and cross-compiler prefix are defined through the `ARCH` and `CROSS_COMPILE` variables in the toplevel Makefile.

- ▶ `ARCH` is the name of the architecture. It is defined by the name of the subdirectory in `arch/` in the kernel sources
 - ▶ Example: `arm` if you want to compile a kernel for the `arm` architecture.
- ▶ `CROSS_COMPILE` is the prefix of the cross compilation tools
 - ▶ Example: `arm-linux-` if your compiler is `arm-linux-gcc`



Specifying cross-compilation (2)

Two solutions to define `ARCH` and `CROSS_COMPILE`:

- ▶ Pass `ARCH` and `CROSS_COMPILE` on the `make` command line:

```
make ARCH=arm CROSS_COMPILE=arm-linux- ...
```

Drawback: it is easy to forget to pass these variables when you run any `make` command, causing your build and configuration to be screwed up.

- ▶ Define `ARCH` and `CROSS_COMPILE` as environment variables:

```
export ARCH=arm
```

```
export CROSS_COMPILE=arm-linux-
```

Drawback: it only works inside the current shell or terminal. You could put these settings in a file that you source every time you start working on the project. If you only work on a single architecture with always the same toolchain, you could even put these settings in your `~/.bashrc` file to make them permanent and visible from any terminal.



Predefined configuration files

- ▶ Default configuration files available, per board or per-CPU family
 - ▶ They are stored in `arch/<arch>/configs/`, and are just minimal `.config` files
 - ▶ This is the most common way of configuring a kernel for embedded platforms
- ▶ Run `make help` to find if one is available for your platform
- ▶ To load a default configuration file, just run `make acme_defconfig`
 - ▶ This will overwrite your existing `.config` file!
- ▶ To create your own default configuration file
 - ▶ `make savedefconfig`, to create a minimal configuration file
 - ▶ `mv defconfig arch/<arch>/configs/myown_defconfig`



Configuring the kernel

- ▶ After loading a default configuration file, you can adjust the configuration to your needs with the normal `xconfig`, `gconfig` or `menuconfig` interfaces
- ▶ You can also start the configuration from scratch without loading a default configuration file
- ▶ As the architecture is different from your host architecture
 - ▶ Some options will be different from the native configuration (processor and architecture specific options, specific drivers, etc.)
 - ▶ Many options will be identical (filesystems, network protocols, architecture-independent drivers, etc.)
- ▶ Make sure you have the support for the right CPU, the right board and the right device drivers.



Device Tree

- ▶ Many embedded architectures have a lot of non-discoverable hardware.
- ▶ Depending on the architecture, such hardware is either described using C code directly within the kernel, or using a special hardware description language in a *Device Tree*.
- ▶ ARM, PowerPC, OpenRISC, ARC, Microblaze are examples of architectures using the Device Tree.
- ▶ A *Device Tree Source*, written by kernel developers, is compiled into a binary *Device Tree Blob*, passed at boot time to the kernel.
 - ▶ There is one different Device Tree for each board/platform supported by the kernel, available in `arch/arm/boot/dts/<board>.dtb`.
- ▶ The bootloader must load both the kernel image and the Device Tree Blob in memory before starting the kernel.



Building and installing the kernel

- ▶ Run `make`
- ▶ Copy the final kernel image to the target storage
 - ▶ can be `uImage`, `zImage`, `vmlinux`, `bzImage` in `arch/<arch>/boot`
 - ▶ copying the Device Tree Blob might be necessary as well, they are available in `arch/<arch>/boot/dts`
- ▶ `make install` is rarely used in embedded development, as the kernel image is a single file, easy to handle
 - ▶ It is however possible to customize the make install behaviour in `arch/<arch>/boot/install.sh`
- ▶ `make modules_install` is used even in embedded development, as it installs many modules and description files
 - ▶ `make INSTALL_MOD_PATH=<dir>/ modules_install`
 - ▶ The `INSTALL_MOD_PATH` variable is needed to install the modules in the target root filesystem instead of your host root filesystem.



