Understanding What Those 250 Million Transistors are Doing

John Goodacre
ARM Processor Division
Director – Program Management

"The Elba Program"

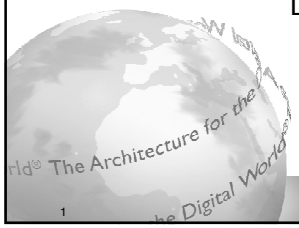

Abstract:

ARM has continuously developed test silicon for its IP to confirm its operation and characteristic. However today's challenges are so much greater, with ARM's latest testchip "Elba" consisting of over 250M transistors there are so many different aspects that need to be understood. In this summary, we'll look at Elba along with some of the design goals and special structures and techniques employed to be able to continue to characterize and understand the operation of today's more complex IP.

“Elba” Program: Background

- Advanced Product Development project in ARM aimed at comparing Cortex™-A9 MPCore™ processor against other high performance CPU implementations.
  - Use of TSMC 40G to allow comparison against other high performance processes
  - Tweaking of logical/physical technology and RTL/Library IP to extract max performance and lowest power.
  - Development of two macros High/Mid Performance to compare against equivalent competitive product.
  - *Develop a device, analyse split lots, and put it on an appropriate board and prove SoC power management strategies*
- Identify and offer program output as various product offerings
  - Enhanced RTL and physical IP, (Artisan® Processor Optimization Pack [POP])
  - Processor hard-macro with validated power management schemes (Osprey)
- Evolve an ecosystem around the support of high-performance “general purpose ARM devices”
  - Build a high performance platform to support the Osprey macro
  - Integrate ARM IP from multiple sources and present an optimized solution
  - Enhance application-specific opportunities by extending the ARM values of low-power across more markets

Elba: High-level Design Target

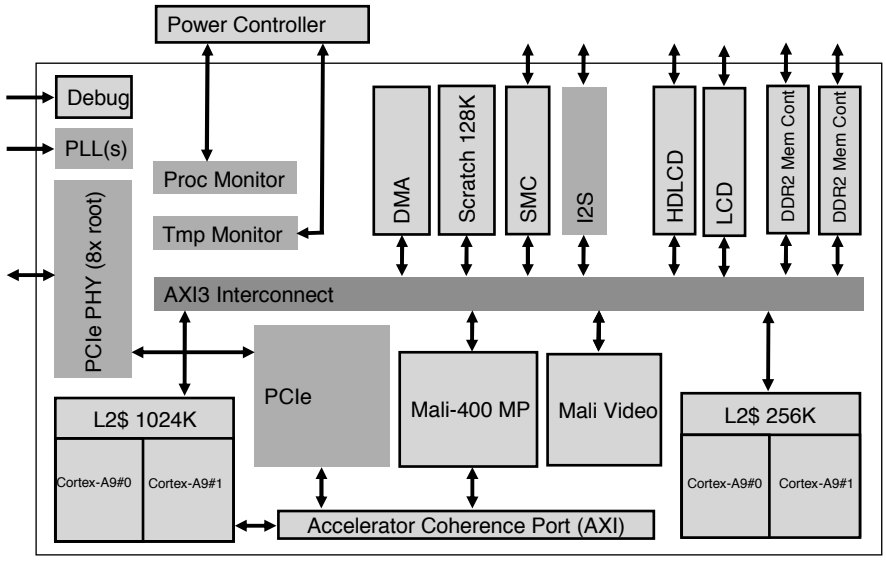

Osprey: The Cortex-A9 Hard Macro

- Designed with simplified SoC integration in mind
  - Only synchronous interfaces are L2 cache RAM and AMBA® AXI™ master ports
  - All GHz+ logic contained within macro
  - Debug, reset, interrupts etc. are all synchronised in the macro
- AMBA AXI interface capable of 2:1 bus to CPU clock ratio
  - Additional asynchronous interface can be added outside of the macro for DVFS operation
  - To achieve 2:1 ratio must be closely coupled to bus fabric
- L2 macro interface is very critical
  - Cache controller supports programmable read and write latency
  - Achievable latency dependant on floorplan and L2 RAM performance
- CoreSight™ debug and trace integration kit shipped within macro

