Laser based localisation using the Riegl and Velodyne laser range scanner
(April – June)

Introduction

- **Aim:** To perform global localisation using a map acquired using the Riegl and localise a vehicle using the Velodyne

2D feature map

Some challenges

- The vehicle does not have odometry feedback
- Have to make some conservative assumptions about the odometry (eg. max distance travelled, and max rotation, between Velodyne readings)

Test sites

- Monash University (corner of Blackburn and Wellington)
- Pomonal (near the Grampians, about 4-5 hours from Melbourne)

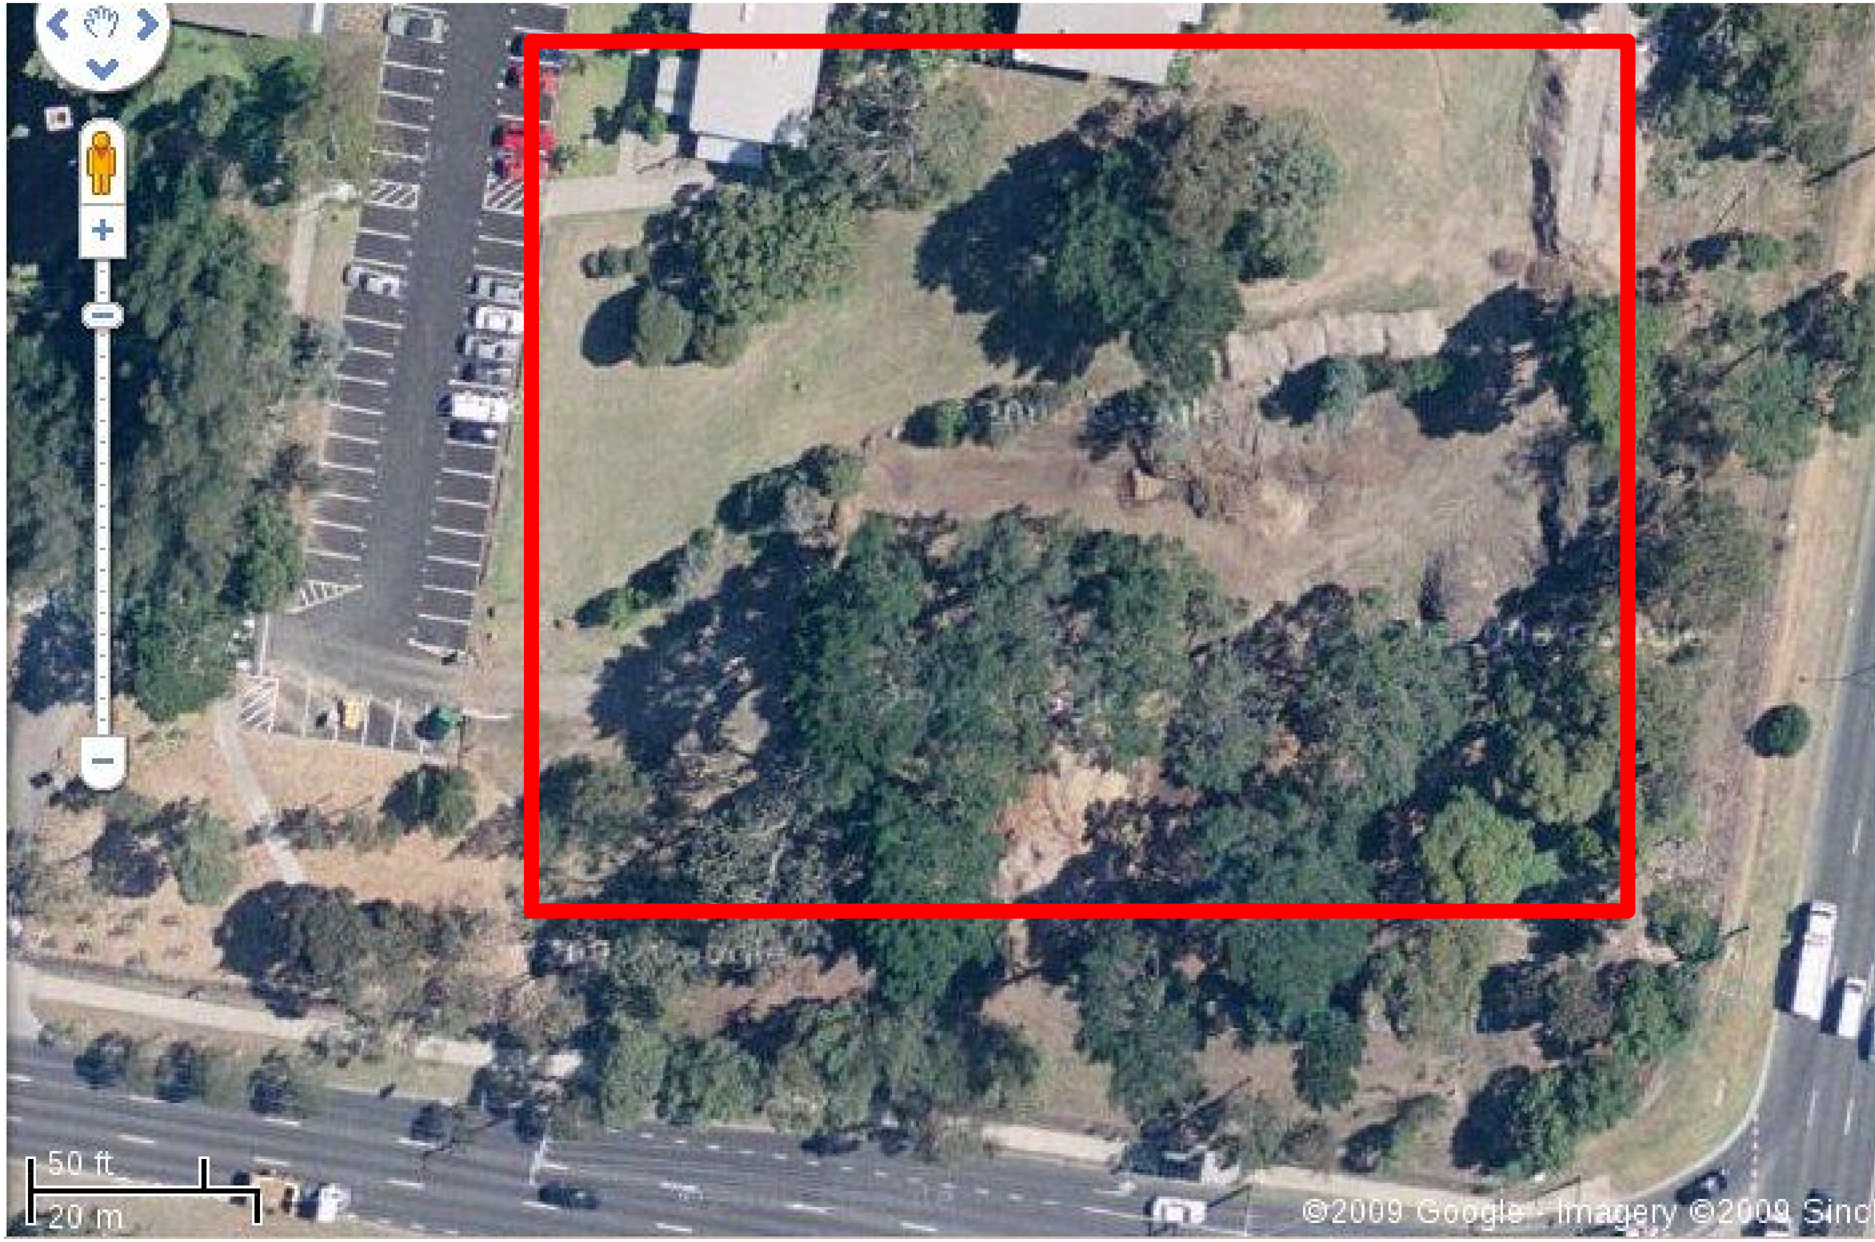

Blackburn

Wellington

