

Developing the Bifrost GPU architecture for mainstream graphics

ARM


Dan Wilson
Senior Product Manager

ARM Technical Symposia China
October 31 – November 4, 2016

©ARM 2016

Graphics processing drivers

- Virtual reality extending into mainstream devices
 - Augmented realities, virtual spaces and immersive videos
- Vulkan showing great momentum
 - Unity releases Vulkan renderer preview
- Demand for high quality screens and cameras
 - Smartphone resolutions are rapidly increasing



ARM Mali GPUs and their architectures

Current

Mali-G71
Bifrost

Mali-G51
Bifrost

Mali-470
Utgard

Previous

Mali-T860/880
Midgard

Mali-T820/830
Midgard

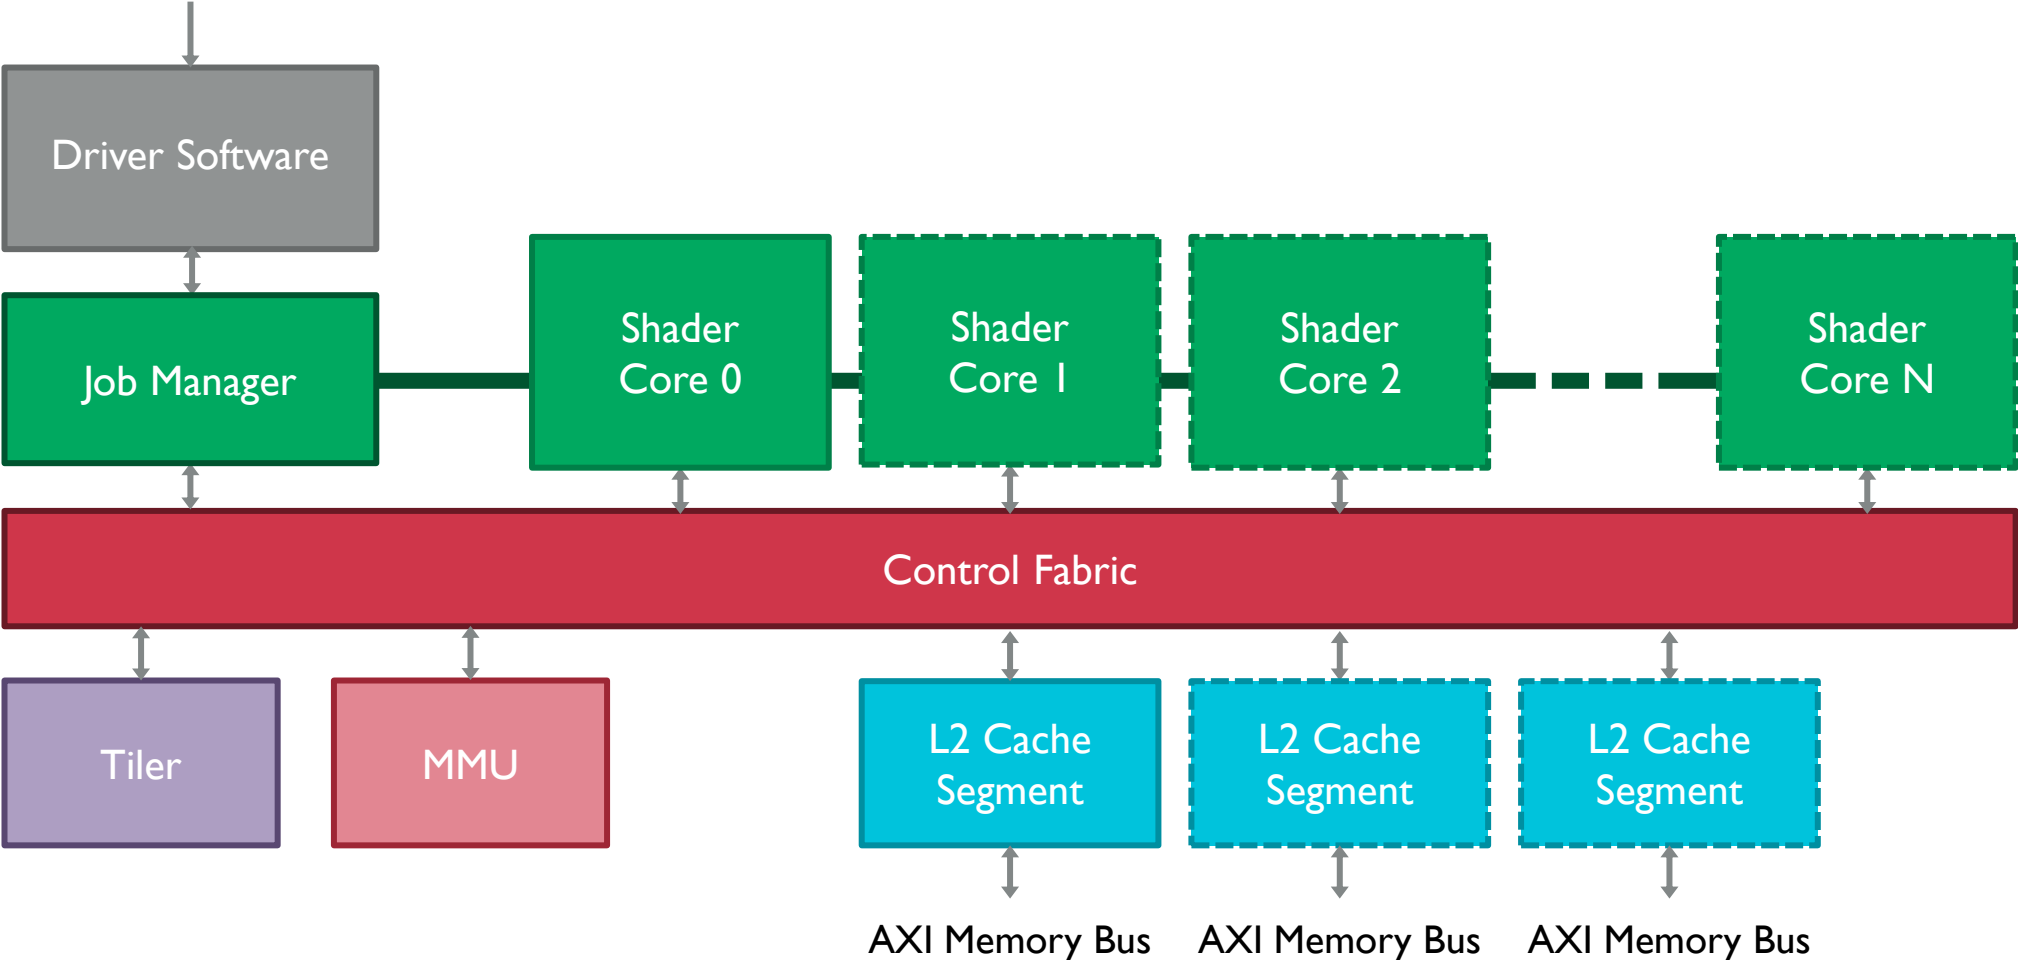
Mali-400/450
Utgard

Performance

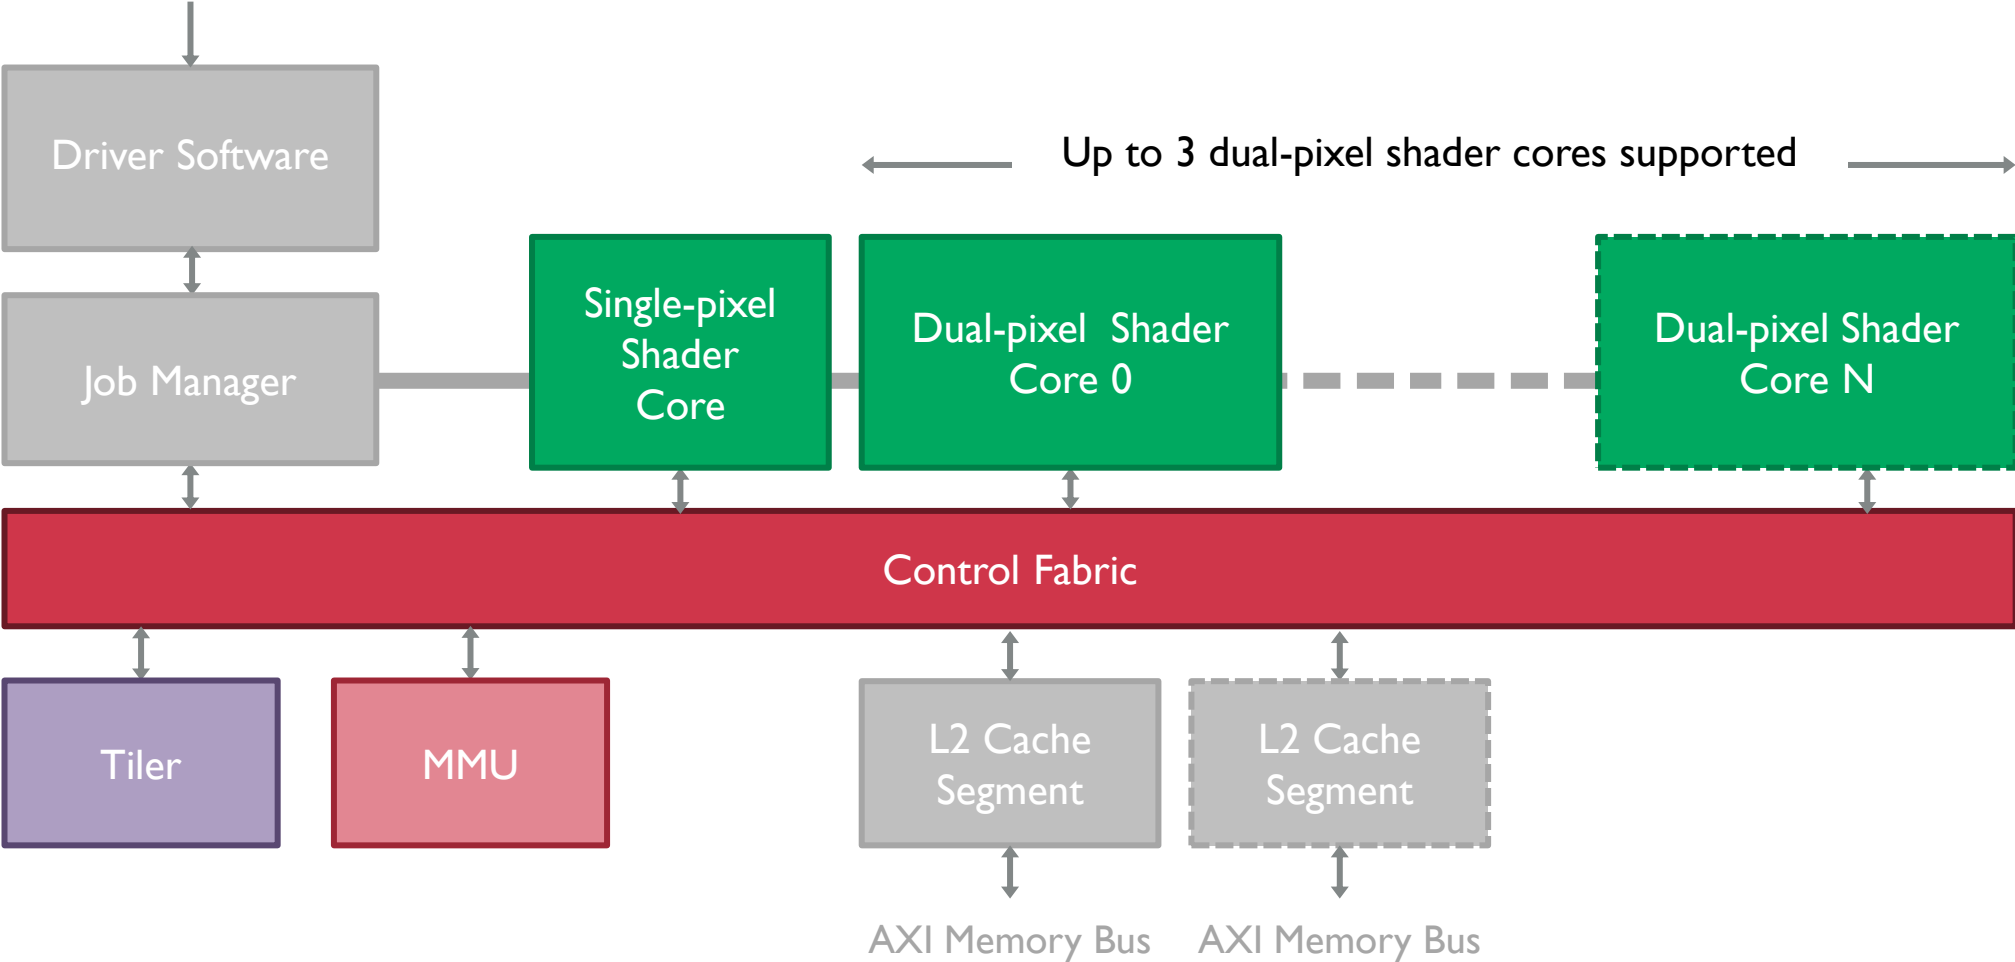
