Person-Tracking Security Camera

Team A5:
Jerry Ding, Nathan Levin, Karthik Natarajan

Application Area

- A security camera to combat package theft, monitor homes/businesses, etc.
- Designed for flexible deployment
  - Babysitter around the house
  - Business is closed for the weekend
- Can cover a wide area without multiple cameras or expensive and image-distorting fisheye lenses.

Solution Approach

- Product in one sentence:
  - A compact and self-contained security camera that automatically tracks and zooms into any suspicious person, and that an average store or homeowner can easily install and use.
- “Automatically tracks and zooms”
  - Mechanical apparatus to follow persons of interest
  - Includes optical zoom, a feature almost unheard of below $800
- “Compact and self-contained”
  - Powered by batteries and easily portable
- “Any suspicious person”
  - Is able to identify humans

Complete Solution (overview)

- The system stays in a powered-down state most of the time, and boots in the presence of people.
- We use deep learning to generate bounding boxes for all people in view. One is selected to be the anchor box, which determine how the camera will move.
- The coordinates of the bounding box corners are used to determine where the camera should point next and whether it should zoom in or not.
- The anchor box is periodically changed to a fresh person.
    - If there are no boxes, the system powers itself down.

Complete Solution

- Human identification accomplished by FPGA accelerated machine learning inference
  - Customized Yolov3 neural net
  - Customized Xilinx Deephi DPU
- Control logic processed on CPU part of SoC
- Motor interfacing handled by Adafruit Featherboard running Arduino C code

