Board File to Device Tree Migration
A War Story

Pantelis Antoniou
panto@antoniou-consulting.com
panto@irc.freenode.net / #beagle

Antoniou
Consulting
Embedded Linux Experts

Or how I stopped worrying and learned to love the Device Tree (grudgingly)

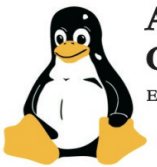

Antoniou Consulting
Embedded Linux Experts

This was originally a BoF

Audience input is welcomed!
Tell us your horror stories!
We can relate!

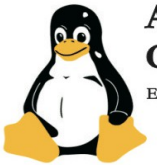

Antoniou Consulting
Embedded Linux Experts

In The Beginning

- There was the original beaglebone.
- Cheap ARM Cortex-A8 OSHW board.
- Linux kernel based on a 3.2 derived heavily patched kernel
- git log --pretty=oneline v3.2..linux-ti33x-psp-3.2.34-r18a+gitr720e07b4c1f687b61b147b31c698cb6816d72f01 | wc -l
  14895
- Pretty hard to submit to the mainline.
- And there came the new beaglebone black.

Antoniou Consulting
Embedded Linux Experts

Trouble in paradise

- Original TI kernel uses a board specific file: `board-am335xevm.c`
- 4K+ lines
- Supports a number of TI am335x boards (note the name).
- Very hard to do anything new and not affect the other boards.

Antoniou Consulting
Embedded Linux Experts

Trouble in paradise (cont)

```
beaglebone_cape_setup()


if (!strncmp("BB-BONE-LCD7-01", cape_config.partnumber, 15)) {
        pr_info("BeagleBone cape: initializing LCD cape\n");
        if (!strncmp("00A2", cape_config.version, 4)) {
                setup_pin_mux(bbtoys7a2_pin_mux);
                gpio_request(BEAGLEBONE_LCD7A2_AVDD_EN,
                             "BONE_LCD_AVDD_EN");
                gpio_direction_output(BEAGLEBONE_LCD7A2_AVDD_EN, 1);
```

Antoniou Consulting
Embedded Linux Experts

The (benign) Dictator has spoken

- http://article.gmane.org/gmane.linux.ports.arm.omap/55060
- “Gaah. Guys, this whole ARM thing is a f*cking pain in the ass.”
- “Somebody needs to get a grip in the ARM community.”
- “Somebody in the ARM community really needs to step up and tell people to stop dicking around.”

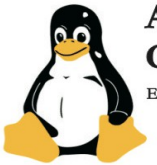

No More New Board Files

Antoniou Consulting
Embedded Linux Experts

The Promised Land

- *"The Device Tree is a data structure for describing hardware. Rather than hard coding every detail of a device into an operating system, many aspects of the hardware can be described in a data structure that is passed to the operating system at boot time."*

The Journey Through the Desert

- *Bye Bye Board File.*
- Board described by foo.dts
  - dts=Device Tree Source
- dtc foo.dts -> foo.dtb
  - dtb=Device Tree binary Blob
- Kernel passed both kernel+foo.dtb
- Kernel unflattens foo.dtb and populates the device busses.
- No More Platform Data to configure drivers!

The Device Tree Commandments

- Thou shalt not use platform data.
- Thou shalt use DT nodes in the boot DTB to configure your board drivers.
- Thou shalt not use board/arch specific drivers in favor of general purpose frameworks
- Thou shalt rework platform data callbacks in a general purpose notification mechanism.

DT is purely data.

Antoniou Consulting
Embedded Linux Experts

What is so special about the BB(B)

- Capes, small add on boards that can stack.
- Direct connection to the (very) muxed pins of the am335x SoC.
- BB Black has to be compatible with capes that re-use the pins of HDMI + eMMC.
- Auto cape detection via an on board EEPROM.
- Device Tree Overlays + capemgr.
- Most advanced usage of Device Tree in a shipping Linux product to date.

Let's talk about The Driver Problem

- Silicon Vendors Hate Updating Drivers
- Significant amount of drivers had no DT bindings.
- Significant amount of drivers had no working DT bindings.
- Some DT bindings just don't make sense.
- No central kernel tree with all the bits for a given SoC.

Antoniou Consulting
Embedded Linux Experts

Booting mainline

- Significant difference in what a customer/user considers 'booting' and what the vendor provides.
- Vendor
  - Boots with serial + initrd
- User
  - All the peripherals (that I use) work.

The Catch 22 of a Mainline kernel

- Vendors can not deal with churn
  - No stable in-kernel interfaces – on purpose
  - No monetary incentive
  - Not enough developers to both support the new SoCs in the pipeline and update the old SoC kernel support to mainline
- Users/community can not deal with old kernel versions
  - Most distros require mainline support
  - Keeping track of thousand of patches undesirable
  - It is just no fun working on old software

Antoniou Consulting
Embedded Linux Experts

Device Tree Criticism

- Device Tree has been sold to the community as a cure for the ARM kernel mess.
- It has failed to clean the mess but:
  - The problem is not technical.
  - Transition handled poorly.
  - Maintainer hostility (not eating your own dogfood)
  - Structure in a place where there was none; bound to draw fire from the old guard.

Device Tree Advantages on BB

- It works – shipping product.
- Hardware hacking without having to:
  - Recompile the kernel.
  - Modify any kernel source file.
  - Know the internals of a driver in order to configure it.
  - DT syntax is (quirky but) human-readable.
  - Same mmc image supports both white + black.
  - git log --pretty=oneline v3.8.13..HEAD | wc -l
    746

Questions? Flames?

Speak Your Mind

Antoniou Consulting
Embedded Linux Experts