



Linaro Ltd.

## News from the ARM architecture

Arnd Bergmann

© 2010 IBM Corporation

- 1 Historic issues with ARM kernel code
- 2 Future directions
- 3 Current workflow
- 4 Ongoing code changes
- 5 Ideal architecture ports
- 6 Working with the arm-soc process

# Historic issues with ARM kernel code

# History of ARM linux code, some time ago

- Strong focus on embedded systems

# History of ARM linux code, some time ago

- Strong focus on embedded systems
- Lots of out of tree code

## History of ARM linux code, some time ago

- Strong focus on embedded systems
- Lots of out of tree code
- Aging code base

## History of ARM linux code, some time ago

- Strong focus on embedded systems
- Lots of out of tree code
- Aging code base
- Losing the war on complexity

# History of ARM linux code, some time ago

- Strong focus on embedded systems
- Lots of out of tree code
- Aging code base
- Losing the war on complexity
- Close to collapse



# Technical problems in the ARM codebase

- Any new hardware mandates code changes
- No common platform model
- Multiple platforms mutually exclusive at compile time
- Infrastructure in hardware specific code

# Management problems

- Independent vendor specific communities
- Little incentive for cooperation and peer review

# Management problems

- Independent vendor specific communities
- Little incentive for cooperation and peer review
- Russell King could no longer keep up  
⇒ Lack of pushback on crap

# Management problems

- Independent vendor specific communities
- Little incentive for cooperation and peer review
- Russell King could no longer keep up
  - ⇒ Lack of pushback on crap
- Linus Torvalds could not keep up either
  - ⇒ Lack of merging new code

# Future directions

# Going beyond embedded

- Desktops
- Servers
- 64 bit CPU implementation

# Going beyond embedded

- Desktops
- Servers
- 64 bit CPU implementation
- One out of the two main linux architectures

# Ideal architecture port

- Generic user space ABI
- Device drivers in subsystems, not subarchitectures
- Boot time hardware detection
- No mutually exclusive build time options
- Minimum kernel changes for new hardware



# Current workflow

## 2011 direction changes

- Flattened device tree
- Single zImage for armv6/v7
- Duplicate code removal
- Moving device drivers to subsystems
- Co-maintaining arm-soc.git

# Changes we did not do

- Converge on a single clean platform
- Move legacy platforms to subdir

# 2011 achievements

- 1 new source tree

# 2011 achievements

- 1 new source tree
- 2 subarch maintainer summits

## 2011 achievements

- 1 new source tree
- 2 subarch maintainer summits
- 3 merge windows using arm-soc.git

## 2011 achievements

- 1 new source tree
- 2 subarch maintainer summits
- 3 merge windows using arm-soc.git
- 4 clean new subarchitectures

# 2012 achievements

- Common clock implementation
- Clock DT bindings



## 2012 achievements

- Common clock implementation
- Clock DT bindings
- Pinctrl subsystem

## 2012 achievements

- Common clock implementation
- Clock DT bindings
- Pinctrl subsystem
- Started removing board files

## Changesets merged

Kernel	rmk	arm-soc	arch/arm total	overall
2.6.39	680	0	1187	11031
3.0	265	22	518	9843
3.1	274	557	980	9380
3.2	286	701	1355	12695
3.3	XXX	XXX	XXX	XXX
3.4	XXX	XXX	XXX	XXX
3.5-rc1	XXX	XXX	XXX	XXX

## Top arm-soc contributors by changesets

174	Arnd Bergmann
78	Olof Johansson
69	Shawn Guo
65	Kevin Hilman
57	Benoit Cousson
53	Kukjin Kim
52	Tony Lindgren
51	Linus Torvalds
50	Jean-Christophe PLAGNIOL-VILLARD
46	Fabio Estevam

Total number of contributors: 236

## Top contributors by non-merge changesets

69	Shawn Guo
65	Kevin Hilman
57	Benoit Cousson
50	Jean-Christophe PLAGNIOL-VILLARD
46	Fabio Estevam
43	Stephen Warren
43	Mark Brown
38	Thomas Abraham
34	Paul Walmsley
32	Kukjin Kim

