FossXR 2022

Free drivers for Oculus Rift headsets

- Jan Schmidt <jan@centricular.com>
- @thaytan

Centricular

Oculus Headsets

- Rift DK2 / CV1
  - “Outside-In” tracking
- Rift S
  - “Inside-out” tracking
- **Not** Quest / Quest 2

Constellation System (CV1)

- Camera sensors see IR
- LED models from firmware
  - headband adjustments, occlusion mean they don’t match
- LEDs are pulsed in sync with the camera
- Track IR blobs… extract poses
- Need to know camera poses

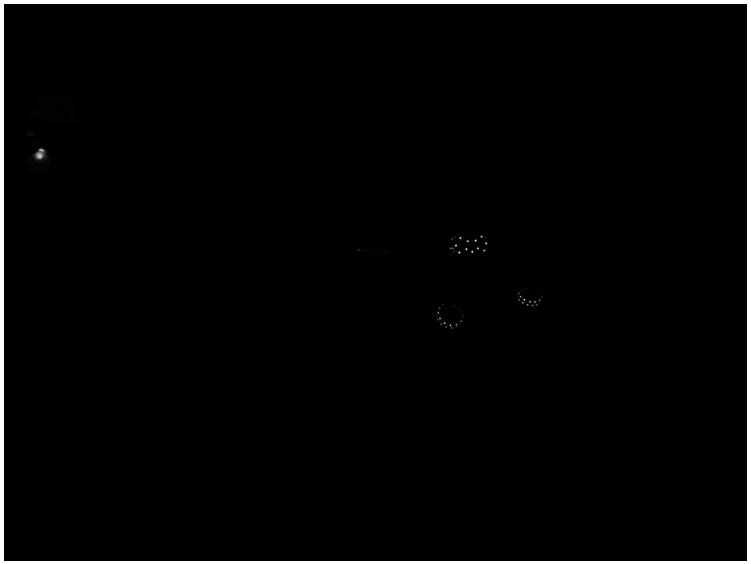

Camera Poses

- Need to know the camera positions
- Work backward from views of the headset
- Run once each time the configuration changes

- Gives an [x,y,z] position+[w,x,y,z] quaternion for each camera
- Can do online estimation

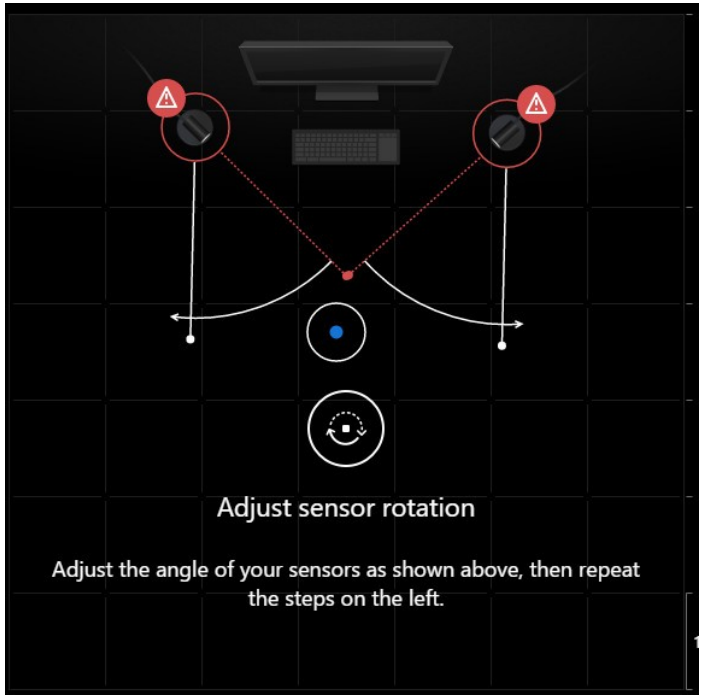

CV1 Details

- Mostly a UVC camera, with quirks
  - 52.0833 FPS (19.2ms / frame)
  - No Linux kernel support for Variable length controls
  - UVC in userspace = scheduling problems
- Frame exposure synchronisation
  - HMD ↔ Controller radio link
- Match up HMD exposures with capture
  - Based on frame arrival times and IMU sample times
  - Frames from different cameras have different arrival times though!

