EFFICIENTLY LEVERAGING ACCELERATED PARALLEL PROCESSING FOR RADAR SIGNAL INTERPRETATION

Latest graphics chips have evolved into highly programmable computing engines with massive parallel processing capability. This additional computing power is also highly attractive for high-end military applications such as radar signal interpretation, but existing COTS (Commercial Off The Shelf) products often do not meet all the requirements of this demanding application area.

For system integrators who want to save valuable engineering resources which are needed to adapt a graphics solution to meet application-specific requirements, vendors such as Tech Source provide both rugged COTS and customised solutions that are based on high-performance and long-term available embedded graphics chips.

Architecture of modern graphics chips

As graphics chips become more powerful, new ways to harness this potential for embedded applications arise, such as utilising the hundreds of shader units in today's GPUs to process hundreds of floating point operations at the same time, also referred to as GPGPU (general-purpose graphics processing unit) computing or AMD Accelerated Parallel Processing (APP) technology. To understand how an application can make use of this architecture, it is necessary to understand how the architecture is laid out. Consider first what a graphics processor is required to do. Every time the computer screen refreshes itself (~60 - 80 times per second) the graphics chip needs to transform a 3-dimensional image and project it onto a 2-dimensional surface. This translation is done for each pixel on the screen. Given a resolution of 1280 x 1024, that is over 1.3 million pixels per screen or 104 million pixels per second. To make this happen as fast as possible, modern graphics chips leverage an architecture made up of hundreds of compute cores. Each core works independently to carry out the computations required for a single pixel. Working in parallel, these pipelines process hundreds of pixels at the same time. This same architecture can be used to perform rapid mathematical calculations for a myriad of applications.

Benefits for military applications

Basically, many applications can benefit from APP technology where extensive integer or floating point operations have to be performed as quickly as possible. Take, for instance, radar signal interpretation. Radar interpretation has to carry out a large number of floating point operations in order to identify targets, movement, etc. Radar data can be collected much faster than it can be processed, making these computations the primary bottleneck. To improve performance, radar system designers must add more processors (or cores) or increase their clock speeds, both of which work in direct opposition to the need to make these systems smaller and consume less power to meet SWaP (Size, Weight and Power) and thermal management requirements. However, leveraging the parallel processing capabilities of a GPU fulfils these requirements as it enables developers to increase performance and at the same time decrease the power consumption and size of their application.

Cambridge Pixel leverages the accelerated parallel processing capabilities of the AMD Radeon E6760 GPU for its radar scan conversion solution which is deployed on the Type 45 Destroyer and QE Class aircraft carriers of the Royal Navy.

Parallel processing architectures

Parallel processing can provide significant increases in performance per watt over a traditional CPU architecture for applications that can leverage it. And it provides a competitive solution to traditional DSP and FPGA solutions. Furthermore APP technology also helps to simplify the embedded system design process, as the GPU plus software programming approach enables developers to replace the functionality of proprietary hardware such as FPGAs and DSPs by programming the processor and GPU to perform those same functions. Therefore, by leveraging broadly available silicon solutions and hardware paired with a programmable software solution, standard hardware can be leveraged to deliver highly flexible solutions that are less complex, leading to an accelerated time-to-market.

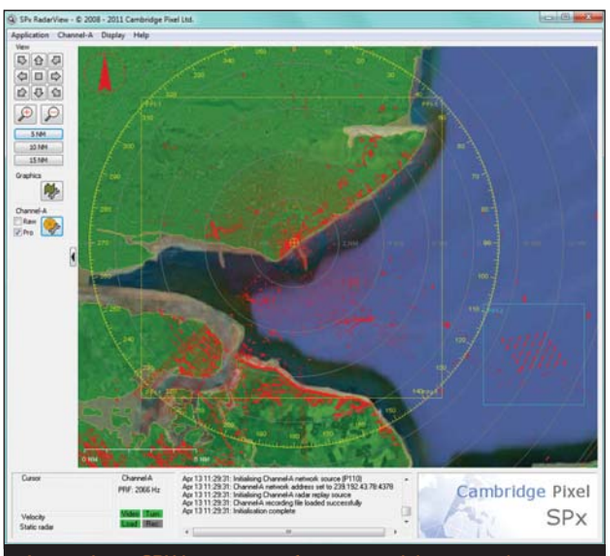

Leveraging a GPU increases performance and decreases the power consumption and size of applications.

Programming across multiple cores

Programming across multiple cores running in parallel requires a different approach than the classical programming of serial applications on a single processor. Early pioneers in parallel processing developed cumbersome applications in which they selected graphics operations in OpenGL that used the same math they needed, and then copied the results from the frame buffer. Proprietary parallel processing languages such as CUDA and Brook+ were a big improvement to the ease-of-use, however, at the expense of portability. Today the mainstream approach is to use an open standard programming language called OpenCL. OpenCL was created by architects from AMD, Apple, IBM, Intel, Nvidia, and Sony. As it has matured, it has become the API of choice for code that is portable across different hardware and software platforms. In this way, programmers can benefit from the fastest processors on the market each year, which helps to encourage competition and innovation between vendors. Another benefit of OpenCL is its ability to leverage multi-core CPUs together with multi-core GPUs. Requiring only one programming platform for both the CPU and GPU and enabling applications to benefit from increases in the number of CPU cores for certain operations.

