

Tech Brief: AMD FirePro™ SDI-Link and AMD DirectGMA Technology

V 1.0, Alexis Mather – 9 September 2011

Introduction

With the increasing prevalence of GPU-based video processing in professional video domains, it is increasingly critical that low-latency (real-time), high-throughput links be established between the various components of video processing systems. This is particularly true for the GPU itself and the interface boards providing video input and output capability from those systems.

This technology brief outlines AMD's approach to resolving the problem of real-time interconnects between GPU and video I/O, in the form of AMD FirePro™ SDI-Link. We start by examining the areas where this type of low-latency, high-throughput video processing is of interest. Next, a full picture of the AMD FirePro SDI-Link platform is painted, describing the key components of this platform. Finally, we review the areas where AMD FirePro SDI-Link can shine, and take a look at the prospects for wide adoption of this approach.

SDI and Graphics Processing in Video Pipelines

Serial Digital Interface (SDI) is a virtually ubiquitous standard for inter-connecting equipment (cameras, displays, recorders, servers, mixers, keyers and other devices) particularly in broadcast but also post-production. With the accelerating use of virtual sets in production and pre-production, SDI has increasing relevance in professional video pipelines.

SDI Input and Output (I/O) in PC-based systems (now the vast majority of systems in broadcast) is typically provided by board-level products that offer direct connection to SDI or connections to break-out-boxes that provide connections on the desktop or in rack-mounted systems. Some of the key vendors for such systems include AJA Video Systems, Blackmagic Design, Bluefish444, Deltacast, DVS, and Matrox.

A key consideration that we will explore is the need for GPU-based rendering and processing to be tightly coupled with SDI I/O to and from systems in these pipelines. Given the real-time nature of many of the workflows, it is critical that the path between the GPU doing the work and the devices that receive and transmit be as fast and unencumbered as possible. Specifically, the latency – defined as the time between the request of an operation and its completion – must be as low as possible.

Broadcast

Broadcast workflows vary greatly, but many (if not most) of them share the need to keep latency to a minimum. An example is virtual sets, often used in election coverage and increasingly in newscasting. More examples include:

- On-air/Live broadcast graphics
- News/Sports/Commentary/Analysis
- Virtual Sets
- Real-time Video/Digital Cinema Production

Production & Post

Post-production has long benefitted from having SDI as a standard interconnect. A key reason for this is that some key stages of post-production (e.g. on-line editing, color grading, finishing) require the viewing of content on 'reference' displays. Such professional displays (which can include 4K projectors for film work) have SDI inputs, benefitting from SDI's use of standardized color spaces. Low latency is critical in these workflows as it can help ensure that PC-displays and the reference monitors used alongside them are kept completely in sync (a lack of sync during playback is extremely distracting).

Increasingly, broadcast techniques (such as virtual sets) are meeting with formerly offline pre-visualization techniques for increasingly interactive on-set, live content creation, leading to GPU-assisted production techniques that also require SDI in this space.

Examples of use cases include:

- On-set 3D VFX pre-viz
- On-set/daily color grading and review
- Real-time video/D-Cinema post-production
- Final conform, color grading, finishing
- Real-time video format transcoding

Today's Challenge

A key issue facing system developers (including integrators and software vendors) today is how to achieve the lowest-latency connection between the GPU and video I/O board, while simultaneously benefiting from the capabilities offered by world-class SDI board vendors (flexible number of I/O ports, robust support for latest 3G-SDI standard).

Today, approaches for GPU-and-Video I/O interoperability fall into two broad categories:

Approaches that allow for low latency but that offer little choice in SDI capabilities.

These are typically proprietary board-level solutions built by integrators or other vendors. While these solutions can meet the latency requirements, the frequency with which they are updated can lead to a lack of flexibility as standards and market needs evolve.

Approaches that are based on video I/O boards from industry-recognized SDI board vendors

Industry veteran vendors have, as a core business, to meet the I/O requirements the market demands: flexible numbers of inputs and outputs, timely support for standards as they emerge, robust compatibility with other SDI devices. Here, real-time considerations have not been core to the design, and latency challenges still typically exist.

Unfortunately, to date, the choice above between the two approaches has left a marked, unmet need for low latency GPU-and-video interoperability with highly flexible, up-to-date SDI I/O support. It is this unmet need that AMD has spent time understanding and looking to address.

The AMD FirePro™ SDI-Link Platform

The AMD FirePro SDI-Link platform was developed in response to the challenges facing the professional video markets described above. The key goal in developing the platform has been to remove the compromise that faces solution providers between low-latency and flexible video I/O products.

AMD has developed key technology to address the latency problem, reached out to leading industry partners for their world-class SDI solutions, and delivered a new product (the AMD FirePro™ V7900 SDI professional graphics card); these come together in the AMD FirePro™ SDI-Link platform to enable solution providers to deliver the very best to their own customers.

A fundamental part of AMD's strategy in implementing AMD FirePro SDI-Link has been to reach out to industry partners who are well established in this space; they both understand customer requirements for SDI Input/Output, and meet them comprehensively with their board-level products.