Area-efficiency

Ultra-low-power


Bifrost GPU design



Mali-G51 scalable and area-efficient design



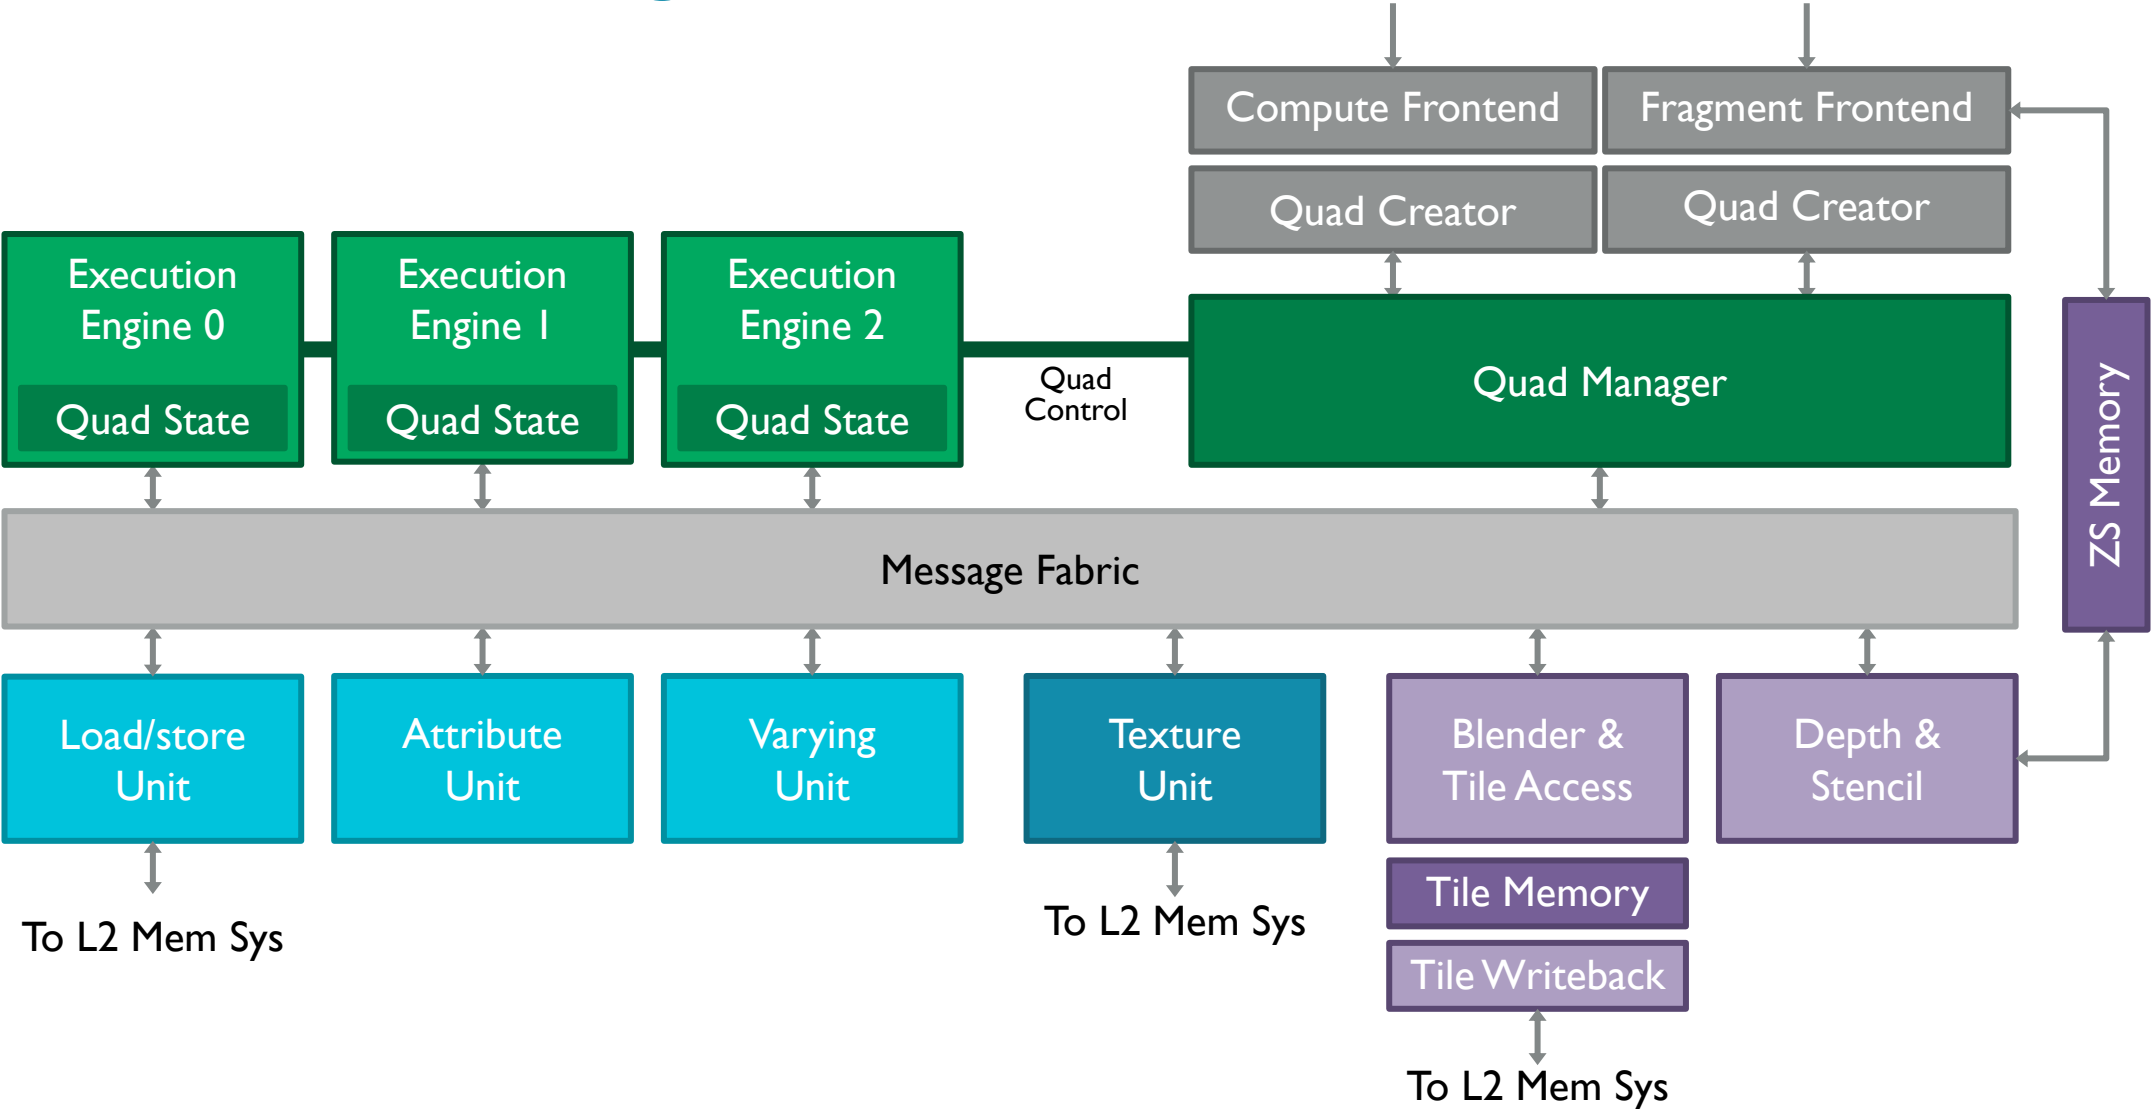
Dual pixel shader core



