WHITE PAPER

NVIDIA'S Tegra K1: A Game-Changer for Rugged Embedded Computing

abaco
SYSTEMS

NVIDIA’S Tegra K1: A Game-Changer for Rugged Embedded Computing

by Dustin Franklin, Abaco Systems

GPU technology is already making significant inroads into applications such as radar, sonar and ISR that can readily benefit from the very high degree of parallelism offered by graphics processors.

The extensive compute capability offered by GPU technology has, however, challenges with power consumption and heat dissipation – a concern in smaller, lighter weight platforms that do not have significant power at their disposal and that are problematic to cool.

But: the limitations of deploying GPU technology in SWaP-constrained environments were largely eliminated when, earlier this year, NVIDIA® announced that Abaco Systems was its preferred provider of rugged solutions, destined for harsh, challenging environments, based on the Tegra® K1.

192 Cores. 327 GFLOPS. Less Than 10W of Power.

NVIDIA’s Tegra K1 (TK1) is the first ARM system-on-chip (SoC) with integrated CUDA®. With 192 Kepler GPU cores and four ARM Cortex-A15 cores (Figure 1) delivering a total of 327 GFLOPS of compute performance, TK1 has the capacity to process lots of data with CUDA while typically drawing less than 6W of power (including the SoC and DRAM). CUDA (Compute Unified Device Architecture) is a parallel computing platform and programming model invented by NVIDIA that gives program developers direct access to the virtual instruction set and memory of the parallel computational elements in CUDA GPUs, and allows them to be used for general purpose processing.

To appreciate the enormity of what NVIDIA has achieved: a 6U VPX multiprocessor single board computer featuring earlier GPU technology consumes ~100 watts to deliver 645 GFLOPS of performance. In other words: twice as much compute horsepower, but with over ten times the power consumption. Such platforms, with their raw compute capability and flexibility, will continue to be at the heart of many leading edge embedded computing deployments. There are, however, a growing number of applications being envisaged for the future that will significantly benefit from ‘only’ half the processing performance if power/cooling can be reduced so substantially.

The Tegra K1, then, brings game-changing performance to low-SWAP and small form factor (SFF) applications in the sub- 10W domain, all the while supporting a developer-friendly Ubuntu® Linux® software environment delivering an experience more like that of a desktop rather than an embedded SoC. Tegra K1 is plug-andplay and can stream high-bandwidth peripherals, sensors, and network interfaces via built-in USB 3.0 and PCIe® gen2 x4/x1 ports. TK1 is geared for sensor processing and offers additional hardwareaccelerated functionality asynchronous to CUDA, like H.264 encoding and decoding engines and dual MIPI CSI-2 camera interfaces and image service processors (ISP). There are many exciting embedded applications for TK1 which leverage its natural ability as a media processor and low-power platform for quickly integrating devices and sensors.