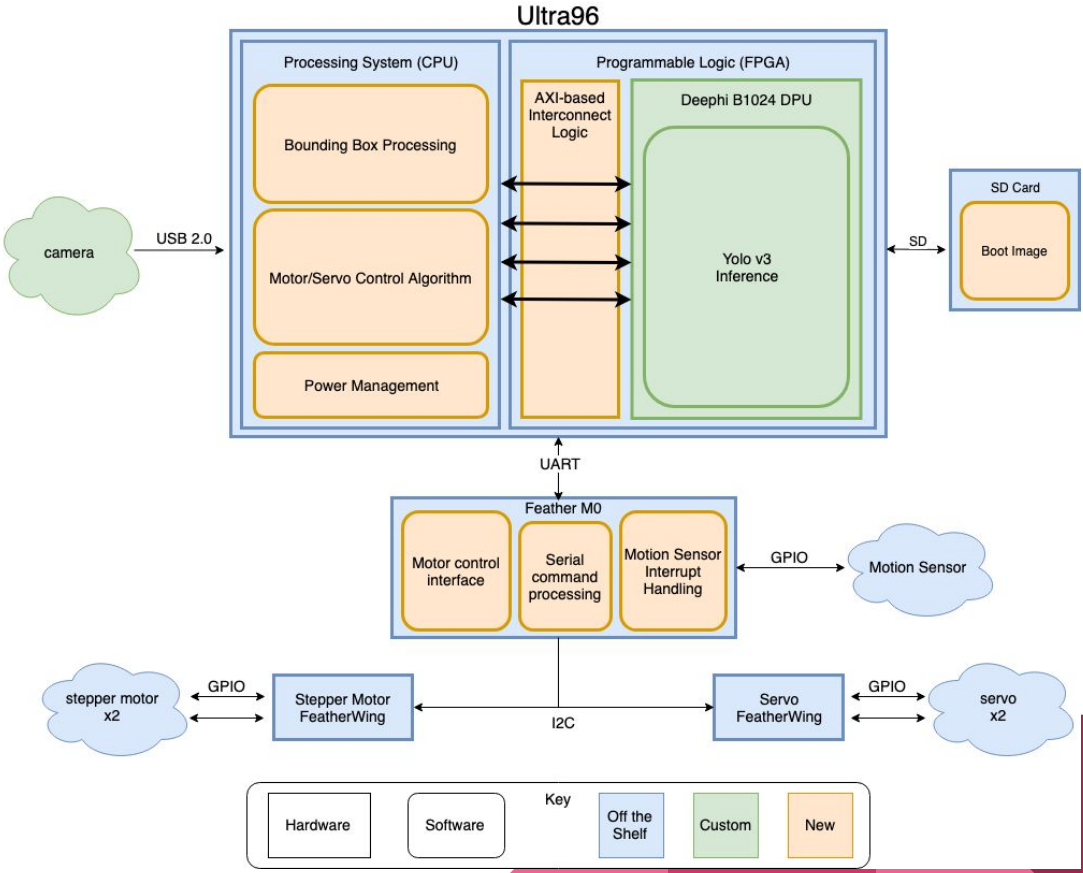

Metrics - Goals

- Runtime
  - 500 minutes active OR
  - 30 days idle

- Battery pack is ~100Wh. This implies:

  Average active power ≤ 12W
  Average passive power ≤ 139mW

- 720p video footage, tracking >10 fps

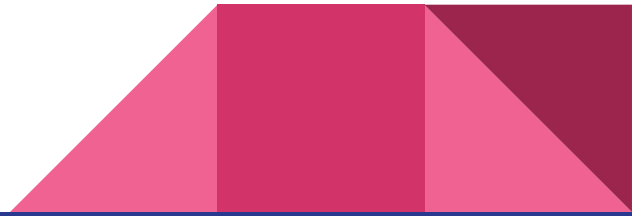

Evaluation - Power

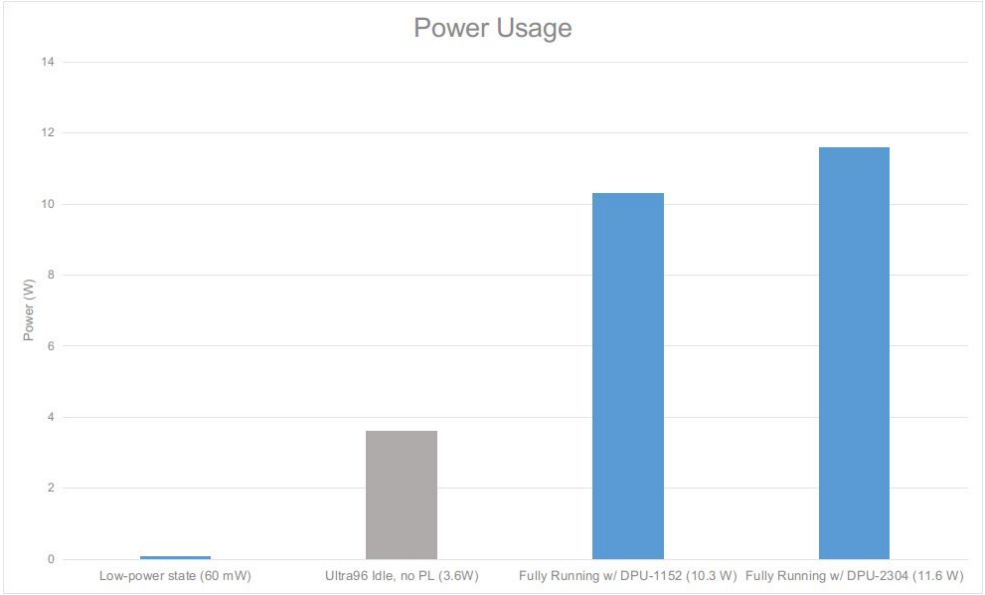

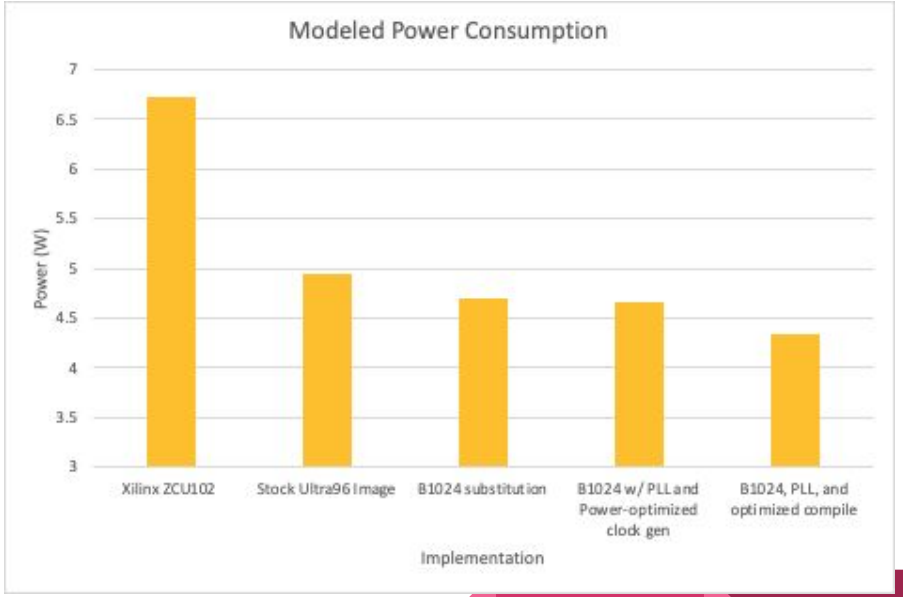

Evaluation - DPU configurations

- Xilinx said their IP couldn’t be used for the Ultra96. **They were wrong.**
- DPU implementation customized to reduce active power draw.
- Maintained adequate performance despite different DPU implementation.

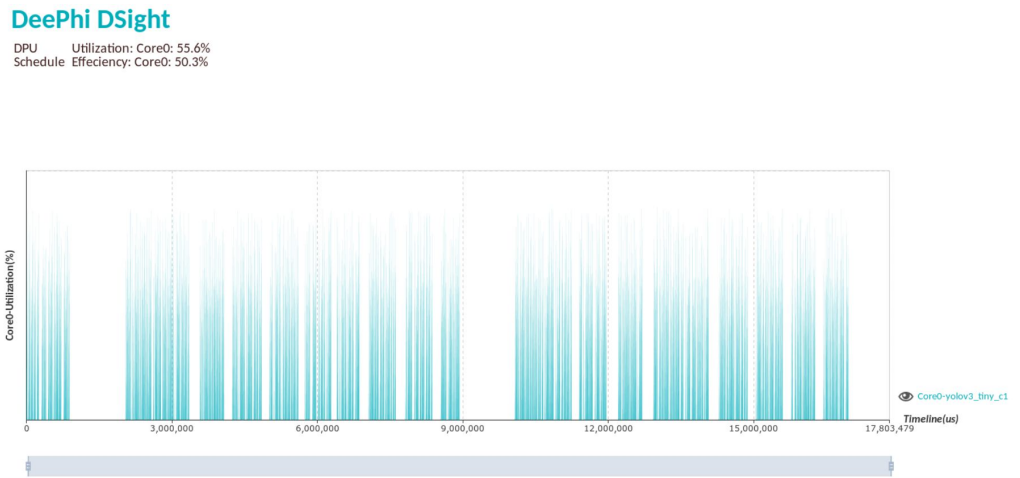

Evaluation - Tracking

Averaged over random frames at tracking:
(Height of person) / (Height of frame) = 86%

Project Management

Changes:

- Changed to 2-phase low power mode
- CPU inference
  - Unnecessary delay and complexity
- Optimization from direct HW & SW instead of HLS

Remaining tasks:

- Enclosure
- Finalize integration of motion detectors
- Final tuning

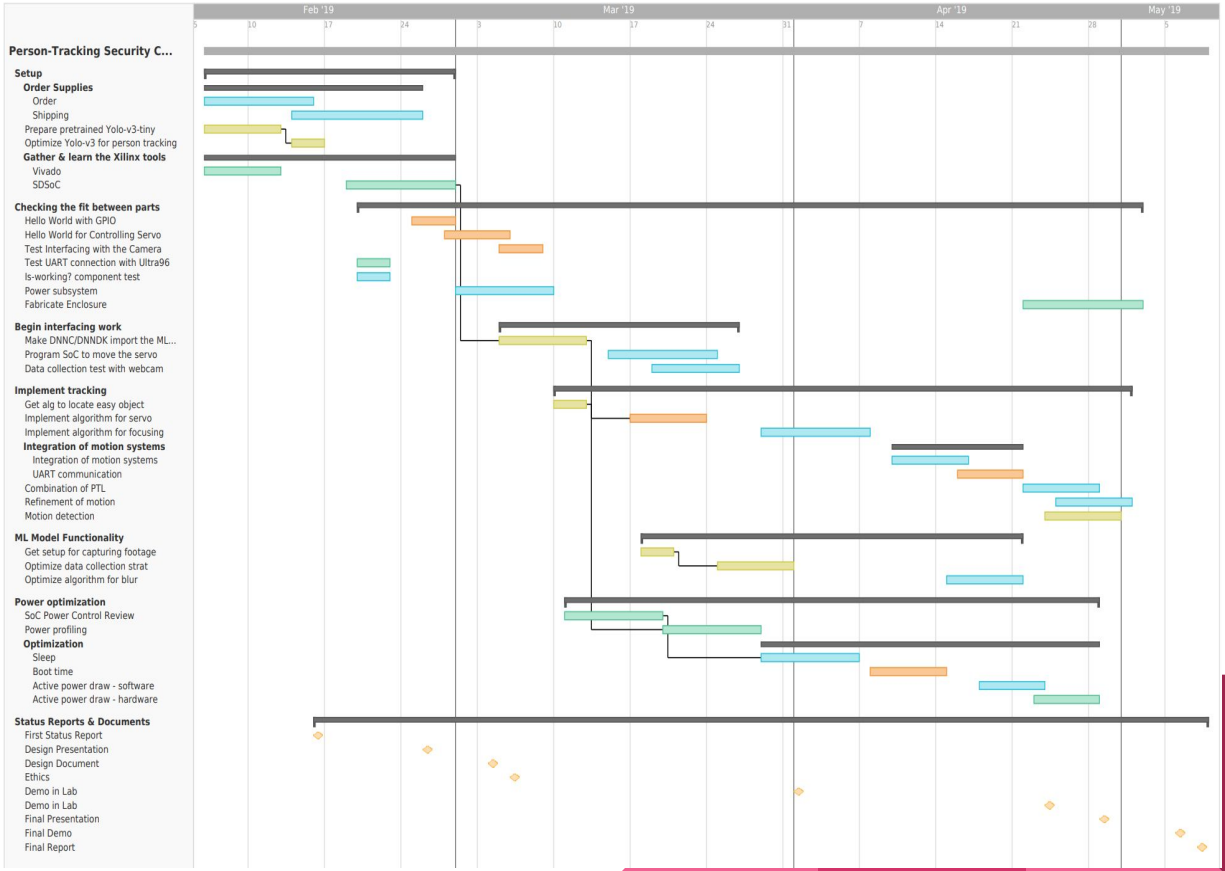

Lessons Learned

- “Plans are worthless, but planning is everything.” — Dwight Eisenhower
  - Plan before acting - understand **what** you are doing.
  - Be able to rationalize every implementation detail.
- “No plan of operations extends with any certainty beyond the first contact with the main hostile force.” — Helmuth von Moltke
  - Use that rationalization to understand **why** and **when** to change.
  - We avoided a lot of trouble by smoothly adjusting to challenges.
- Hardware is hard, software is soft.
  - Don’t rely too much on hardware for inflexible details.

Conclusion

- Successfully created a robust tracking and zoom system
- Capable of handling a number of people within the frame
- Able to quickly adjust the zoom to track targets at various depths
- Fits within target battery life parameters
- Offers features above and beyond the competition

Questions?