Booting with U-Boot

- ▶ Recent versions of U-Boot can boot the `zImage` binary.
- ▶ Older versions require a special kernel image format: `uImage`
 - ▶ `uImage` is generated from `zImage` using the `mkimage` tool. It is done automatically by the kernel `make uImage` target.
 - ▶ On some ARM platforms, `make uImage` requires passing a `LOADADDR` environment variable, which indicates at which physical memory address the kernel will be executed.
- ▶ In addition to the kernel image, U-Boot can also pass a *Device Tree Blob* to the kernel.
- ▶ The typical boot process is therefore:
 1. Load `zImage` or `uImage` at address `X` in memory
 2. Load `<board>.dtb` at address `Y` in memory
 3. Start the kernel with `bootz X - Y` or `bootm X - Y`
The `-` in the middle indicates no *initramfs*



Kernel command line

- ▶ In addition to the compile time configuration, the kernel behaviour can be adjusted with no recompilation using the **kernel command line**
- ▶ The kernel command line is a string that defines various arguments to the kernel
 - ▶ It is very important for system configuration
 - ▶ `root=` for the root filesystem (covered later)
 - ▶ `console=` for the destination of kernel messages
 - ▶ Many more exist. The most important ones are documented in `Documentation/kernel-parameters.txt` in kernel sources.
- ▶ This kernel command line is either
 - ▶ Passed by the bootloader. In U-Boot, the contents of the `bootargs` environment variable is automatically passed to the kernel
 - ▶ Built into the kernel, using the `CONFIG_CMDLINE` option.



Practical lab - Kernel compiling and booting



1st lab: board and bootloader setup:

- ▶ Prepare the board and access its serial port
- ▶ Configure its bootloader to use TFTP

2nd lab: kernel compiling and booting:

- ▶ Set up a cross-compiling environment
- ▶ Cross-compile a kernel for an ARM target platform
- ▶ Boot this kernel from a directory on your workstation, accessed by the board through NFS



Using kernel modules



Advantages of modules

- ▶ Modules make it easy to develop drivers without rebooting: load, test, unload, rebuild, load...
- ▶ Useful to keep the kernel image size to the minimum (essential in GNU/Linux distributions for PCs).
- ▶ Also useful to reduce boot time: you don't spend time initializing devices and kernel features that you only need later.
- ▶ Caution: once loaded, have full control and privileges in the system. No particular protection. That's why only the `root` user can load and unload modules.



Module dependencies

- ▶ Some kernel modules can depend on other modules, which need to be loaded first.
- ▶ Example: the `usb-storage` module depends on the `scsi_mod`, `libusual` and `usbcore` modules.
- ▶ Dependencies are described in
`/lib/modules/<kernel-version>/modules.dep`
This file is generated when you run `make modules_install`.



Kernel log

When a new module is loaded, related information is available in the kernel log.

- ▶ The kernel keeps its messages in a circular buffer (so that it doesn't consume more memory with many messages)
- ▶ Kernel log messages are available through the `dmesg` command (**d**iagnostics **m**essage)
- ▶ Kernel log messages are also displayed in the system console (console messages can be filtered by level using the `loglevel` kernel parameter, or completely disabled with the `quiet` parameter).
- ▶ Note that you can write to the kernel log from user space too:

```
echo "<n>Debug info" > /dev/kmsg
```



Module utilities (1)

- ▶ `modinfo <module_name>`

`modinfo <module_path>.ko`

Gets information about a module: parameters, license, description and dependencies.

Very useful before deciding to load a module or not.

- ▶ `sudo insmod <module_path>.ko`

Tries to load the given module. The full path to the module object file must be given.



Understanding module loading issues

- ▶ When loading a module fails, `insmod` often doesn't give you enough details!
- ▶ Details are often available in the kernel log.
- ▶ Example:

```
$ sudo insmod ./intr_monitor.ko
insmod: error inserting './intr_monitor.ko': -1 Device or resource busy
$ dmesg
[17549774.552000] Failed to register handler for irq channel 2
```



Module utilities (2)

- ▶ `sudo modprobe <module_name>`

Most common usage of `modprobe`: tries to load all the modules the given module depends on, and then this module. Lots of other options are available. `modprobe` automatically looks in `/lib/modules/<version>/` for the object file corresponding to the given module name.

- ▶ `lsmod`

Displays the list of loaded modules

Compare its output with the contents of `/proc/modules`!



Module utilities (3)

- ▶ `sudo rmmod <module_name>`

Tries to remove the given module.