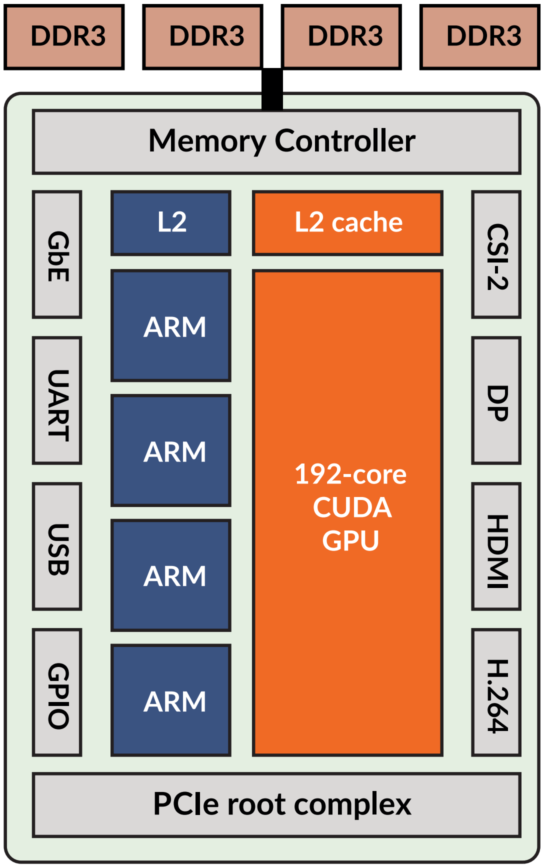

Figure 1 Simple TK1 block diagram

Figure 2 The launch of Tegra K1 sees the possibility of a range of GPU-based solutions with substantially different capabilities – but with code compatibility enabling one development for multiple deployments

Case Study #1 Robotics/Unmanned Vehicle Platform

Embedded applications commonly require elements of video processing, digital signal processing (DSP), command and control, and so on. In this example architecture with Tegra K1, CUDA is used to process imagery from highdefinition GigEVision gigabit cameras and simultaneously perform world mapping and obstacle detection operations on a 180° LIDAR scanning rangefinder. (Figure 3) Additionally devices such as GPS receivers, inertial measurement unit (IMU), motor controllers, and other sensors are integrated to demonstrate using TK1 for autonomous navigation and motion control of a mobile platform (such as a robot or unmanned vehicle). Teleoperation capability is provided by applying Tegra's hardware-accelerated H.264 compression to the video and streaming over WiFi, 3G/4G, or satellite downlink to a remote ground station or other networked platform. This architecture provides an example framework for perception modeling and unmanned autonomy using TK1 as the system's central processor and sensor interface. It's Tegra's low power consumption and minimal heat dissipation that make it the ideal processor for confined environments such as robots or small unmanned vehicles, giving them a local processing capability that would previously have been unthinkable.

