Porting FreeBSD/arm to Marvell SoC

Rafał Jaworowski
raj@semihalf.com

BSDCan 2008, Ottawa

Introduction – TOC

- ARM architecture basics
- Highlights of contemporary FreeBSD/arm support
- Marvell SoC family overview
- Notes on porting FreeBSD/arm to Orion and Kirkwood
  - Host system, reference source code
  - Scope of porting activities (rework, new development)
  - Challenges and issues
  - Debugging, testing
- Current state of support
- Upcoming development

What is ARM?

- 32-bit RISC architecture designed by ARM Ltd.
  - Widely used in embedded applications
  - Recognized for low power consumption
  - From simple machines to advanced systems
- Very popular
  - Estimated 75% of all RISC CPUs on the market
- Licensed
  - A lot of incarnations from various vendors
- History
  - Origins at Acorn Computers Ltd. (80's)
  - Acorn RISC Machine ›› Advanced RISC Machine ›› ARM architecture
  - More significant: ARM6 (90's)

ARM nomenclature

- Multitude of variations, source of confusion...
- Family
  - ARM{*family*}{*features:* MMU, caches, Thumb, VFP, pipeline depth, synthesizable etc.}
  - ARM9TDMI, ARM9E, ARM10E, ARM11, ...
  - A group of processor implementations sharing the same H/W characteristics
- Architecture version
  - ARMv4T, ARMv5TEJ, ARMv6, ARMv6KZ, ...
  - Designates instruction set (Thumb first introduced in ARMv4T)
- Core
  - ARM926EJ-S, ARM1022E, ARM1156T2(F)-S, ...

ARM technical highlights

- RISC
  - Reduced number of instructions
  - Pipelines: processing of instructions broken down to smaller units executed in parallel
  - Registers: large set of general purpose registers
  - Load-store architecture: CPU operates on data in registers, dedicated instructions for transfers between registers and memory
- ARM specific (non-RISC)
  - Variable execution time for some instructions
  - Thumb (effectively a second 16-bit instruction set)
  - Enhanced instructions (DSP, Jazelle, VFP etc.)

ARM technical highlights – cont'd

- Bus architecture
  - Von Neumann (ARM7): unified L1 cache
  - Harvard (ARM9 and beyond): split L1 cache
- L1 Caches
  - Virtual (ARM7-10)
    - Software maintained coherency
    - Flush/invalidate on every context switch
    - DMA issues
  - Physical (ARM11)
    - Hardware enforced coherency

ARM technical highlights – cont'd

- Memory mangement
  - MMU (demand paging)
  - MPU (simpler schemes)
- Usually ARM core is part of a System-On-Chip
- Soft-cores
  - synthesizable and netlist
- Hard-cores
- Not microprogrammed
- Unique features (Thumb, Jazelle, DSP, VFP)

Terminology

- Platform
  - Complete system
  - Analogous to a motherboard in a PC, only needs a PSU
  - Built around SoC
- SoC (System-On-Chip)
  - Highly integrated circuit
  - Local buses
  - Memory controllers (DRAM, FLASH etc.)
  - Peripherals (Ethernet, UART, PCI/PCIE bridges, USB, DMA engines)
  - CPU
- Core CPU
  - Main processing unit(s)
  - Integral part of the SoC