AMD Accelerated Parallel Processing (APP) SDK

Supported by the industry standard OpenCL programming language, GPGPU application software development can be sped up using the AMD Accelerated Parallel Processing (APP) SDK. The SDK includes developer tools such as compiler, debugger, code profiler and math libraries. With the new Accelerated Parallel Processing SDK v2.3 developers can build applications that take advantage of the processing power for AMD CPUs, GPUs and APUs.

Challenges on the hardware side

Application developers are still faced with some challenges when it comes to harnessing the full potential of modern graphics chips for their embedded applications: Existing COTS products often do not meet all the requirements of military applications such as rugged, long-lasting designs suitable for extremely rough conditions and extended temperature ranges. One unique concern that many embedded designers have is ensuring that their designs can be built and supported for several years which requires their suppliers to provide extended product longevity.

To address the challenge of long term availability, AMD's embedded graphics processor and APU solutions for instance have a planned availability of five years of active manufacturing, plus another two years of possible "end of life" commitment. And to meet the environmental requirements with commercial-off-the-shelf solutions, vendors such as Tech Source offer products with long life cycles of 7 years or more (10 years or more with customer forecasts) and overall support for the product of 15 years or longer. With 20 years of experience in the development of high performance graphics hardware for the ATC and defense markets, they can enable a vendor to cut out time-consuming and costly hardware development work. Being experts on the requirements for these markets they can provide fully featured embedded graphics solutions for the most common embedded form factors such as VME, VPX, CompactPCI, AdvancedTCA, and other single board computers (SBCs) utilising an XMC mezzanine card interface. Tech Source provides rugged and long-term available COTS solutions and can also offer fully customised solutions (with specific I/O), software solutions and direct access to support from the design engineers themselves, not tech support or customer service personnel.

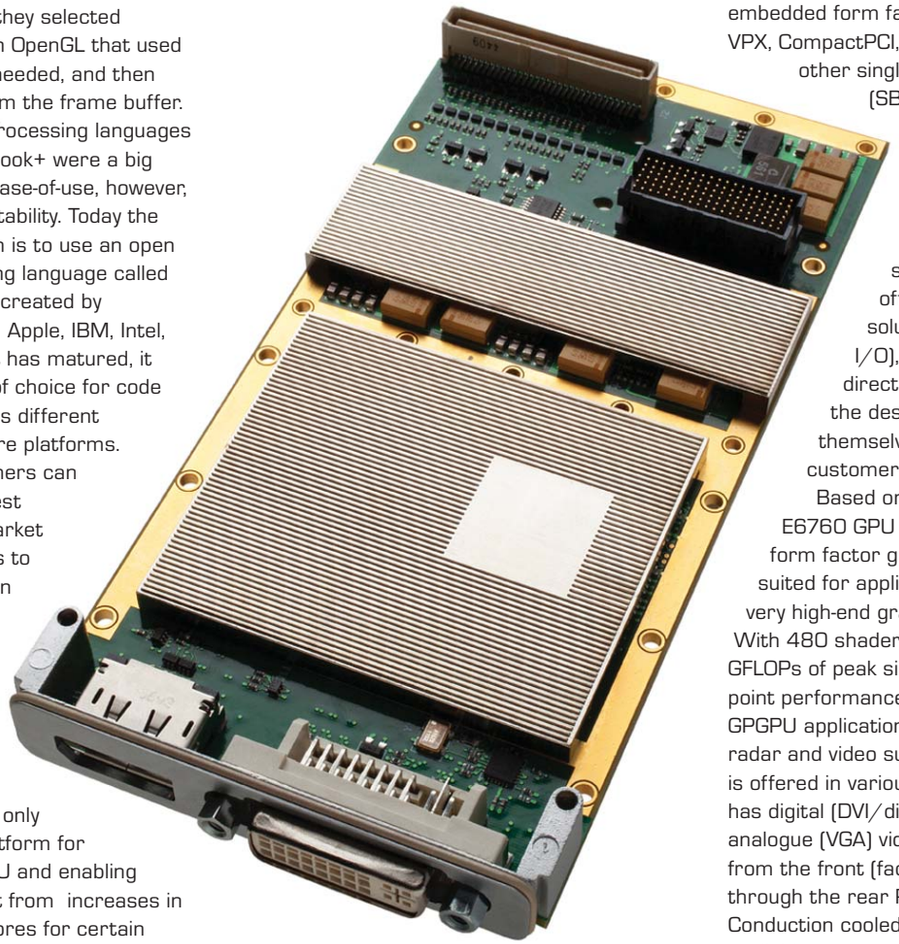

Based on the AMD Radeon E6760 GPU the Condor 3000x XMC form factor graphics/video card is suited for applications that require very high-end graphics and computation. With 480 shader units delivering 576 GFLOPs of peak single precision floating point performance, it is also ideal for GPGPU applications such as ultra-sound, radar and video surveillance. The product is offered in various ruggedised levels and has digital (DVI/display port) and analogue (VGA) video outputs available from the front (face plate) of the card or through the rear Pn4/P16 connectors. Conduction cooled versions are also available.

Conclusion

The parallel processing potential of today's graphics solutions is highly attractive for performance hungry floating point applications in markets such avionics, military and air traffic control. Application developers who want to harness this potential can save valuable engineering resources by trusting in an experienced supplier such as Tech Source to receive a graphics solution that is built to meet their individual application- and market-specific requirements. Application developers can thus concentrate on their core competence and improve time-to-market for their graphics- and computing-intense embedded applications.

By CAMERON SWEN, Marketing Manager, AMD Embedded Solutions