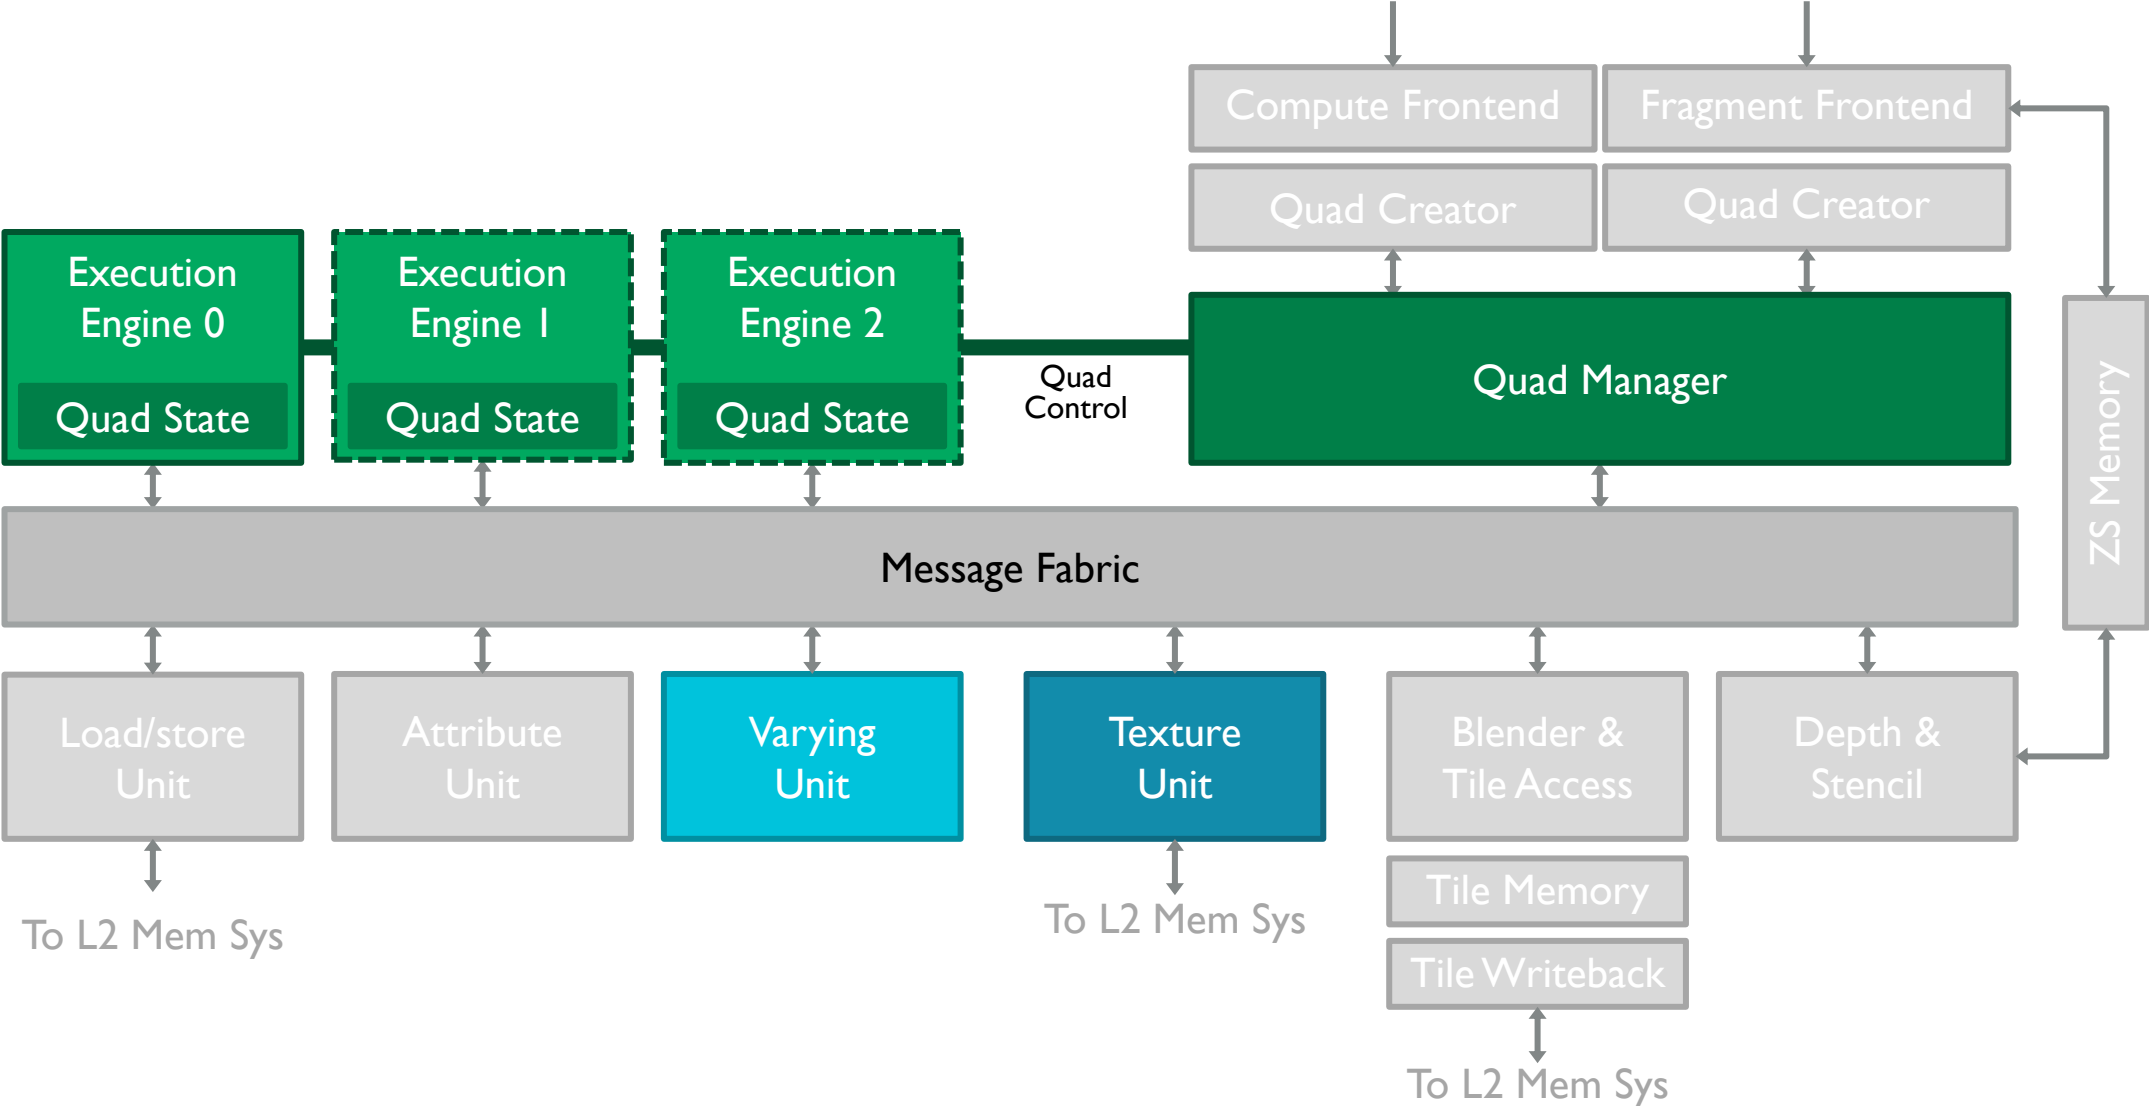
- MP2 and higher configurations use a dual-pixel shader core
 - Increases area efficiency
 - Delivers 2 textured pixels per clock cycle (2-tppc)