Constellation System (Rift S)

- 5 cameras on the headset
- Fewer LEDs = easier
  - Only the controllers

- But the cameras move
  - The SLAM/VIO is the hard part
  - (later)

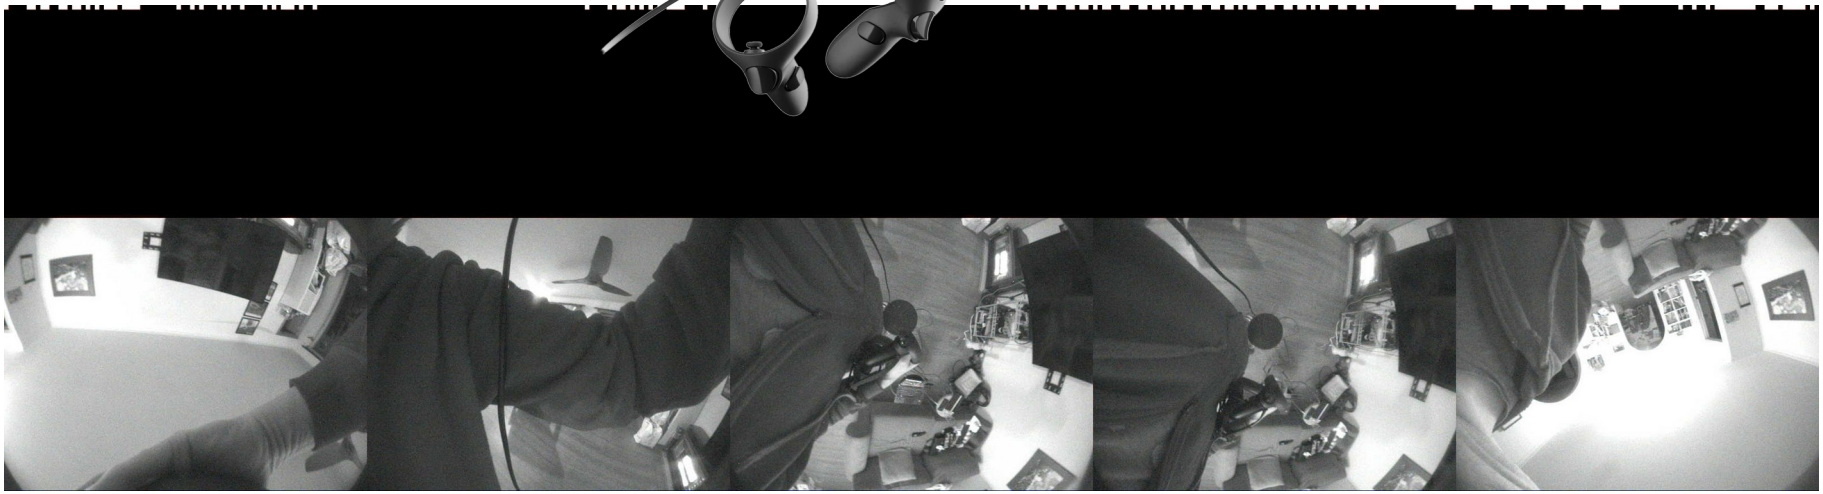

Correspondences

- Matching which blob is which LED. Home-brew depth-first search
- Pre-sort LED model positions based into lists of proximate neighbours
- Sort observed blobs by proximity
- Match groups of 4 LEDs to 4 blobs. Extract pose using LambdaTwist P3P and validate 4th point – then assess ‘pose score’
- Score based on expected matches in the bounding box / visibility of LEDs
- Two-pass strategy for big speed increase (test only nearest LEDs first)

Using the IMU

- 3DOF tracking
  - Let’s us align gravity vectors
  - Reduces the viable correspondences
- Can do better – 2-point correspondence + gravity
- Can do even better… 6DOF fusion