Will only be allowed if the module is no longer in use (for example, no more processes opening a device file)

- ▶ `sudo modprobe -r <module_name>`

Tries to remove the given module and all dependent modules (which are no longer needed after removing the module)



Passing parameters to modules

- ▶ Find available parameters:

```
modinfo snd-intel8x0m
```

- ▶ Through `insmod`:

```
sudo insmod ./snd-intel8x0m.ko index=-2
```

- ▶ Through `modprobe`:

Set parameters in `/etc/modprobe.conf` or in any file in `/etc/modprobe.d/`:

```
options snd-intel8x0m index=-2
```

- ▶ Through the kernel command line, when the driver is built statically into the kernel:

```
snd-intel8x0m.index=-2
```

- ▶ `snd-intel8x0m` is the *driver name*
- ▶ `index` is the *driver parameter name*
- ▶ `-2` is the *driver parameter value*



Check module parameter values


How to find the current values for the parameters of a loaded module?

- ▶ Check `/sys/module/<name>/parameters`.
- ▶ There is one file per parameter, containing the parameter value.



Linux Kernel in a Nutshell, Dec 2006

- ▶ By Greg Kroah-Hartman, O'Reilly
<http://www.kroah.com/1kn/>
- ▶ A good reference book and guide on configuring, compiling and managing the Linux kernel sources.
- ▶ Freely available on-line!
Great companion to the printed book for easy electronic searches!
Available as single PDF file on
<http://free-electrons.com/community/kernel/1kn/>
- ▶ Our rating: 2 stars






Developing Kernel Modules

Free Electrons

© Copyright 2004-2014, Free Electrons.
Creative Commons BY-SA 3.0 license.
Corrections, suggestions, contributions and translations are welcome!





Hello Module 1/2

```
/* hello.c */
#include <linux/init.h>
#include <linux/module.h>
#include <linux/kernel.h>

static int __init hello_init(void)
{
    pr_alert("Good morrow to this fair assembly.\n");
    return 0;
}

static void __exit hello_exit(void)
{
    pr_alert("Alas, poor world, what treasure hast thou lost!\n");
}

module_init(hello_init);
module_exit(hello_exit);
MODULE_LICENSE("GPL");
MODULE_DESCRIPTION("Greeting module");
MODULE_AUTHOR("William Shakespeare");
```



Hello Module 2/2

- ▶ `__init`
 - ▶ removed after initialization (static kernel or module.)
- ▶ `__exit`
 - ▶ discarded when module compiled statically into the kernel, or when module unloading support is not enabled.
- ▶ Example available on
<http://git.free-electrons.com/training-materials/plain/code/hello/hello.c>



Hello Module Explanations

- ▶ Headers specific to the Linux kernel: `linux/xxx.h`
 - ▶ No access to the usual C library, we're doing kernel programming
- ▶ An initialization function
 - ▶ Called when the module is loaded, returns an error code (0 on success, negative value on failure)
 - ▶ Declared by the `module_init()` macro: the name of the function doesn't matter, even though `<modulename>_init()` is a convention.
- ▶ A cleanup function
 - ▶ Called when the module is unloaded
 - ▶ Declared by the `module_exit()` macro.
- ▶ Metadata information declared using `MODULE_LICENSE()`, `MODULE_DESCRIPTION()` and `MODULE_AUTHOR()`




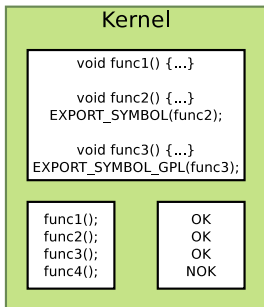


Symbols Exported to Modules 1/2

- ▶ From a kernel module, only a limited number of kernel functions can be called
- ▶ Functions and variables have to be explicitly exported by the kernel to be visible to a kernel module
- ▶ Two macros are used in the kernel to export functions and variables:
 - ▶ `EXPORT_SYMBOL(symbolname)`, which exports a function or variable to all modules
 - ▶ `EXPORT_SYMBOL_GPL(symbolname)`, which exports a function or variable only to GPL modules
- ▶ A normal driver should not need any non-exported function.



