RADEON PRO

Solid State Graphics (SSG) Technical Brief

Introduction

Modern GPUs are continually evolving to keeping pace with the growing demands of users. Methods and algorithms such as ray tracing for high-end rendering and non-linear video editing place high demands on the GPU's graphics compute capabilities. These demands extend to the system level when applications require large data sets to be used, for example in cases such as 8k/16k video playback or media servers where multimedia contents are pushed to the next level of photorealism.

The GPU is designed to handle enormous amounts of data. However, at the system level, large data sets are not instantaneously available to the GPU. The Radeon™ Pro SSG (Solid State Graphics) solution strives to move one step closer to providing instantaneous access to these large data sets.

RADEON PRO SSG

Exploiting Local NVM to Access Large Data Sets

In general, non-volatile memory (NVM) is connected either through the SATA or PCIe® system interfaces with the system processor (CPU) acting as an intermediary.

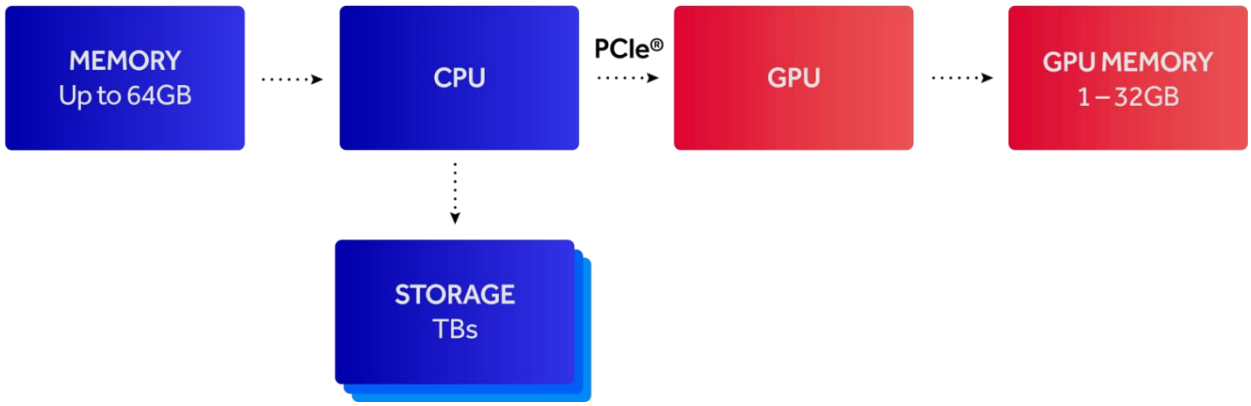

Figure 1: Traditional system architecture for mass storage.

To ensure that the requests from all storage devices are serviced, all communication must traverse the PCIe® root complex with proper acknowledgement from the CPU. Data transfer between two endpoints is typically performed through CPU memory in three stages:

1. Data is read from one PCIe® endpoint to operating system file cache (CPU memory)
2. Data is copied from operating system file cache to process memory (CPU memory)
3. Data is written from process memory to the other PCIe® endpoint

If data exists in one endpoint device but needs to be processed in another endpoint device, then the rate of completion may be limited by transport speed, rather than processing speed. Data that has to be transported must also be copied in multiple locations before reaching its final destination as shown in Figure 2 below.

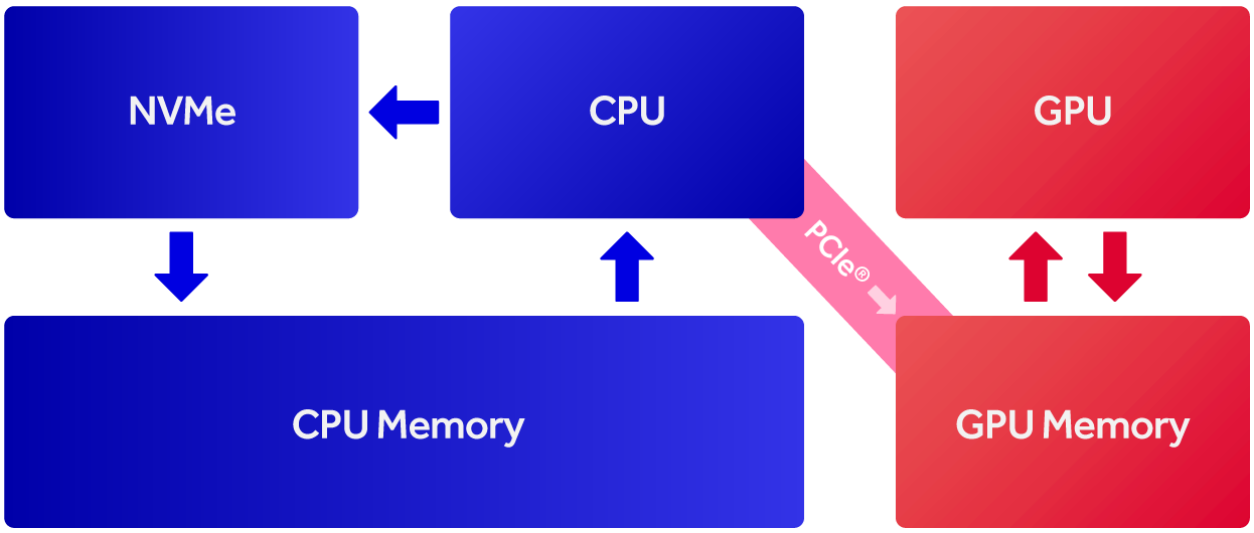

Figure 2: Data path for typical graphics transaction necessitates many copy transactions that are expensive to process.