Total number of contributors: 235

## arm-soc maintainer role

- Sponsored by Linaro and Google
- Vendor neutral
- Define and enforce common rules for everyone
- Merging code upstream to Linus
- Coordination with Russell
- Pushing back on crap

## arm-soc workflow

- Pulling many per soc topic branches
- Pushing cross-soc topic branches
- One for-next branch, rebasing
- Early staging branches
- Not quite clean allowed when doing cleanups

## arm-soc workflow

- Pulling many per soc topic branches
- Pushing cross-soc topic branches
- One for-next branch, rebasing
- Early staging branches
- Not quite clean allowed when doing cleanups
- DEMO



# Ongoing code changes

# Multiplatform: Single zImage status

- Taking longer than expected

# Multiplatform: Single zImage status

- Taking longer than expected
- Breakthrough!

## Multiplatform: Single zImage status

- Taking longer than expected
- Breakthrough!
- Progress largely invisible
- Focus on ARMv6 and ARMv7

## Multiplatform: Single zImage status

- Taking longer than expected
- Breakthrough!
- Progress largely invisible
- Focus on ARMv6 and ARMv7
- Building vexpress+imx+omap2+ux500 now

## Multiplatform: Single zImage status

- Taking longer than expected
- Breakthrough!
- Progress largely invisible
- Focus on ARMv6 and ARMv7
- Building vexpress+imx+omap2+ux500 now
- Booting on vexpress so far

## Device tree status

- Taking longer than expected

## Device tree status

- Taking longer than expected
- Breakthrough!



## Device tree status

- Taking longer than expected
- Breakthrough!
- Pinmux bindings in 3.4 (Linus Walleij)

## Device tree status

- Taking longer than expected
- Breakthrough!
- Pinmux bindings in 3.4 (Linus Walleij)
- Clock bindings in 3.4 (Grant Likely/Mike Turquette)

## Device tree status

- Taking longer than expected
- Breakthrough!
- Pinmux bindings in 3.4 (Linus Walleij)
- Clock bindings in 3.4 (Grant Likely/Mike Turquette)
- New platforms: prima2, zynq, highbank, picocell, spear13xx
- Largely converted: exynos, i.mx, tegra, at91, vexpress, ux500
- In progress: omap, shmobile, lpc32xx, msm

## Device tree status

- Taking longer than expected
- Breakthrough!
- Pinmux bindings in 3.4 (Linus Walleij)
- Clock bindings in 3.4 (Grant Likely/Mike Turquette)
- New platforms: prima2, zynq, highbank, picocell, spear13xx
- Largely converted: exynos, i.mx, tegra, at91, vexpress, ux500
- In progress: omap, shmobile, lpc32xx, msm
- Lots of simple device driver bindings

# Randconfig builds

- Taking longer than expected

# Randconfig builds

- Taking longer than expected
- Delayed-by: Arnd Bergmann <arnd@arndb.de>
- Integration into Linaro CI loop

## Randconfig builds

- Taking longer than expected
- Delayed-by: Arnd Bergmann <arnd@arndb.de>
- Integration into Linaro CI loop
- about 150 patch series
- Regression testing
- Submitted by Mathieu Poirier

# Memory management changes

- Continuous memory allocator
- dma-buf infrastructure
- generic iommu handling in dma-mapping API



## Device driver subsystems to use

- MTD
- DRM (graphics)
- ASoC (sound)
- IndustrialIO (ADC, amplifiers, ...)
- PWM
- LED

# Framework drivers

- iommu
- dma-engine
- regulator
- clock
- gpiolib
- pinctrl

# Platform drivers

- irqdomain
- sparse irq
- clocksource
- devicetree

# Ideal file layout

- No board files
- Drivers in subsystems
- platform data in include/linux/platform\_data/
- register definitions in drivers
- IRQ/GPIO/MMIO/... definitions in DT

# Good Examples

- Highbank
- SPEAr13xx

# Infrastructure under construction

- DMA bindings
- System controller subsystem
- drivers/irqchip
- early console abstraction
- pwm subsystem
- USB host controller
- SMP
- PCI I/O space

