Person-Tracking Security Camera

Team A5:
Jerry Ding, Nathan Levin, Karthik Natarajan

Use Case

- Security cameras often lack the resolution to clearly display the face of an intruder. Even with 720p footage, faces can be hard to see at the end of a 20 feet driveway.
- Multi-camera systems are expensive and difficult to position
- Pedestrian tracking technology exists, and current neural CV approaches can have very high accuracy.
- We plan to use this to make a security camera physically track one person with a zoomed-in view.
- ECE Areas: Hardware, Software, Systems

Case Study - Package Theft

Limited field of view and low resolution for important features.

Subject still at large.

Requirements

- Must be able to track people at a rate of 10 fps
  - Able to track Usain Bolt with meter accuracy
- Data recorded to external media at 720p30
- Passively cooled up to 80°F external temperature
- Able to keep subject detected and in frame 80% of the time while subject within the range of the camera

Components - Hardware (Pessimistic Pricing)

Item	Cost
Ultra96 FPGA SoC board	$250.00
Logitech C920 Camera	$50.00
Cables	$20.00
PSU	$20.00
Motor + controller	$40.00
Poster	$100.00
Misc.	$50.00
Total	$530.00

Solution Approach: Hardware

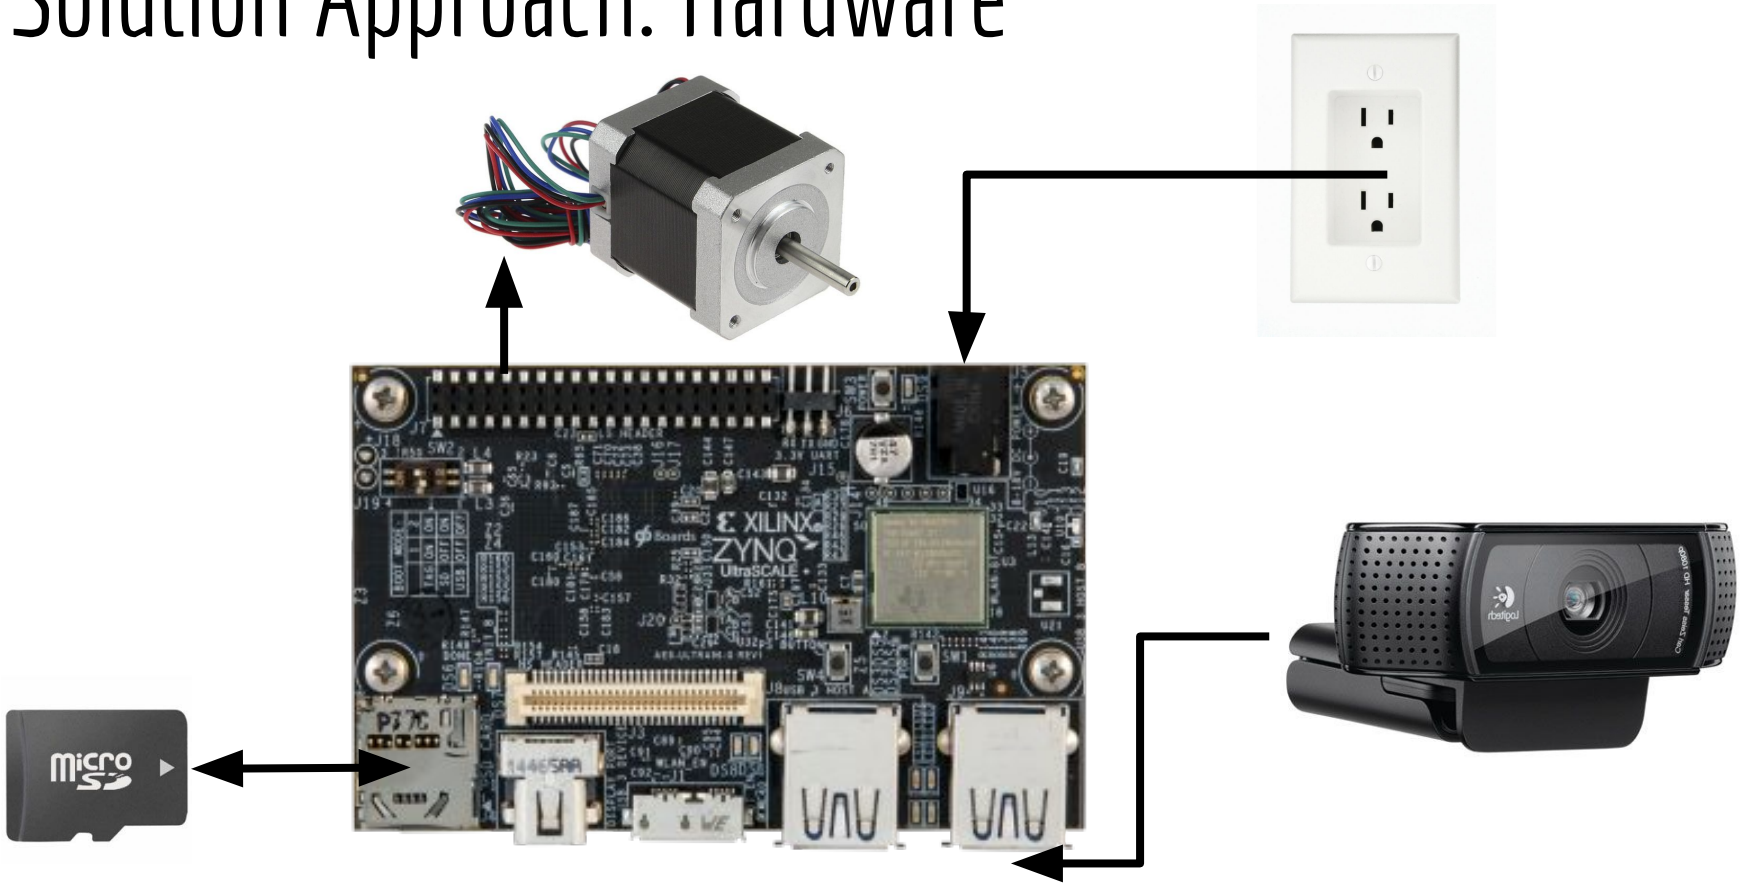

Solution Approach: Software

- Xilinx Tools
  - reVISION software suite
  - NNNDK (Deep Neural Network Development Kit)
  - Power estimation and system monitoring
  - Vivado
- Gstreamer 1.2
  - Video recording
- Neural network
  - TensorFlow
  - YOLO-v3-tiny

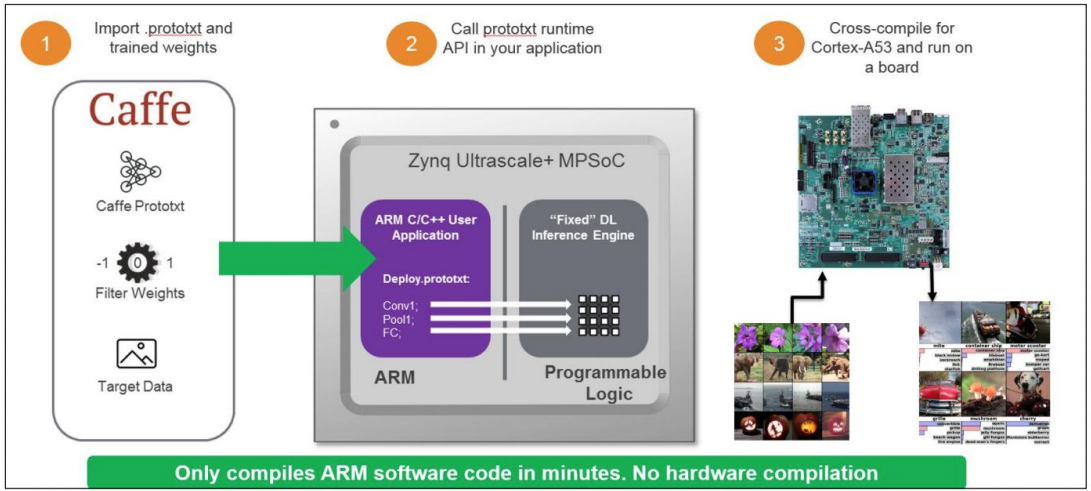

Solution Approach: Machine Learning

Xilinx provides reVISION, a tool to enable HW acceleration for ML models.

We plan to download the pretrained Yolo-v3-tiny network and use existing tools to convert it into a TensorFlow model which can be used by reVISION.

A fallback is to train a model using the Caltech Pedestrian Detection dataset.

