Mil Tech Insider

Graphics Processing Unit accelerates scan conversion

By John Wemekamp

Many examples abound as to why the adoption of COTS technologies has been such a success story for the military. The maximum benefit is obtainable where massive commercial technology investment can be leveraged directly to solve a similar technology problem in another domain. Nowhere is this level of investment more evident than in the commercial video gaming market, providing the opportunity for the technology being developed for real-time, dynamic, and highly realistic gaming to be applied with an innovative twist to military applications. Examples of this include radar scan conversion or image processing on live sensor data streams.

The majority of surveillance radars still use rotating antennae, whether they are in fixed ground positions or mounted on ships or large aircraft. Once digitized, radar video output is in polar form (range and azimuth), which must be converted to raster format for display on an operator's screen. This scan conversion process is a very intense and repetitive algorithm that has traditionally been performed by a host processor or, more recently, with FPGA assistance to a host processor. The slow decay of data in the display buffer results in trails forming behind moving objects, making them clearly visible.

Off-the-shelf embedded computers are used extensively for the display and processing elements of radar systems. These computers can range from embeddable, industrial PCs for benign environment applications to VMEbus or its newer sibling VPX for harsher, more critical environments. Using current technology, the embedded host computer in the form of an SBC would be supported by a scan converter for displaying and windowing the radar image, plus a separate GPU to overlay graphics onto the radar and present a broader operational picture to the operator. These functions are typically implemented as PCI or PMC modules attached to the host SBC.

The demands of increasingly sophisticated, and hence more expensive to produce, video games are driving the investment in GPUs and the tools to support them. These GPUs are part of continuously evolving families of products necessarily offering application portability across generations and manufacturers. In the search for greater performance, the latest offerings from vendors such as NVIDIA and ATI incorporate large memory space and very high-performance, general-purpose processing capability. These support single-thread processes used in fragment and vertex manipulation required by real-time, three-dimensional (3D) animation. This on-GPU processing allows much faster image rendering directly into display memory than would be possible using a host processor. Until recently, it was difficult to make use of this latent capability; however, the introduction of shader languages – such as OpenGL's Shader Language (GLSL) or the High Level Shader Language (HLSL) supported by Microsoft – makes for much simpler programming of these functions using C code.

By loading the incoming radar in polar form into GPU memory, an algorithm running on the GPU can directly convert it into raster format in display memory in real time. Used in this way, the latest generations of GPUs offer an order of magnitude improvement in performance over FPGA-assisted scan converter solutions. For example, it is normally anticipated that an increase in display resolution will impact scan converter performance as more pixel locations are manipulated. However, when using the GPU even with display resolutions of 2k x 2k pixels – often used for air traffic control or naval surveillance radars – there is no impact on overall performance. Additionally, displaying the traditional plan view of the radar image, the GPU allows the image to be easily distorted in real time to show projections from different origins and angles. This is illustrated in Figure 1, which shows a high-resolution, scan-converted and projected radar image created by the patent-pending SoftScan developed by Curtiss-Wright Controls Embedded Computing (CWCEC) using GLSL.

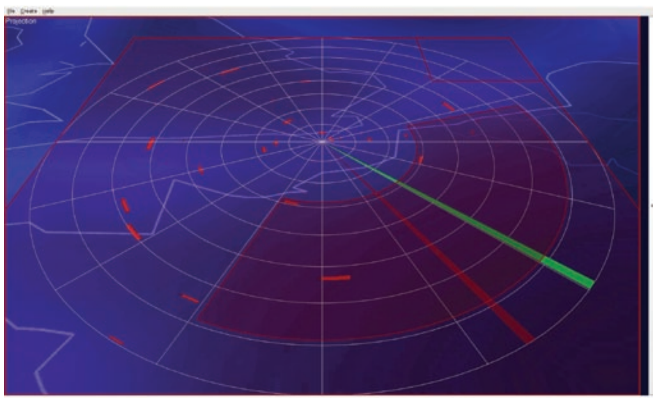

Figure 1

Two-dimensional radar scan conversion is one example of how new levels of system performance can be unleashed using the latent power of the GPU. The same principles can be applied to 3D radar (range, azimuth, and height) where the GPU also provides the ideal platform for displaying the images in many projections and visualizations, able to display a number of them simultaneously on the same screen as well, if required. The use of shader language and the inevitable growth in capability of GPUs generated purely by commercial market pressures appear set to translate into a host of new military battlefield visualization options and performance benefits, truly reinforcing the value of the COTS proposition.

To learn more, e-mail John at john.wemekamp@curtisswright.com.