Configuration	Architecture			
	Shader core configuration	Execution engines	Texels / Clk	Pixels / Clk
MP1	1x Single core	1	1	1
MP2	1x Dual core	3	2	2
MP3	1x Dual core + 1x Single core	4	3	3
MP4	2 x Dual core	6	4	4
MP6	3 x Dual core	9	6	6

Bifrost core design



Mali-G51 development focus




Texture unit




- The texture unit has been re-designed for Mali-G51
 - Delivers 1 texel/clock for Single core and 2 texels/clock for Dual core
- Designed for increased tolerance to high memory system latency
- Pipeline extensively re-designed to reduce effective pipeline length
 - Reduces die area and power consumption
- YUV texture sampling performance increased
 - Up to 3x the performance

Focus on mainstream content




- Instruction set rebalanced for graphics workloads
 - Added new instructions for shader functions used by more recent graphics apps
- Tailoring capabilities to mainstream configurations
 - Tiler
 - Memory management unit (MMU)

AFBC 1.2



- ARM frame buffer compression (AFBC) has been updated to version 1.2
- DDR-optimized data order for headers and body buffers
 - Improving GPU performance in bandwidth limited scenarios
 - Improving DPU performance in rotation use cases
- Optimized compression for constant color blocks
- Full backwards compatibility with AFBC 1.1 specification

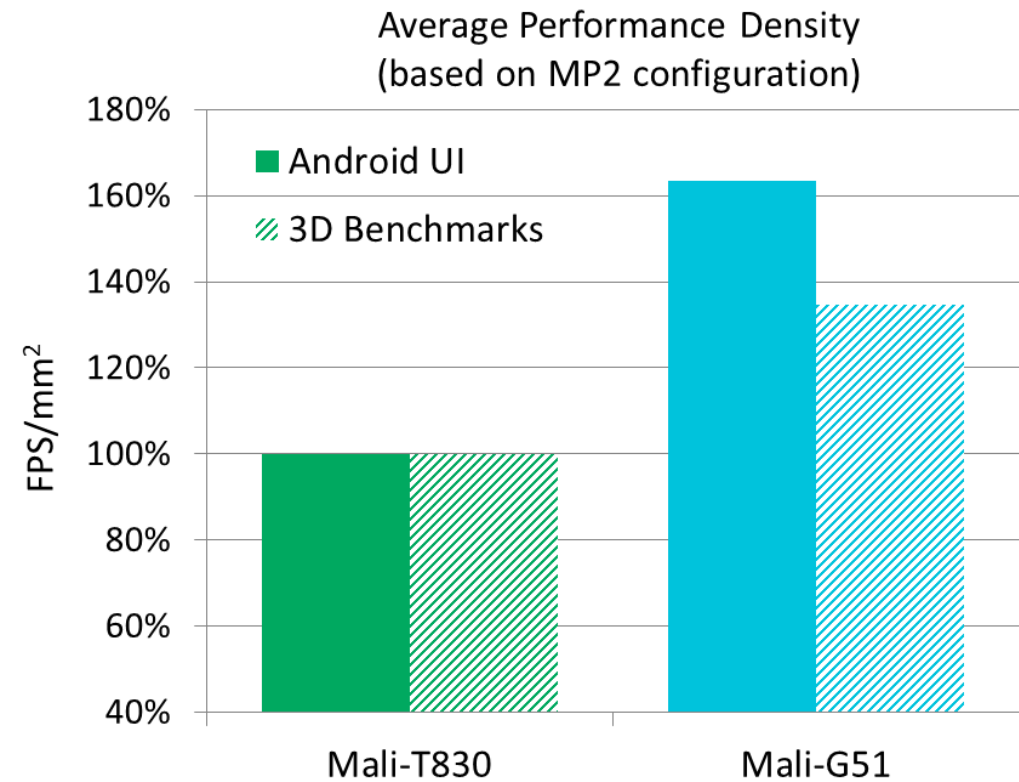
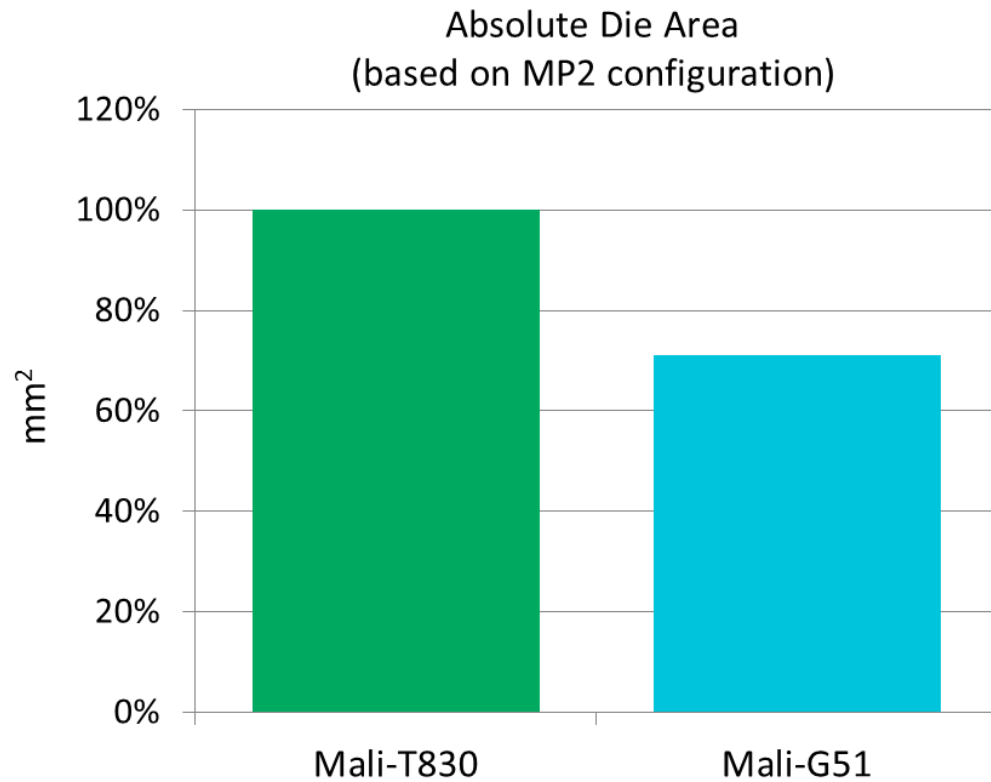
Message preload



