Actions

- Turn Event notifications on
- Download
- Rss

Engagement

Author: tariq.ahmad

5739 views

Attendees

Total Confirmed Attendees: 134

- cyb_at
- cghaba
- dactacc4f
- fdu13652
- aqsanche
- twincat
- olicou
- floyda1
- fusionpower
- ram894000

1 2 3 4 5

This event occurred in the past.

Using Python Overlays to Experiment with Neural Networks

Watch On-Demand

When

8 Apr 2020 11:00 AM to 12:00 PM Central Time (US & Canada)

Where

On24

Event Type

6 (On24)

AVNET Reach Further™

Utilizing PYNQ (Python Productivity for Zynq) from Xilinx and the Avnet Ultra96-V2 Single Board Computer

Summary

Python Productivity for Zynq, or PYNQ, has the ability to present programmable logic circuits as hardware libraries called overlays. These overlays are analogous to software libraries. A software engineer can select the overlay that best matches their application. The overlay can be accessed through an application programming interface (API). Using existing community overlays, this course will examine how a beginner can experiment with neural networks using PYNQ on Ultra96. The course will then point you to the best resources to begin your own PYNQ journey!

What You Will Learn In This Webinar:

- Understand what PYNQ and PYNQ overlays are
- Learn how to set up Ultra96-V2 to boot and use PYNQ See both binary and quantized neural network examples running with PYNQ on Ultra96-V2 Explore the best avenues to get started with PYNQ and Ultra96-V2
- See both binary and quantized neural network examples running with PYNQ on Ultra96-V2
- Explore the best avenues to get started with PYNQ and Ultra96-V2

```
class:  great white shark    probability: 79.08%
class:      tiger shark      probability: 18.42%
class:     hammerhead        probability:  1.41%
class:        sturgeon       probability:  0.01%
class:        stingray       probability:  0.01%
```

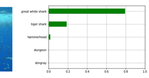

About the Ultra96-V2:

Like Ultra96, the Ultra96-V2 is an Arm-based, Xilinx Zynq UltraScale+ ™ MPSoC development board based on the Linaro 96Boards Consumer Edition (CE) specification. Ultra96-V2 is available in more countries around the world as it has been designed with a certified radio module from Microchip. Engineers have options of connecting to Ultra96-V2 through a Webserver using integrated wireless access point capability or to use the provided PetaLinux desktop environment which can be viewed on the integrated Mini DisplayPort video output. Multiple application examples and on-board development options are provided as examples.

Ultra96-V2 provides four user-controllable LEDs. Engineers may also interact with the board through the 96Boards-compatible low-speed and high-speed expansion connectors by adding peripheral accessories such as those included in the MikroE Click Mezzanine for 96Boards (available as an accessory). Micron LPDDR4 memory provides 2 GB of RAM in a 512M x 32 configuration. Wireless options include 802.11b/g/n Wi-Fi and Bluetooth 5 Low Energy. The radio module is Agency Certified in over 75 countries. UART6 are accessible on a header as well as through the expansion connector. JTAG is available through a header. The convenient JTAG/UART Pod (available as an accessory) provides both JTAG and UART connections via USB. I2C is available through the expansion connector. Two Microchip USB3320 USB 2.0 ULPI Transceivers and one Microchip USB5744 4-Port SS/HS USB Controller Hub enable multiple USB connections. Ultra96-V2 provides one upstream (device) and two downstream (host) USB 3.0 connections. A USB 2.0 downstream (host) interface is provided via the high-speed expansion. An IDT VersaClock 6E clock generator provides timing for USB 3.0, USB 2.0, DisplayPort, and the Xilinx MPSoC primary clock input. The integrated Infineon programmable power regulators generate all on-board voltages from an external 12V supply (available as an accessory) as well as providing access to power telemetry through PMBus connectivity.