Symbols exported to modules 2/2





Module License

- ▶ Several usages
 - ▶ Used to restrict the kernel functions that the module can use if it isn't a GPL licensed module
 - ▶ Difference between `EXPORT_SYMBOL()` and `EXPORT_SYMBOL_GPL()`
 - ▶ Used by kernel developers to identify issues coming from proprietary drivers, which they can't do anything about ("Tainted" kernel notice in kernel crashes and oopses).
 - ▶ Useful for users to check that their system is 100% free (check `/proc/sys/kernel/tainted`)
- ▶ Values
 - ▶ GPL compatible (see `include/linux/license.h`: GPL, GPL v2, GPL and additional rights, Dual MIT/GPL, Dual BSD/GPL, Dual MPL/GPL)
 - ▶ Proprietary



Compiling a Module

- ▶ Two solutions
 - ▶ *Out of tree*
 - ▶ When the code is outside of the kernel source tree, in a different directory
 - ▶ Advantage: Might be easier to handle than modifications to the kernel itself
 - ▶ Drawbacks: Not integrated to the kernel configuration/compilation process, needs to be built separately, the driver cannot be built statically
 - ▶ Inside the kernel tree
 - ▶ Well integrated into the kernel configuration/compilation process
 - ▶ Driver can be built statically if needed



Compiling an out-of-tree Module 1/2

- ▶ The below `Makefile` should be reusable for any single-file out-of-tree Linux module
- ▶ The source file is `hello.c`
- ▶ Just run `make` to build the `hello.ko` file

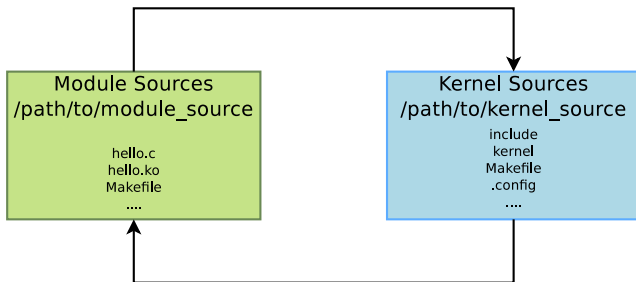
```
ifneq ($(KERNELRELEASE),)
obj-m := hello.o
else
KDIR := /path/to/kernel/sources

all:
<tab>$(MAKE) -C $(KDIR) M=$$PWD
endif
```

- ▶ For `KDIR`, you can either set:
 - ▶ full kernel source directory
(configured + `make modules_prepare`)
 - ▶ or just kernel headers directory (`make headers_install`)



Compiling an out-of-tree Module 2/2



- ▶ The module Makefile is interpreted with `KERNELRELEASE` undefined, so it calls the kernel Makefile, passing the module directory in the `M` variable
- ▶ The kernel Makefile knows how to compile a module, and thanks to the `M` variable, knows where the Makefile for our module is. The module Makefile is interpreted with `KERNELRELEASE` defined, so the kernel sees the `obj-m` definition.



Modules and Kernel Version

- ▶ To be compiled, a kernel module needs access to the kernel headers, containing the definitions of functions, types and constants.
- ▶ Two solutions
 - ▶ Full kernel sources
 - ▶ Only kernel headers (`linux-headers-*` packages in Debian/Ubuntu distributions)
- ▶ The sources or headers must be configured
 - ▶ Many macros or functions depend on the configuration
- ▶ A kernel module compiled against version X of kernel headers will not load in kernel version Y
 - ▶ `modprobe` / `insmod` will say `Invalid module format`



New Driver in Kernel Sources 1/2

- ▶ To add a new driver to the kernel sources:
 - ▶ Add your new source file to the appropriate source directory.
Example: `drivers/usb/serial/navman.c`
 - ▶ Single file drivers in the common case, even if the file is several thousand lines of code big. Only really big drivers are split in several files or have their own directory.
 - ▶ Describe the configuration interface for your new driver by adding the following lines to the `Kconfig` file in this directory:

```
config USB_SERIAL_NAVMAN
    tristate "USB Navman GPS device"
    depends on USB_SERIAL
    help
        To compile this driver as a module, choose M
        here: the module will be called navman.
```



New Driver in Kernel Sources 2/2

- ▶ Add a line in the `Makefile` file based on the `Kconfig` setting:
`obj-$(CONFIG_USB_SERIAL_NAVMAN) += navman.o`
- ▶ It tells the kernel build system to build `navman.c` when the `USB_SERIAL_NAVMAN` option is enabled. It works both if compiled statically or as a module.
 - ▶ Run `make xconfig` and see your new options!
 - ▶ Run `make` and your new files are compiled!
 - ▶ See `Documentation/kbuild/` for details and more elaborate examples like drivers with several source files, or drivers in their own subdirectory, etc.



