



**Linaro
connect**
San Francisco 2017

(e)MMC/SD/SDIO State of Affairs

Linus Walleij
Ulf Hansson



MMC (MultiMedia)



SD (Secure Digital)



eMMC (embedded-)



SDIO (SD I/O)



First MMC stack by Russell King in kernel 2.6.9

<rmk@flint.arm.linux.org.uk>

[MMC] Add MMC core support

This patch adds core support to the Linux kernel for driving MMC interfaces found on embedded devices, such as found in the Intel PXA and ARM MMCI primecell. This patch does not add support for SD or SDIO cards.

It is vaguely based upon the handhelds.org MMC code, but the bulk of the core has been rewritten from scratch.



ENGINEERS AND DEVICES
WORKING TOGETHER

Lean, Mean and Clean

```
$ ls -al
total 108
 280 Oct 18 2004 .
 960 Oct 18 2004 ..
1346 Oct 18 2004 Kconfig
 292 Oct 18 2004 Makefile
10690 Oct 18 2004 mmc_block.c
20362 Oct 18 2004 mmc.c
 510 Oct 18 2004 mmc.h
14611 Oct 18 2004 mmci.c
 4912 Oct 18 2004 mmci.h
 5298 Oct 18 2004 mmc_queue.c
 832 Oct 18 2004 mmc_queue.h
 5695 Oct 18 2004 mmc_sysfs.c
13644 Oct 18 2004 pxamci.c
 2184 Oct 18 2004 pxamci.h
```



Everything Grew by Feature Creep



Block Device Usage Violation

```
dev ls -al mmc*
```

```
brw-rw---- 1 root root 179, 0 Jan 1 02:11 mmcblk0
brw-rw---- 1 root root 179, 1 Jan 1 02:11 mmcblk0p1
brw-rw---- 1 root root 179, 2 Jan 1 02:11 mmcblk0p2
brw-rw---- 1 root root 179, 5 Jan 1 02:11 mmcblk0p5
brw-rw---- 1 root root 179, 8 Jan 1 02:11 mmcblk3
brw-rw---- 1 root root 179, 16 Jan 1 02:11 mmcblk3boot0
brw-rw---- 1 root root 179, 24 Jan 1 02:11 mmcblk3boot1
brw-rw---- 1 root root 179, 9 Jan 1 02:11 mmcblk3p1
brw-rw---- 1 root root 179, 32 Jan 1 02:11 mmcblk3rpbm
```

Backup?

```
dd if=/dev/mmcblk3 of=/tmp/card.img
```



The BIG MMC LOCK

```
int __mmc_claim_host(struct mmc_host *host, atomic_t
*abort)
{
    DECLARE_WAITQUEUE(wait, current);
    unsigned long flags;
    int stop;
    bool pm = false;

    might_sleep();

    add_wait_queue(&host->wq, &wait);
    spin_lock_irqsave(&host->lock, flags);
    while (1) {
        set_current_state(TASK_UNINTERRUPTIBLE);
        stop = abort ? atomic_read(abort) : 0;
        if (stop || !host->claimed ||
host->claimer == current)
            break;
        spin_unlock_irqrestore(&host->lock,
flags);
    }
}
```

```
schedule();
        spin_lock_irqsave(&host->lock, flags);
    }
    set_current_state(TASK_RUNNING);
    if (!stop) {
        host->claimed = 1;
        host->claimer = current;
        host->claim_cnt += 1;
        if (host->claim_cnt == 1)
            pm = true;
    } else
        wake_up(&host->wq);

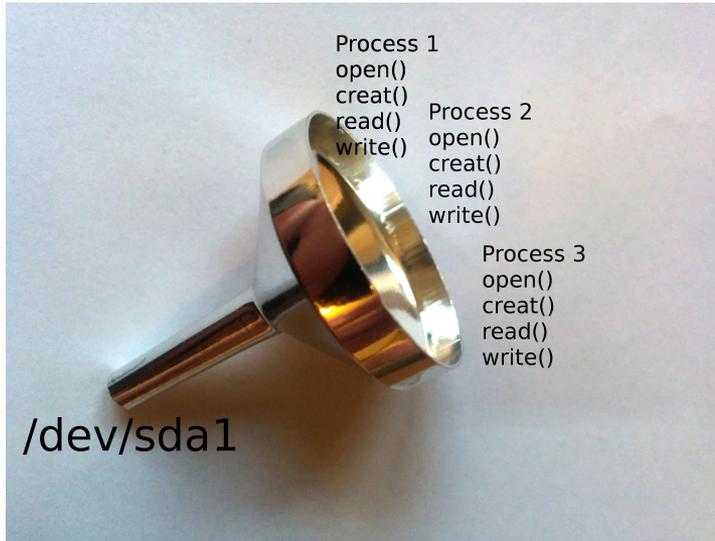
    spin_unlock_irqrestore(&host->lock, flags);
    remove_wait_queue(&host->wq, &wait);

    if (pm)
        pm_runtime_get_sync(mmc_dev(host));

    return stop;
}
EXPORT_SYMBOL(__mmc_claim_host);
```



Add Multiqueue (blk-mq) Support



Steps to Success

- Get rid of all uses of the "big MMC lock":
 1. Move RPMB partitions to a character device - not really a block device - working with the Intel people on generic RPMB support too also helps their case, patch exists today
 2. Move boot partitions and generic partitions to become Linux partitions
Upside: makes it possible to back up the whole mmcblkN (except RPMB) with the dd command
 - ∨ Alt 1: do not make the partition core aware of these partitions "special status", just move them to the end of the default partition, give them bogus block numbers beyond the end of the default partition back there, register partitions directly with the block layer at this place
Downside: it is not really true
 - ∨ Alt 2: do a "deep awareness" of the special partitions in the block layer partition code, so it is not just addressing "some blocks somewhere on a device" but actually handling addressed-from-0 special partitions
Upside
 4. Delete dual-mode card support, we only support eMMC/SD or SDIO not both at the same time.
 5. Split eMMC/SD block access and SDIO host access in two worlds
 6. The block layer can use a smaller lock, like a mutex or semaphore for its use cases alone
 7. Voila: no big MMC host lock



Other Developments in MMC: Command Queuing





**Linaro
connect**
San Francisco 2017

Thank You

#SFO17

SFO17 keynotes and videos on: connect.linaro.org

For further information: www.linaro.org