Features

- Xilinx Zynq UltraScale+ MPSoC ZU3EG A484
- Micron 2 GB (512M x32) LPDDR4 Memory
- Delkin 16 GB microSD card + adapter
- PetaLinux environment available for download
- Microchip Wi-Fi / Bluetooth
- IDT VersaClock 6E clock generator
- Infineon high efficiency power management
- Aavid/Boyd Heatsink
- Mini DisplayPort (MiniDP or mDP
- 1x USB 3.0 Type Micro-B upstream port
- 2x USB 3.0, 1x USB 2.0 Type A downstream ports
- 40-pin 96Boards Low-speed expansion header
- 60-pin 96Boards High-speed expansion header
- 85mm x 54mm form factor
- Linaro 96Boards Consumer Edition compatible

Target Applications

- Artificial Intelligence
- Machine Learning
- IoT/Cloud connectivity for add-on sensors
- Embedded Computing
- Robotic
- Entry level Zynq UltraScale+ MPSoC development environment
- Training, prototyping and proof of-concept demo platform
- Wireless design and demonstrations using Wi-Fi and Bluetooth

The Presenters:

Bryan Fletcher , Technical Marketing Director, Avnet

I've been a fan of programmable logic technology since early in my career at Hewlett-Packard and Evans & Sutherland. I work now in Avnet's Products & Emerging Technologies group, currently working on Ultra96-V2, after working on many development boards over the years and as a Field Applications Engineer.

Fred Kellerman, Wireless Edge Communication Systems Architect, Avnet

I have extensive experience with software system development for Xilinx RFSoC, MPSoC, DSP, FPGA and multi-processor real-time embedded and digital wireless communication systems. I have had many opportunities to give technical presentations and teach at RIT for the EE department. I program with C / C++ / Python / Matlab and VHDL. Currently I am working with Avnet and Xilinx to integrate their ACAP, MPSoC and RFSoC ICs into 5G cellular beamforming and other communication systems. I worked with Avnet and Xilinx to help nurture their open source Python productivity for ZYNQ framework (PYNQ) specifically for the Avnet / Xilinx Ultra96 platform.

#_xilinx_webinar #_python #_xilinx #_5ggs_webinar #_machine learning #_ai #_pynq #_recording_available #_avnet_webinar

tahmad78 — 22 comments — 1 member is here

Leave a comment...

Edit • Insert • Format • Tools • Comment

Top Comments

tariq.ahmad *over 5 years ago* +3

Hi Fred27 , there was a problem with the registrations but it should now be fixed. I think neuromodulator tried as well. It should work for both of you now. Can you try again and let me know if everything ...

bhfletcher *over 5 years ago in reply to weiwei2* +3

I wasn't sure, so I fired up PYNQ to take a look. Once I got into the QNN folder, then I selected the tiny-yolo-image notebook. Here is the description. After reading through the rest of the descriptions ...

dubbie *over 5 years ago* +2

This seems to be a 'doorway' to an interesting approach to using artificial neural networks. I could follow most of it conceptually, especially the parts about classifying images, training and evaluating...

Older | Best | Newest

tariq.ahmad *over 5 years ago*

Hi Fred27 , there was a problem with the registrations but it should now be fixed. I think neuromodulator tried as well. It should work for both of you now. Can you try again and let me know if everything is working the way it should?

Fred27 *over 5 years ago in reply to tariq.ahmad*

Seems to work now. Thanks.

dubbie *over 5 years ago*

This seems to be a 'doorway' to an interesting approach to using artificial neural networks. I could follow most of it conceptually, especially the parts about classifying images, training and evaluating the results of using the ANN (PYNQ). However, it also looks quite complicated and thee seems to be a great deal to learn, in order to achieve a good understanding and to be able to exploit this approach personally.

A pity the boards cost (relatively) so much, it is not something that I would be able to invest in, just to find out a little bit more.

Dubbie

Problemchild *over 5 years ago in reply to dubbie*

True but this is a method of running Networks in/on a optimised FPGA environment, you could do this on a PC especially with say CUDA/OPENCL support and in fact if you are training a network then this would be preferable as the network would be run on the CPU side of things which isnt that powerful on the Zinq.

