

# Security Video Analytics on Xilinx Spartan-3A DSP

The processing requirements of video analytics take advantage of Xilinx FPGA parallelism, embedded and DSP processing.

by Csaba Rekeczky  
co-CTO and Vice President  
Eutecus, Inc.  
[rcsaba@eutecus.com](mailto:rcsaba@eutecus.com)

Joe Mallett  
Senior Product Line Manager  
Xilinx, Inc.  
[jmallett@xilinx.com](mailto:jmallett@xilinx.com)

Akos Zarandy  
co-CTO and Vice President  
Eutecus, Inc.  
[zarandy@eutecus.com](mailto:zarandy@eutecus.com)

The processing bandwidth requirements for a wide range of security analytics applications are forcing companies to reconsider their approach to system hardware. A single video and imaging DSP processor is insufficient for performing some of the computationally intensive analytics operations at acceptable data rates. Also, no reliable and robust solution has been demonstrated that handles high-definition (HD) resolution at full video frame rates. This has forced systems engineers to consider either a multichip or an alternative single-chip system. Both solutions have advantages and disadvantages.

A multichip system comprised of multiple DSPs generally offers designers a more familiar design flow, but has added PCB costs, takes up board/system space and can create system performance issues. A single-chip solution, on the other hand, would seemingly have cost, footprint and power advantages—but it could potentially present designers with a steeper learning curve, adding complexity and engineering cost to the design project and potentially delaying the product release.

That was the dilemma we faced here at Eutecus, Inc., a video analytics company based in Berkeley, Calif., during the system specification phase of our next-generation analytics product, the Multi-core Video Analytics Engine (MVE™).

We had implemented our first-generation product on Texas Instruments' DaVinci Digital Media System-on-Chip platform. But for our second generation,

we needed a bit more processing power and system integration. We quickly decided that a multidevice, DSP solution wasn't cost- or system-effective. We needed a single-chip solution that would allow us to easily port the IP developed in our earlier product and add more to it for the MVE.

With a bit of research, we found the Xilinx® Spartan®-3A DSP 3400A. The device provided 126 dedicated XtremeDSP™ DSP48A slices, had more than enough performance to accommodate our system requirements and came in at an attractive price.

Further, our migration fears were quickly laid to rest when we realized that the Spartan-3A DSP was supported by the Xilinx Embedded Development Kit. The EDK allowed us to implement a dual-processor hardware architecture, based on the Xilinx MicroBlaze™ embedded processor, similar to the dual-processor

hardware architecture we had been using on Texas Instruments' DaVinci platform.

With our device selected, we set out to create a single-chip analytics design by porting our existing DaVinci code base to the Xilinx dual-processor embedded system. We then created just the right set of accelerator blocks in the FPGA fabric to meet our exact performance requirements, which included processing high-definition video at full frame rates. The result was the MVE, which is sold into the aerospace/defense, machine vision and surveillance markets.

**Video Analytics Product Overview**

The Multi-core Video Analytics Engine relies on our InstantVision Embedded™ software and a specialized Cellular Multi-core Video Analytics (C-MVA™) coprocessor equipped with many advanced features and capabilities.

The latest version of the MVE/C-MVA is capable of handling HD resolution at video frame rates. It consumes less than 1 watt and executes multiple event-detection and classification algorithms fully in parallel. Figure 1 shows the output of a video analytics traffic-monitoring example, which classifies different types of vehicles, flow direction, lane changes and lane violations—all concurrently and marked by different colors.

We designed the C-MVA coprocessor in such a way that we can significantly expand the complexity of its operations to support the analytics functions in the dense-object space, which is particularly challenging because it requires analysis of overlapping and incomplete objects/events. Application-specific DSPs offer extremely poor support for this type of feature as well as for processing scaling. Both are much more flexible within FPGAs.

The 126 XtremeDSP DSP48A slices within the Xilinx Spartan-3A DSP 3400A FPGA are capable of 30 GMACs of DSP performance, so the device was well suited to the demanding cost and performance requirements of video analytics. The Xilinx FPGA also allowed us to add future video analytics functions and the associated event-detection examples, based on our customers' needs. We've summarized them in Table 1.