Large Dataset Access in Radeon Pro SSG

Certain data operations require interaction from the CPU/root complex in order to remain coherent; however, if the integrity datasets can be managed locally, then peer to peer (P2P) communication can be employed removing the need for the data to traverse the CPU. The main advantage of this is the ability to utilize a low latency data path to effectively move large data sets from a memory endpoint to a processing endpoint (GPU). The Radeon™ Pro SSG solution takes this concept one step further by physically co-locating high performance NVMe with a high performance GPU to create a fully optimized path for transferring and processing of large data sets.

The following list describes the general steps required for peer-to-peer read operation:

1. Reserve and pin some part of the PCIe® exposed video memory
2. Instruct the SSD to perform a read operation, supplying the bus address of the exposed video memory
3. Process the data on the GPU.

These steps permit traditionally bandwidth-intensive transactions to be handled within the SSG storage architecture (figure 3).

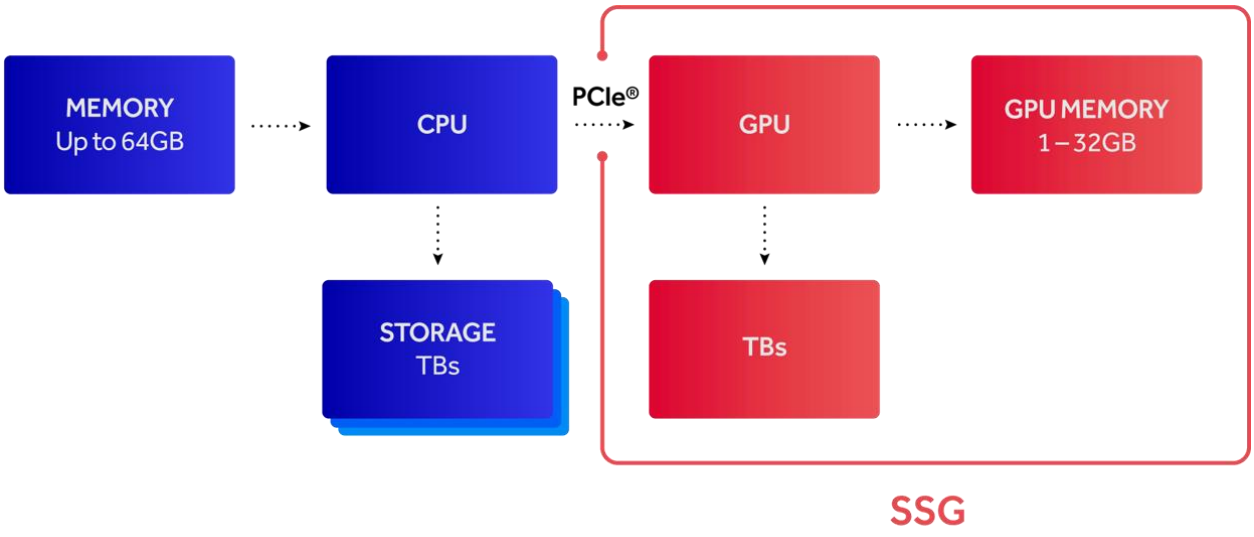

Figure 3: The system storage architecture with Radeon™ Pro SSG.

Once the path between the GPU and NVMe is established, an additional benefit of the Radeon Pro SSG solution is revealed. Because the CPU is removed from the path of the data being transferred, it is free to service other tasks, thus helping improve the overall efficiency of the system.

Radeon™ Pro SSG impact on User Experience

Raw 8K Video playback

Raw 8K video content is currently the most advanced standard for video playback. High-end cameras have the capability of capturing reality in 8K frames, but current playback solutions are often unable to display these frames at full frame rate (>30 FPS) to ensure a smooth or real-time experience. As a result, video playback or editing is often jerky and slow, leading to longer editing times and user frustration.

With the Radeon Pro SSG, video playback performance can not only meet, but exceed, real-time performance requirements because of the card's ability to process and present this same data at speeds that far exceed typical solid state mass storage. In addition, smooth scrubbing through raw 8K content now becomes possible as the user has almost instant access to any position in the video stream.

By being able to handle these large data streams, the AMD Radeon Pro SSG card can empower the user to play back, manipulate, stitch and post-process raw 8K content similar to how current solutions handle 4K content today.