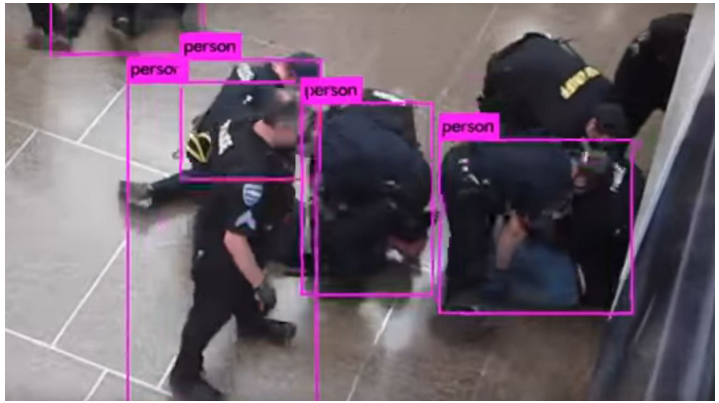

Testing, Verification and Metrics

- Our camera should be able to track a person within 20 ft of your door
- Measure mean of distance between subject and center of the frame
- Run camera for 5 hours with passive cooling and monitor for overheating
- Accuracy = (# frames identified) / (# of frames subject is within range)

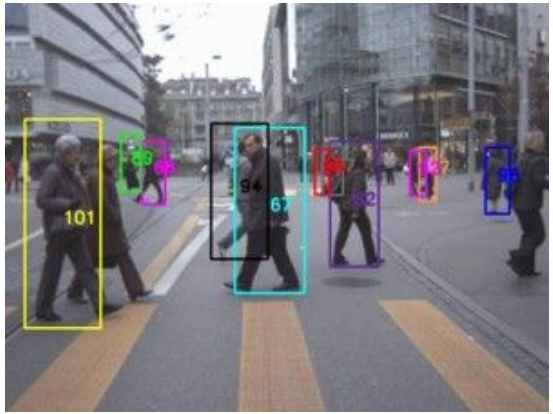

Some useful measurements

- Track the power output of the board with Xilinx's SysMon
- Track the percentage of the picture the person takes up when zoomed
- Track the miss rate, a.k.a. % of times where a person walking through camera's field of vision was not detected.

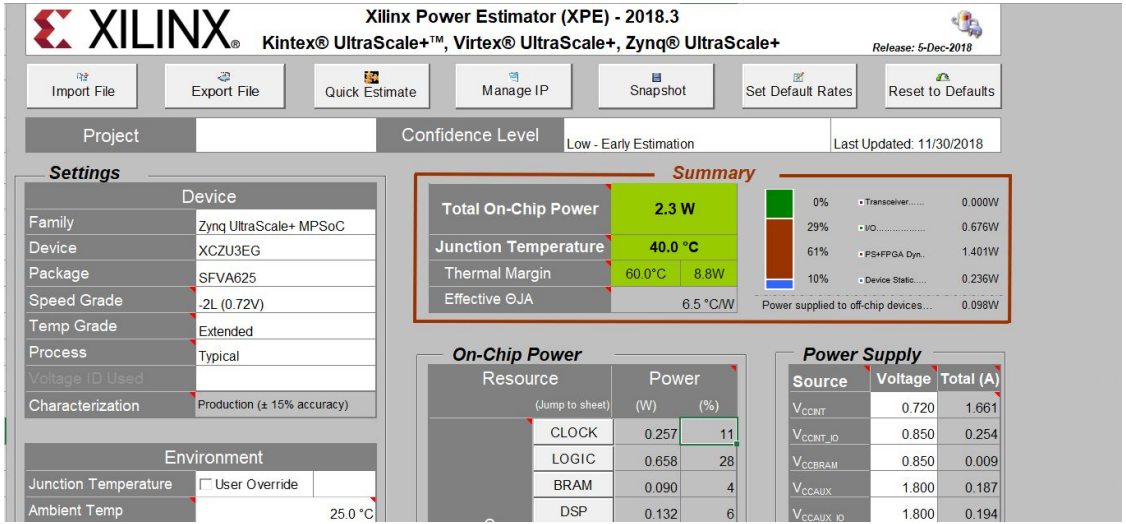

Tasks and Division of Labor

	Hardware	Mechatronics	Machine Learning	I/O
Nathan	■	■		
Jerry	■		■	
Karthik		■	■	■

Schedule


Feb '19
Mar '19
Apr '19
May '19
5
10
17
24
3
10
17
24
31
7
14
21
28
5
Person-Tracking Security C...
Setup
Order Supplies
Order
Shipping
Prepare pretrained Yolo-v3-tiny
Gather & learn the Xilinx tools
Vivado
reVISION
Checking the fit between parts
Test signalling to external devices
Is-working? component test
Make a mount for the servo
Begin interfacing work
Make reVISION import the ML model
Program SoC to move the servo
Data collection test with webcam
Implement tracking
Get alg to locate easy object (red dot)
Implement control algorithm for servo
ML Model Functionality
Get setup for capturing footage
Optimize data collection strat
Power optimization
HLS Review
Power profiling
Optimization
Status Reports & Documents
First Status Report
Design Presentation
Design Document
Ethics
Demo in Lab
Demo in Lab
Final Presentation
Final Report