



Author:

ISSUED TO:

COPY NO.

GFXIP_9x SX

Micro-Architecture Specification

Rev 1.0 – Last Edit: 12-Feb-14


THIS DOCUMENT CONTAINS CONFIDENTIAL INFORMATION THAT COULD BE SUBSTANTIALLY DETRIMENTAL TO THE INTEREST OF AMD THROUGH UNLICENSED USE OR UNAUTHORIZED DISCLOSURE.

1. **Preserve this document's integrity:**
 - ⇒ **Do not reproduce any portions of it.**
 - ⇒ **Do not separate any pages from this cover.**
2. **This document is issued to you alone. Do not transfer it to or share it with another person, even within your organization.**
3. **Store this document in a locked cabinet accessible only by authorized users. Do not leave it unattended.**
4. **When you no longer need this document, return it to AMD. Please do not discard it.**

"Copyright 2011, Advanced Micro Devices, Inc. ("AMD"). All rights reserved. This work contains confidential, proprietary to the reader information and trade secrets of AMD. No part of this document may be used, reproduced, or transmitted in any form or by any means without the prior written permission of AMD."

AMD, the AMD Arrow Logo and combinations thereof are trademarks of Advanced Micro Devices, Inc. PCIe is a registered trademark of PCI-SIG. HDMI is a trademark of HDMI Licensing, LLC.

AMD (NYSE: AMD) is a semiconductor design innovator leading the next era of vivid digital experiences with its ground-breaking AMD Fusion Accelerated Processing Units (APU). AMD's graphics and computing technologies power a variety of devices including PCs, game consoles and the powerful computers that drive the Internet and businesses. For more information, visit <http://www.amd.com>.

 AMD	ORIGINATE 29-Nov-11	EDIT DATE 12-Feb-14	DOCUMENT-VER. NUM. 1.0	PAGE 2 of 35
--	------------------------	------------------------	---------------------------	-----------------

Revision History

Date	By	Revision	Description
2/12/2014	Dong Yiao	0.1	Initial Copy from gfx8
2/12/2014	Dong Yiao	0.1	First edit for out-of-date contents



