GPGPU COTS Platforms

High-Performance Computing Solutions.

abaco SYSTEMS

WE INNOVATE. WE DELIVER.
YOU SUCCEED.

Rugged GPGPU COTS Boards for Military and Aerospace

GPGPU platforms deliver new levels of performance for size, weight and power (SWaP) constrained mission payloads.

The world of high-performance computing is undergoing a revolution, thanks to advances in General Purpose computing on Graphics Processing Units (GPGPU). The idea behind GPGPU is to use a GPU, which typically handles computation for computer graphics only, to perform parallel computation in applications that have traditionally been handled by the CPU.

A multi-GPU platform hosted by one or more CPUs is able to perform heterogeneous computing, harnessing the parallel computing power of the many-core GPUs to provide very large increases in performance with minimal programming complexity.

PARALLEL COMPUTING

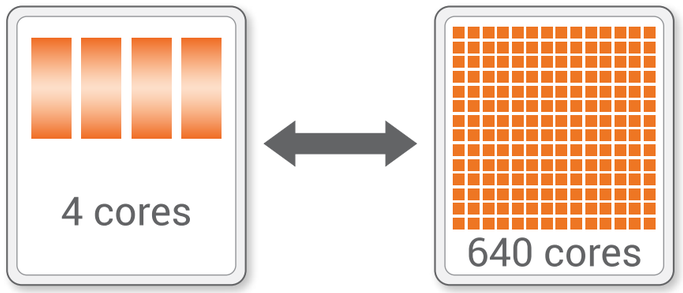

Optimized for throughput computing

Additionally, programmers are helped by software development environments such as Compute Unified Device Architecture (CUDA) and OpenCL, which allow them to harness the many-core, parallel processing capabilities of the GPGPU platforms.

While greatly increasing functional capability, the GPGPU platform also delivers the performance with far less size, weight and power (SWaP). This results in significant savings in cost, risk, and time-to-market.

Lab-proven technologies ruggedized for the harsh environment of military applications

Now these benefits are fully available to rugged military and aerospace applications. With a full range of Abaco rugged GPGPU boards and systems, the advantages of GPGPU are no longer confined to controlled environments at universities, research centers and hospitals.

The unique partnership between Abaco and NVIDIA allows for new product development using NVIDIA GPUs based on the award-winning CUDA architecture, for military and aerospace applications.

Rugged GPGPU COTS Boards for Military and Aerospace

Development Ease

Increases in performance will be obtained in application areas such as Software Defined Radio, sonar, and medical imaging. But what is less obvious is the change in development strategy offered by GPGPU technology. The only other technology currently offering massively parallel processing capability is Field Programmable Gate Arrays (FPGAs).

Although FPGAs provide very high performance data processing, developing high-performance FPGA cores requires a very specialized skill set built on a hardware engineering background, whereas developing code for GPGPU processors is much more about software expertise. For companies with a background in multi-processor GPP/ DSP-based system architecture, the move to GPGPU will be much less disruptive than a move to FPGA processors. The processing power, system size and power consumption enabled by GPGPU are compelling factors, but the addition of programming ease makes such a system tough to match.

Typical GPGPU applications

Radar

One of the biggest challenges for today's radar systems is to provide more capability—range, number of targets, speed, and so on —while meeting ever more stringent SWaP constraints. The extra speed offered by GPGPU platforms translates directly to more area coverage and more security for the operating team.

One rack containing 72 conventional processors (18 6U boards) and producing a peak capability of 576 GFLOPS can take up four cubic feet, weigh over 105 pounds and consume over 2,000 watts. GPGPU technology can allow system designers to fit an unprecedented amount of processing power into a very compact package. The use of three 3U VPX boards can yield peak processing power of 3916.8 GFLOPS in less than 0.4 cubic feet.

Data Encryption/Decryption

There are several standards for encryption of data, including the Advanced Encryption Standard (AES). AES is the first publicly accessible and open cipher approved by the U.S. National Security Agency for top secret information, typically requiring 256-bit keys at this level. The time to encrypt a block of data increases linearly with the size of the key.