Osprey: Cortex-A9 Dual-core Hard Macro

- Combines Cortex-A9 cluster with L2 cache controller and high speed part of the debug interface
  - Includes additional synchronisation logic to ease SoC integration
  - Level 2 cache memories are outside the macro for flexibility (256K-8MB)

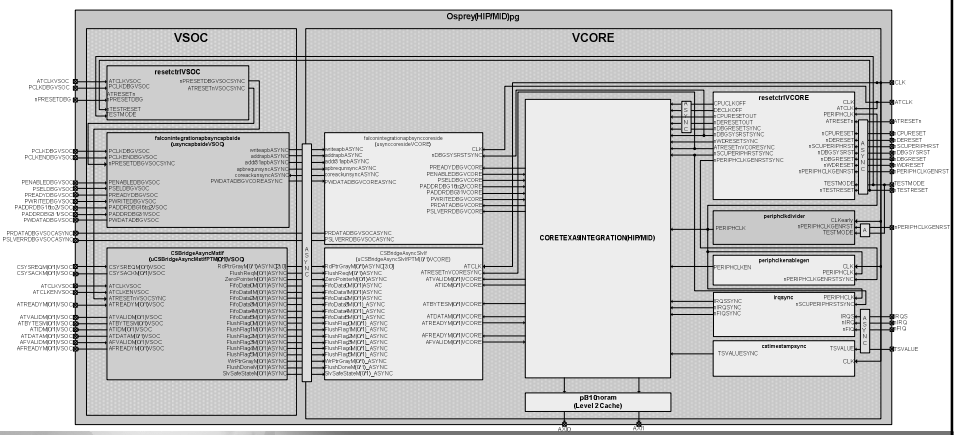

Osprey: Processor Power Gating

Eight Independent Regions

- CPU-0
- CPU-1
- NEON-0
- NEON-1
- PTM-0
- PTM-1
- Debug
- Memory BIST Controller
- CoreLink™ L2C-310 Level 2 Cache and the Cortex-A9 Snoop Control Unit

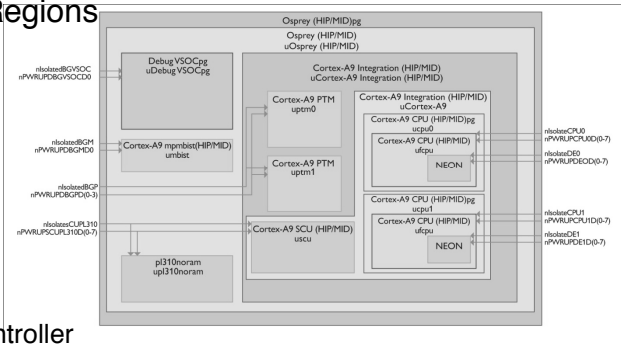

- Controlled either during boot (eg Debug / MBIST), or at run-time (eg NEON™ or CPU)

Osprey: In-Rush Analysis

- Need to limit in-rush current when coming out of power gated mode
  - Ensuring neighbouring power domains do not brown out
- Use small number of “little switches” to charge up the rails
  - limits current to manageable level
- Use more number of “big switches” to supply current in mission mode
- 50L RVT switches being used to reduce the leakage in OFF state
  - Provides 4x reduction compared to 40L RVT
- Guideline used for signoff (ff 125c corner)
  - Inrush current < Mission mode current
  - Wakeup time within 100ns per domain

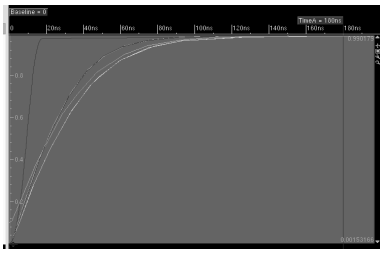

Power Gated Region Control

- Each CPU, NEON and the SCU/L2 logic have 8 groups of Power Switches
  - Analysis shows ~100ns power up with a 10ns delay between each group
- Power up order is not in sequence due to more switches on chain 0 (switches in chain 0 would be reduced if we did the implementation again)