Figure 1 – Road map of the Multi-core Video Analytics Engine (MVE) and example application

Event Detection Examples Enabled by C-MVA	Ver. 1	Ver. 2	Ver. 3
Detect and count all moving persons, vehicles and other objects	Green	Green	Green
Detect, count and differentiate people, vehicles and other objects	Green	Green	Green
Detect objects moving in line with the specified motion flow/counterflow	Green	Green	Green
Detect and classify lost or stolen objects, entering-loitering in forbidden area, lane violation	Green	Green	Green
Detect and classify gestures, characteristic motion paths	White	Green	Green
Detect and classify related behaviors within the crowd	White	White	Green

Table 1 – Supported video analytics functions for typical event detection applications

Further, the Xilinx FPGA and ISE® Design Suite tools gave our analytics design teams more flexibility in customizing solutions for end customers. We can tailor the video analytics engines and system-on-chip (SoC) solutions quickly, by rapidly prototyping for both standard- and high-definition video processing. This allows us to efficiently use the available resources in the Spartan-3A DSP 3400A or the lower-cost Spartan-3A DSP 1800A FPGA device based on the customer's needs.

An FPGA solution has the added benefit of allowing us to create a variety of derivative end products that use the same hardware platform. Since we have designed multiple analytics accelerator engines using VHDL, we can integrate specific cores into the C-MVA coprocessor. This approach

allowed our engineers to reuse our dual-MicroBlaze embedded system to create a different FPGA programming file, resulting in an extremely scalable solution that we can easily tailor to a wide variety of analytics applications.

**Migrating from DaVinci to Xilinx FPGA**

Our previous-generation video analytics products were based on the TI DaVinci Digital Media System-on-Chip TMS320DM6446, which included both the ARM9x processor and the C64x+ DSP coprocessor. Our design used the ARM9x for communications and control and the C64x+ for the DSP processing for the analytics algorithms. However, that combined system could not address the processing requirements our second-generation

product would need. Thus, we turned to the Spartan-3A DSP FPGA family.

We simplified the task of design migration by creating a Xilinx embedded system that included two MicroBlaze v7 soft-core processors running independently. This architecture allowed us to port the ARM and DSP processor code separately, which greatly simplified design migration. Figure 2 shows a block diagram of the Euteucus hardware system and the MVE-based reference SoC design.

Our MVE engine consists of the InstantVision Embedded software running on the MicroBlaze (MB0), system control and communications on the MicroBlaze (MB1) and the C-MVA coprocessor, which we designed as a modular chain of hardware accelerator IP cores running in the FPGA fabric.

Migrating our ARM and DSP code proved to be straightforward using the Xilinx ISE Design Suite and the MicroBlaze soft cores. One of the distinct advantages of our InstantVision cross-platform environment is that it was written in high-level, standard C/C++ language and required little modification.

Once we ported the code, we validated that it had the correct functional behavior and identified any performance bottlenecks. Accelerating the C/C++ code that we initially developed for the TI processors proved to be the critical challenge, as we used several of the DaVinci C64x+ coprocessor accelerator blocks during the assembly-level optimization for this platform. We followed a series of steps in this transition, starting with initially replacing these blocks with high-level C functions. Eventually, we replaced the majority of these functions with equivalent accelerator blocks running on the FPGA fabric.

From a functional point of view, our solution has three layers that comprise the MVE, which receives the standard-/high-definition video flow as input data and then generates the event-detection metadata. This resultant metadata provides the object/event-tracking and classification results, along with several image flows for debugging purposes as the output of the analysis. We implemented functional

blocks as either embedded software running on a MicroBlaze processor or specialized IP cores. We placed these specialized hardware accelerators into the FPGA fabric, and the complete chain of these accelerators comprises the C-MVA analytics coprocessor.

As shown in Figure 3, the three algorithmic layers of the MVE video analytics engine consist of several main functional blocks, most of which we can significantly

accelerate by using specialized IP cores that rely on dynamic configuration of the resources available in the FPGA. We designed the C-MVA coprocessor based on these IP cores, so as to accelerate the processing front end and midlayer (see Figure 4) of the entire analytics algorithm. This modular approach, supported by Xilinx's ISE Design Suite, allowed us to scale the system in terms of both performance and power consumption.




Figure 2 – Dual-MicroBlaze™ System-on-Chip (SoC) architecture MVE Engine coprocessor block diagram