The computation load required to maintain encryption of a real-time data stream can be prohibitive. With the advent of CUDA and the addition of crucial arithmetic, bitwise logical and shift operations as well as the ability to use texture caches to index tables, GPUs are now a viable alternative to general purpose processors for data encryption/ decryption. Performance gains up to 10x have been demonstrated.

Situational Awareness

Surveillance of large areas has historically been achieved by using an array of sensors connected to a bank of monitors, with separate or multiplexed displays for each video stream. Such arrangements present the operator with a confusing array of disparate video feeds, require a great deal of space, and consume a large amount of power. In a dynamic, real-time scenario, there is also a danger of information overload for an operator attempting to interpret such large volumes of imagery.

Interrelationships between sensors is not always obvious, and important contextual visual information can be overlooked. Many such systems rely on the operator for "event" detection, but large volumes of information, coupled to the effects of stress and fatigue, can significantly reduce operator effectiveness. Abaco's image processing subsystem overcomes these issues and greatly improves the performance of surveillance assets and their operators. We offer a previously unattainable level of situational awareness to platforms such as armored vehicles, aircraft, remote unmanned platforms and security and surveillance systems

IED Detection

Improvised Explosive Devices (IEDs) are a major cause of injuries and fatalities among ground troops. A number of techniques for automated detection of IEDs are used, and all of these require processing a high volume of data. The effectiveness of the solution depends on how fast the algorithm can reliably operate on that data. GPGPU technology is proving to be a highly effective solution for such high throughput computations.

Ground Change Detection relies on realtime image processing, and may be applied to sensors mounted on ground vehicles or UAVs. The system needs to apply image registration and stabilization, and moving object extraction, before comparison with normalized geo-referenced data, all while dealing with lighting and legitimate scene changes.

Ground Penetrating Radar (GPR) allows construction of a 3D model of the ground, identifying any suspicious objects or changes from normalized data. GPR can be applied to ground mobile or airborne systems. In addition, behavioral modeling— based on live sensor imagery, possibly combined with wide-area surveillance data—can be used to identify hostile intent and potential threats, giving operational forces time to assess the risk and take appropriate defensive actions.

Target Tracking

GPGPU-based video trackers and image processors are at the heart of target tracking systems where they provide the highest performance solutions in the smallest, fully ruggedized hardware packages.

Target detection and target acquisition processes identify objects within an area of the video image display that meets the user-defined target criteria. A range of detection algorithms is built into the system to meet situational requirements.

When one or more targets have been detected, the tracking system can enter automatic or manual tracking mode. Automatic target acquisition may be prioritized by using several different factors, such as target nearest to the boresight or the largest target. If a system is in auto-track mode, the video tracker automatically tracks the selected target and can control almost any type of pan-and-tilt or gimbal system to track the target.

GPGPU-based automatic video trackers feature a wide range of proven high performance detection and tracking algorithms that can be tailored to the operational scenario, including centroid, phase correlation and edge detection. Algorithms may be combined to meet particularly demanding tracking scenarios.

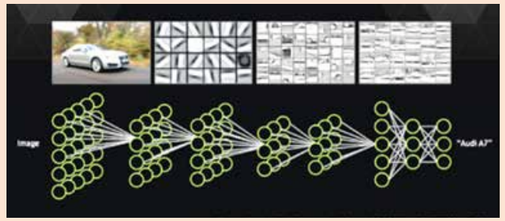

Autonomous Systems

Autonomous capabilities are transforming the defense industry, and robust GPU-enabled graphics/vision/compute solutions applying artificial intelligence are critical. The massively parallel architecture of GPUs naturally supports the neural network layers required for efficient machine learning, as well as autonomous command and control functions. GPUs are capable of supporting the entire autonomous stack:

- > Target classification/tracking
- > Sensor Fusion
- > Perception
- > Situational Awareness
- > Decision-making
- > Execution

