Something about NAND

Andrew Turner
andrew@FreeBSD.org

FreeBSD Developer Summit
Meeting Plaza
Maarssen, The Netherlands
October 6 – 8, 2011

NAND Flash structure (1)

- Split into blocks and pages
- Erase blocks, read/write pages
- Pages split into main area and extra area (OOB)
- For example:
  - 8640 byte page (8192 main bytes + 448 OOB bytes)
  - 256 pages per block
  - 2048 blocks on device

FreeBSD

NAND Flash structure (2)

- Bits can spontaneously fail
- Need ECC to detect and correct failed bits
- Blocks have a limited number of erase cycles
- Need a bad block table or mark blocks as bad

Historical Drivers

- John Birrell's Driver
  - Written in 2003/2004
  - Host and device halves in a single driver
  - Only works with NAND flash attached to the memory bus
- Andrew Turner's Driver
  - Initially based on John Birrell's Driver, later rewritten
  - Split into host and device halves with a NAND simulator
- Marcel Moolenaar's Driver
  - Only a NAND flash simulator

FreeBSD

NAND Framework

Contains 4 layers

- NAND Flash Controller – CPU interface
  - e.g. Marvell, S3C2410, ...
- NANDbus – Bus Access
- NAND Chip – NAND command set
  - e.g. ONFi, Samsung large page, Samsung small page
- GEOM device

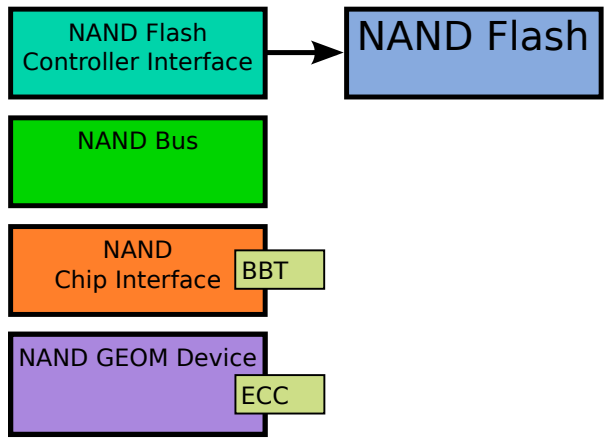

FreeBSD

NAND Filesystem

We need one

There are rumors of one in development.

FreeBSD

Changes I would like

- Rework ECC and bad block handling
- Pass both block and sector size into GEOM
- Flash transition framework

Future considerations

- Number of bits ECC needs to fix is increasing
  Need to correct 24+ bits per page, new ECC schemes
- EZ-NAND – NAND flash with built in ECC (ECC Zero)
- High-Speed NAND flash – 166+ MT/s
- Multiple planes and logical units (LUN) to increase parallelism

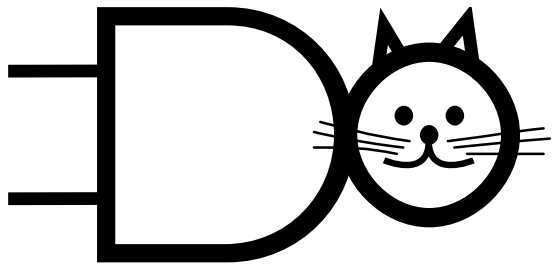

Questions?

FreeBSD