Generic System-On-Chip

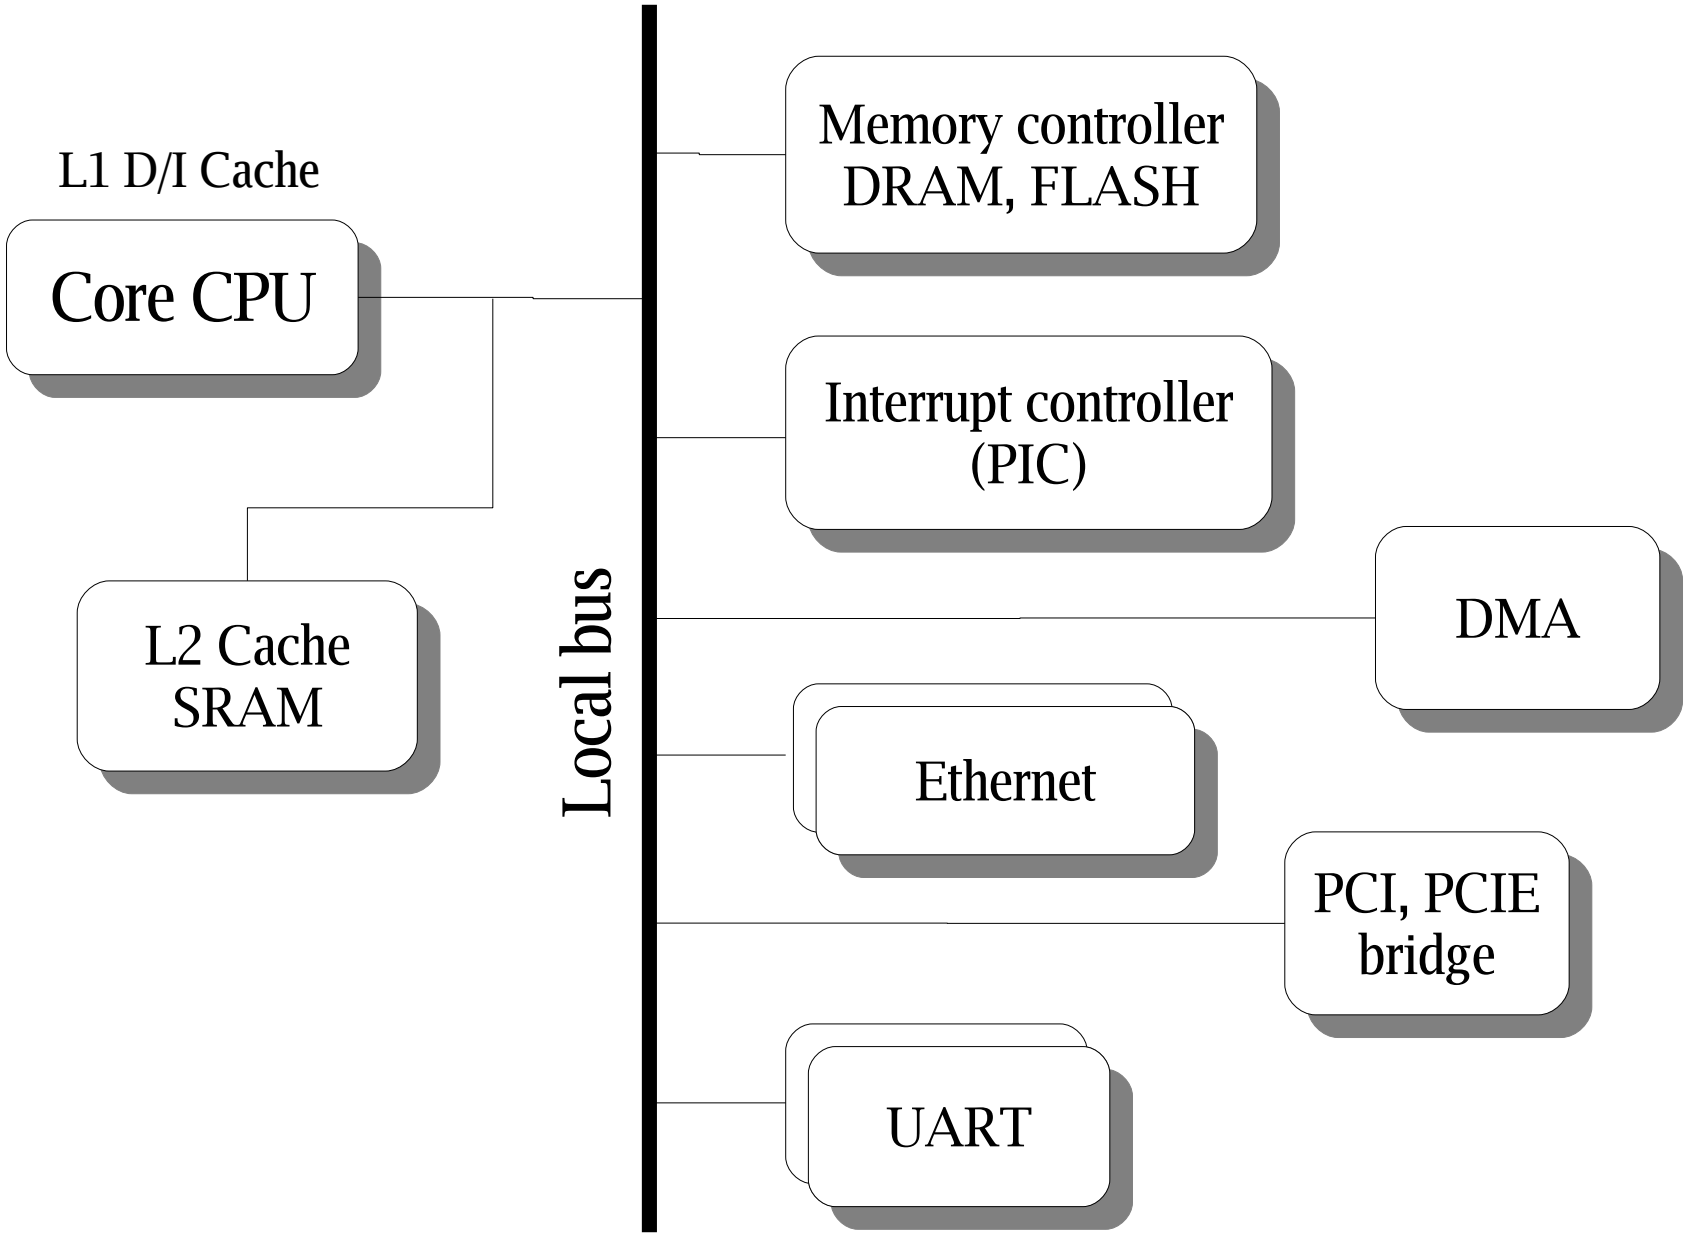

Marvell SoC devices

- Common features
  - Based on the core compliant with ARMv5TE architecture spec
  - Integrated with many peripherals and offloading engines into one chip
  - Endian mode selected at boot time
- Extensions to ARM architecture, non-typical features
  - Out-of-order execution
  - Branch prediction
  - Super scalar pipeline
  - DSP extensions

Systems families

- Orion - 88F5181, 88F5182
  - Feroceon 88FR331 CPU Core (150nm)
  - up to 500MHz
  - Single issue
  - Integrated I/D L1 cache (32KB/32KB)
    - 4-way, direct mapped
- Orion - 88F5281
  - Feroceon 88FR531 Core (150nm)
  - up to 500MHz
  - Dual issue
  - Integrated L1 cache
    - Direct mapped 32KB I-cache
    - 4-way, set associative, 32KB D-cache
    - Write-allocate
  - Branch Target Buffer (BTB)
  - Vector Floating Point Unit (VFP)
- Kirkwood - 88F6180, 88F6192, 88F6281
  - Feroceon 88FR131 CPU Core (65nm)
  - up to 1.5GHz
  - Single issue
  - Integrated I/D L1 cache (16KB/16KB)
    - 4-way, set associative
  - Unified L2 cache (256KB)
    - 4-way, set associative
    - Write-back or write-through
    - Physically tagged
  - Vector Floating Point Unit (VFP)