Raytracing

Raytracing renders a 3D image with photorealistic results by fully simulating the interaction of light as it encounters a virtual object; this technique creates large data sets that must be transformed as the viewer's perspective changes. These data sets can grow exponentially as additional lights, objects and perspectives are accounted for.

The nature of raytracing means that it is best suited for applications that operate with some pre-rendering stages in their workflow. Prior to the Radeon Pro SSG graphics solution, pre-rendering of complex objects or scenes was sluggish as a rule. Real time interactions with these scenes was not possible due to the high latencies involved in prefetching the associated data sets. With Radeon Pro SSG technology, the transport time of these data sets can be minimized to improve latency and responsiveness during manipulation of the scene.

Performance Advantages

As an example of the performance increases which are possible with the Radeon Pro SSG card, consider the example of 8k Video playback below. Here you can see that the best case for a workstation equipped with an incredibly fast PCIe® SSD is 17 FPS. In contrast, the Radeon Pro SSG architecture takes that to 92 FPS—a 5.4X performance increase (figure 4).[1]

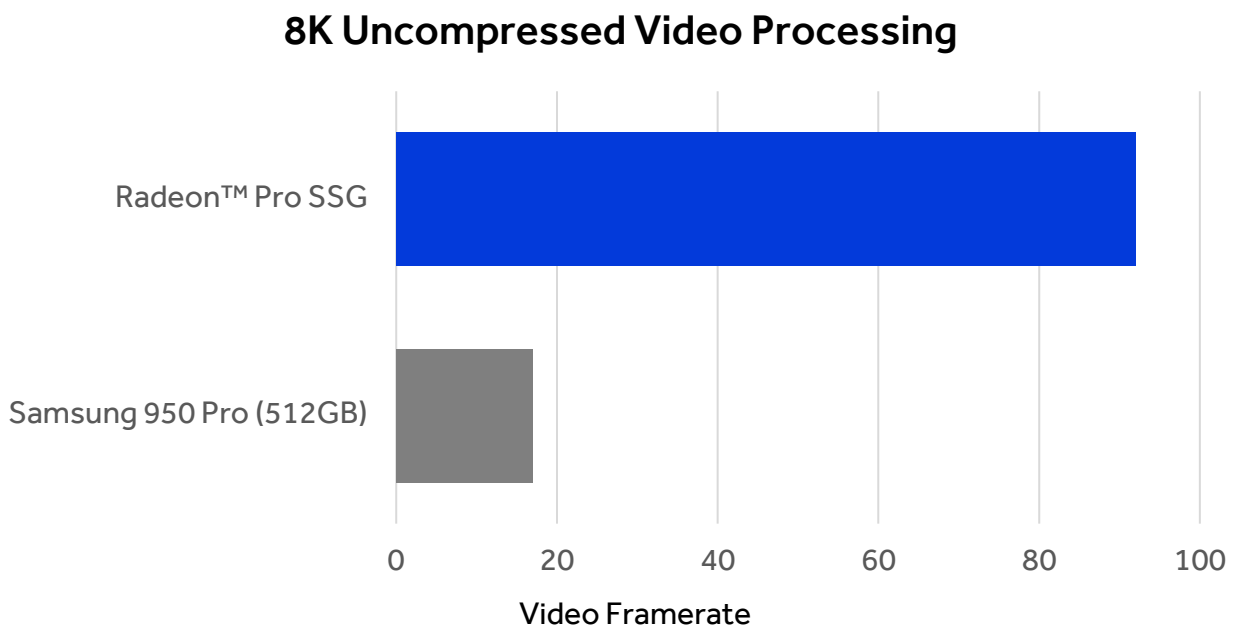

Figure 4: *Radeon™ Pro SSG enables over 5X the performance of traditional SSDs in raw 8K video playback.*[1]

Summary

The ability to process data quickly is at the root of virtually every engineering and design problem. As datasets grow in size and complexity, the task of processing this data falls on the GPU. GPUs have been largely keeping up with this challenge, but external factors like platform topology restricts the GPU's ability to access large data sets effectively.

The Radeon™ Pro SSG Developer Kit gives developers a solution without these restrictions. The large data sets and the powerful GPU now reside within close proximity of each other. The path between these two endpoints is further optimized to create a dedicated, fast lane for information, removing bottlenecks when presented with enormous data sets and ensuring that developers have the power to easily work with the ever increasing data sets their applications demand.

Legal Attributions

---

[1] Based on AMD internal testing as of 7/25/2016. System configuration: Core i7-4790K, 16GB DDR3-1600, Gigabyte GA-Z97X-Gaming5, Samsung 850 Evo (OS drive), Windows 10 x64. Software configuration: Radeon Software 16.30.2010-160629n-304071E, AMD LiquidFlashPlayer media player. Raw 8K video framerate from M.2 NVMe Samsung 950 Pro 512GB SSD: 17 FPS @ 0.85 GB/s sustained data rate. Raw 8K video framerate from Radeon Pro SSG: 92 FPS at 4.61GB/s sustained data rate. RPW-1