

Design Choices for FPGA-based SoCs When Adding a SATA Storage

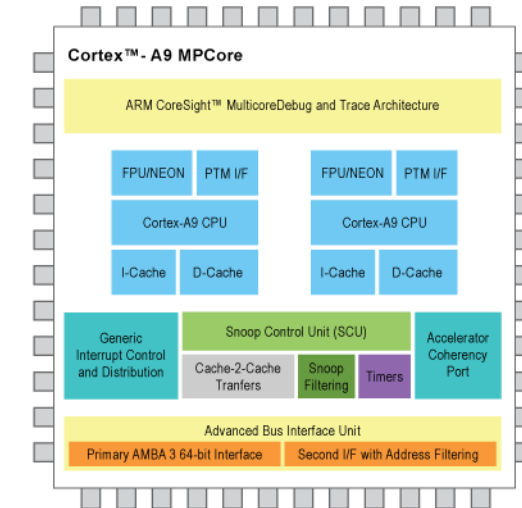
- ▶ Lorenz Kolb & Endric Schubert, Missing Link Electronics
Rudolf Usselmann, ASICS World Services



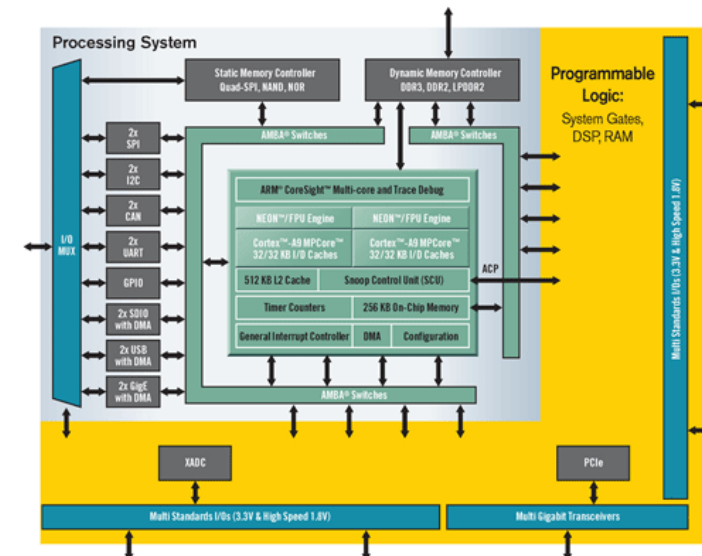
Motivation for SATA Storage for FPGA based SoCs

- ARM Cortex A9MP CPUs: fast enough to run rich OS and software
- Wide range of applications require I/O and signal processing flexibility
 - Data-Logging
 - Test & Measurement
 - Advanced Driver Assist Systems (ADAS)
 - Telematics
 - Machine Visioning
 - Broadcasting Applications
- FPGA based SoC a flexible and cost efficient alternative to embedded PCs

“Leave Your PC Behind!”





Altera SoC



Xilinx ZYNQ

Storage Solutions for FPGA based SoCs

	USB Thumbdrives	Compact Flash	SD Cards	SSD
FPGA Design Requirements	USB 2.0 OTG built-in into most devices	3 rd party IP core with parallel I/O; needs many FPGA pins	SDIO built-in into most devices	3 rd party IP core, needs Gigabit transceivers
Flexibility	Consumer driven, plenty of devices available	Less often used, past it's prime	Consumer driven, plenty of devices available	Consumer driven, plenty of devices available, future proof
Capacity	Typ. 16 GB	Typ. 64 GB	Typ. 64 GB	>256 GB
Performance	30 MB/s	133 MB/s	< 100 MB/s	> 400 MB/s
Design Cost	No extra cost	Extra cost of 3 rd party IP core, needs extra FPGA resources	No extra cost	Extra cost of 3 rd party SATA AHCI IP core, needs extra FPGA resources