SoC features overview

- 88F5281
  - Gigabit Ethernet
  - GPIO/MPP
  - IDMA engine
  - Interrupt controller
  - PCI/PCI-X, PCI-Express
  - TWSI (I2C), Timers, Watchdog
  - UART
  - USB 2.0
- 88F6281
  - Crypto engine
  - Gigabit Ethernet
  - GPIO/MPP
  - Interrupt controller
  - NAND
  - PCI-Express
  - SATA
  - SD/SDIO/MMC
  - S/PDIF, I2S
  - SPI
  - TDM
  - TWSI (I2C), Timers, Watchdog
  - UART
  - USB 2.0
  - XOR engine (DMA)

Development environment

- FreeBSD 8.0-CURRENT source tree
  - Started with mid-Nov 2007 snapshot
  - Keeping in sync, some 2-3 weeks behind latest HEAD
  - In-tree toolchain
- Development platform
  - Marvell reference designs: DB-88F5182, DB-88F5281, DB-88F6281, RD-88F6281
  - Complete systems
- Firmware
  - U-Boot derived from 1.1.4 version, extended with Feroceon support
  - Additional features with regards to stock U-Boot (diagnostics subsystem, extra dozen commands)

State of the FreeBSD/arm

- CVS
  - AT91
  - StrongARM 11x0
  - XScale: I80321, I8134x, IXP425
  - Released in FreeBSD 7.x
- Perforce
  - EP93xx
  - Orion 5x
  - S3C2xx0
  - XScale: PXA2x0
- Originates from NetBSD/arm port
- Most of the infrastructure in place
  - pmap(9), bus_dma(9), bus_space(9)

Scope of the Marvell SoC port

- ARM9E startup
  - Rework and extensions within existing ARM startup area
  - Optimizations for bigger KVA
- Machine-dependent
  - Flexible RAM size recognized when setting up pmap (page tables): more and more RAM
  - Decode windows, GPIO/MPP setup
- Device drivers hierarchy (*newbus*)
  - Internal bus model (mbus)
  - Critical drivers: timers, interrupt controller
  - Successful reuse example: uart(4), ehci(4), only *mbus* attachment required