# Cleanup vs. new code development

- Reasonable proportion of cleanups

# Cleanup vs. new code development

- Reasonable proportion of cleanups
- Reducing the amount of crap
- More crap getting removed than added



# Cleanup vs. new code development

- Reasonable proportion of cleanups
- Reducing the amount of crap
- More crap getting removed than added
  - Coding style
  - Non-portable code
  - Bloat
  - Private infrastructure
  - Missing infrastructure
  - Excessive Macro use

# Cleanup vs. new code development

- Reasonable proportion of cleanups
- Reducing the amount of crap
- More crap getting removed than added
  - Coding style
  - Non-portable code
  - Bloat
  - Private infrastructure
  - Missing infrastructure
  - Excessive Macro use
- Some crap is very well written

# Platform specific problems: ux500

# Platform specific problems: ux500

- out of tree code

## Platform specific problems: ux500

- out of tree code
- common clock

# Platform specific problems: omap

# Platform specific problems: omap

- hwmod

## Platform specific problems: omap

- hwmod
- absolute code size



# Platform specific problems: samsung

## Platform specific problems: samsung

- fragmentation: exynos, s5p, s3c64xx, s3c24xx

## Platform specific problems: samsung

- fragmentation: exynos, s5p, s3c64xx, s3c24xx
- premature DMA DT binding

# Platform specific problems: i.mx

## Platform specific problems: i.mx

- macros in arch/arm/plat-mxc/devices/

## Platform specific problems: i.mx

- macros in arch/arm/plat-mxc/devices/
- not much else really

# Platform specific problems: vexpress

# Platform specific problems: vexpress

- clock support



# Platform specific problems: vexpress

- clock support
- device tree support partial

# Working with the arm-soc process

- Topic branches
- Separate cleanups and bug fixes
- Timing

## Topic branches

- Independent development
- Based on a -rc version
- Can include dependencies

## Topic branches

- Independent development
- Based on a -rc version
- Can include dependencies
- Too many dependencies:
  - arm-soc cannot linearize them: Game over

## Topic branches

- Independent development
- Based on a -rc version
- Can include dependencies
- Too many dependencies:
  - arm-soc cannot linearize them: Game over
- Too many conflicts
  - cannot resolve: Game over

# Typical topic branches

- fixes (urgent)
- fixes (non-urgent)
- cleanups
- dt
- board
- pm
- drivers

## Typical topic branches (continued)

- new-soc
- defconfig
- maintainers
- clock
- pinctrl
- new subsys

# Cleanup branches

- Usually first
- Features based on cleanups
- Large but nonintrusive



# Cleanup branches

- Usually first
- Features based on cleanups
- Large but nonintrusive
- Your ticket for other changes

# Changeset comments

- Always too short
- Explain why a change is done, not how
- shortlog style "ARM: omap: do something"
- Only bug fixes after -rc1
- Use "cc: stable@vger.kernel.org" for backports

# Pull requests

- Send to [arm@kernel.org](mailto:arm@kernel.org)
- Signed tags
- Get your gpg keys signed
- Make branches bisectable
- Tag descriptions

# Pull timing

- As early as possible after -rc1
  - ...but no earlier
- Staging branches
- As early as possible after -rc1

# Summary

- Lots of work getting done:  
device tree, consolidation, cleanup
- Not running out of work anytime soon:  
new platforms, ARMv8, multiplatform kernels

# Summary

- Lots of work getting done:  
device tree, consolidation, cleanup
- Not running out of work anytime soon:  
new platforms, ARMv8, multiplatform kernels
- Linux on ARM world domination coming,  
still busy with the details

# Legal Statement

This work represents the view of the author and does not necessarily represent the view of IBM.

IBM, IBM (logo), e-business (logo), pSeries, e (logo) server, and xSeries are trademarks or registered trademarks of International Business Machines Corporation in the United States and/or other countries.

Linux is a registered trademark of Linus Torvalds.

Other company, product, and service names may be trademarks or service marks of others.

# Questions?