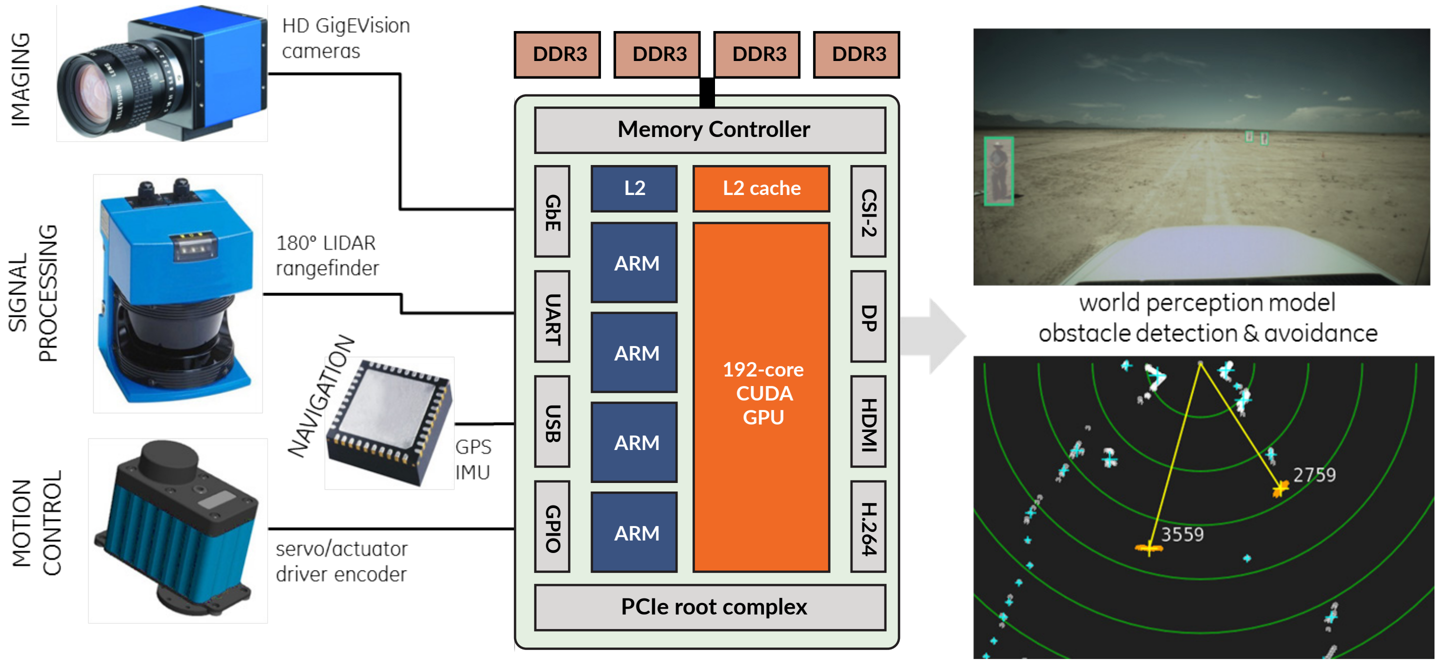

Figure 3 Sensor processing pipeline implemented using Tegra K1 for autonomous navigation.

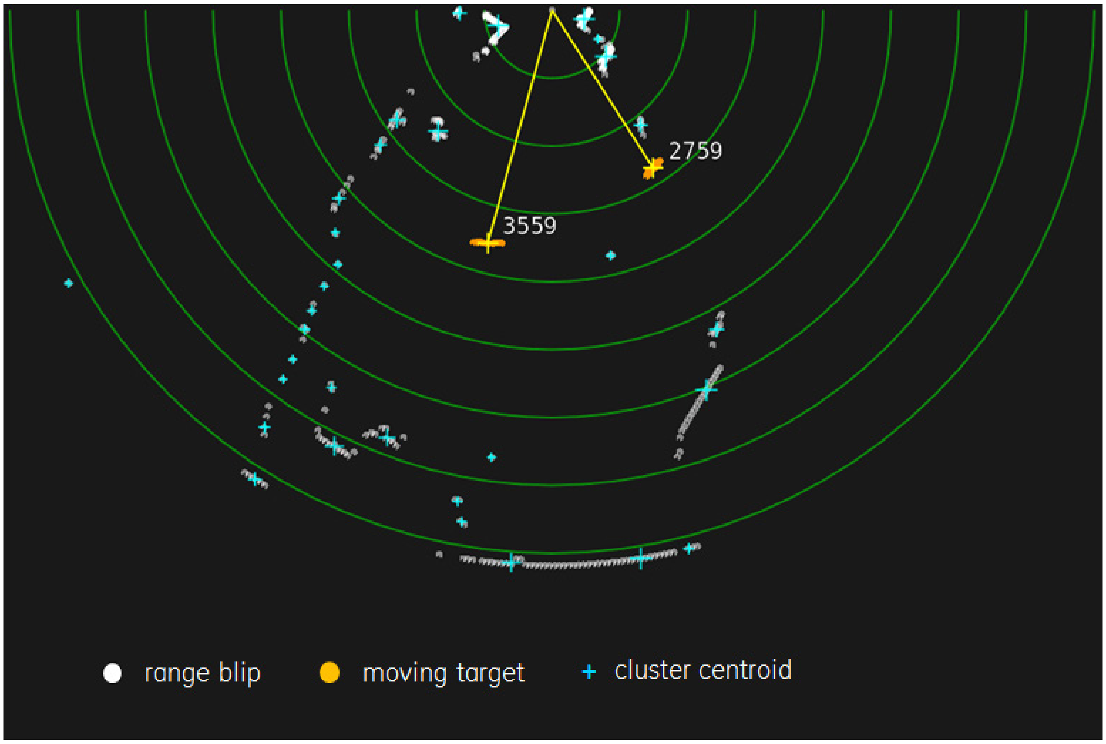

Figure 4 *LIDAR-driven PPI display visualizes static and moving obstacles in the platform's environment.*

