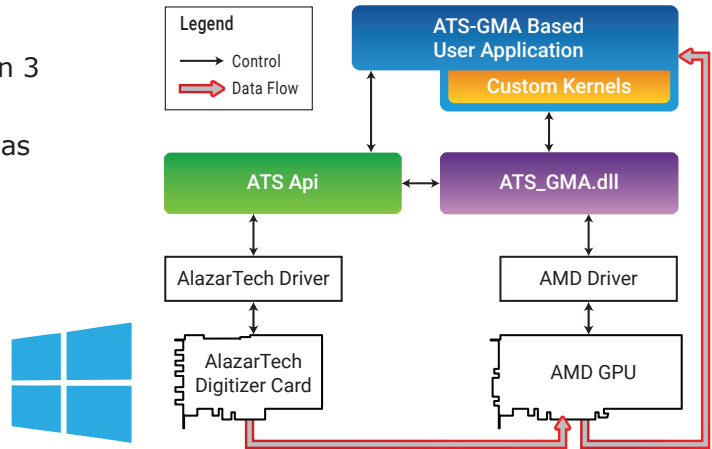


- DMA A/D data to GPU at high speed
- Up to 6.9 GB/s transfer rate for PCIe Gen 3 digitizer boards
- Stream data to GPU with latency as low as 100 μ s
- Supports AMD Radeon™ Pro GPUs and AMD SDK version 2.9 and higher, which support OpenCL™
- Write your own OpenCL kernels
- Designed to work with AlazarTech PCI Express waveform digitizers
- Compatible with Windows® 10



Product	GPU Compatibility	Operating System	Kernel Programming Language	Throughput to GPU	Latency
ATS-GMA-BASE	AMD Radeon Pro	Windows 10 64-bit	OpenCL	Up to 6.9 GB/s	~100 μ s

Overview

Finally, you can DMA A/D data from a PCIe based waveform digitizer to a GPU without any CPU involvement.

ATS-GMA-BASE is a software library that allows users to DMA data from AlazarTech PCI Express waveform digitizers to AMD Radeon Pro Graphical Processing Units (GPUs) at rates up to 6.9 GB/s with a latency as low as 100 μ s. ATS-GMA-BASE does not use any host memory buffers for temporary storage.

With this very low latency and high speed data transfer, it is now possible to acquire signals at 4 GS/s sample rate and 12 bit resolution and DMA the gapless stream to a GPU for complex signal processing

Modern GPUs include very powerful processing units and a very high speed graphical memory bus. This combination makes them perfectly suited for signal processing applications.

Unfortunately, it is not easy for other hardware devices, such as waveform digitizers, to DMA data directly to the GPU's on-board memory. This forces users to manually copy data from the buffer returned by the waveform digitizer to the GPU. This copying process is relatively slow and causes a drastic reduction in overall data throughput.

A commonly used alternate approach with NVIDIA® GPUs is to use CUDA streams to overlap multiple buffer transfers. This method increases the transfer speed to approximately 5 GB/s, but suffers from large data latency.

ATS-GMA-BASE solves this problem by allowing users to DMA data directly from AlazarTech PCI Express waveform digitizers to AMD Radeon Pro GPUs. Rates up to 6.9 GB/s have been achieved.

Method	Transfer Speed	Latency
Manual memcopy	~900 MB/s	~ 3.5 ms
CUDA streams	~5 GB/s	~ 3.5 ms
ATS-GMA	Up to 6.9 GB/s	~ 100 μ s

ATS-GMA-BASE includes an example program that demonstrates how to use the ATS-GMA-BASE library to DMA data from a waveform digitizer to a GPU. The example also shows how to do simple data processing on the GPU using OpenCL kernels, and how to transfer the processed data to host memory (RAM). Users can use this example program as a starting point to create their own kernels to do GPU-based DSP.

ATS-GMA-BASE also includes user-defined flags that enable users to unpack data to 16 bits per sample and to deinterleave data from multi-channel acquisitions.