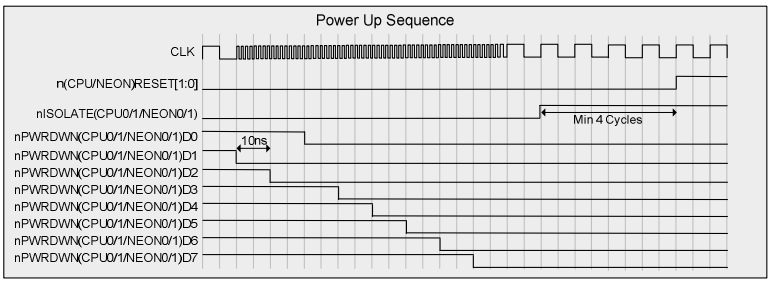

Selecting Different Transistor Sizes

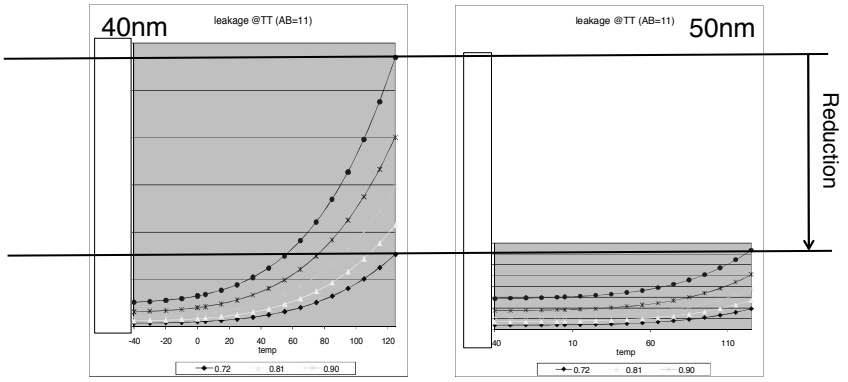

- 50nm cells behave “like HVt”
  - Significant leakage reduction at higher temperatures
  - No additional mask cost, better low voltage operation
  - Can still use HVt if desired

Elba: Designing a Osprey “Testchip”

- Different from most SoC developments as this is a development platform and not a product
  - A silicon qualification vehicle for the Cortex-A9 hard macros
  - An IP proving vehicle for new cores
  - A software development vehicle
  - Demonstrate diverse applications from server to compute platform
  - Support future R&D activities around power management etc.
- Design/implementation decisions are sometimes counter-intuitive
  - Packaging cost is not an issue, silicon area is
  - An unusual mixture of IP and connectivity
    - Required to support multiple and complex clocking schemes
  - Optimized across multiple applications
  - DFT/DFM yield decisions may be compromised (not for hard macro)