Hello Module with Parameters 1/2

```
/* hello_param.c */
#include <linux/init.h>
#include <linux/module.h>

MODULE_LICENSE("GPL");

/* A couple of parameters that can be passed in: how many
   times we say hello, and to whom */

static char *whom = "world";
module_param(whom, charp, 0);

static int howmany = 1;
module_param(howmany, int, 0);
```



Hello Module with Parameters 2/2

```
static int __init hello_init(void)
{
    int i;
    for (i = 0; i < howmany; i++)
        pr_alert("(d) Hello, %s\n", i, whom);
    return 0;
}

static void __exit hello_exit(void)
{
    pr_alert("Goodbye, cruel %s\n", whom);
}

module_init(hello_init);
module_exit(hello_exit);
```

Thanks to Jonathan Corbet for the example!

Source code available on:

http://git.free-electrons.com/training-materials/plain/code/hello-param/hello_param.c



Declaring a module parameter

```
module_param(  
    name, /* name of an already defined variable */  
    type, /* either byte, short, ushort, int, uint, long, ulong,  
           charp, bool or invbool. (checked at run time!) */  
    perm /* for /sys/module/<module_name>/parameters/<param>,  
           0: no such module parameter value file */  
);
```

```
/* Example */  
static int irq=5;  
module_param(irq, int, S_IRUGO);
```

Modules parameter arrays are also possible with
`module_param_array()`.



Practical lab - Writing Modules




- ▶ Create, compile and load your first module
- ▶ Add module parameters
- ▶ Access kernel internals from your module



Useful general-purpose kernel APIs

Free Electrons

© Copyright 2004-2014, Free Electrons.
Creative Commons BY-SA 3.0 license.
Corrections, suggestions, contributions and translations are welcome!





Memory/string utilities

- ▶ In `linux/string.h`
 - ▶ Memory-related: `memset()`, `memcpy()`, `memmove()`, `memscan()`, `memcmp()`, `memchr()`
 - ▶ String-related: `strcpy()`, `strcat()`, `strcmp()`, `strchr()`, `strrchr()`, `strlen()` and variants
 - ▶ Allocate and copy a string: `kstrdup()`, `kstrndup()`
 - ▶ Allocate and copy a memory area: `kmemdup()`
- ▶ In `linux/kernel.h`
 - ▶ String to int conversion: `simple_strtoul()`, `simple_strtol()`, `simple_strtoull()`, `simple_strtoll()`
 - ▶ Other string functions: `sprintf()`, `sscanf()`



Linked lists

- ▶ Convenient linked-list facility in `linux/list.h`
 - ▶ Used in thousands of places in the kernel
- ▶ Add a `struct list_head` member to the structure whose instances will be part of the linked list. It is usually named `node` when each instance needs to only be part of a single list.
- ▶ Define the list with the `LIST_HEAD()` macro for a global list, or define a `struct list_head` element and initialize it with `INIT_LIST_HEAD()` for lists embedded in a structure.
- ▶ Then use the `list_*`() API to manipulate the list
 - ▶ Add elements: `list_add()`, `list_add_tail()`
 - ▶ Remove, move or replace elements: `list_del()`, `list_move()`, `list_move_tail()`, `list_replace()`
 - ▶ Test the list: `list_empty()`
 - ▶ Iterate over the list: `list_for_each_*`() family of macros



Linked Lists Examples (1)

```
► From include/linux/atmel_tc.h

/*
 * Definition of a list element, with a
 * struct list_head member
 */
struct atmel_tc
{
    /* some members */
    struct list_head node;
};
```



Linked Lists Examples (2)

► From `drivers/misc/atmel_tclib.c`

```
/* Define the global list */
```

```
static LIST_HEAD(tc_list);
```

```
static int __init tc_probe(struct platform_device *pdev) {
```

```
    struct atmel_tc *tc;
```

```
    tc = kzalloc(sizeof(struct atmel_tc), GFP_KERNEL);
```

```
    /* Add an element to the list */
```

```
    list_add_tail(&tc->node, &tc_list);
```

```
}
```

```
struct atmel_tc *atmel_tc_alloc(unsigned block, const char *name)
```

```
{
```

```
    struct atmel_tc *tc;
```

```
    /* Iterate over the list elements */
```

```
    list_for_each_entry(tc, &tc_list, node) {
```

```
        /* Do something with tc */
```

```
    }
```

```
    [...]
```

```
}
```