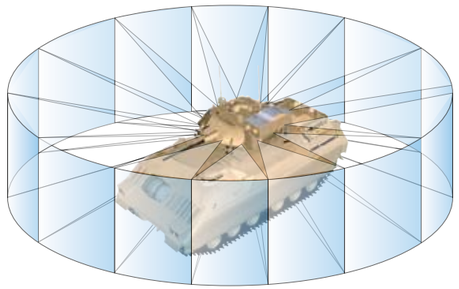

Abaco GPGPU Platforms

GVC1001
Graphics/Vision/AI computing redefined

- Powerful graphics, vision and image processing platform
- Industry-leading performance with NVIDIA® Jetson AGX Xavier™ SoM
- More comprehensive I/O for graphics, vision and image processing applications
- 256 GB NVMe SSD bulk storage as standard
- Abaco AXIS ImageFlex for accelerated image processing and manipulation
- Interoperability with the RES3000 family for multiple GigE camera aggregation

GRA115Q
Graphics & GPGPU Output Board

- Available with NVIDIA Quadro RTX™ 5000 or Quadro RTX™ 3000 GPU which can result in a 4X performance improvement over NVIDIA Pascal™ class GPUs
- Supports new customers with latest DisplayPort ™ output technology
- Provides technology insertion path for existing GRA112D/3D, GRA112Q/3Q and GR5 customers with single link DVI ports
- Decreases system bottle necks for overall increased system performance
- Accelerates image processing and manipulation with Abaco AXIS ImageFlex

IPN254
High Performance Computing Multiprocessor SOSA™ Alignment, technology insertion

- By combing the latest NVIDIA® Quadro RTX™ 3000 GPU with the latest 9th generation Intel® Xeon® E CPU, the IPN254 delivers maximum processing performance
- Alignment with the SOSA™ Technical Standard, an alternative version of the IPN254 provides a form/fit/function-compatible technology insertion solution for IPN252 users
- Supports new customers with the latest DisplayPort ™ output technology

GR5
3U VPX NVIDIA Quadro P2000 Graphics, Video & GPGPU Card

- Features most recent NVIDIA chip-down Pascal P2000 technology: do more with less
- Supports new customers with latest DisplayPort output technology
- Provides technology insertion path for existing GRA112D/3D customers with single link DVI ports
- Decreases system bottle necks for overall increased system performance
- Accelerates image processing and manipulation with Abaco AXIS ImageFlex

NVP2000
XMC Graphics/Video/GPGPU Board

- Most recent NVIDIA chip-down Pascal technology: do more with less
- Significant GPU performance in a small form factor
- Increased frequency and resolution for improved operator displays
- Lower SWaP COTS graphics output solution
- Reduced system bottle necks for overall increased system performance
- Decreased development costs through provision of key board level drivers and integration

NVP2102
Ultra-High Performance Graphics/ Video Capture XMC

- Lower SWaP COTS video ingest and graphics solution
- Most recent NVIDIA chip-down Pascal technology: do more with less
- Significant video ingress and egress capability
- Direct video capture to GPU memory and output to lower system latency
- Potential to lower overall system, spares and maintenance costs
- Accelerated image processing and manipulation with Abaco AXIS ImageFlex

GR2
Dual Channel Graphics Output Board

- 3U VPX graphics output card based on the NVIDIA Pascal Quadro P5000/P3000 GPUs
- Cost-effective and low risk alternative for customers only requiring video output and ideal for artificial intelligence, machine learning, autonomous systems and high performance embedded computing applications
- Enables cost-effective and low risk technology refresh cycles with newer GPUs

GR4
Quad Channel Video Capture Board

- 3U VPX video capture and processing card with increased performance based on NVIDIA Pascal Quadro P5000/P3000 GPUs
- 3G-SDI input/output (four channels each) supports modern sensor and display requirements
- Designed for military and defense applications as well as challenging commercial/industrial applications

Software

AXIS ImageFlex

ImageFlex is an image processing and visualization toolkit enabling rapid development of high performance image processing, visualization and autonomy applications aimed at size, weight and power (SWaP) sensitive applications. It is focused on high performance GPU processing and graphics.