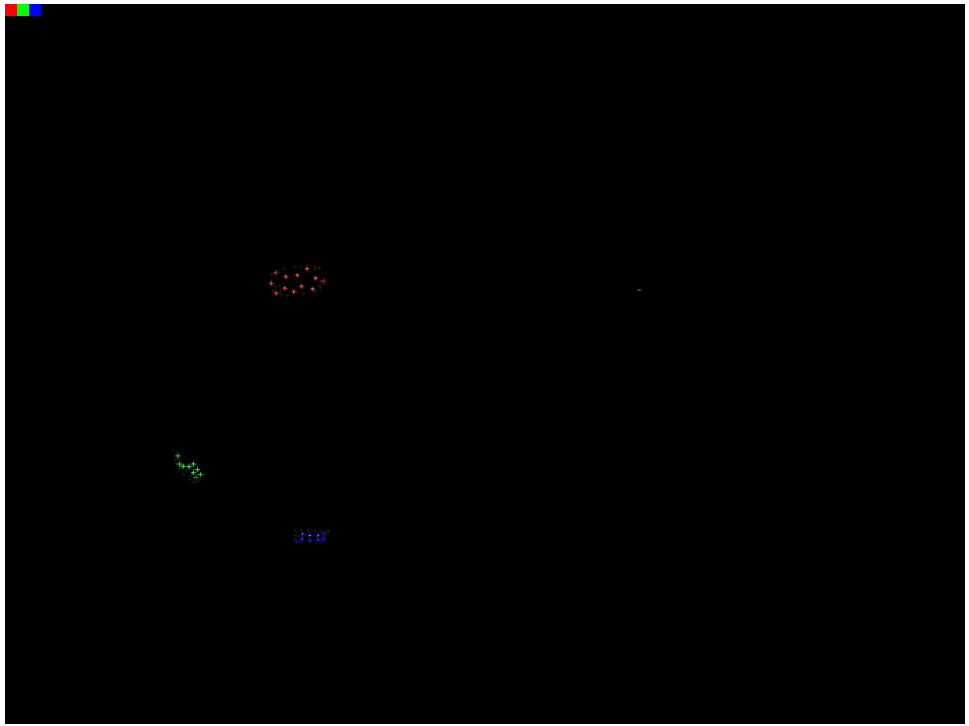

IMU+Vision Fusion

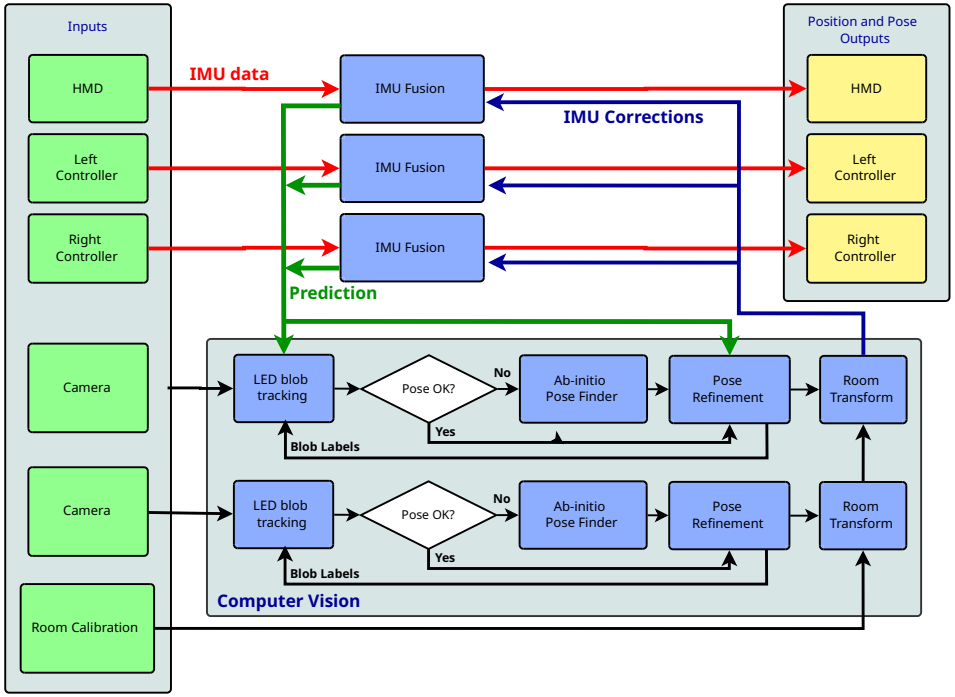

Centricular

Latency

- Frames start arriving every 19.2ms
  - USB transfer time ~17-18ms
- Image processing takes time
  - JPEG decode (for USB 2.0), 2-3ms
  - Blob extraction, RANSAC, 1-10ms
  - Correspondence search – **can be over 100ms**
    - (but more often < 40ms)
- IMU fusion is very quick
  - has to be less than 1ms

Kalman Filtering

- Improved sensor fusion
  - Unscented Kalman Filter
  - Tracks position, rotation, extracts IMU biases
  - “Slots” for lagged position updating.
- Pretty expensive
  - Runs every 1ms for the headset, 2ms for controllers
  - Could perhaps run at camera rate and predict in between?