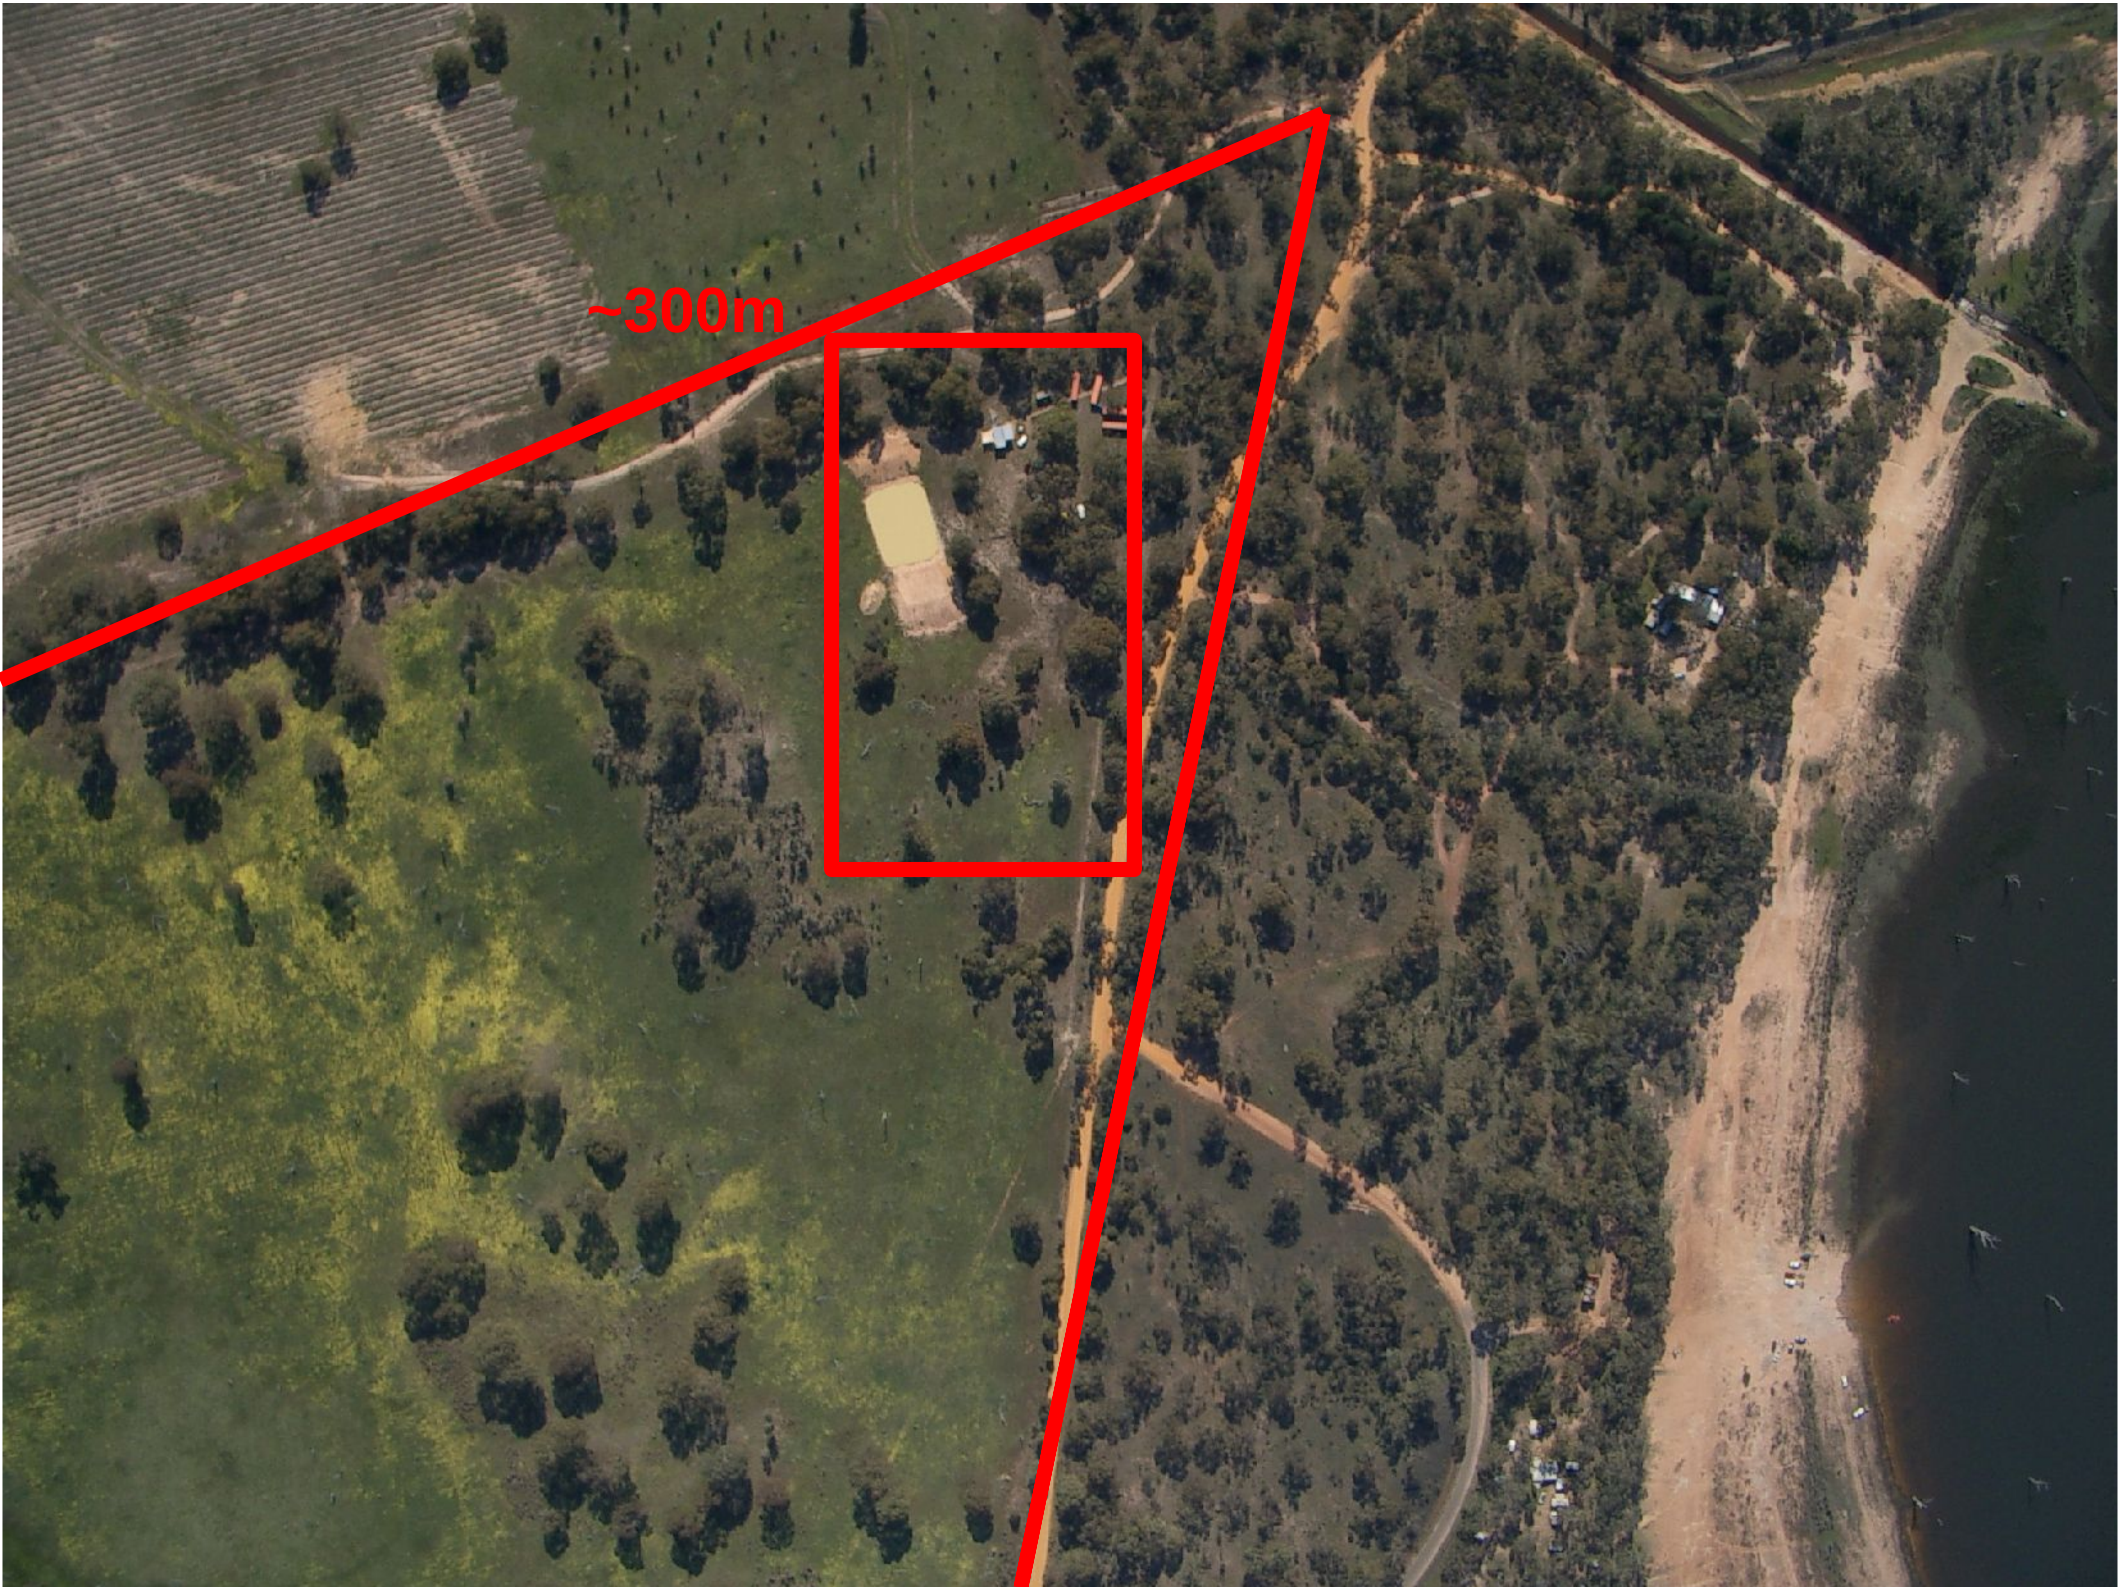
~300m

General approach

- Build a 3D model of the environment using the Riegl (same method used in my thesis)
- Generate a 2D feature map from the 3D model
- Cache the laser signatures to speed up scan matching (cached every 1 metre)
- Perform scan matching with the Velodyne data and the cached laser signatures
- Refine estimated pose using ray tracing

Building the 3D model

- 4 scans collected at Monash University
- 13 scans collected Pomonal
- Scans were manually registered
- (see demo)

2D feature map

- Project the 3D points onto a 2D image (1 pixel = 10cm)
- Keep points with Z height between 0.5 and 1.0 metres relative to the ground plane
- Grow each feature by 1 pixel (improves scan matching reliability)
- Don't need to apply any noise filtering eg. 1 pixel features. This makes it worse in fact. What appears to look like noise to us in the image is actually a feature in real life eg. thin tree branches.