- Work is requested from fixed-function units by sending messages across the message fabric
- Messages are normally generated by specialized instructions in the execution engines
- The quad manager can now send messages directly to the varying unit instead of via the execution engines
 - Occurs when the early instructions are varying or varying-texture loads
 - Increases the performance of common fragment shaders



Area and performance

- Measured on Mali-G51 Beta candidate RTL

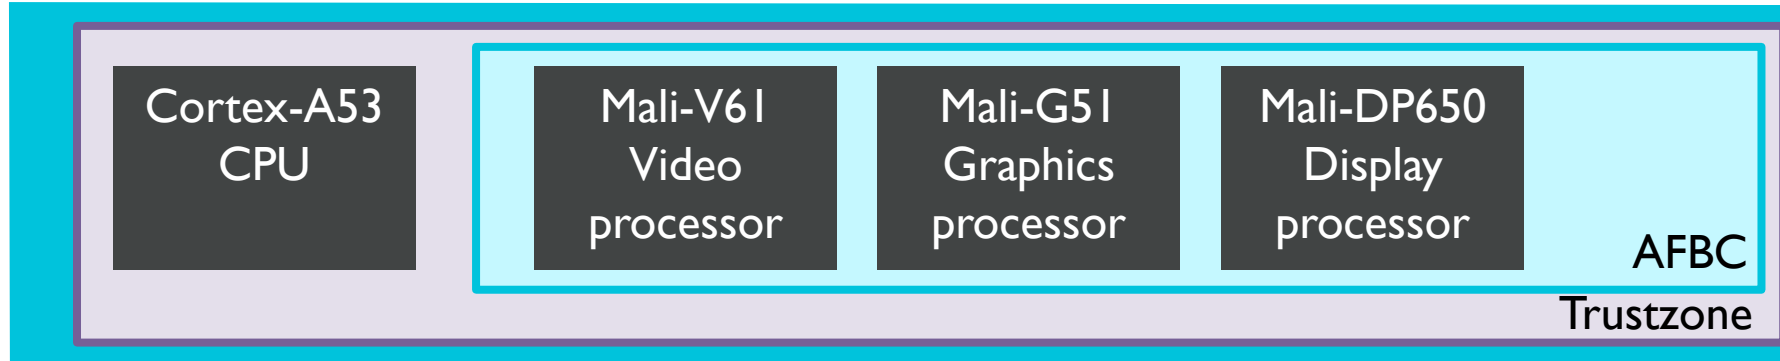


Power


- Measured on Mali-G51 Beta candidate RTL



System level performance advantages




The premium quality, protected, low-latency, low-power multimedia pipeline



Mali-G51 status

- RTL released to lead partners
- Working towards production-quality RTL in December
 - SoCs expected in 2017, devices expected in 2018
- Mali-G51 publically launched at 2016 ARM Tech Symposia
 - The first Bifrost GPU for mainstream devices
 - Up to 1.6x more area and energy efficient
 - The smallest Vulkan enabled Mali GPU
 - For more information [read the launch blog](#)



Thank you

ARM

The trademarks featured in this presentation are registered and/or unregistered trademarks of ARM Limited (or its subsidiaries) in the EU and/or elsewhere. All rights reserved. All other marks featured may be trademarks of their respective owners.

Copyright © 2016 ARM Limited

©ARM 2016