GPU-Based Signal Processing

Graphical Processing Units (GPUs) were originally designed for rendering high-quality video for gaming applications, which required being able to perform massive amount of real-time calculations. The highly parallel architecture of modern GPUs also makes them an ideal platform for digital signal processing (DSP) and high performance computing (HPC) systems.

In the past, complex real-time signal processing, such as FFT, correlation, FIR filtering etc., could only be achieved using dedicated DSP processors or by implementing the algorithms inside and FPGA or an ASIC. All these methods are non-trivial, expensive, time consuming and require highly specialized engineering skills.

Using GPUs, users can implement any algorithm that can be parallelized in a GPU using well known software techniques and gain a better than 10-fold improvement over CPU based signal processing. The reason why GPUs perform so well for DSP applications is that they contain hundreds of processing cores (kernels) running in parallel, while sharing a very high speed graphical memory bank.

Benchmarks

All ATS-GMA benchmarks in this datasheet were measured using an ATS9373 in a system using an Asus® Prime X299-A machine using an Intel i9 7900X @ 3.3 GHz, DDR4 memory (32 GB RAM) and Radeon Pro WX 9100 GPU.

Data Throughput

The following throughput was measured in continuous streaming mode:

Buffer Size	Throughput*
1 MB	6.9 GB/s
2 MB	6.9 GB/s
4 MB	6.9 GB/s
8 MB	6.9 GB/s
16 MB	6.9 GB/s


* Based on benchmarks done in March 2018

A Typical ATS-GMA-BASE Application

A typical user application that uses ATS-GMA-BASE consists of the following minimum sections:

- 1) User application sets up waveform digitizer hardware (sample rate, input range, trigger parameters etc.).
- 2) User application allocates buffers in GPU memory and sets up the GPU.
- 3) User application starts data capture.
ATS-GMA-BASE starts DMA'ing data to GPU, one buffer at a time.
- 4) User-written GPU kernels do the following:
 - Process a buffer;
 - Copy result buffer to user memory, if required;
 - Get next buffer, and repeat.
- 5) User application running on CPU consumes result buffer, if required.
For highest performance, make sure data consumption is faster than the rate at which result buffers are supplied by GPU kernels.
- 6) This continues until the application has to be closed.

ATS-GMA Data Flow



ATS-GMA-BASE is supplied with an example user application in source code. The application includes GPU kernels that use ATS-GMA-BASE to receive data, do very simple signal processing (data inversion), and, if required, copy the processed (inverted) data to a user buffer. All this is done at the highest possible data transfer rate.

Programmers can replace the data inversion code with application-specific signal processing kernels to develop custom applications.

Computer Power Supply

GPUs are power hungry. The Radeon Pro WX 7100 requires a power supply that can provide at least 650 Watts of power. As such, users must make sure their computer's power supply has sufficient capacity.

Compatible GPUs

ATS-GMA is designed to be compatible with AMD Radeon Pro GPUs using AMD APP SDK version 2.9 and higher. Testing was done using AMD APP SDK version 2.9 and Radeon Pro WX 7100 and Radeon Pro WX 9100.

Kernel programming must be done in OpenCL. AMD GPUs do not support CUDA®, consequently, ATS-GMA does not support CUDA. If you need to program in CUDA, AlazarTech offers [ATS-GPU](#).

It should be noted that the current version of ATS-GMA supports only one GPU at a time. If you have multiple GPUs installed in your computer, ATS-GMA will let you select one of them for use.

Compatible Waveform Digitizers

All AlazarTech PCI Express waveform digitizers are compatible with ATS-GMA. Only single-board configurations are supported at this time.

AlazarTech's PCI bus waveform digitizers are not supported, as the host CPU is more than capable of handling data rates generated by PCI bus boards.

ATS-GMA cannot directly be interfaced with non-AlazarTech waveform digitizers.