Simplified virtual memory view

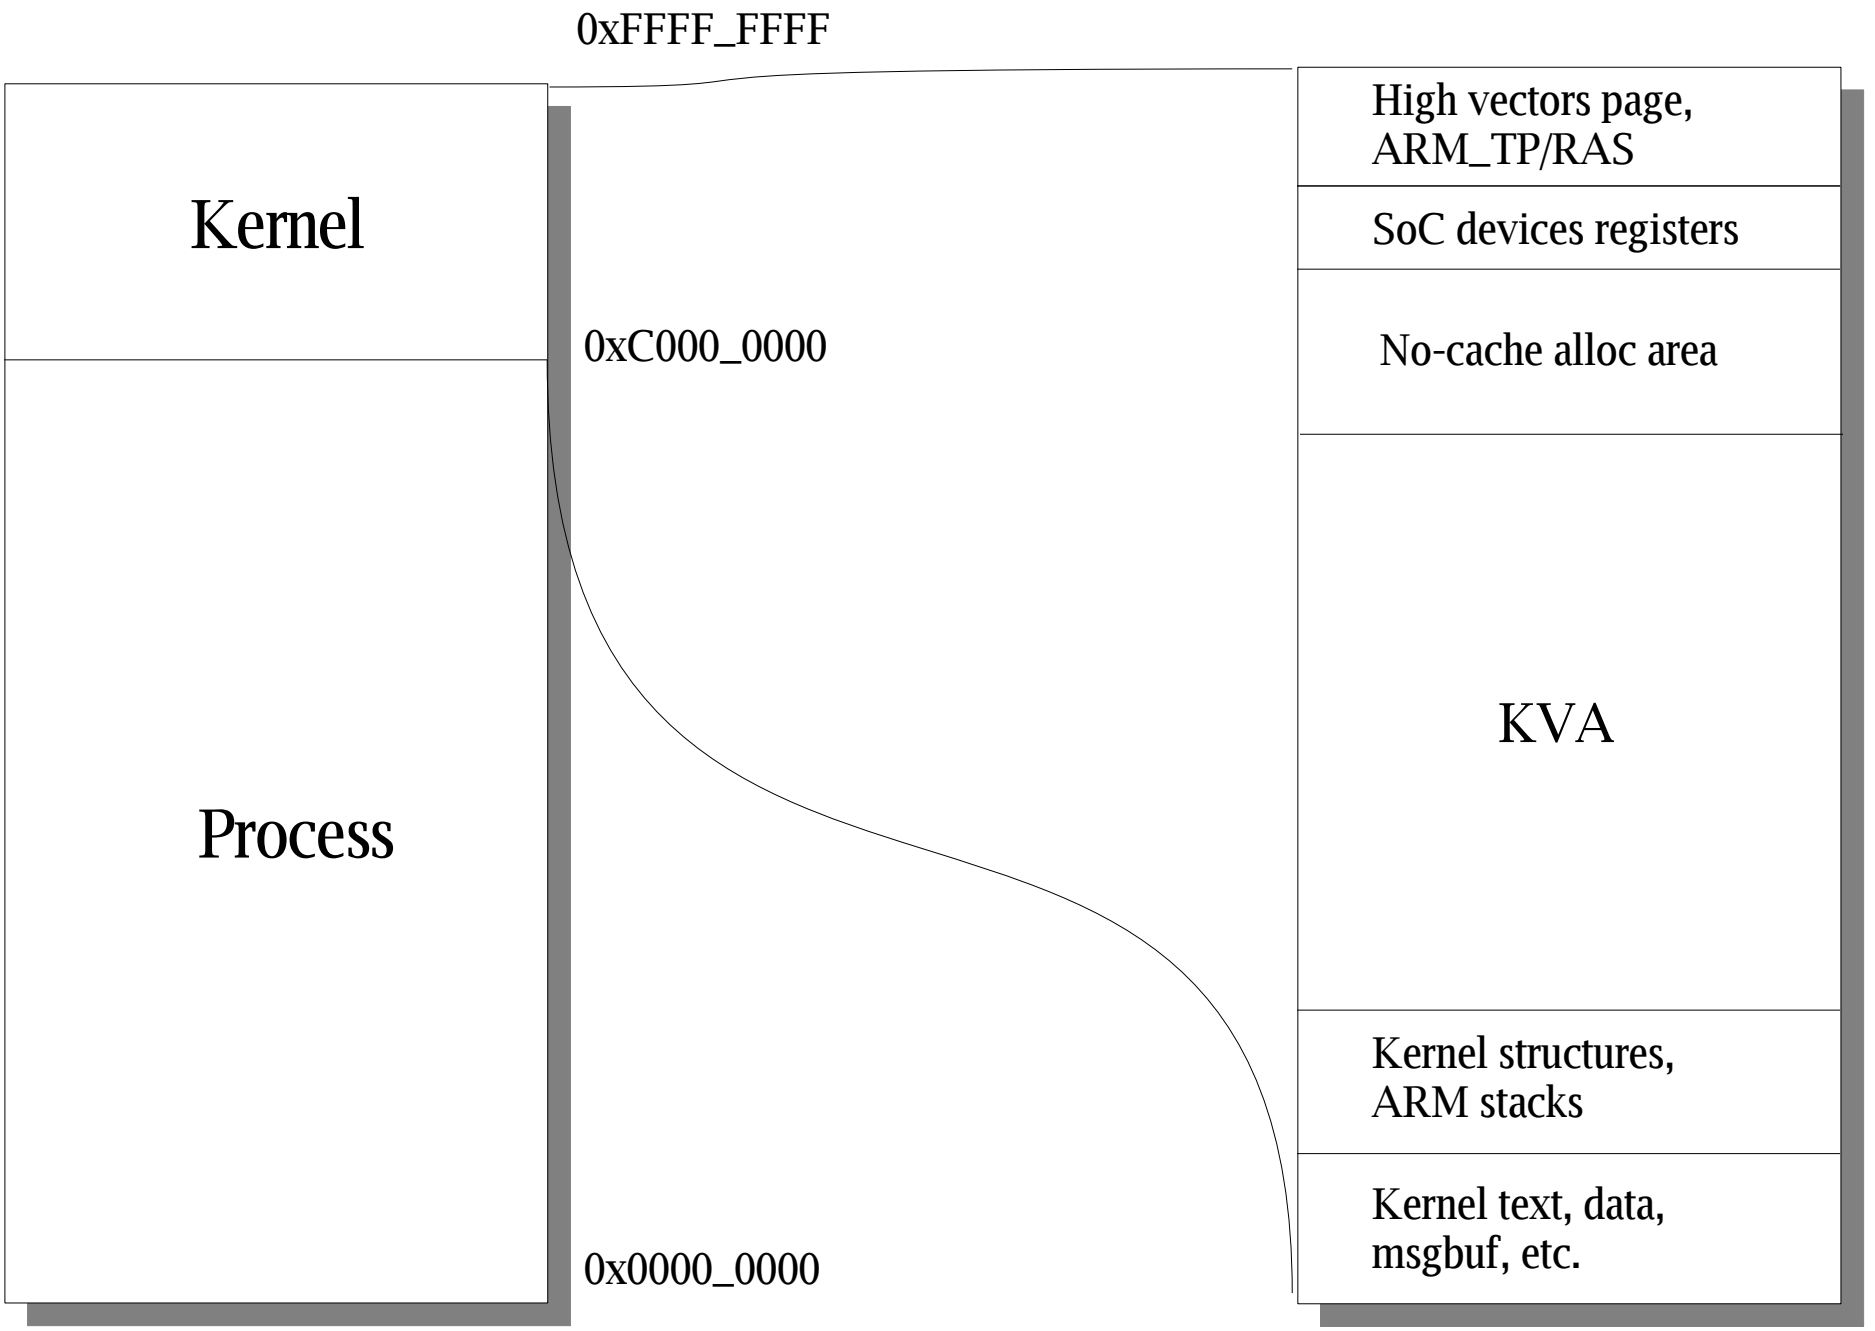

New development highlights

- Gigabit Ethernet controller
  - mge(4)
- IDMA
  - General purpose DMA engine
  - idma(4) driver, test application
- TWSI (I2C)
  - Controller driver for iicbus(4) framework, test application
- PCI, PCI-Express
- Infrastructure
  - Decode windows setup
  - Low-level core operations (L2 cache support in newer Feroceons)
  - Design towards multi-SoC/multi-platform support