The four key components of AMD FirePro SDI-Link are:

- AMD DirectGMA technology
- A software development kit (SDK), with access to AMD's engineering support team
- Support from 3rd party SDI I/O board vendors and certification for their products¹
- Supported AMD FirePro™ professional graphics cards, starting with the AMD FirePro™ V7900 SDI

AMD DirectGMA Technology

AMD DirectGMA (for **D**irect **G**raphics **M**emory **A**ccess) technology is built upon the PC industry's full adoption of the PCI-Express standard. PCI-Express allows for a rapid interconnect between the 'host' (that encompasses a CPU, main memory, and general I/O control) and additional devices (including GPUs, storage controllers, network controllers and the like). In addition to the host-device relationship, PCI-Express allows for devices to communicate directly with one another over PCI-Express.

AMD DirectGMA technology implements PCI-Express 'peer-to-peer' transfers. The 'peers' are the AMD FirePro™ professional graphics card and the 3rd party's SDI I/O board; they are peers relative to the host system. Peer-to-peer transfers are extremely interesting in the context of achieving low latency because these transfers completely bypass any need to traverse the host's main memory, utilize the CPU, or incur additional redundant transfers over PCI-Express.

¹ Certification program expected to launch in Q4, 2011.

AMD DirectGMA technology includes two key pieces: the hardware capability to drive the PCI-Express transfers through a GPU engine dedicated to the purpose, and the driver infrastructure to make sure the capability is exposed to software developers.

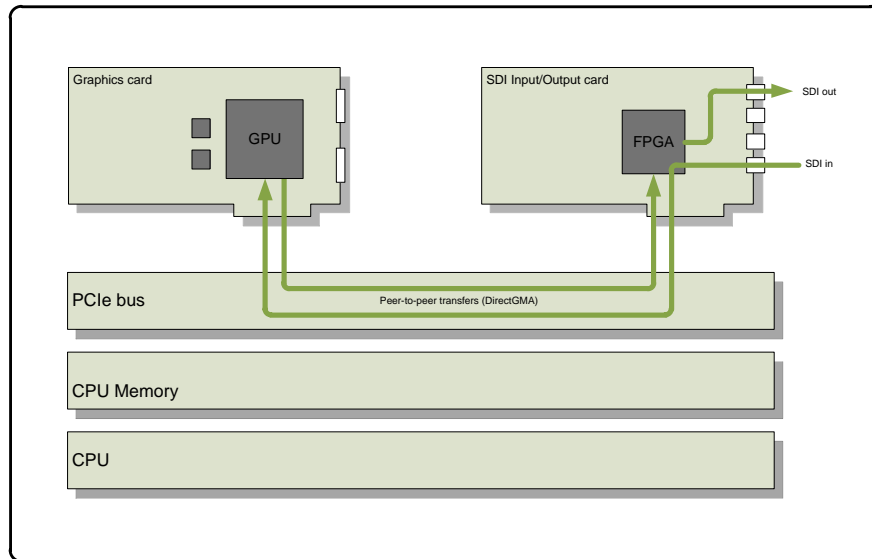


Diagram illustrating the AMD DirectGMA Technology in action. Note that redundant passes through system memory are avoided.

Software Development Kit (SDK)

The AMD FirePro™ SDI-Link SDK is a comprehensive set of documentation and samples to help software developers through the process of adopting AMD FirePro SDI-Link and AMD DirectGMA technology in their pipelines.

Control over AMD DirectGMA technology is done in two parts: handling incoming video data (from SDI I/O board) and transferring GPU-rendered content to the SDI I/O board. Both of these directions are handled by simple graphics API extensions. Some portions of the transfers also need to be handled via the SDKs of the 3rd party I/O board vendors.

For a software developer who already has an OpenGL pipeline using a preferred SDI I/O board, the changes required to move over to AMD DirectGMA technology are extremely simple; initial implementations have been complemented in a matter of a mere few days of effort.

The AMD FirePro™ SDI-Link SDK will support OpenGL, DirectX® and OpenCL™ APIs².

SDI I/O Board Certification

In order to help give integrators and software partners a starting point for their own solutions, AMD will maintain a certification program to document validated configurations of AMD products with 3rd party products and to track validated platforms and motherboards, along with software and firmware revisions.

AMD understands the need for rigorous validation of solutions in broadcast and production, making AMD DirectGMA technology and its support from AMD FirePro™ SDI-Link technology partners, an extremely attractive for low-latency solutions.

AMD FirePro™ SDI-Link Capable Graphics Cards

The AMD FirePro™ SDI-Link platform represents a significant innovation for customers in the broadcast and production space. A key part of delivering the capability to these demanding customers is creating products specifically tailored to them.

The *AMD FirePro™ V7900 SDI* professional graphics card, launched at the same time as AMD FirePro™ SDI-Link, is the first such product from AMD. It is anticipated that other products will join the AMD FirePro™ line of products supporting the AMD FirePro SDI-Link platform.

The New GPU-SDI Landscape

We saw earlier that the key challenge the broadcast, production and post-production industry has faced to date is the lack of SDI solutions that are both flexible and offer low latency interoperability with GPUs. Having now taken a look at the AMD FirePro™ SDI-Link platform, and particularly the AMD DirectGMA technology that underpins the core capability, it should be apparent that this collaboration-based approach to professional video pipelines:

- Meets the need for low latency interoperability between GPUs and SDI I/O boards
- Allows customers and integrators the flexibility to pick their SDI I/O vendor and the capabilities they need
- Is well supported from an SDK and board certification point of view.

² OpenCL support implement via OpenGL-OpenCL interoperability already in AMD's graphics drivers.

In short, the AMD FirePro™ SDI-Link platform introduces an era where GPU-accelerated video pipelines no longer have to make a compromise between low latency and flexible SDI I/O connectivity.