Standards Body <http://www.sata-io.org>

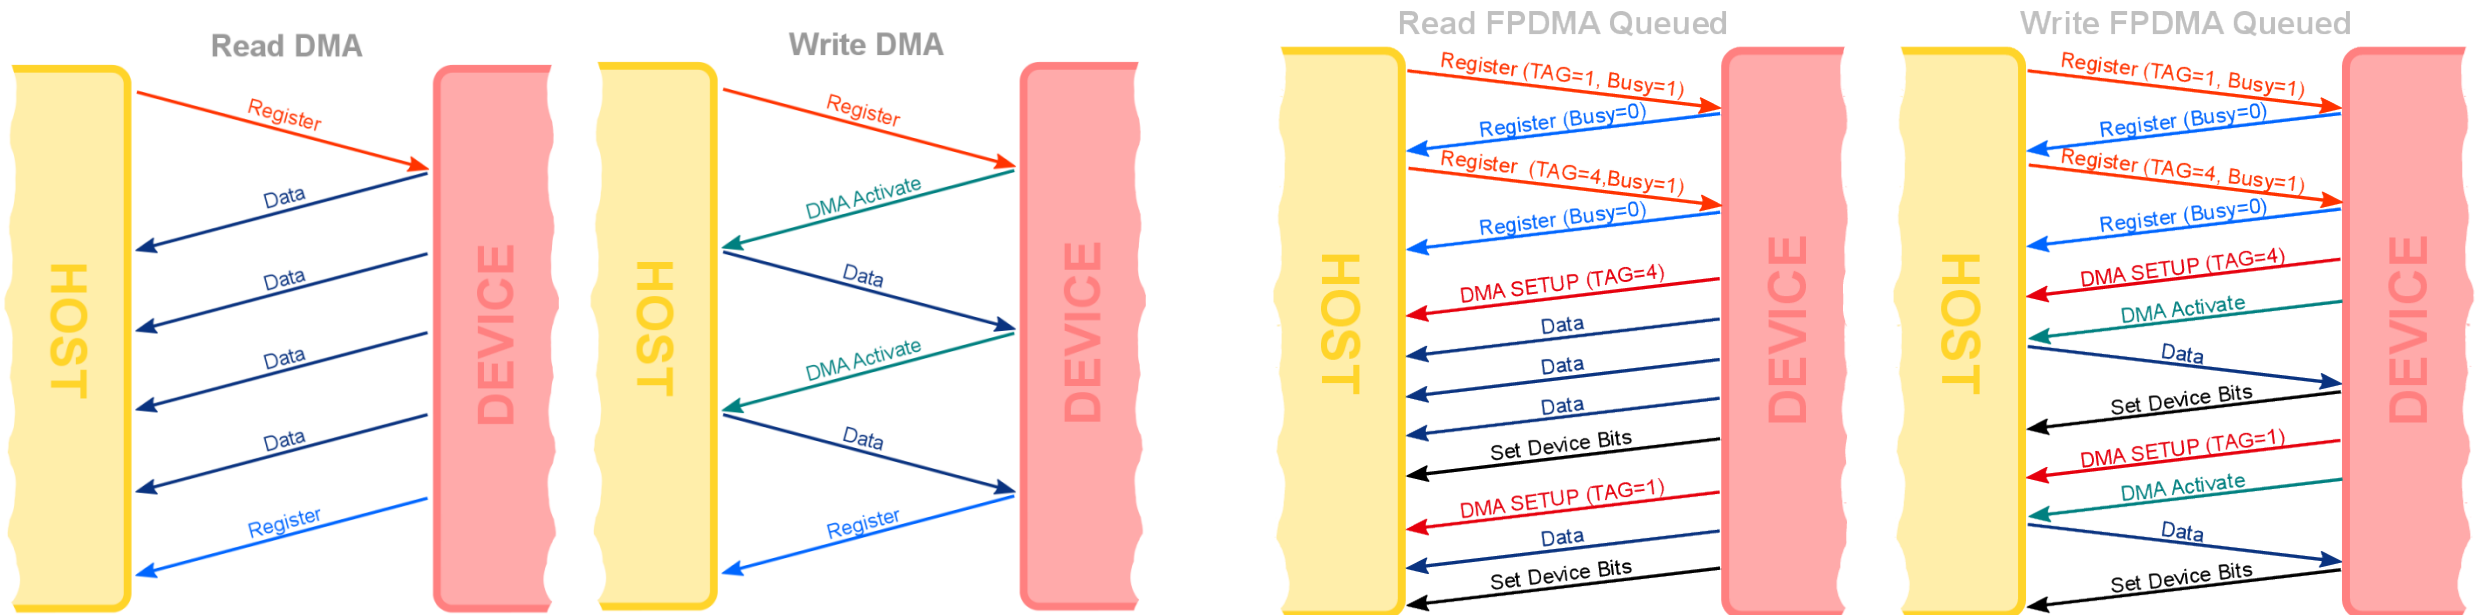
Official naming	In-official naming	Netto-Datarate
Serial ATA 1.5 Gbit/s	SATA I	150 MB/s
Serial ATA 3.0 Gbit/s, SATA Revision 2.x	SATA II, SATA-300	300 MB/s
Serial ATA 6.0 Gbit/s, SATA Revision 3.x	SATA III, SATA-600	600 MB/s