Newbus hierarchy

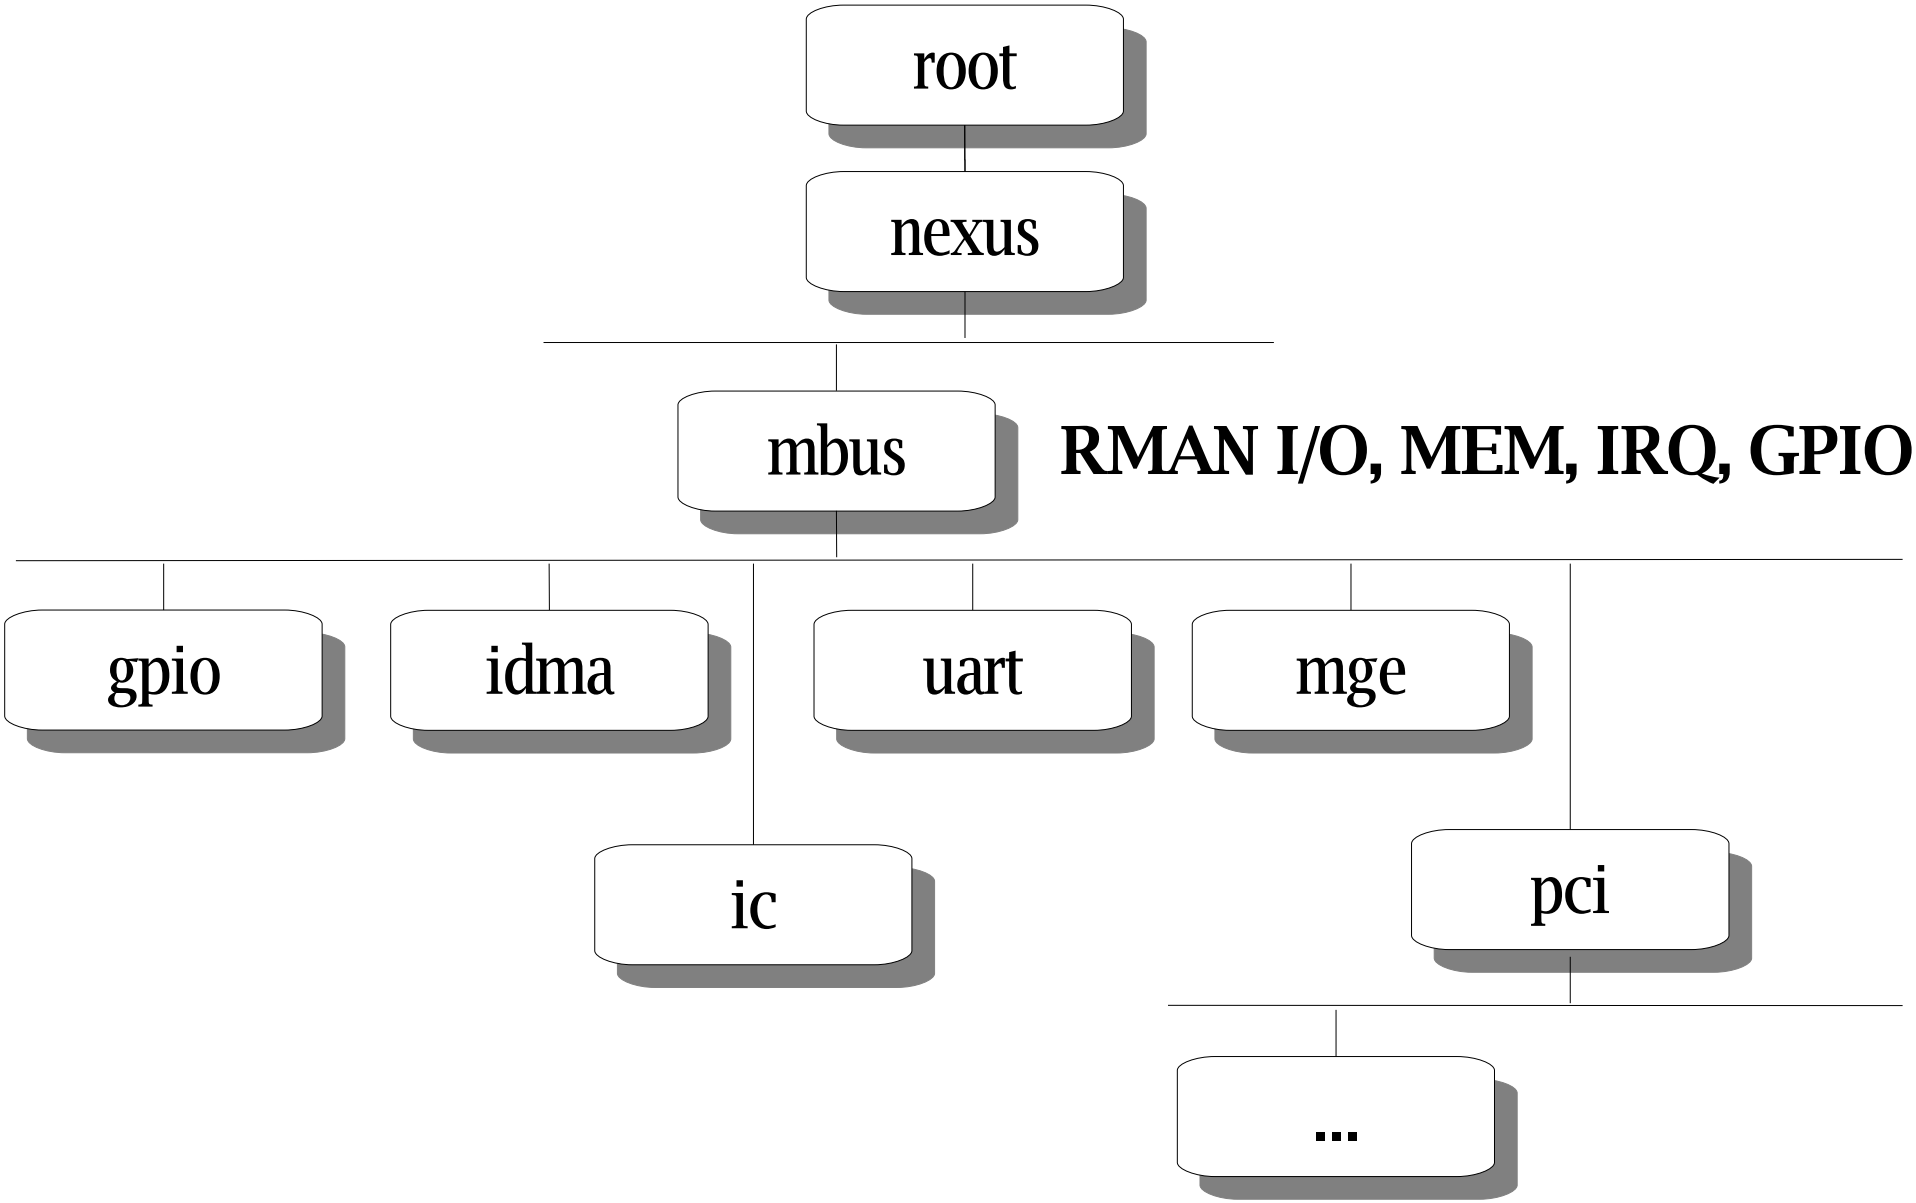

More challenging aspects – virtual cache

- ARM9
  - Virtually-tagged cache
- Coherency
  - Maintained by software
  - Issues not seen on systems with physically-tagged caches
  - Some drivers and code not well prepared: USB, pmap required additional/better handling or bus_dma(9) operations
- Write-allocate
  - Optional feature of write-back mode
  - Complicates virtual cache management
  - Important during “write miss” scenario