Avoiding Glitches

- Extracted poses aren’t always right
  - Mis-identified LEDs
  - Room for improvement
  - RANSAC PnP flakiness
- Prediction time limited when tracking is lost
- 1€ exponential filter for smoothing reported pose

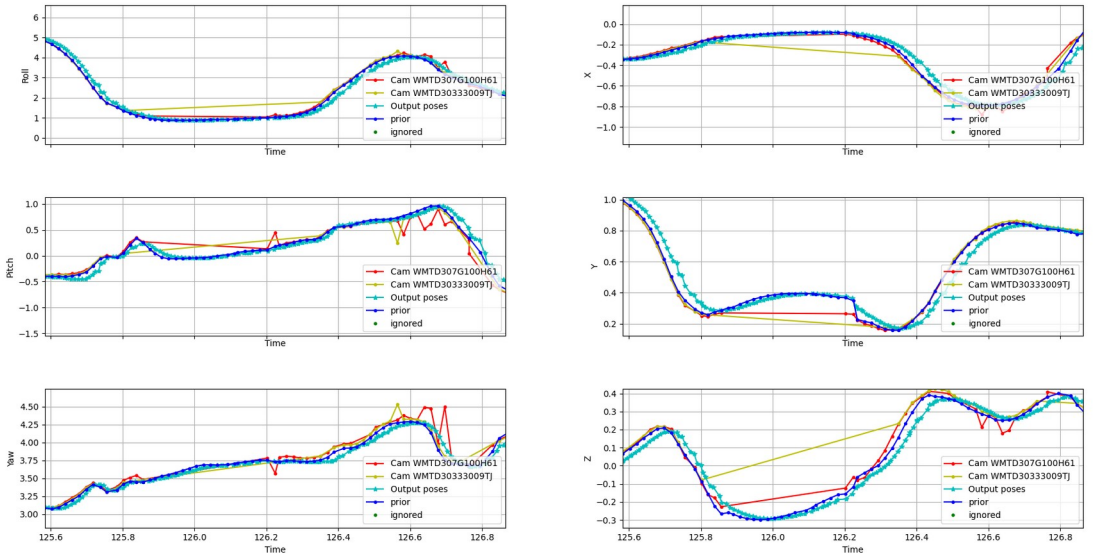

Good Tracking

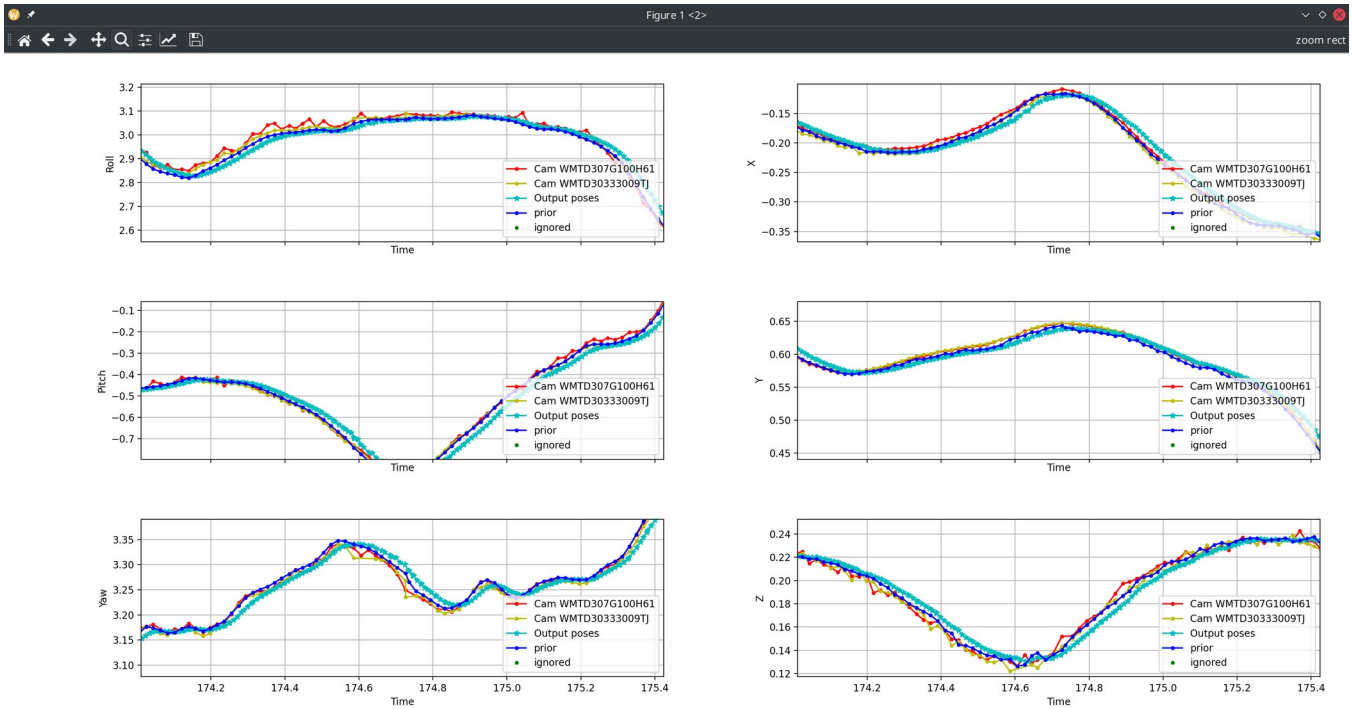

Centricular

Rift S Inside-out SLAM/VIO

- Monado, Basalt
- Exposure compensation
- Distortion compensation
  - Native “Fisheye62” model
  - Basalt conversion
- Attach the calculated pose to each frame

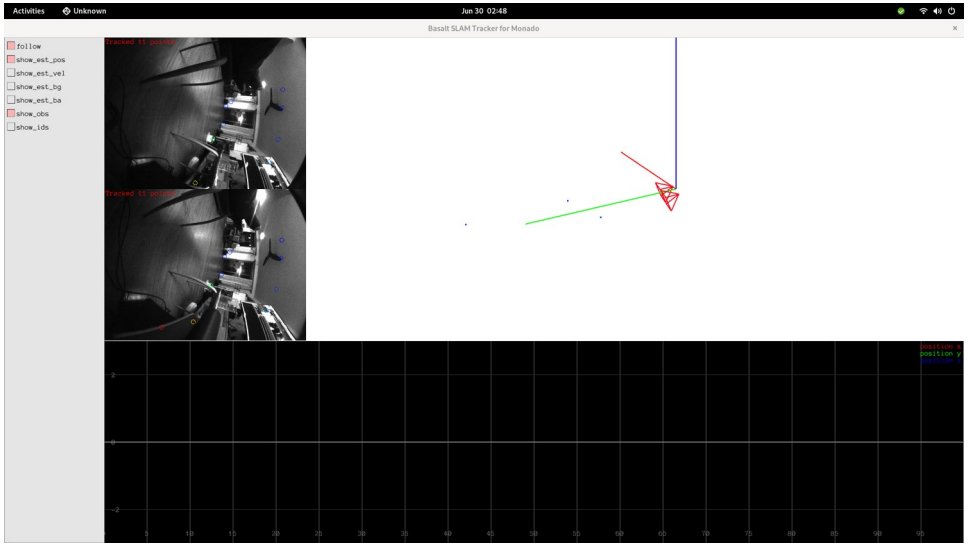

Centricular

Rift S Controllers

- Need the camera pose to predict controller LEDs
  - From the previous interleaved frame + prediction
  - SLAM better keep up (prediction error directly affects controller jitter)
- Controllers might cross view boundaries
- Unlike CV1, camera frames all arrive together

Future Directions

- Fusion performance improvements
  - IMU integration, fusion at camera rates
  - Explore optimisation approaches
- Improve pose extraction
  - Better blob position refinement
  - Figure out OpenCV ransac glitches
  - ML approaches to correspondence?
- Continue simulator / replay work
- Controller tracking for Rift S + WMR

Protocol Reversing

Protocol Reversing

- Sources of information:
  - Code decompiling
  - USB packet captures
- Either could be a breach of the EULA
- But still might be legal in your jurisdiction
- Log files from the official software can be enlightening

USB packet capture

- Wireshark + USBpcap on Windows
- Find which USB root device the port is on first
- Great to capture the first connect
  - Usually a firmware update
  - Capture without fw update too