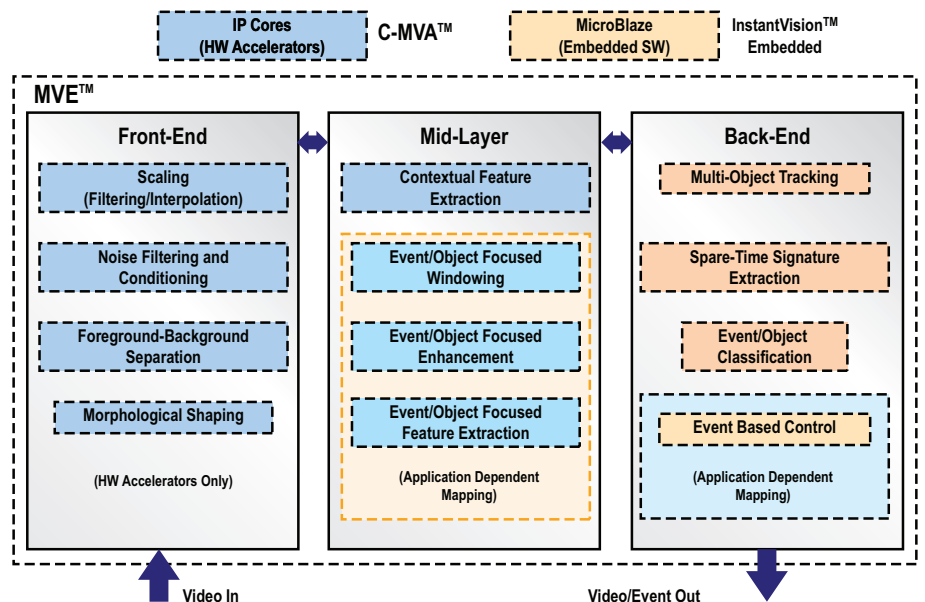


Figure 3 – Block diagram of the video analytics algorithm organization




Figure 4 – MVE analytics engine, InstantVision and driver software

**Turbo-charging using FPGA Accelerator Blocks**

To truly realize the full potential of an FPGA-based video analytics system, we needed to design and integrate the video accelerator engines into the embedded base system. We anticipated several of the performance bottlenecks, so our design team had begun early development of a set of accelerators using VHDL. The code profiler, included as part of the Xilinx ISE Design Suite and the Embedded Development Kit, proved instrumental in helping us identify further performance bottlenecks and develop all the accelerator blocks we required for this design. Table 2 provides a comprehensive list of IP core families.

Our development team, like those at many other companies, consisted of separate hardware and software developers. It was critical to the success of this project to maintain developer productivity by preserving sufficient abstraction between these two design domains. We streamlined this task using a feature in Xilinx Platform Studio, Create IP Wizard, which generates RTL templates and software driver files for hardware accelerator blocks. These templates include the interface logic the design required to access registers, DMA logic and FIFOs from the embedded system. Once we used the template to create the RTL, we placed the RTL into the embed-

ded IP catalog, where a developer can further modify it as needed.

Our IP core development procedure includes a generic, modular peripheral block development flow for the PLB46-MPMC-OPB-based backbone. These peripherals consist of both single- and multi-I/O prototypes (SIMO, MIMO, MISO models), allowing us to flexibly create a multithread coprocessor pipeline for demanding image flow processing algorithms. We achieved this by combining the IP cores in almost arbitrary order and configuring them during the design and customization of various analytics engines.

The MVE analytics engine consists of the InstantVision Embedded software modules and the hardware accelerators that make up the C-MVA analytics coprocessor. We prototyped the MVE in a Xilinx Spartan-3A-DSP 3400A FPGA and created our SoC reference design. It includes all the required I/O functions for communication and data streaming (see Figure 2 for the complete hardware-firmware block diagram). This complete SoC reference design, encompassing not only the MVE analytics engine but also all the supporting I/O modules, uses 91 percent of the logic slices, 81 percent of the block RAMs and 32 percent of the DSP slices.

Separating out the MVE analytics engine (excluding the MPMC-PLB part of the backbone and specialized I/O components) uses only 46 percent of the logic slices, 44 percent of the block RAMs and 23 percent of the DSP slices, thus making a migration path to the lower-cost Spartan-3A-DSP 1800A FPGA device feasible.

We designed all the IP cores of the C-MVA coprocessor to complete their associated processing within a single clock cycle. This feature, combined with the asynchronous FSL interfaces, in turn allows the system integrator to drive the C-MVA coprocessor with a different clock domain from the rest of the system. Doing so allows the C-MVA to run at the lower pixel clock frequency while driving the backbone at a higher-frequency internal system clock, greatly reducing power consumption while maintaining the system's performance requirements.