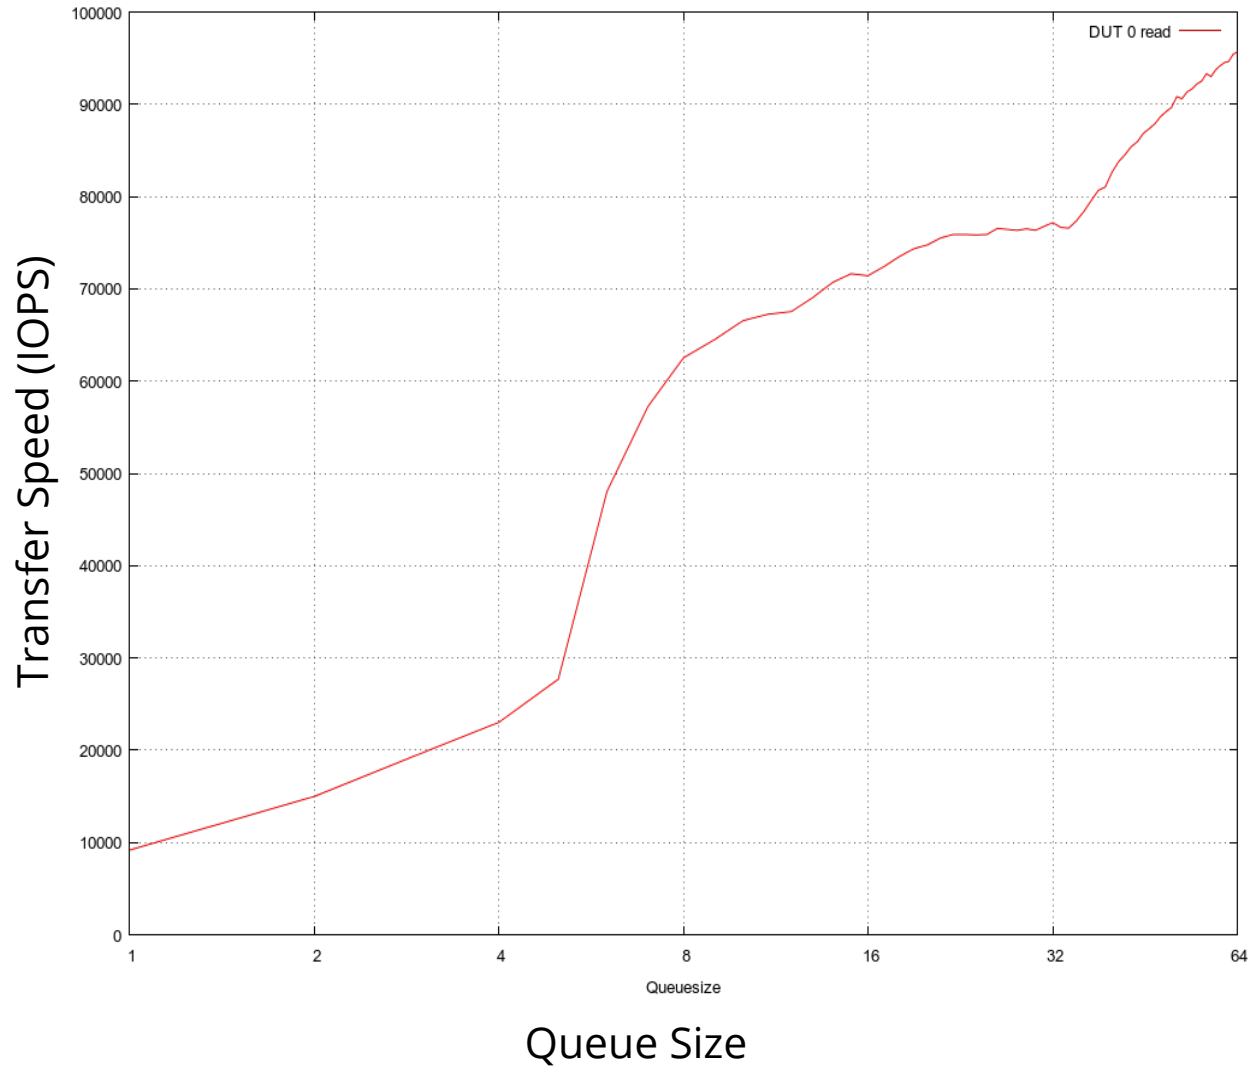
Performance Aspects

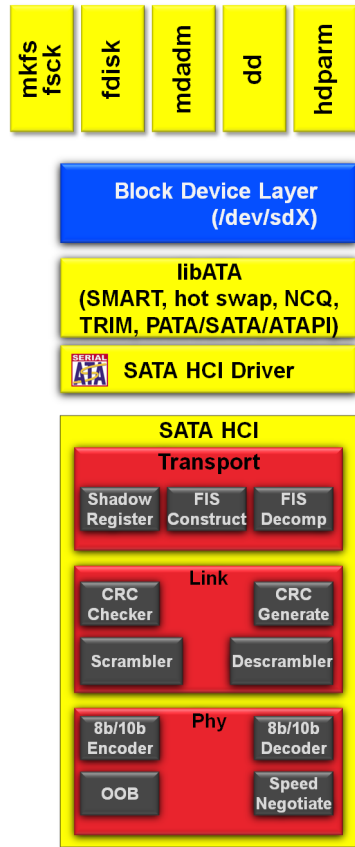
- Bandwidth (MB per second)
 - read vs. write, sequential vs. random access, compression or not
- IOPs (I/O operations per second)
 - Defined by Flash cell types, host and device controller