Virtual cache + write-allocate

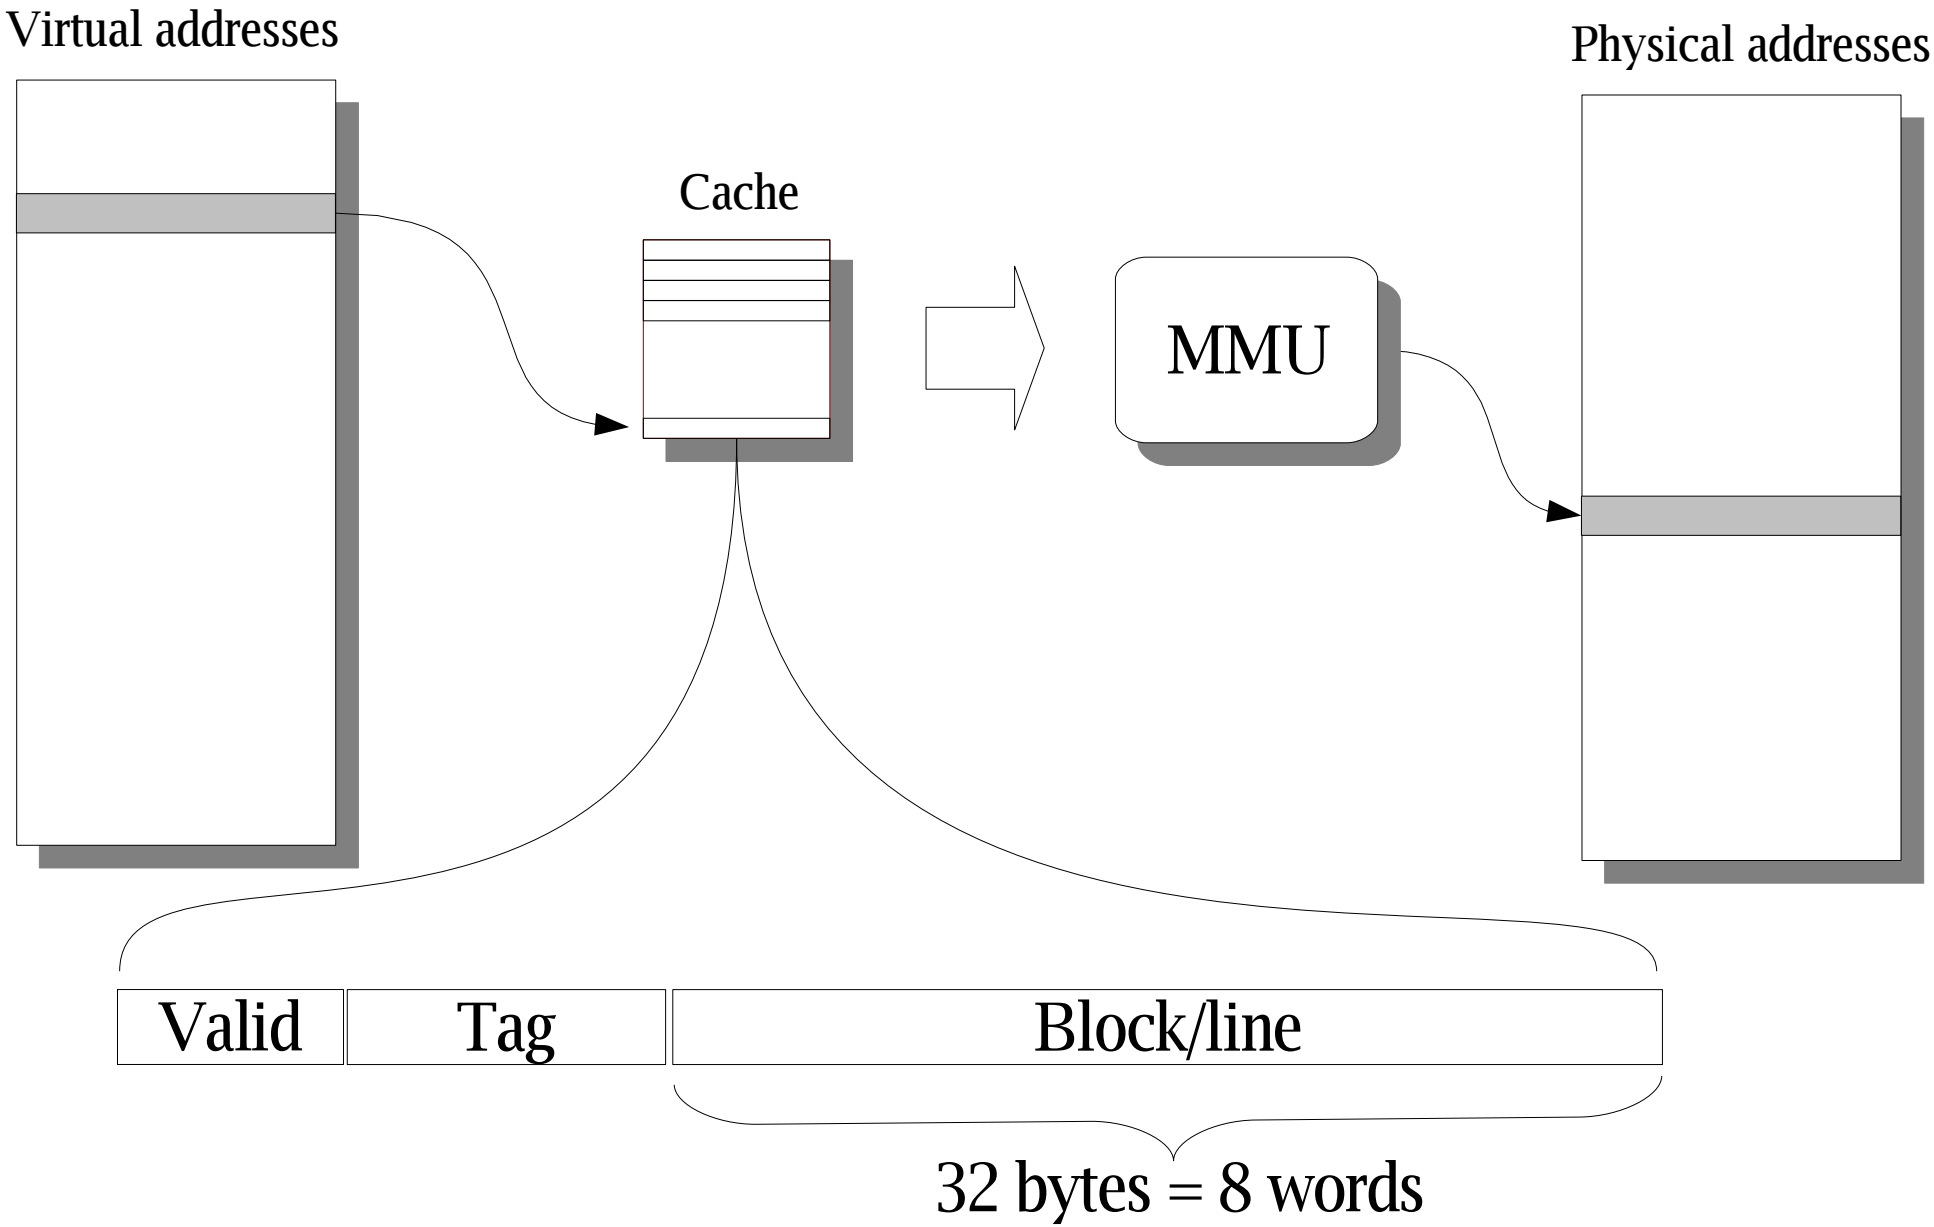

Virtual cache + multiple mappings

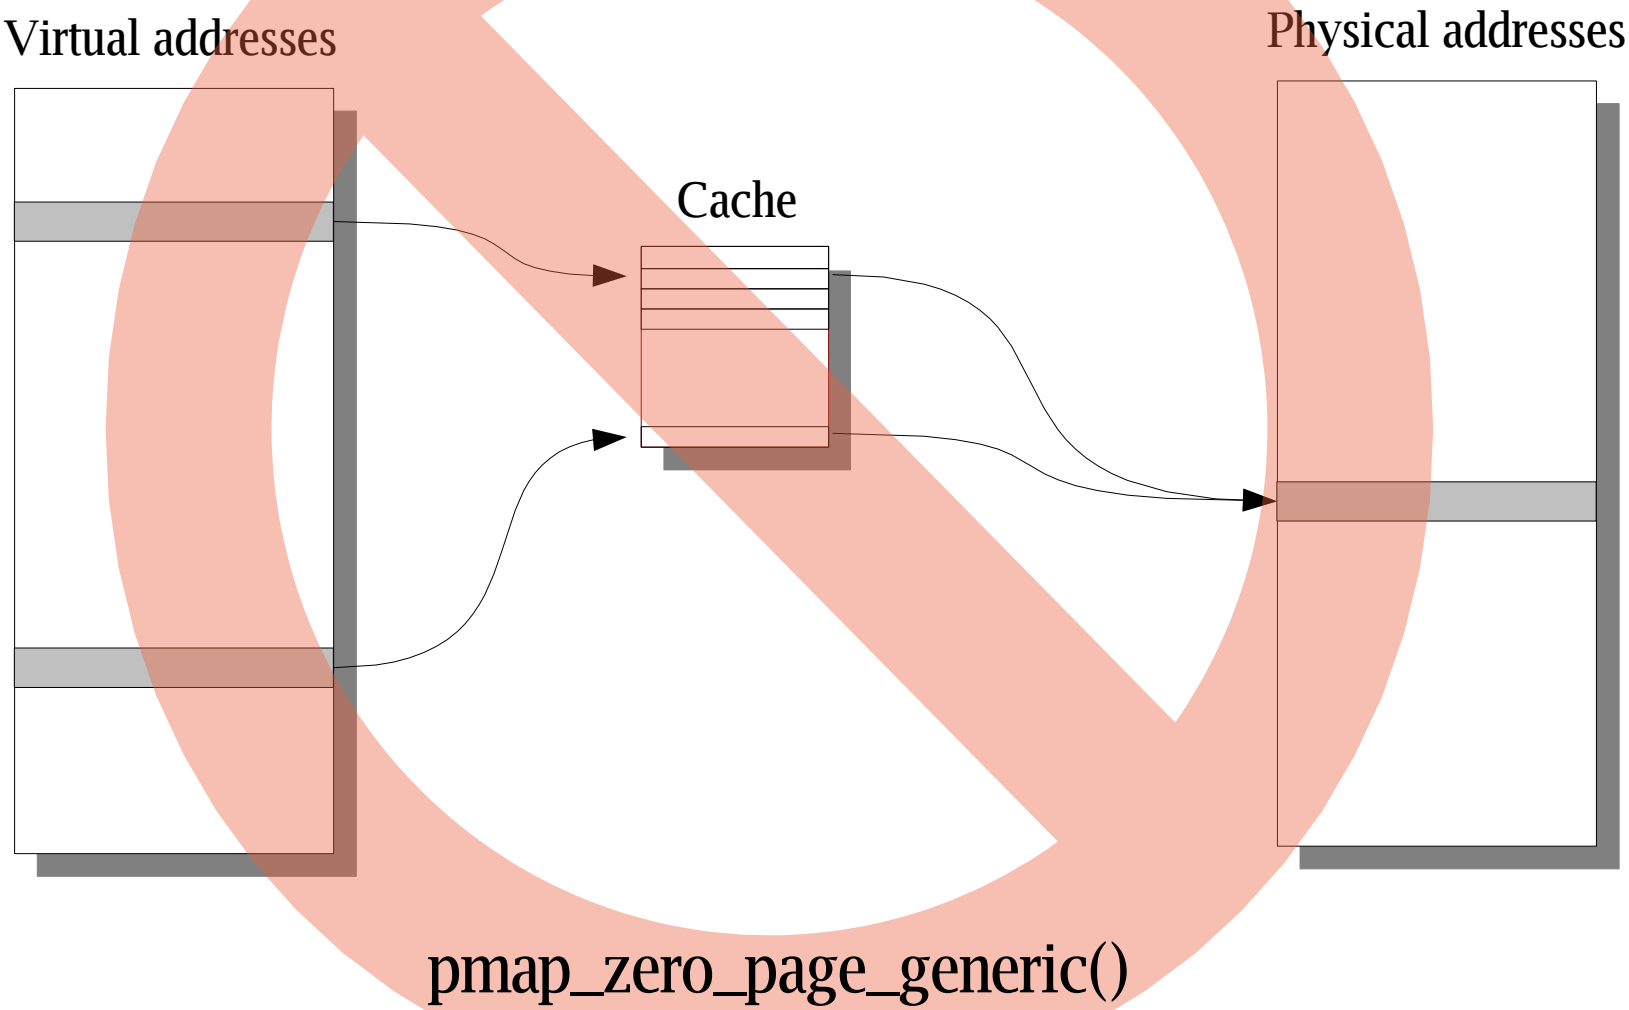

Methodology notes

- Working with real targets, not simulator
- Reviewing code from all other FreeBSD architectures, other operating systems for reference
- Using available FreeBSD tools and infrastructure
  - MFS before networking or storage operational
  - GEOM_UZIP, mkuzip
- Debugging
  - JTAG H/W debugger
  - KDB/DDB: bring as early as possible
  - 7-seg LED are fun and help