As they said you can run some of the stuff on smaller boards like the MiniZed if you are willing to take it's limitations

neuromodulator *over 5 years ago*

It worked now!

weiwei2 *over 5 years ago*

after watching the recorded webinar, i have a question.

during the webinar, i see that there is yolo example in folder. Is the yolo example also make use of the PL for acceleration ?

bhfletcher *over 5 years ago in reply to weiwei2*

I wasn't sure, so I fired up PYNQ to take a look. Once I got into the QNN folder, then I selected the tiny-yolo-image notebook. Here is the description.

Object detection demo

This notebook shows an example of object detection of an image. The network that is used for inference is a variant of Tiny Yolo, whose topology is illustrated in the following picture. The grey-colored layers have been quantized with 1 bit for weights and 2 bit for activations, and will be executed in the PYNQ accelerator, while the other layers are executed in software.

The image processing is performed within darknet by using python bindings.

After reading through the rest of the descriptions, it appears everything in pink/purple below is done in the PL with a hardware accelerator. Here is the description for Step 4 that describes the accelerated part:

Here's a look at the classified image with bounding boxes:

Finally, here is the performance analysis that shows how much was done in SW on the PS vs. the HW on the PL. You can see that the PL performed 4000+ operations in about 1/3rd the time it took the PS to perform 100 operations.

Good luck!

Bryan

Problemchild *over 5 years ago in reply to bhfletcher*

so in this example the PL worked out 2 orders of magnitude faster than the CPU!

Not bad at all... What sort of improvements would we expect in general ?

ankur508 *over 5 years ago in reply to bhfletcher*

Is object classification (ref. Tiny Yolo)performed for static image or from realtime stream.

weiwei2 *over 5 years ago in reply to bhfletcher*

thanks a lot for the information. this is useful 😊

is a new example not found when ultra96 first launched if i remember correctly. back then i remember only have the bnn

bhfletcher *over 5 years ago in reply to Problemchild*

John,

If you compare a software-only algorithm to a PL accelerator, the performance improvement is highly dependent on the algorithm itself. Every one of them is different. It also depends on how much effort you put into the coding. There is no rule of thumb. However, to give you a single example, the Matrix Multiply used in Avnet's SDSoC course based on Ultra96 demonstrates a 24x performance gain. While Vitis has now replaced SDSoC, the principles in that course are all still relevant, including the material around Vivado HLS, which has not changed.

http://avnet.me/ftp_on_demand --> A Practical Guide to Getting Started with Xilinx SDSoC 2018.3

Bryan

bhfletcher *over 5 years ago in reply to weiwei2*

Cheah,

I don't have a timeline for what was added when. However, as I mentioned in the webinar, these Overlays are on Xilinx GitHub and support Ultra96:

- Quantized Neural Networks -https://github.com/Xilinx/QNN-MO-PYNQ
- Binarized Neural Networks -https://github.com/Xilinx/BNN-PYNQ
- OpenCV vision pipelines -https://github.com/Xilinx/PYNQ-Computer-Vision
- Extended Kalman Filter on GPS data -https://github.com/dilpa14/pynq-ekf
- PYNQ HelloWorld image resizer -https://github.com/Xilinx/PYNQ-HelloWorld

You know also have the Vitis AI effort, which is unrelated to PYNQ. That effort is also putting accelerators into the PL, although not necessarily within the PYNQ framework. Avnet is working with Xilinx to show more examples of Vitis AI on Ultra96. For now, you can read more about it at Xilinx:

https://developer.xilinx.com/en/get-started/ai.html

You can also check out Avnet's Deep Learning course for Ultra96:

http://avnet.me/ftp_on_demand --> Introduction to Deep Learning with Xilinx SoCs

Bryan

View More

tariq.ahmad *over 5 years ago in reply to weiwei2*

bhfletcher Can you help with this?

Fred27 *over 5 years ago*

Sounds interesting. ~~However, I can't seem to register my attendance. Anyone else getting this?~~ [Now fixed]

View More

An Avnet Company © 2025 Premier Farnell Limited. All Rights Reserved.
Premier Farnell Ltd, registered in England and Wales (no 0879487), registered office: Farnell House, Forge Lane, Leeds LS12 2NE.
ICP 备案号 10230364