Elba: Block Diagram

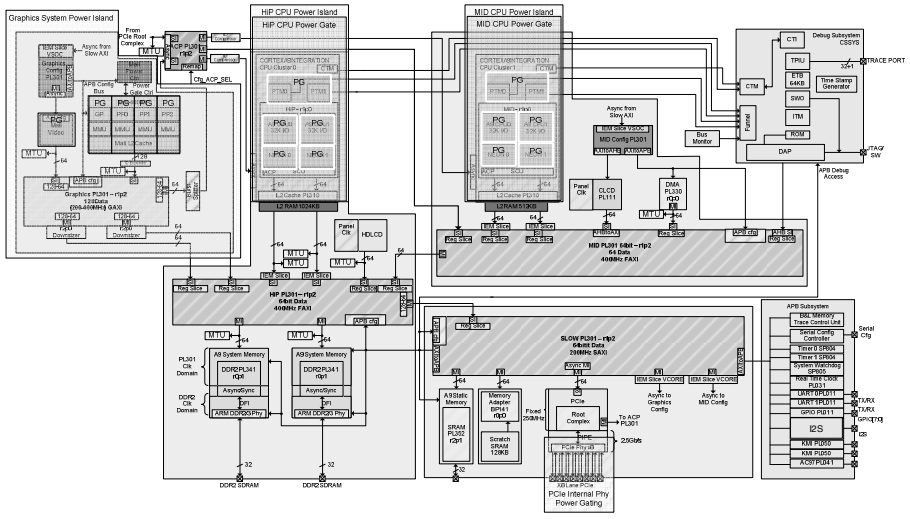

Elba: Floorplan

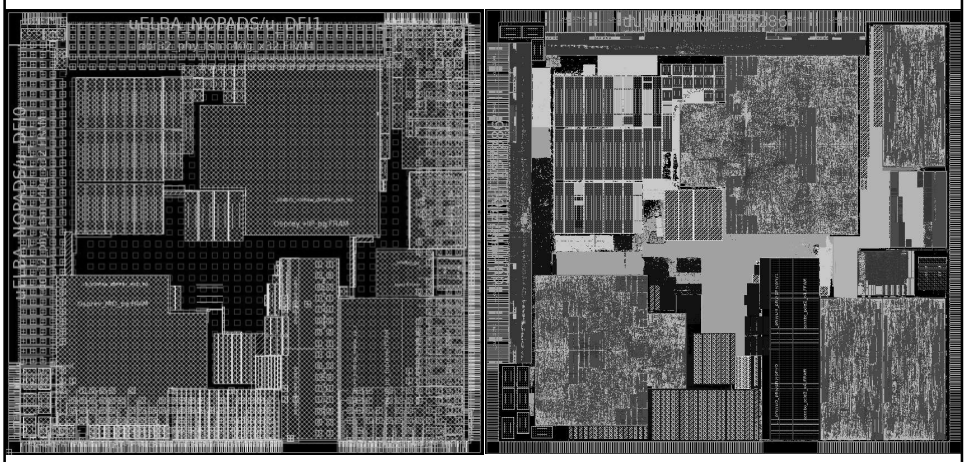

Elba Clocks

- Multiple Synchronous and Asynchronous clock domains
  - 40 clock definitions (84 for CTS)
  - 26 Asynchronous Clock Groups
  - 118 Multi-Cycle Path definitions
  - 300 False Paths

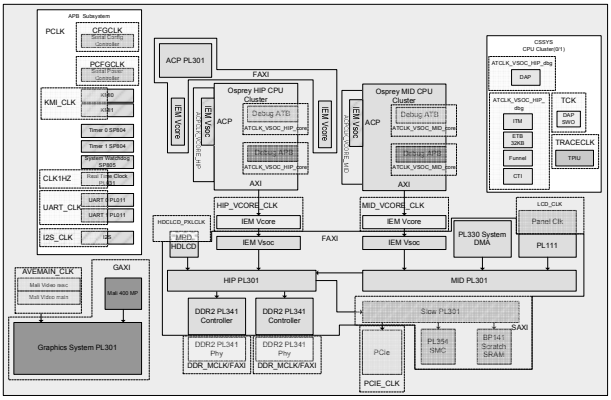

Elba Voltage Islands

- **SoC Power** – This is an always on domain and powers the System AXI, Memory controllers (excluding phys), APB peripherals, System LCD Controllers, DMA engine and scratch RAM
- **HIP VCore Power** – This is a variable domain and powers the high performance macro. Within this domain are a number of power-gated regions discussed later
- **MID VCore Power** – This is a variable domain and powers the MID performance macro. Within this domain are a number of power gated regions discussed later
- **Graphics Power** – This powers the full graphics subsystem. Including Mali™-400 MP, Mali Video and CoreLink NIC-301 Interconnect. This whole domain can be powered down when not required.