FEATURE SUMMARY

Visualization framework API:
- Image creation and management
- CPU to GPU data movement
- 2D “overlay” drawing

Image processing API:
- Image manipulation
- Lens distortion correction.
- Complex image morphing
- Image fusion
- Image stabilization

Interoperability API:
- OpenCL interoperability
- CUDA interoperability

Custom extendibility:
- Easy creation of custom OpenGL “shader”
- 2D & 3D Matrix computation functions

Tools for Deploying Neural Nets:
- Image annotator GUI tool
- Caffe to TensorRT converter

Abaco quick start application examples:
- “Basics” example, showing all key functionality
- AI based object detection
- LIDAR point-cloud display
- “SkyBox” example for 360 situational awareness
- Image fusion example
- Image stabilization and target tracking examples
- OpenCV and OpenVX interoperability examples

CUDA

CUDA is a parallel computing language created by NVIDIA that exploits the massively parallel characteristics of NVIDIA’s ubiquitous silicon. CUDA is taught in universities worldwide and used in many R&D labs, so a large number of programmers are available and there is a wealth of web-based resources.

Integrated Software Modules

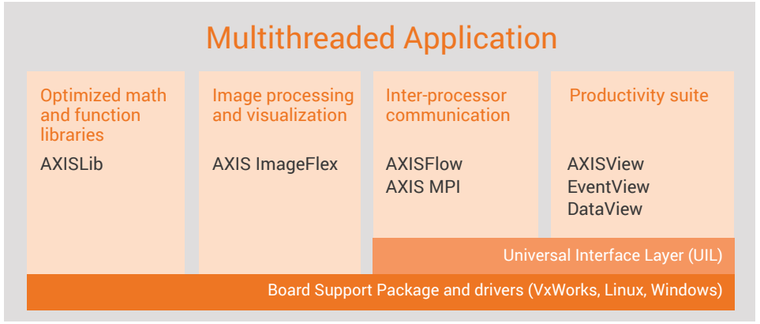

CUDA software development tools:
- NVIDIA C Compiler and debugger for parallel GPU code
- CUDA Visual Profiler
- CUDA SDK with examples of best-practice guides
- Parallel Nsight® IDE

Advanced libraries that include:
- NVIDIA Performance Primitives (image and video)
- Image processing: ArrayFire and OpenVPX
- Math and signal processing: cuFFT, cu BLAS, cuSPARSE, cuSOLVER, CUDA Math Library
- Deep learning: cuDNN, TensorRT

C for CUDA extends C by allowing the programmer to define C functions, called kernels, that when called are executed N times in parallel by N different CUDA threads, as opposed to only once like regular C functions. A kernel is defined using the __global__ declaration specifier and the number of CUDA threads for each call is specified using a new <<>> syntax.

OpenCL

Abaco GPGPU products also support Open Computer Language (OpenCL), the first open language for writing programs that execute across heterogenous architectures such as CPUs, GPUs, and FPGAs.. It includes a C-type language for writing kernels, defines APIs, and provides parallel computing using task-based and data-based parallelism. There are now a wide variety of open source and third party libraries and tools for OpenCL development.

OpenCL development tools:
- Compilers and debuggers from NVIDIA, Intel, AMD and Xilinx
- GPUSTATUSMonitor

Open source libraries:
- Image processing: ArrayFire and OpenVPX
- Math and Signal processing: clFFT, clBLAS, clSPARSE

OpenGL

OpenGL is supported on all Abaco GPU platforms. It is a standard specification defining a cross-language, cross-platform API for writing applications that produce 2D and 3D graphics. The shader language, introduced in OpenGL 1.4, also allows for OpenGL to be used for a level of compute particularly tailored to image processing applications such as image fusion

The AXIS Multiprocessing software suite facilitates the development of complex applications over multiple clusters of GPU platforms.

AXISFlow provides a communications API for multi-threaded/multi-core/multiprocessor communications.

AXISView provides a set of GUI tools enabling system visualization, application instrumentation, debug and monitoring.

CUDA exploits the massively parallel characteristics of NVIDIA’s ubiquitous silicon. CUDA is taught in universities worldwide and used in many R&D labs, so a large number of programmers are available and there is a wealth of web-based resources.

