

## Course Description

This course covers the advanced features of the Versal™ ACAP AI Engine, including debugging an application in the Vitis™ unified software platform, using filter intrinsics, implementing a system design in hardware, and optimizing an AI Engine kernel program.

The emphasis of this course is on:

- Describing the advanced features of the Versal ACAP AI Engine architecture
- Debugging applications using the Vitis unified software platform
- Using AI Engine filter intrinsics and programming a FIR filter using filter intrinsics
- Implementing system integration in hardware
- Optimizing AI Engine kernels using compiler directives, programming style, and efficient movement of data

### Level – ACAP 4

### Course Details

- 2 days live instructor led training (in person or online)
- 14 lectures
- 4 labs

**Price** – \$1,600 or 16 Xilinx Training Credits

**Course Part Number** – ACAP-AIE3

**Who Should Attend?** – Software and hardware developers, system architects, and anyone who needs to accelerate their software applications using Xilinx devices

### Prerequisites

- Comfort with the C/C++ programming language
- Software development flow
- Vitis software for application acceleration development flow
- [Designing with Versal AI Engine 1](#)
- [Designing with Versal AI Engine 2](#)

### Software Tools

- Vitis unified software platform 2020.2

### Hardware

- Architecture: Xilinx Versal ACAPs

Check with [Morgan Advanced Programmable Systems, Inc.](#) for the specifics of the in-class lab board or other customizations.

After completing this comprehensive training, you will have the necessary skills to:

- Debug an application using the simulation debugging methodology and event traces
- Identify and debug the various problems that arise in application development
- Utilize various AI Engine kernel optimization techniques, such as compiler directives, software pipelining, coding for performance, and core utilization
- Apply C coding guidelines for performance improvement, including function inlining, pointer restricting, and code shuffling
- Implement an AI Engine kernel using intrinsics for a symmetric FIR with `mul4_sym` and `mac4_sym`
- Implement an AI Engine kernel using a non-symmetric FIR with `mul4_nc` and `mac4_nc`
- Develop and generate an AI Engine kernel using Model Composer

## Course Outline

### Day 1

- **Introduction to AI Engine Architecture (Review)**  
Introduces the architecture of the AI Engine and describes the AI Engine interfaces that are available, including the memory, lock, core debug, cascaded stream, and AXI-Stream interfaces. {Lecture}
- **Versal AI Engine and Data Movement**  
Describes the memory module architecture for the AI Engine and how memory can be accessed by the AI Engines in the AI Engine arrays. {Lecture}
- **Debugging AI Engine Applications 1**  
Describes the application simulation debugging methodology and debugging with event traces, such as AI Engine events, DMA events, lock events, and stream events. Also demonstrates how to visualize these events in the Vitis unified software platform. {Lecture, Lab}
- **Debugging AI Engine Applications 2 (Use Cases)**  
Reviews various use cases of problems that arise, such as memory conflicts and deadlock analysis. Also, covers performance analysis (profiling) in hardware. {Lecture}
- **Overview of AI Engine Kernel Optimization**  
Explains the various AI Engine kernel optimization techniques, such as compiler directives, software pipelining, coding for performance, and core utilization. {Lecture}
- **AI Engine Kernel Optimization – Compiler Directives**  
Describes the usage of compiler directives for loop unrolling, loop flattening, and software pipelining to help improve the performance of AI Engine kernels. {Lecture}
- **AI Engine Kernel Optimization – Coding Style**  
Covers the C coding guidelines for performance improvement, including function inlining, pointer restricting, and code shuffling. {Lecture}
- **AI Engine Kernel Optimization**  
Illustrates kernel optimization techniques, such as the restrict keyword, custom pragmas, and code restructuring. Also covers using additional accumulators for improved scheduling and calculating AI Engine core utilization for the kernels to help improve performance. {Lab}

### Day 2

- **Advanced C++ Kernel Programming**  
Provides an overview of C++ kernel template functionality and the different types of states and kernel instance states using C++ classes. Also covered are kernel instance states with scalar parameters in a constructor and kernel instance state with array parameters in a constructor. {Lecture, Labs}
- **Data Types - Review**  
Provides an AI Engine functional overview and identifies the supported vector data types and high-width registers for allowing single instruction, multiple data (SIMD) instructions. {Lecture}
- **AI Engine Symmetric Filter Implementation**  
Describes advanced MAC intrinsic syntax, including the intrinsics for symmetric FIR implementation, such as `mul4_sym` and `mac4_sym`. Also provides guidelines for choosing the right fixed-point intrinsics for a FIR filter. {Lecture, Lab}

- **AI Engine Non-Symmetric Filter Implementation**  
Describes the intrinsics for non-symmetric FIR implementations, such as mul4\_nc and mac4\_nc. Also provides guidelines for choosing the right intrinsics for a FIR filter. {Lecture}
- **Floating-point Operations**  
Explains the floating-point operations fpmul, fpmac, and fpmisc as well as the fully configurable, floating-point intrinsics fpmac\_conf. {Lecture}
- **Model Composer for AI Engine Development**  
Introduces Model Composer and how it helps with kernel development for the AI Engine and modeling a heterogeneous device. {Lecture}

- Even better, and upon request, you can use these computers after hours on training days to experiment with labs. This is not possible for in-person training.
- Additionally, just like in-person training, the laptops and devCards, tools, OS, and licensing are setup in advance.
- In some ways, live online-training is better than in-person...for example, you can grant the instructor permission to look at your Vivado, PetaLinux terminal, or Vitis for extended periods of time if your lab is not going exactly as planned to a missed step.
- This is often more comfortable than two engineers crowding around a laptop screen.
- Taking remote training also allows you to learn some tips and tricks for working remote. Whether your devCard is in the lab down the hall, or across the world via VPN, you can control your Xilinx based device quickly and efficiently.

## Register Today

Morgan Advanced Programmable Systems, Inc. (Morgan A.P.S.) delivers public and private courses in locations throughout the central US region; including Iowa, Illinois, Kansas, Minnesota, Missouri, Nebraska, North Dakota, South Dakota and Wisconsin.

Visit [morgan-aps.com/training](http://morgan-aps.com/training), for full course schedule and training information.



You must have your tuition payment information available when you enroll. We accept credit cards (Visa, MasterCard, or American Express) as well as purchase orders and Xilinx training credits.

## Student Cancellation Policy

- Students cancellations received more than 7 days before the first day of class are entitled to a 100% refund. Refunds will be processed within 14 days.
- Student cancellations received less than 7 days before the first day of class are entitled to a 100% credit toward a future class.
- Student cancellations must be sent [here](#).

## Morgan A.P.S. Course Cancellation Policy

- We regret from time-to-time classes will need to be rescheduled or cancelled.
- In the event of cancellation, live on-line training may be offered as a substitute.
- Morgan A.P.S. may cancel a class up to 7 days before the scheduled start date of the class; all students will be entitled to a 100% refund.
- Under no circumstances is Morgan A.P.S. responsible or liable for travel, lodging or other incidental costs. Please be aware of this cancellation policy when making your arrangements.
- For additional information or to schedule a private class contact us [here](#).

## Online training with real hardware

During the Covid-19 period, some companies do not allow their staff to participate in live in-person training.

- Consequently, Morgan Advanced Programmable Systems, Inc. has set up a training VPN where engineer participants can take classes online using the same computers and devCards used during in-person training.