Overall Power Control System

- Off-Chip Cortex-M3 microcontroller
- Power change request interrupt (Mail Box Mechanism)
- UART communications for high level messaging
- Serial Power Control interface (SPI like)
  - External control by Micro or Cortex-A9 can access via APB
- Cortex-A9 Interrupt, STANDBYWFI and PWRCTL monitoring

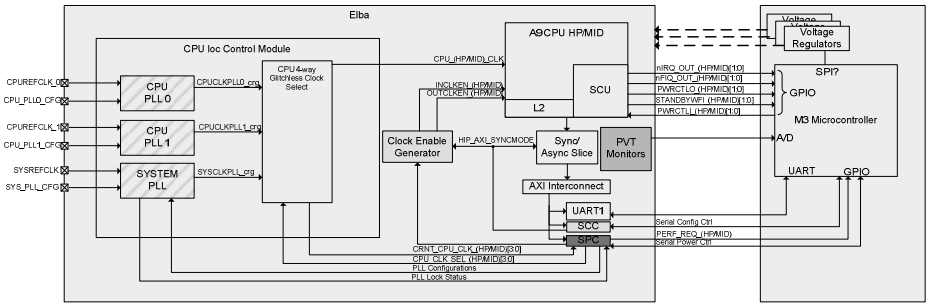

Mali-400 MP Power Gating

- Mali-400 MP Configuration
  - 3 Fragment Processors
  - 1 Vertex Processor
  - 128KB L2 Cache
- Power Gating
  - Graphics - independent PP & GP
  - Mali video engine
- Includes Mali PMU
  - Automatic power down when IDLE
  - Implementation has 2 logic and 4 RAM power control chains

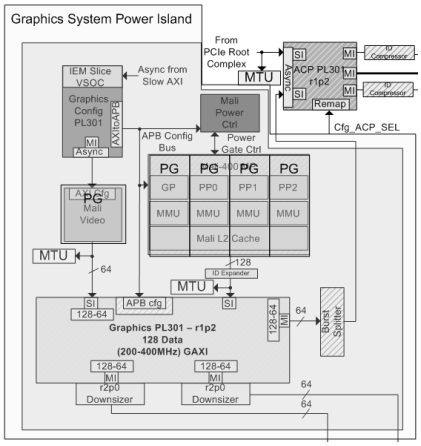

Mali-400 MP PMU

- For Mali GPU Pixel Processors and Geometry Processor
- Elba custom power control state machine
- Uses Mali PMU power-up request and acknowledge signals to start sequence
- Transparent to Mali driver

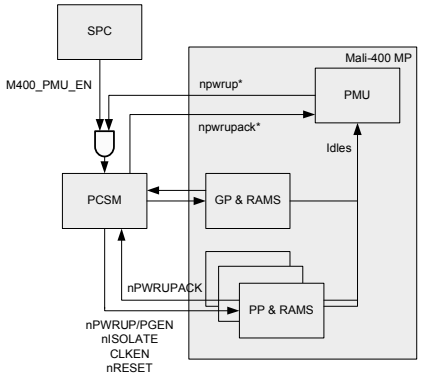

Processor Deep Sleep Modes

- SCU + L2 Active
  - Master via ACP port can snoop L2 or be forwarded to L3
  - GIC live, so wake-up from standard SoC interrupt sources (nIRQ_OUT/FIQ_OUT from A9)
- L2 Dormant
  - L2 and SCU logic powered down
  - Uses internal retention mode
  - GIC live
- Deep-Shutdown
  - L2 data lost
  - GIC powered-down
    - On-chip interrupt wake-up sources or external power-up request
    - Example UART, KMI, WatchDog and Timers
  - Software state retention/restore (hibernation) typically used

Power State	SCU/L2C-310 Logic	CPU0	NEON0	CPU1	NEON1	L2 RAMS
CPUs active	ON	OFF/ON	OFF/ON	OFF/ON	OFF/ON	ON
acp_l2active	ON	OFF	OFF	OFF	OFF	ON
l2dormant	OFF	OFF	OFF	OFF	OFF	RETENTION
deep_shutdown	OFF	OFF	OFF	OFF	OFF	OFF

Dynamic Voltage and Frequency Scaling

- 3 PLLs
  - 1 Fixed for 100% performance and SoC Clocks
    - Also provides 50% performance
  - 2 Variable
    - Configured to any frequency
- Clock Switch
  - Provide glitchless switching between all clocks (If PLLs are locked)
  - Controlled via the serial SPC or directly via CPU
- Independent for both HIP and MID
  - Not a primary operating mode but both Cortex-A9 clusters can operate together
- Overall, more complex than needed, but good for R&D !