The scanning LIDAR used produces range samples every 0.5° over 180 degrees, which are grouped into clusters using mean shift and tracked when motion is detected. CUDA was used to process all range samples simultaneously and perform change detection versus the octree-partitioned 3D point cloud built from previous georeferenced LIDAR scans, producing a list of static and moving obstacles refreshed in realtime for collision detection and avoidance. A radar-like Plan Position Indicator (PPI) is then rendered on the OpenGL side. (Figure 4) This particular LIDAR was connected via RS232 to a serial port; other LIDARs support Gigabit Ethernet as well. The open-source SICK Toolbox library was used for connecting to the sensor, which compiles and runs out of the box on TK1. Having easy access to LIDAR sensors provides TK1 with submillimeter accurate readings to exploit with CUDA for realtime 3D environment mapping and parallel path planning.

On the imaging side, Tegra K1 has a number of interfaces for streaming highdefinition video, such as CSI-2, USB 3.0, and Gigabit Ethernet. Framegrabbers for other mediums like HD-SDI, CameraLink, LVDS, and others can be integrated with TK1 via its PCIe gen2 x4 port. For this case study, testing was carried out with multiple Gigabit Ethernet cameras from GigEVision-compliant vendors, with resolutions ranging from 1920x1080 up to 2448x2048, and found an individual ARM CPU core sufficient per Gigabit Ethernet port for handling network

Figure 5 *MMTI and trainable HoG pedestrian/vehicle detectors extract dynamic obstacles from HD video at runtime.*

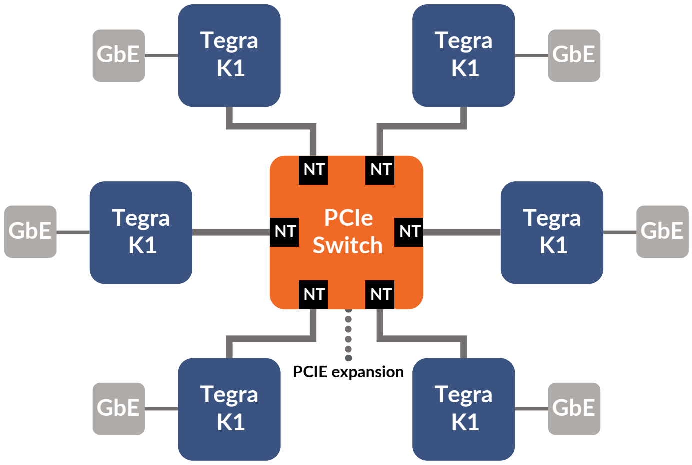

Figure 6 SWaP-optimized tiled architecture, six Tegra K1's interconnected with non-transparent PCIe switching and RDMA.

protocols and packetization using the sockets API. Using the cudaMallocManaged feature new to CUDA 6, the video stream is depacketized by the CPU into a buffer shared with the GPU, requiring zero copies to get the video "into GPU memory" (in the case of TK1, it's physically all the same memory).

Using freely-available libraries like OpenCV, NVIDIA NPP, and VisionWorks, users have the ability to run a myriad of CUDA-accelerated video filters on-thego including optical flow, SLAM, stereo disparity, robust feature extraction and matching, mosaicking, and multiple moving target indicator (MMTI). (Figure 5) Trainable pedestrian and vehicle detectors can run in realtime on TK1 using available Histogram of Oriented Gradients (HoG) implementations. There are many existing CUDA codes available which previously ran on discrete GPUs and are now able to be deployed on Tegra.