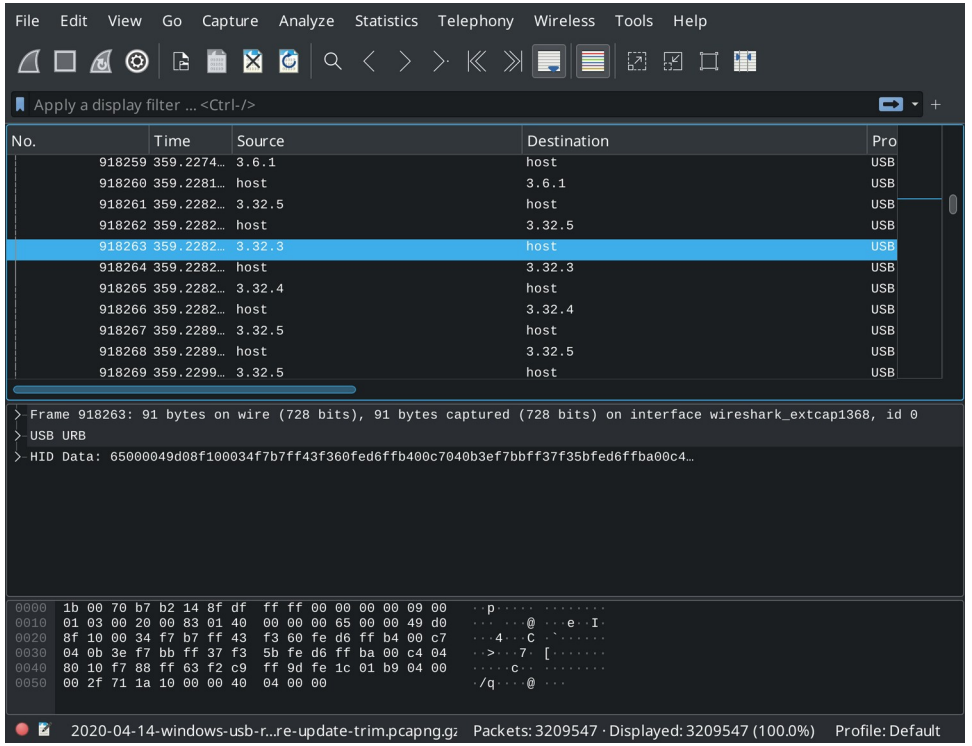

Centricular

USB packets

- Setup via HID GET/SET reports is normal.
- IMU on USB interrupt
- Controllers / radio traffic
  - Often on another USB interrupt endpoint
- Isochronous in for camera
- Isochronous out for audio

Looking for patterns

- Common operations
  - Turn on the screen
  - Enable IMU
- Known values in hex dumps
  - Screen resolutions, physical dimensions
  - Floating point values
- Take lots of notes

```
6 (0x06) - Get the display configuration
  display info:
     06 a0 05 00 0a 01 00 50 fe 03 ef 01 d2 00 5a 00 | .......P......Z.
     0c 00 01 00 02 00                               | ......
    Vertical: 1440 (0x5a0) Horiz: 2560 (0xa00) Hz: 80
    Unk0: 1 Unk1: 32441342 Unk2: 5898450 Unk3: 65548 Unk4: 2

      0x5a0  = 1440     0xa00 = 2560
      0x001              0x50 = 80 Hz
      03fe = 1022        01ef = 495
      00d2 = 210         005a = 90
      000c = 12          0001           0002
```

```
 9 (0x09) = GET REPORT ID 9 (response len 21)
   - imu config

0010                                     09 e8 03 00              ....
0020   00 6f 12 03 42 00 00 80 45 66 66 a3 43 00 00 c8   .o..B...Eff.C...
0030   41

   - Sent right after enabling the HMD with 02 01
   - 0x000003e8 = 1000 (us between accelerometer readings, or 1000Hz?)
   - 0x4203126f = 32.768 (gyro scale)
   - 0x45800000 = 4096.0 (accel scale)
         - 0x43a36666 = 326.8 (temperature scale)
   - 0x41c80000 = 25.0 (temperature offset)
```

Simple tests

- Replay earliest HID packets
  - This is when it’s useful to know which packets might modify firmware
- Omit or reorder packets, see what happens
- Try modifying values in the packets
- Pay attention to inter-packet timing or repetitions
  - Maybe something is polled until a completion value
  - Some operations take time

Rift S Radio Report

Centricular

Questions:

- Jan Schmidt <jan@centricular.com>
- @thaytan