By default, one Frame Information Structure (FIS) packet gets transferred after the other using standard Direct Memory Access (DMA).

With Native Command Queuing (NCQ), FIS packets can be transferred in an interleaving fashion using First-Party DMA (FPDMA).







Functionality Description

Design Aspect

User Programs

Application software development

Device layer

GNU/Linux Operating System

Application layer

Custom Device Driver

Transport layer

SATA AHCI IP Core

Link layer

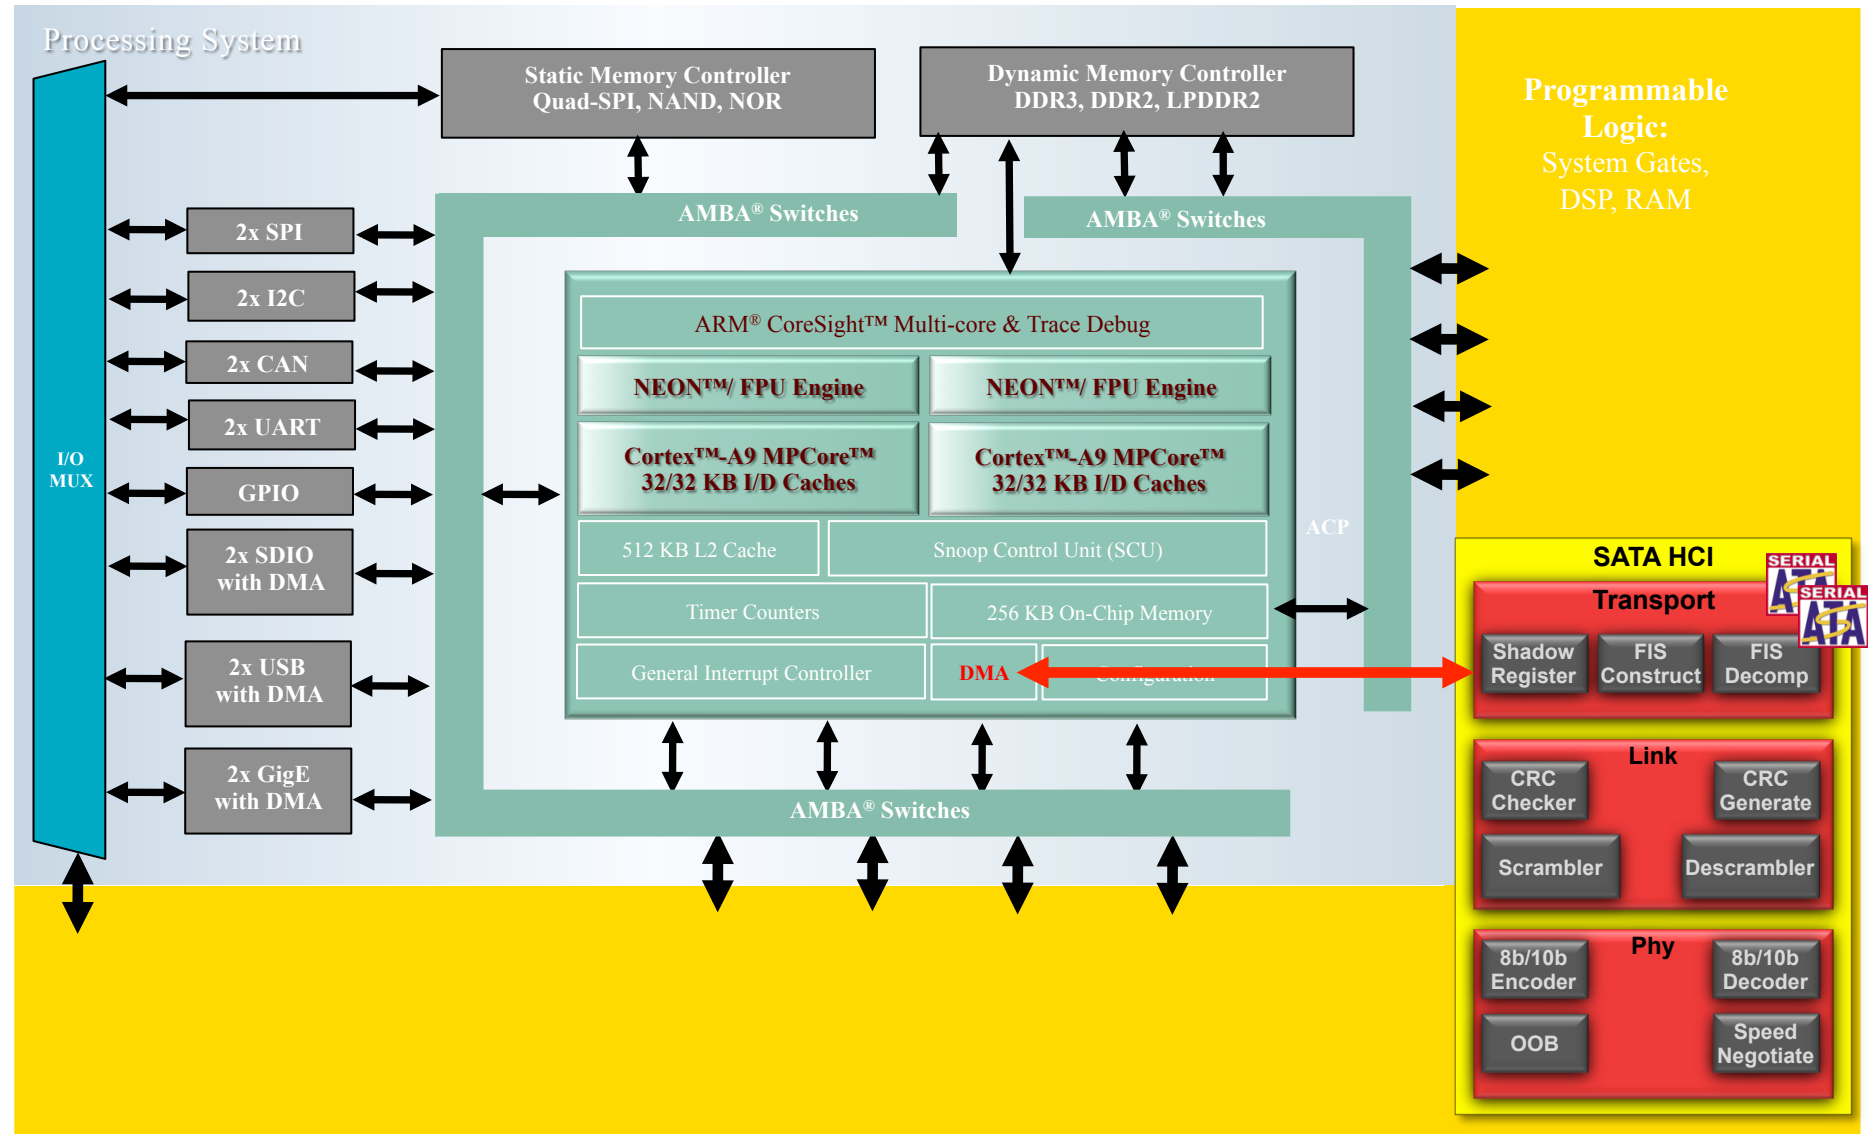
Phy layer

