

# EE / CprE / SE 491

## Machine Learning Heterogeneous Computing

### Bi-Weekly Report #5

Time: Oct. 24 - Nov.8  
Client: JR Spidell  
Faculty Advisor: Diane Rover

#### Team Members:

Sandro Panchame  
Rudolph Nahra  
Alek Comstock  
Jeffery Kasper

#### Biweekly Summary

These weeks, the team began development to use the Xilinx DPU to accelerate inference. We made some changes to our development environment and are making plans to containerize our environment. Made progress on trimming the dataset to hopefully allow for higher quality training time.

#### Past Bi-Weekly Accomplishments

**Jeffery K:** Completed petalinux docker image and distributed it among the team. Created documentation for the docker image to help explain what it does and how to use it.

**Sandro P:**

Tested out the Structural Similarity index on a set of frames from one video. Trimmed down the test sample to roughly 11% of its original size. Discovered negative positions were not just applied to blinks. Had mixed results where it was somewhat effective of removing blinks and keeping the most “unique” frames.

While not all blinks had a negative position, not all non-blinks had a positive position. The resulting dataset did not have all blinks removed, however there was, what I could best describe as, uniqueness. Used .90 as a threshold.

Referenced the part of the data that classified movements, used that to trim out the frames from the dataset where blinks and noise were classified as movements. Further trimmed the dataset by removing rows where negative positions were within the dataset. We will call this the ‘clean’ set.

Improved code where the index was used and now trimming is much faster. Went from 10 hours to 3 minutes on a sample set. Using a .94 threshold for now.

Tried implementing a way for processing the frames, having the model make a prediction, have remodnav classify the movement, and display the classification in the top right corner of the 'video feed.' Since remodnav needs two positions to make a classification, a circular buffer of size two was used.

**Alek C.** I worked on our inter-process communication program. Currently awaiting testing; The RPU side is ready to send messages, and the APU is ready to receive and echo them back. My future steps i'm working towards is having the APU modify the data it receives before echoing it back.

**Rudolph:** Synthesized hardware description for the DPU (twice), exported bitstream. Read about VART (Vitis AI Runtime) libraries. Added VART libraries to our petalinux build. Resolved issue from a long time ago where our CNN crashed the Vitis AI Quantizer. Quantized and compiled our model into a \*.xmodel for the DPU. Reorganized our build process because Vitis didn't want to compile with VART libraries included - now compiling directly on the device. Loaded DPU design into FPGA fabric and confirmed that it is added to linux device tree. Wrote several scripts to accelerate development environment on the Kria board. Confirmed that interprocessor communication still works after all these changes. Finally, wrote program to perform inference using VART libraries on the DPU.

## Pending Issues

**Sandro P:** Implementation of frame extraction, prediction, and classification had an issue tied to remodnav. I have a hunch that the issue may be a perspective projection problem, where remodnav has a variable connected to converting pixels to degrees.

**Alek C.** I need help from teammates to test the program on the board; I do not have access to it, and it's a complicated process that several members have spent a long time working on. I also will need help setting up the code so it only sends to 1 APU at a time, and does so sequentially, to properly batch our work. I may need to delve deeper to figure out how to differentiate writing to different APUs, as currently it seems they all have the same write to/read location.

**Rudolph:** Need assessment of memory usage of program to run on DPU

**Jeffery K.** Create documentation for our project so that understanding what it is and how it works will be clear. Assist Alek as needed with running the software on the Kria board.

Individual Contributions

Team Member	Contribution	Weekly Hours	Total Hours
Sandro Panchame	Implemented the circular buffer. Obtained mixed results from trying to trim the dataset. Looked at the movement classification in the dataset and used that to trim out the blinks and noise from it. We now have a 'clean' set for training. Have some ground work setup in python to extract frames, make predictions, and classify movements as a video is being run. Created a trim set.	13	104
Rudolph Nahra	DPU development	40	199
Alek	Modify bits of the openAMP code	10	91
Jeffery Kasper	Making a docker image to house the petalinux tool set. Documenting the docker image.	19	119

## Plans for Coming Weeks

**Jeffery K** : Assist Alec with running his application code on the hardware and document our project more thoroughly.

**Sandro P**: Train the model on the 'clean set' and then a second model on the 'trim set' and compare performance. Resolve the issue with using remodnav.

**Rudolph**: Develop interprocessor communication method to allow passing of images from RPU to APU to DPU.

**Alek C.**: I hope to further edit the IPC of our code; I hope to have the APU edit the data it receives from the RPU before echoing it back, so it aligns with our end goal of how we want our processes to continue. As pointed out in pending issues, I need to make sure that we send to each core sequentially, as of writing I am not sure how to do that.