CUDA software development tools:
- NVIDIA C Compiler for parallel GPU code
- CUDA Debugger
- CUDA Visual Profiler
- SDK with best-practice guides
- Parallel Nsight®

Advanced libraries that include:
- NVIDIA Performance Primitives (image and video)
- Basic Linear Algebra Subprograms
- VisionWorks, OpenCV
- cuDNN, TensorRT for deep learning

Abaco Systems GPGPU products also support Open Computer Language (OpenCL), the first open language for writing programs that execute across CPUs, GPUs, and other processors. It includes a language for writing kernels, defines APIs, and provides parallel computing using task-based and data-based parallelism.

Open Graphics Library (OpenGL) is a standard specification defining a cross-language, cross-platform API for writing applications that produce 2D and 3D computer graphics. This is used in the graphics output processes.

PROCESSING FLOW ON CUDA

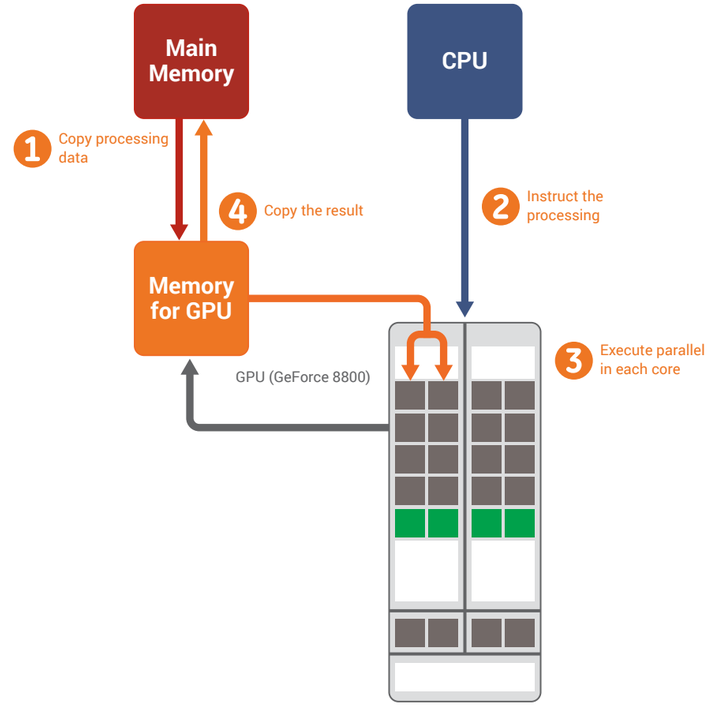

C for CUDA extends C by allowing the programmer to define C functions, called kernels, that when called are executed N times in parallel by N different CUDA threads, as opposed to only once like regular C functions.

A kernel is defined using the __global__ declaration specifier and the number of CUDA threads for each call is specified using a new <<<…>>> syntax:

GPGPU COTS Platforms

Abaco GPGPU-empowered platforms can be easily implemented to either adapt to your legacy applications or to accommodate your new applications.

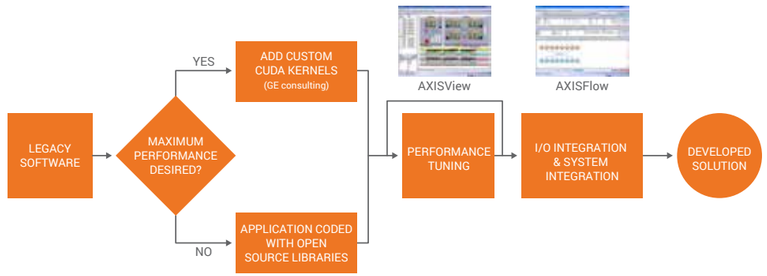

abaco SYSTEMS

WE INNOVATE. WE DELIVER.
YOU SUCCEED.

abaco.com

Americas: **866-OK-ABACO or +1-866-652-2226**
Europe, Africa, Middle East, & Asia Pacific: **+44 (0) 1327-359444**
Locate an Abaco Systems Sales Representative visit: **abaco.com/products/sales**