- Testing
  - In-house developed test framework (expect)
  - Existing tools (iperf etc.)
- Lab

Current state of Orion support

- Single- and multi-user operation
  - Booting with root filesystem mounted from MFS, USB, NFS
  - Little Endian
- Single kernel image for 5x family, separate for 6x
- 5x already in Perforce
  - //depot/projects/arm/...
  - make buildkernel TARGET_ARCH=arm KERNCONF=DB-88F5XXX
- Unsupported
  - VFP (vector floating point coprocessor)
- Coming soon to a theatre near you!
  - Kirkwood 88F6281
  - Discovery MV-78100

Concluding remarks

- References
  - //depot/projects/arm/
  - sys/arm
  - sys/arm/orion
  - sys/dev/mge
- Acknowledgements
  - NetBSD Project
  - FreeBSD/arm team: Olivier Houchard, Sam Leffler, Kevin Lo (thanks for testing on the Linkstation system!), Warner Losh
  - Maen Suleiman @Marvell
  - Grzegorz Bernacki, Bartłomiej Sięka, Jan Sięka @Semihalf

Porting FreeBSD/arm to Marvell SoC

Rafał Jaworowski
raj@semihalf.com

http://www.semihalf.com/pub/bsdcan/2008_marvell_freebsd.pdf

BSDCan 2008, Ottawa