Built-in FPGA high-speed Gigabit Transceivers

Architecture Option 1: Built-in PL330 DMA Controller

Full functionality but very limited performance without hardware support for NCQ

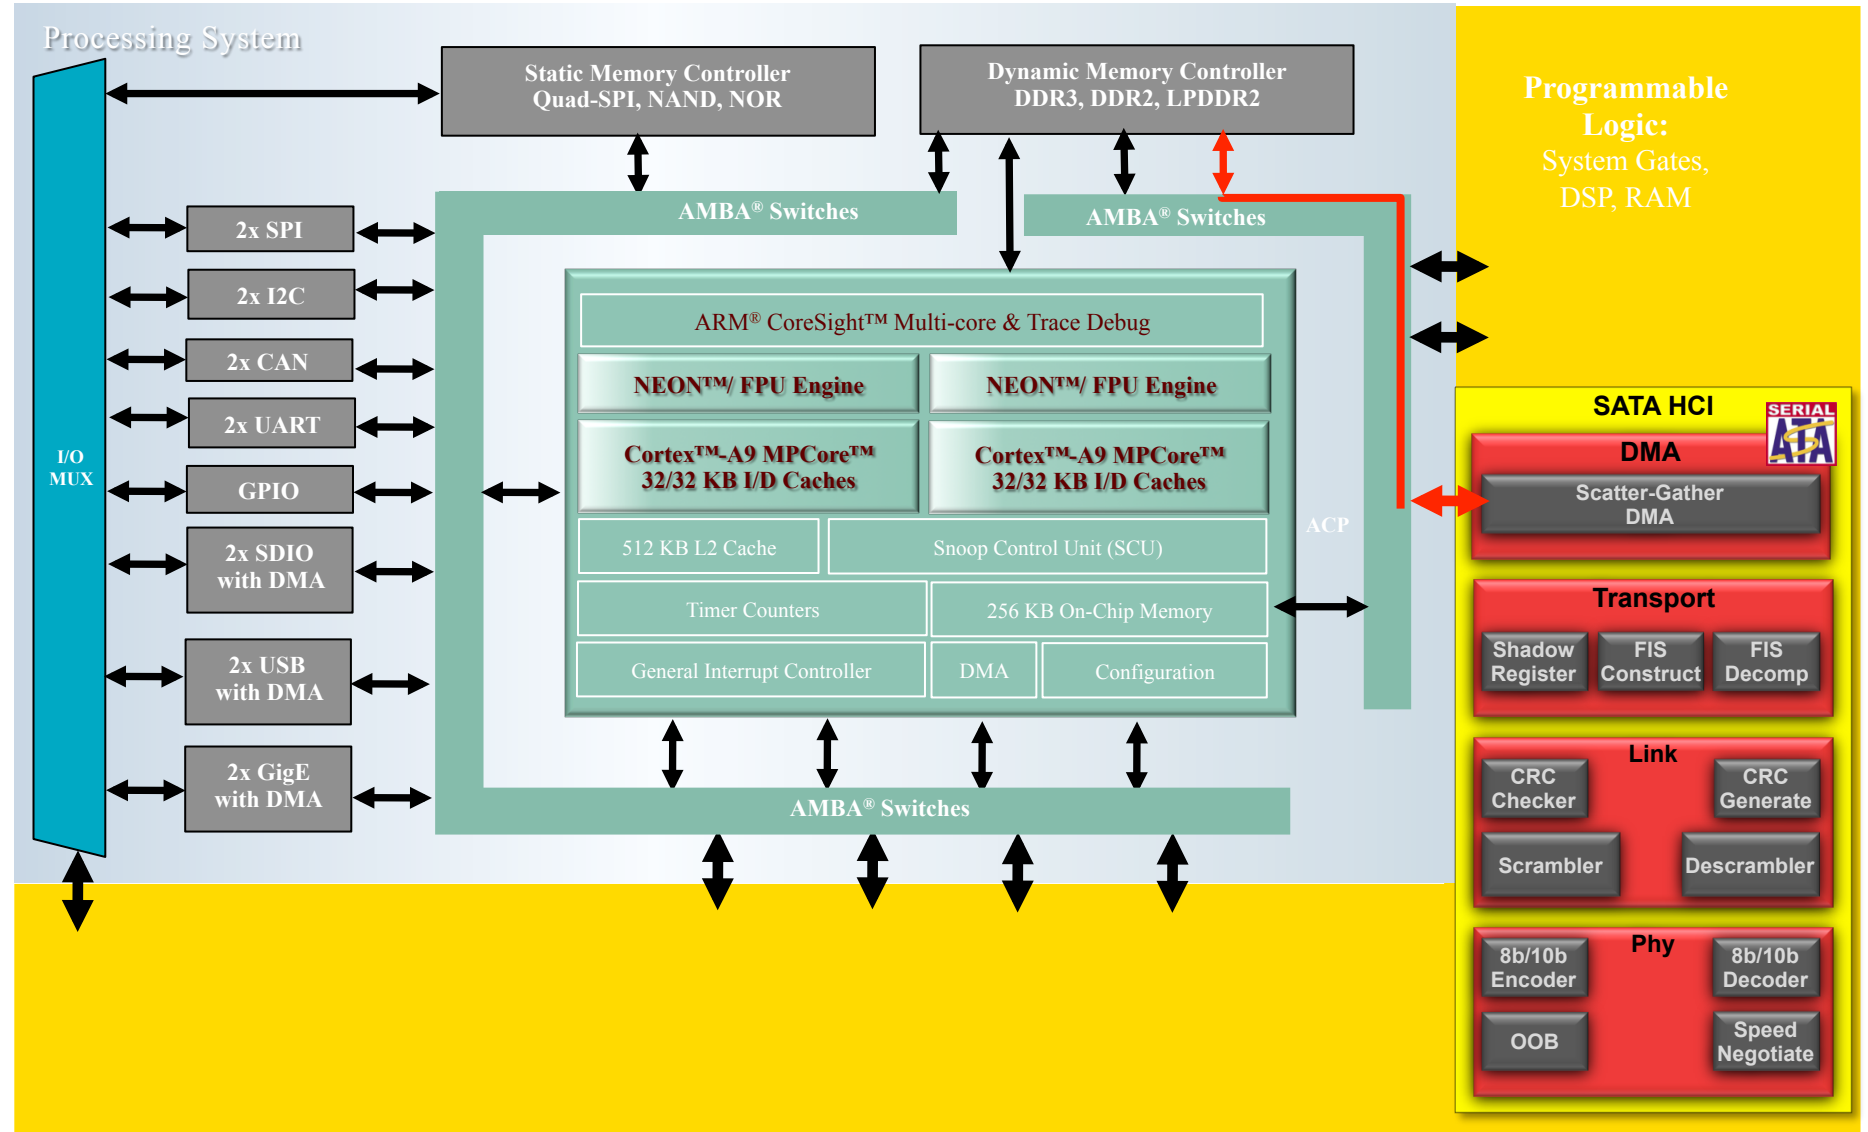
Scatter-Gather support needs software work for Linux