DVFS cont….

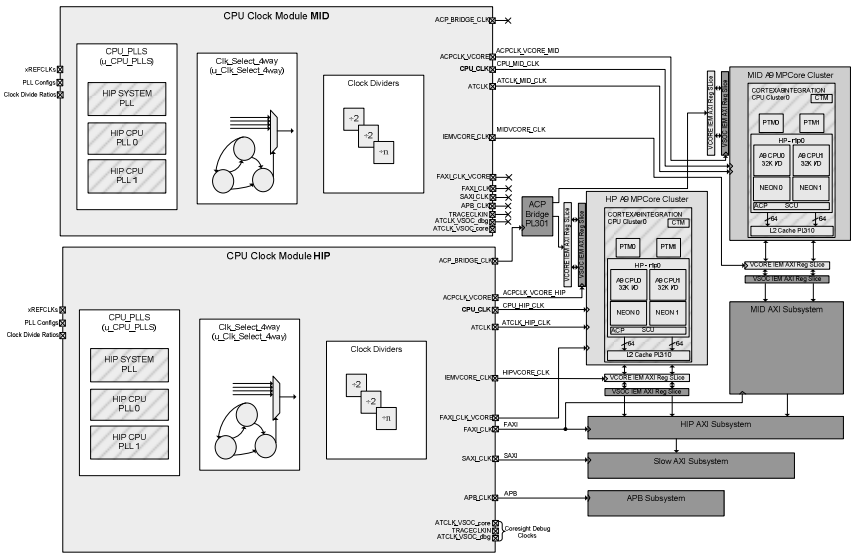

DVFS example switch

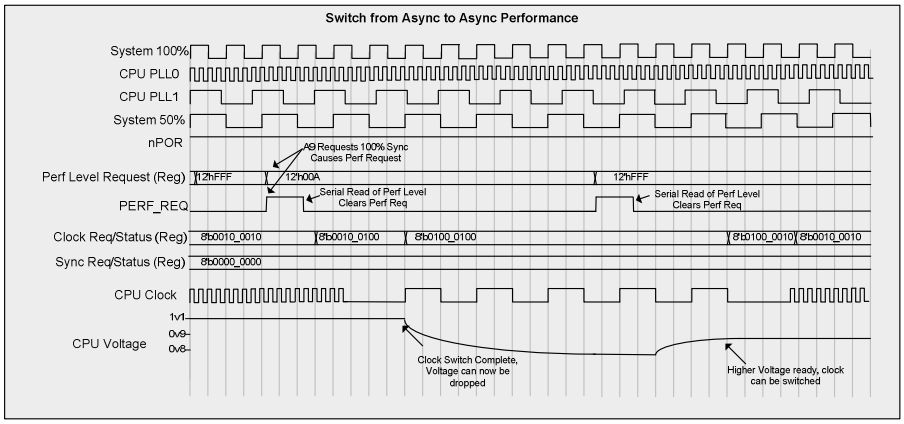

Ongoing Silicon Investigations

- Split lot silicon providing process to implementation correlation
  - The multi-GHz dual-core Cortex-A9 processor is something!
  - The Power-optimized macro is well within mobile power budgets

- Bandwidth monitors on most AMBA AXI bus providing detailed ‘at-speed’ profile data
  - Along with processor performance monitoring and CoreSight components provide significant SoC visibility
- Providing the basis of a “Standard Compute Subsystem”
  - Extending SoC platform standardization to simplify OS portability
  - Providing standard software routines over boot and power management

Video Demo