Data Throughput to GPU

The data transfer rate to GPU is dependent on the generation of PCI Express digitizer board used:

PCIe Link Speed	Transfer Rate
Gen 3: ATS9373, ATS9371	Up to 6.9 GB/s
Gen 2: ATS9360, ATS9416	Up to 3.5 GB/s
Gen 1: ATS9870, ATS9350, ATS9351, ATS9120, ATS9625, ATS9626, ATS9440, ATS9462	Up to 1.6 GB/s (Exact rate is limited by digitizer sample rate)

Add-On ATS-GMA libraries

AlazarTech will be releasing additional libraries for use with ATS-GMA-BASE. Features such as FFT, Zero Padding, and Windowing Function will be available with an upcoming OCT Signal Processing library.

Software Licensing Policy

Users are allowed to freely distribute the ATS-GMA-BASE libraries as long as they have purchased one ATS-GMA-BASE license and there is an AlazarTech PCI Express waveform digitizer present in the same computer. If an AlazarTech PCI Express waveform digitizer is not present in the computer, users must purchase a separate license for each computer on which ATS-GMA is installed.

In no case is the user allowed to distribute or share the source code of ATS-GMA-BASE with other users.

Annual Subscriptions

The purchase of an ATS-GMA-BASE license includes a subscription that provides customers with the following benefits for a period of 1 year on ATS-GMA:

- Download ATS-GMA-BASE updates from the AlazarTech website;
- Receive new example programs as they become available;
- Receive technical support on ATS-GMA-BASE.

Additional add-on libraries for ATS-GMA-BASE are not covered by the annual subscription, i.e. holders of an annual subscription will have to purchase subscriptions for additional libraries separately.

It is recommended that customers renew their ATS-GMA-BASE subscriptions in order to maintain their product up-to-date, and receive new features and technical support.



ATS-GMA-BASE

Real Time Signal Processing Software

Note that support is provided for product bugs, and not for writing custom GPU kernels or for learning GPU programming.

Writing Custom GPU Kernels

ATS-GMA-BASE includes an example program in C/C++ source code, which implements very simple GPU kernels that invert data and write it back to a buffer in computer memory.

Users who need to write their own kernels should start with the included source code, add OpenCL code in the appropriate place, and compile their libraries.

The example program is provided with a Visual Studio project and a CMake build file. We use more recent C++ features, and Visual Studio 2015 and later is required.

Writing, testing, and debugging modified kernels will be the sole responsibility of the user and AlazarTech will not be responsible for assisting the user with such custom modifications.

Users must have expert programming knowledge of OpenCL development in order to customize ATS-GMA kernels.

ORDERING INFORMATION

ATS-GMA-BASE: GPU Streaming Library License + 1 Year Subscription ATSGMA-001

ATS-GMA-BASE-1YR: 1 year extended support & maintenance for ATS-GMA-BASE ATSGMA-002

ATS-GMA main API functions

```
ATS_GMA_AbortCapture
ATS_GMA_AllocBuffer
ATS_GMA_FreeBuffer
ATS_GMA_GetBuffer
ATS_GMA_GetComputeDevice
ATS_GMA_PostBuffer
ATS_GMA_QueryDeviceCount
ATS_GMA_QueryDeviceName
ATS_GMA_Setup
ATS_GMA_SETUP_FLAG
ATS_GMA_SETUP_FLAG_DEINTERLEAVE
ATS_GMA_SETUP_FLAG_UNPACK
ATS_GMA_SETUP_FLAG_USER_DEFINED_CONTEXT
```

† Radeon is a trademark of Advanced Micro Devices, Inc.
OpenCL is a trademark of Apple Inc.
Windows and the Windows logo are trademarks and/or registered trademarks of Microsoft Corporation in the U.S. and/or other countries.
ASUS is either a US registered trademark or trademark of Asustek Computer Inc. in the U.S. and/or other countries.
CUDA is a trademark and/or registered trademark of NVIDIA Corporation in the U.S. and/or other countries.
All other trademarks are the property of their respective owners.

Manufactured By:

Alazar Technologies, Inc.

6600 TRANS-CANADA HIGHWAY, SUITE 310
POINTE-CLAIRE, QC, CANADA H9R 4S2

TOLL FREE: 1-877-7-ALAZAR OR 1-877-725-2927
TEL: (514) 426-4899 FAX: (514) 426-2723

E-MAIL: sales@alazartech.com