Architecture Option 2: 3rd Party Scatter-Gather DMA Controller

Significantly better performance, but does not "max out SSD"!

SSDs internal structure demands NCQ for full R/W performance

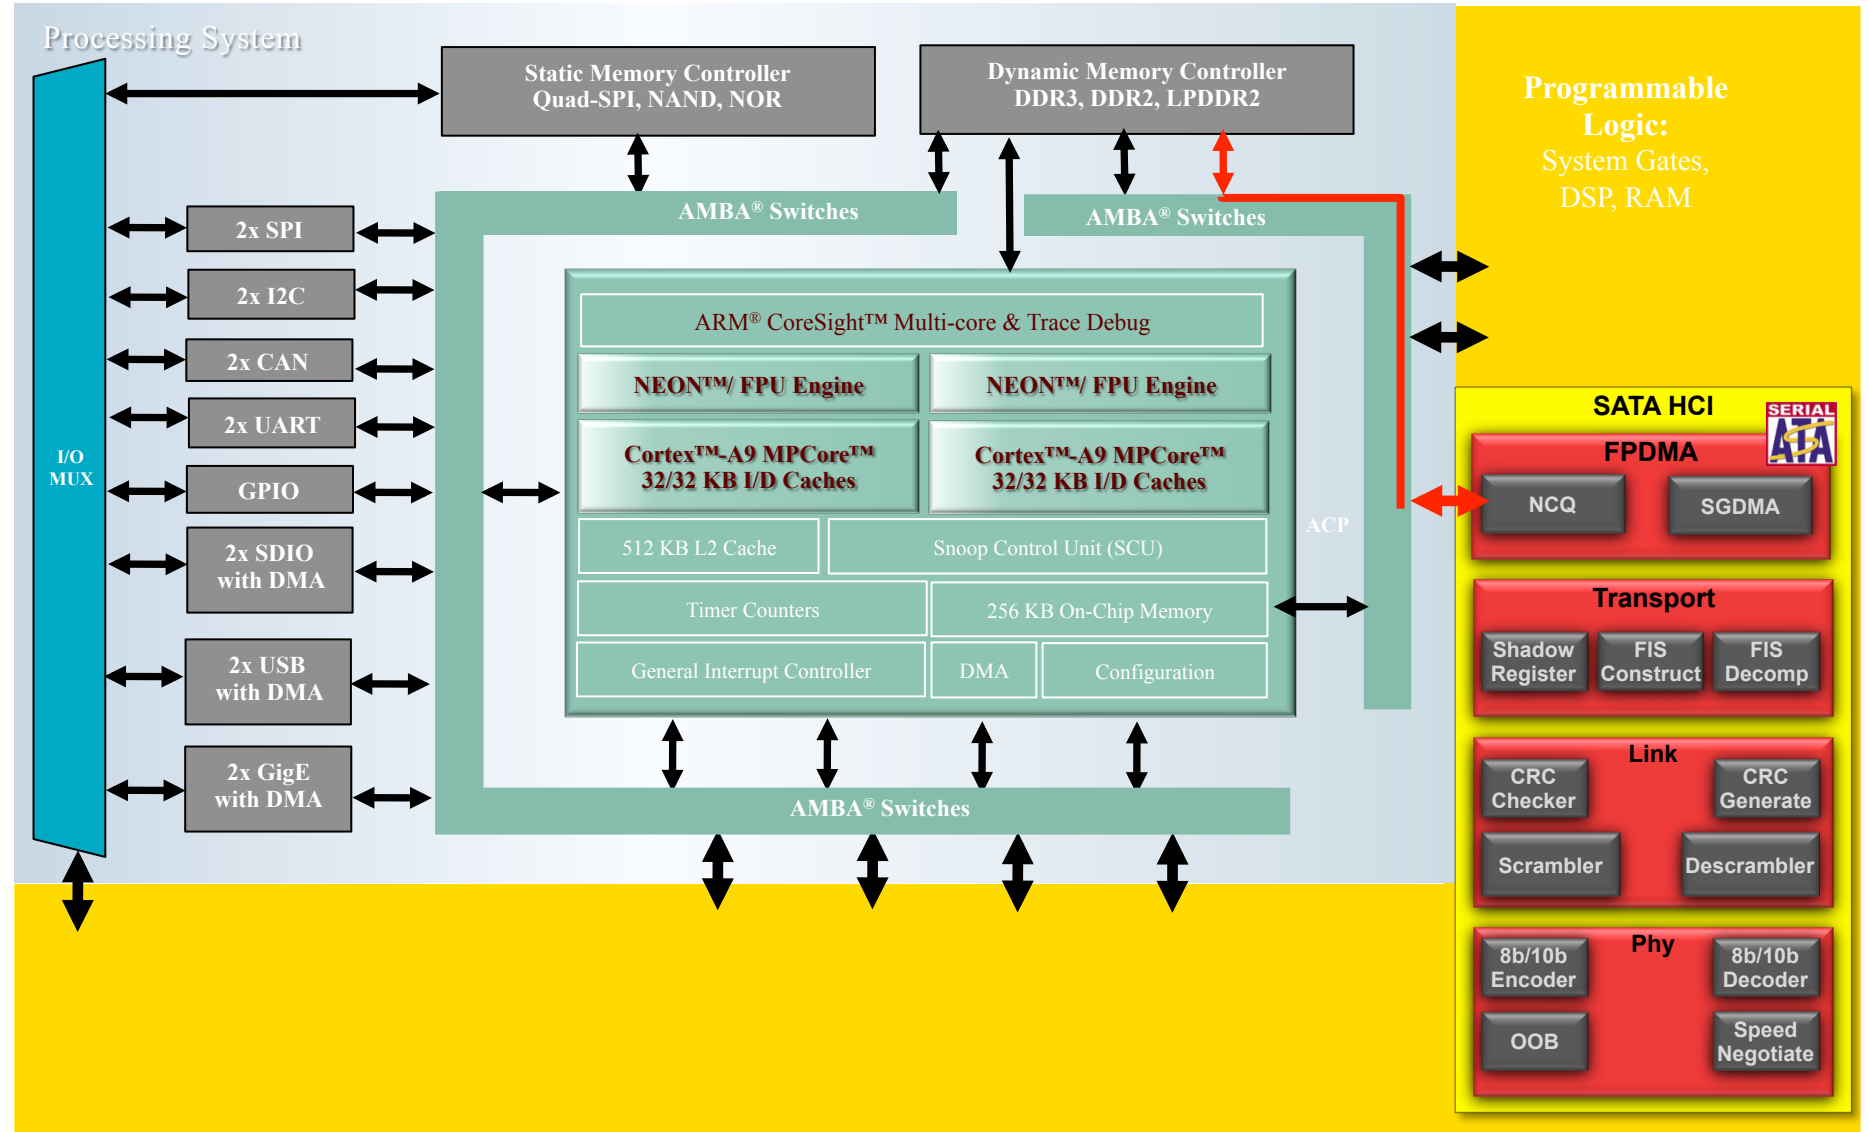


Architecture Option 3: 3rd Party FPDMA Controller with NCQ

Close to "max out SSD"!

Depending on application, bottleneck is in software (IOPS)

This can be caused by too many IRQs

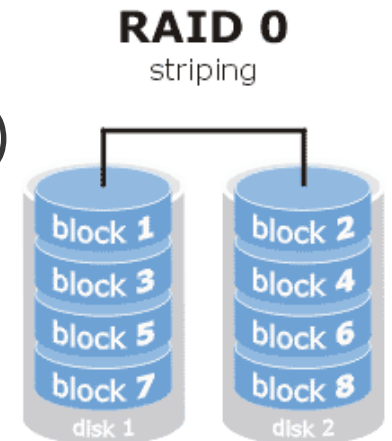


Command Completion Coalescence (CCC) to reduce interrupt and command completion overhead in heavily loaded systems [Serial ATA AHCI 1.3]

- Frees up CPU
- Increases #IOPs
- → Functionality of 3rd party SATA AHCI IP core

Striping over multiple SATA links into multiple SSDs (a.k.a. RAID-0)

- Software-RAID will not work due to CPU load
- → Extra functionality in 3rd party SATA AHCI IP core



- Building performance storage solutions for FPGA-based SoCs is more than an IP core!
- Requires a fine-tuned micro-architecture properly integrated into the software system.
- Benefits of pre-validated FPGA subsystems.

→ **See us at Xilinx booth**
Hall 1/1-205