C-MVA IP Core Family	Ver. 1	Ver. 2	Ver. 3	Function
IPC-WSC				Image flow, up/down scaling and windowing
IPC-CNF				Image flow conditioning and noise filtering, including gain control and contrast modification
IPC-FBS				Foreground-background separation
IPC-BMF				Binary morphological filtering, with size classification and contour-structure shaping
IPC-SFE				Multi-event/object signature and/or feature extraction
IPC-EFE				Event/object-focused enhancement
IPC-EBC				Application-specific event/object-based control
InstantVision™ Embedded				Algorithmic framework and specific modules for video flow analytics

Table 2 – IP core families developed as special hardware accelerator blocks for three generations of MVE / C-MVA

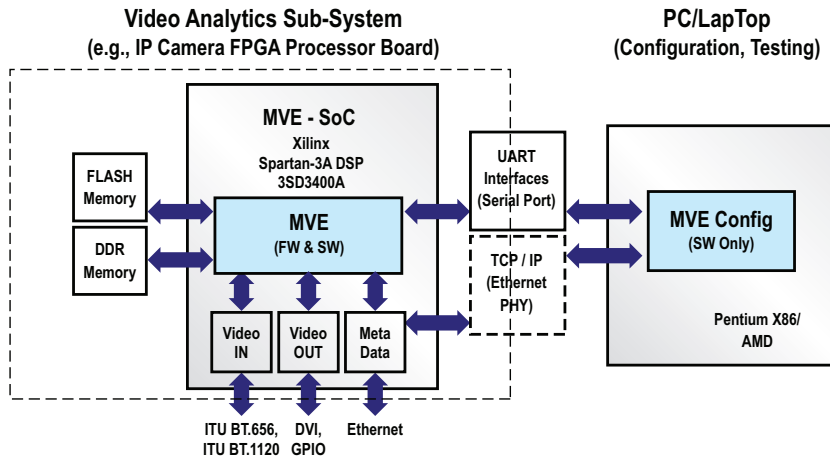


Figure 5 – Complete hardware-firmware-software reference design

**Customization, Packaging and System Integration**

To prove out and further develop the system, we created a security/surveillance demonstration along with all software layers, which allows users to rapidly integrate our product in their systems at various layers (see sidebar). The high-level block diagram of the complete SoC design, which encompasses hardware IP cores, firmware and software in a single reference design, is shown in Figure 4.

We can combine system integration with flexible customization at varying levels within the hardware, firmware and software components. The server-level customization can include tailor-made SoC designs in FPGA, while at the client (configuration) level, modifications are applied

to the WIN32 or .Net API layers. This scheme allows us and our customers to rapidly prototype various configuration and test interfaces.

Users can build client-server communication on UART or TCP/IP to provide flexible configuration management, performance fine-tuning, status monitoring and firmware updating.

Even though we've just finished our second-generation product, we've already begun to look at requirements for our third generation. Judging from our experience with this project, we'll strongly consider Xilinx for the new one, especially as the company introduces reliable, newer and more advanced devices and DSP capabilities on the most advanced process technologies.

**Accelerating Development Using the XtremeDSP Video Starter Kit, Spartan-3A DSP Edition**

As part of our development and demonstration strategies, Euteucus created an MVE Video Analytics Development Kit to give users a rapid development and prototyping platform for FPGA-based video systems. Our development kit is built upon the XtremeDSP Video Starter Kit–Spartan-3A DSP Edition ([http://www.xilinx.com/vsk\\_s3](http://www.xilinx.com/vsk_s3)), which includes an FMC video I/O daughtercard, CMOS camera, cables and Xilinx development software.

After migrating our MVE analytics engine, we were able to leverage this development platform and provide our MVE analytics solution to an existing community of video systems developers for evaluation and purchase with no added hardware costs. Developers who don't already have a Video Starter Kit can easily buy one from a Xilinx distributor. Once programmed into the FPGA, the VSK will boot and begin performing the Euteucus analytics operations. The result is to give developers a quick and easy way to evaluate the performance, capabilities and cost of an FPGA-based video analytics system.

Supporting Your Future  
**HUNT ENGINEERING**

USB connected Programmable FPGA systems

**V-II Pro PowerPC®**

- Virtex®-II Pro XC2VP7
- 256 Mbytes DDR Memory
- Configurable digital I/Os
- PowerPC® boot FLASH
- USB 2 or Standalone

**Software Defined Radio**

- Virtex®-II FPGA 1M gates
- 2 ch 125Mps A/D and D/A
- TI C6203 DSP
- 32Mbytes SDRAM
- Configurable Digital I/O
- USB 2 or Standalone

**Imaging with Virtex®-4FX**

- Virtex®-4 FX12 FPGA
- 128Mbytes DDR Memory
- CameraLink connection
- VHDL Imaging Library
- USB 2 or Standalone

Programmable hardware with cables, device drivers, loading tools, examples and Power Supply. Systems can be used connected to a PC using USB, or can function standalone (without USB) using the Initialisation PROMs.

sales@hunteng.co.uk  
+44 (0)1278 760188  
[www.hunt-rtg.com](http://www.hunt-rtg.com)