Table of Contents

PREFACE 4

1 INTRODUCTION 5

1.1 DEFINITIONS / GLOSSARY OF TERMS 5

1.2 TOP LEVEL DIAGRAM 5

Figure 1: SX in chip context 6

1.3 **R7.x REQUIRED FEATURES** 7

 1.3.1 SX support for 2 PC/GDS redirect busses instead of 1 in 1 SE configurations (16 PIX PER SH) 7

 1.3.2 SX support for 4 RBs. (16 PIX PER SH) 7

 1.3.3 SX support for deeper Position buffer, Color buffers and position alloc storage (16 PIX PER SH) 7

 1.3.4 Streaming of performance counters (STREAM PERF CNTRS (GPU.01 & GPU.02)) 7

1 DETAILED CHANGE DESCRIPTION 8

1.1 SX SUPPORT FOR 2 PC/GDS REDIRECT BUSES INSTEAD OF 1 IN 1 OR 2 SE CONFIGURATIONS (1 SH PER SE ONLY) 8

 1.1.1 Testing 8

1.2 SX SUPPORT FOR 4 RBs. 8

 1.2.1 Testing 9

1.3 SX SUPPORT FOR DEEPER POSITION BUFFER, COLOR BUFFERS AND POSITION ALLOC STORAGE 9

2 PERFORMANCE 9

3 POWER 9

4 HARDWARE IMPLEMENTATION: TOP LEVEL 10

4.1 TOP LEVEL DRAWING 10

Figure 2: SX Top Level Diagram 10

4.2 ADDRESSING THE BUFFERS 10

4.3 FORMAT OF THE DATA 11

4.4 CONTROLS OF AN EXPORT 11

4.5 COLOR EXPORT 12

 4.5.1 The color scoreboard 12

 4.5.2 Export buffer address computation 12

Figure 3: An example of addressing the export buffer 13

4.6 POSITION EXPORTS 13

4.7 REDIRECT EXPORTS 13

5 HARDWARE IMPLEMENTATION: INTERFACE 14

5.1 SHADER CORE INTERFACES (SPI/SQ/SP) 14

 5.1.1 SH_SPI_EXPREQ 15

 5.1.2 SPI_SH_EXPGRANT 16

 5.1.3 SQ_SX_EXPCMD 16

 5.1.4 SPI_SX_EXPADDR 18

 5.1.5 SX_SPI Free Signals 18


 5.1.6 SH_SX_VDATA 19

5.2 SX TO PC INTERFACES 20

 SX_PC_EXPCMD 20

5.3 SX TO PA INTERFACES 21

5.4 SX TO DB INTERFACES 23

 AMD	ORIGINATE 29-Nov-11	EDIT DATE 12-Feb-14	DOCUMENT-VER. NUM. 1.0	PAGE 4 of 35
--	------------------------	------------------------	---------------------------	-----------------


Preface

A micro-architecture (or block level) specification serves many purposes.

- The micro-architecture specification is used by design verification teams to build block and system level verification environments (test benches, test scenarios, test cases and testing methodologies/strategies). Although, a document is clearly not the ONLY vehicle used by design verification teams, it is a critical piece of the verification planning process in that it provides valuable data to support the test plan meetings between design and verification teams.
- The micro-architecture specification is also a useful tool for peer and block design reviews. Interfacing blocks require explicit details of control and handling of data being transferred, transformed between blocks. The specification is the obvious resource that peer teams go to for this type of information.
- This specification is also used by post-silicon verification teams and in the creation of documentation (such as programming guidelines, etc.) that must be prepared for external customers.
- This documentation is also useful when designs are transferred to other teams. For example, derivatives of a graphics core are used in many other products that include an integrated core, hand held devices, etc.

Each new major architecture should contain a micro-architecture specification for each block in the subsystem. In the case where a design is derived from a previous project, that previous project specification would be updated and checked into the new project revision control documentation area. All the delta features would be described in the feature section of the micro-architecture specification and the document in general should be updated to match the new project updated block design.

A template is provided as a means of describing the detail required by all teams that use the micro-architecture specification and to drive consistency between the specifications from one block to another. This template was created by the Design Verification Workgroup which is comprised of representatives across all AMD Business Units. This template has been and will be distributed to other Business Units and Design, Design Verification and Architecture teams for review and feedback. This template was created from a review of many of the existing micro-architecture specifications. Examples are pulled from these documents and presented here to illustrate the type of information that is required. To distinguish between the descriptions of content and examples, all examples appear in *italics*.

 AMD	ORIGINATE 29-Nov-11	EDIT DATE 12-Feb-14	DOCUMENT-VER. NUM. 1.0	PAGE 5 of 35
--	------------------------	------------------------	---------------------------	-----------------

1 Introduction

The SX (shader export) block is responsible to receive and re-order color and position exports from the shader core and forward them to the correct client: PA for position, and the correct DB for color. The SX is also a conduit for parameter and GDS exports and in this role forwards the data unchanged to the PC (parameter cache) block.

The main input for the SX block is the shader output bus which is 2 busses each 16x32 bits wide. The SX output busses are:

- 1) 128 bits wide bus to the PA (primitive assembler) block supporting 1 position per clock
- 2) 256 bits to each DB (supporting up to 4 DB per SX) thus supporting 4 "compressed" pixels per clock (64 bpp) or 2 uncompressed pixels per clock (128 bpp).
- 3) 16x32 bit bus to PC/GDS.

1.1 Definitions / glossary of terms

***Thread** – one instance of a shader program being executed on a vector of pixels, vertices, or primitives. Each thread has its own state which is unique from any other thread.*

***Clause** – a group of instructions all of the same type (all ALU, all texture-fetch, etc.) executed as a group; part of a thread.*

***Wave** – one instruction operates on a wave of pixels/vertices/primitives over 4 clock cycles. This is the basic unit of work. The size of a vector depends on the system configuration, but is always 4 clock cycles.*

***SIMD** – Single Instruction, Multiple Data. Here, a SIMD refers to one slice of the SP machine which all receives the same instruction and operates on a vector of data. Most implementations will have multiple SIMDs. Each SIMD receives a separate instruction from the SQ.*

1.2 Top Level Diagram

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.