In addition to LIDAR devices and cameras, TK1 supports navigational sensors such as GPS and IMU for improved autonomy. These are commonly available as USB or serial devices and can easily be integrated with TK1. One quick way to make a GPS-enabled application is to use libgps/gpsd, which provides a common software interface and GPS datagram for a wide class of NMEA-compliant devices. Meanwhile IMU sensors are integrated to provide accelerometer, gyro, and magnetometer readings at refresh rates of up to 100Hz or more. TK1 fuses the rapid IMU and GPS data using highquality Kalman filtering to deliver realtime interpolated platform positions in 3-space, and then uses these interpolations to further refine visual odometry from optical flow. While less standardized than the NMEA-abiding GPS units, IMU devices commonly ship with vendor-supplied C/ C++ code intended to link with libusb, a standard userspace driver interface for accessing USB devices on Linux. Such userspace drivers leveraging libusb require little effort to migrate from x86 to ARM and enable developers to quickly integrate various devices with TK1 like MOSFET or PWM motor controllers for driving servos and actuators, voltage and current sensors for monitoring battery life, gas/atmospheric sensors, ADCs /DACs, and so on depending on the application at hand. Also Tegra K1 features six GPIO ports for driving discrete signals, useful for connecting switches, buttons, relays, and LEDs.

This case study accounts for common sensory and computing aspects typically found in robotics, machine vision, remote sensing and so on. TK1 provides a developer-friendly environment which takes the legwork out of integration and makes deploying embedded CUDA applications easy while delivering superior performance.

Case Study #2 Tiled Tegra

Some applications may require multiple Tegras working in tandem to meet their requirements. Clusters of Tegra K1 SoCs can be tiled and interconnected with PCIe or Ethernet. (Figure 6) The size, weight, and power advantages gained from implementing such a tiled architecture are substantial and extend the applicability of TK1 into the datacenter and highperformance computing (HPC). Densely distributed topologies with 4, 6, 8 or more K1 SoCs tiled per board are possible and provide scalability beneficial for embedded applications and HPC alike. Consider this example based on an existing embedded system, employing six Tegra K1s:The six K1s are interconnected via PCIe gen2 x4 and a 32-lane PCIe switch with nontransparent (NT) bridging and DMA offload engines. This, along with a lightweight userspace RDMA library, provides low-overhead inter-processor communication between TK1's. Meanwhile sensor interfaces are provided by a Gigabit Ethernet NIC/PHY connected to each Tegra's PCIe gen2 x1 port. There's also a spare PCIe x8 expansion brought out from the PCIe switch for up to 4GB/s of off-board connectivity to user-determined I/O interfaces.

A tiled solution like this is capable of nearly 2 TFLOPS of compute performance while drawing less than 50W, and represents a large increase in the efficiency of low-power clustered SoCs over architectures that utilize higher-power discrete components. The decrease in power enables the placement and routing of all components onboard, resulting in connector-less intercommunication and improved signal integrity and ruggedization. Useful for big data analytics, multi-channel video and signal processing, and machine learning, distributed architectures with TK1 offer substantial performance gains for those truly resource-intensive applications requiring computational density while minimizing SWaP.

A new generation of embedded devices

The ground-breaking computational performance of Tegra K1, driven by NVIDIA's low-power optimizations and the introduction of integrated CUDA, leads a new generation of embedded devices and platforms that leverage TK1's SWaP density to deliver advanced features and capabilities. NVIDIA and Abaco have partnered to bring rugged SFF modules and systems powered by TK1 to the embedded space. Applications in robotics, medical and manwearable devices, software-defined radio, security, surveillance, and others are prime candidates for acceleration with Tegra K1. Beyond this, TK1's ease-of-use promotes scalable, portable embedded systems with shortened development cycles, only furthered by the wealth of existing CUDA libraries and software that now run on Tegra.

The powerful GPU-based multiprocessing platforms of today will continue to be favored in deployments in which the maximum possible pure processing capability is an absolute requirement. There can be little doubt, though, that Tegra K1 offers significant opportunity to bring powerful, rugged embedded computing to places and applications where it was previously impossible.

WE INNOVATE. WE DELIVER. YOU SUCCEED.

Americas: 866-OK-ABACO or +1-866-652-2226 **Asia & Oceania:** +81-3-5544-3973
Europe, Africa, & Middle East: +44 (0) 1327-359444
Locate an Abaco Systems Sales Representative visit: abaco.com/products/sales

abaco.com @AbacoSys

abaco SYSTEMS