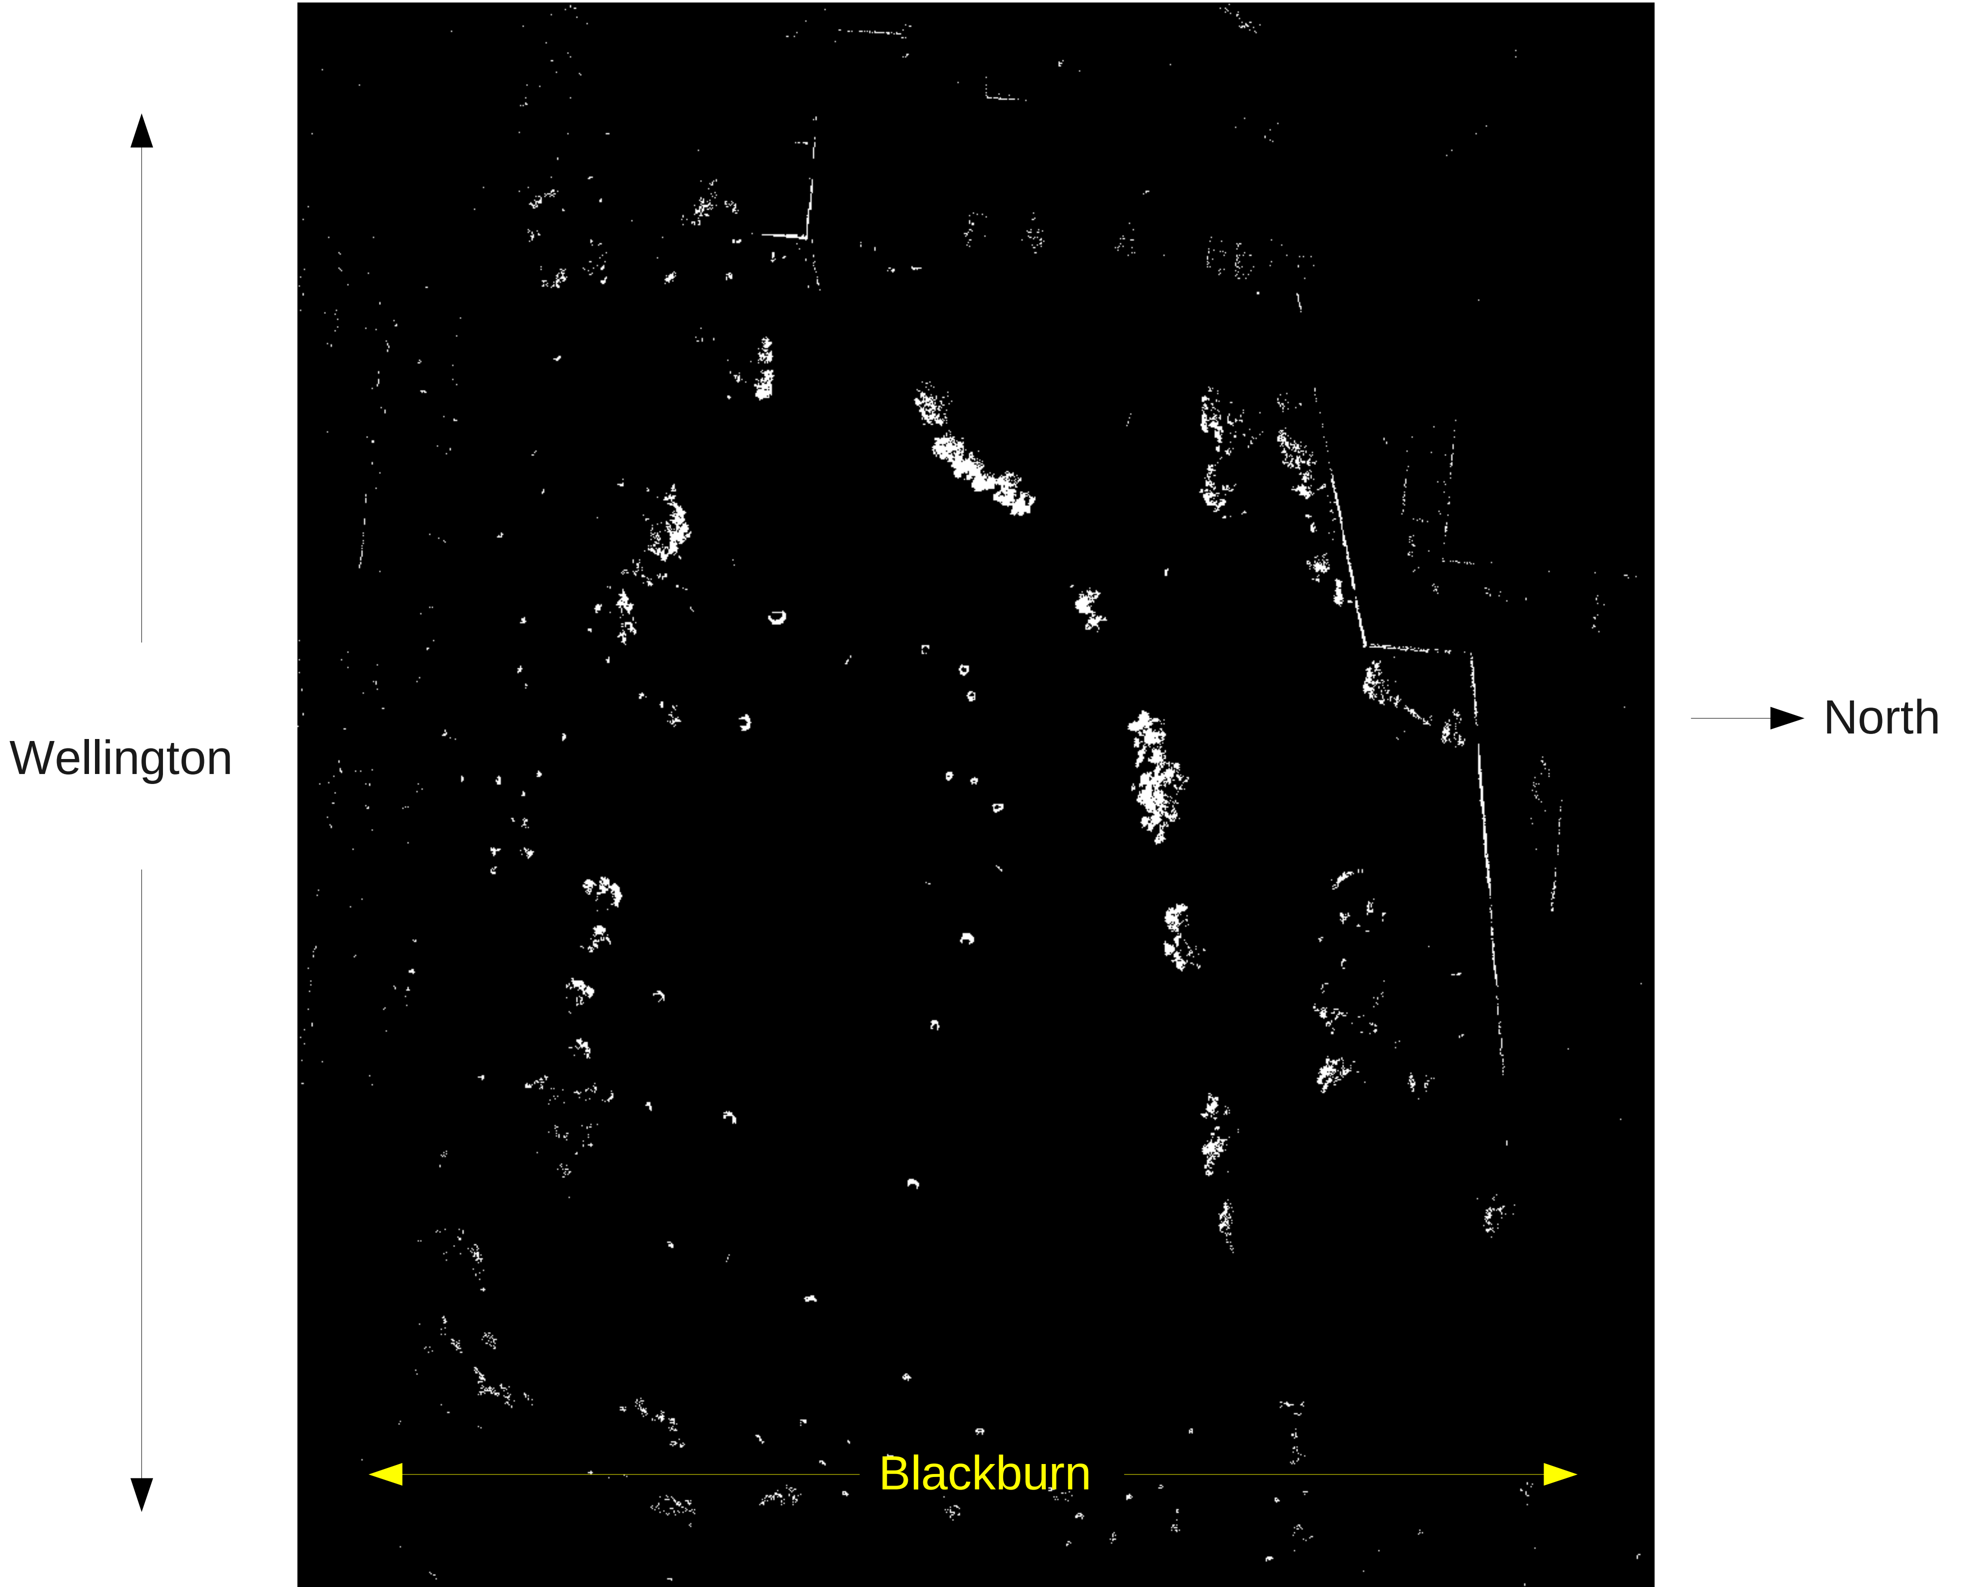

Wellington
North
Blackburn

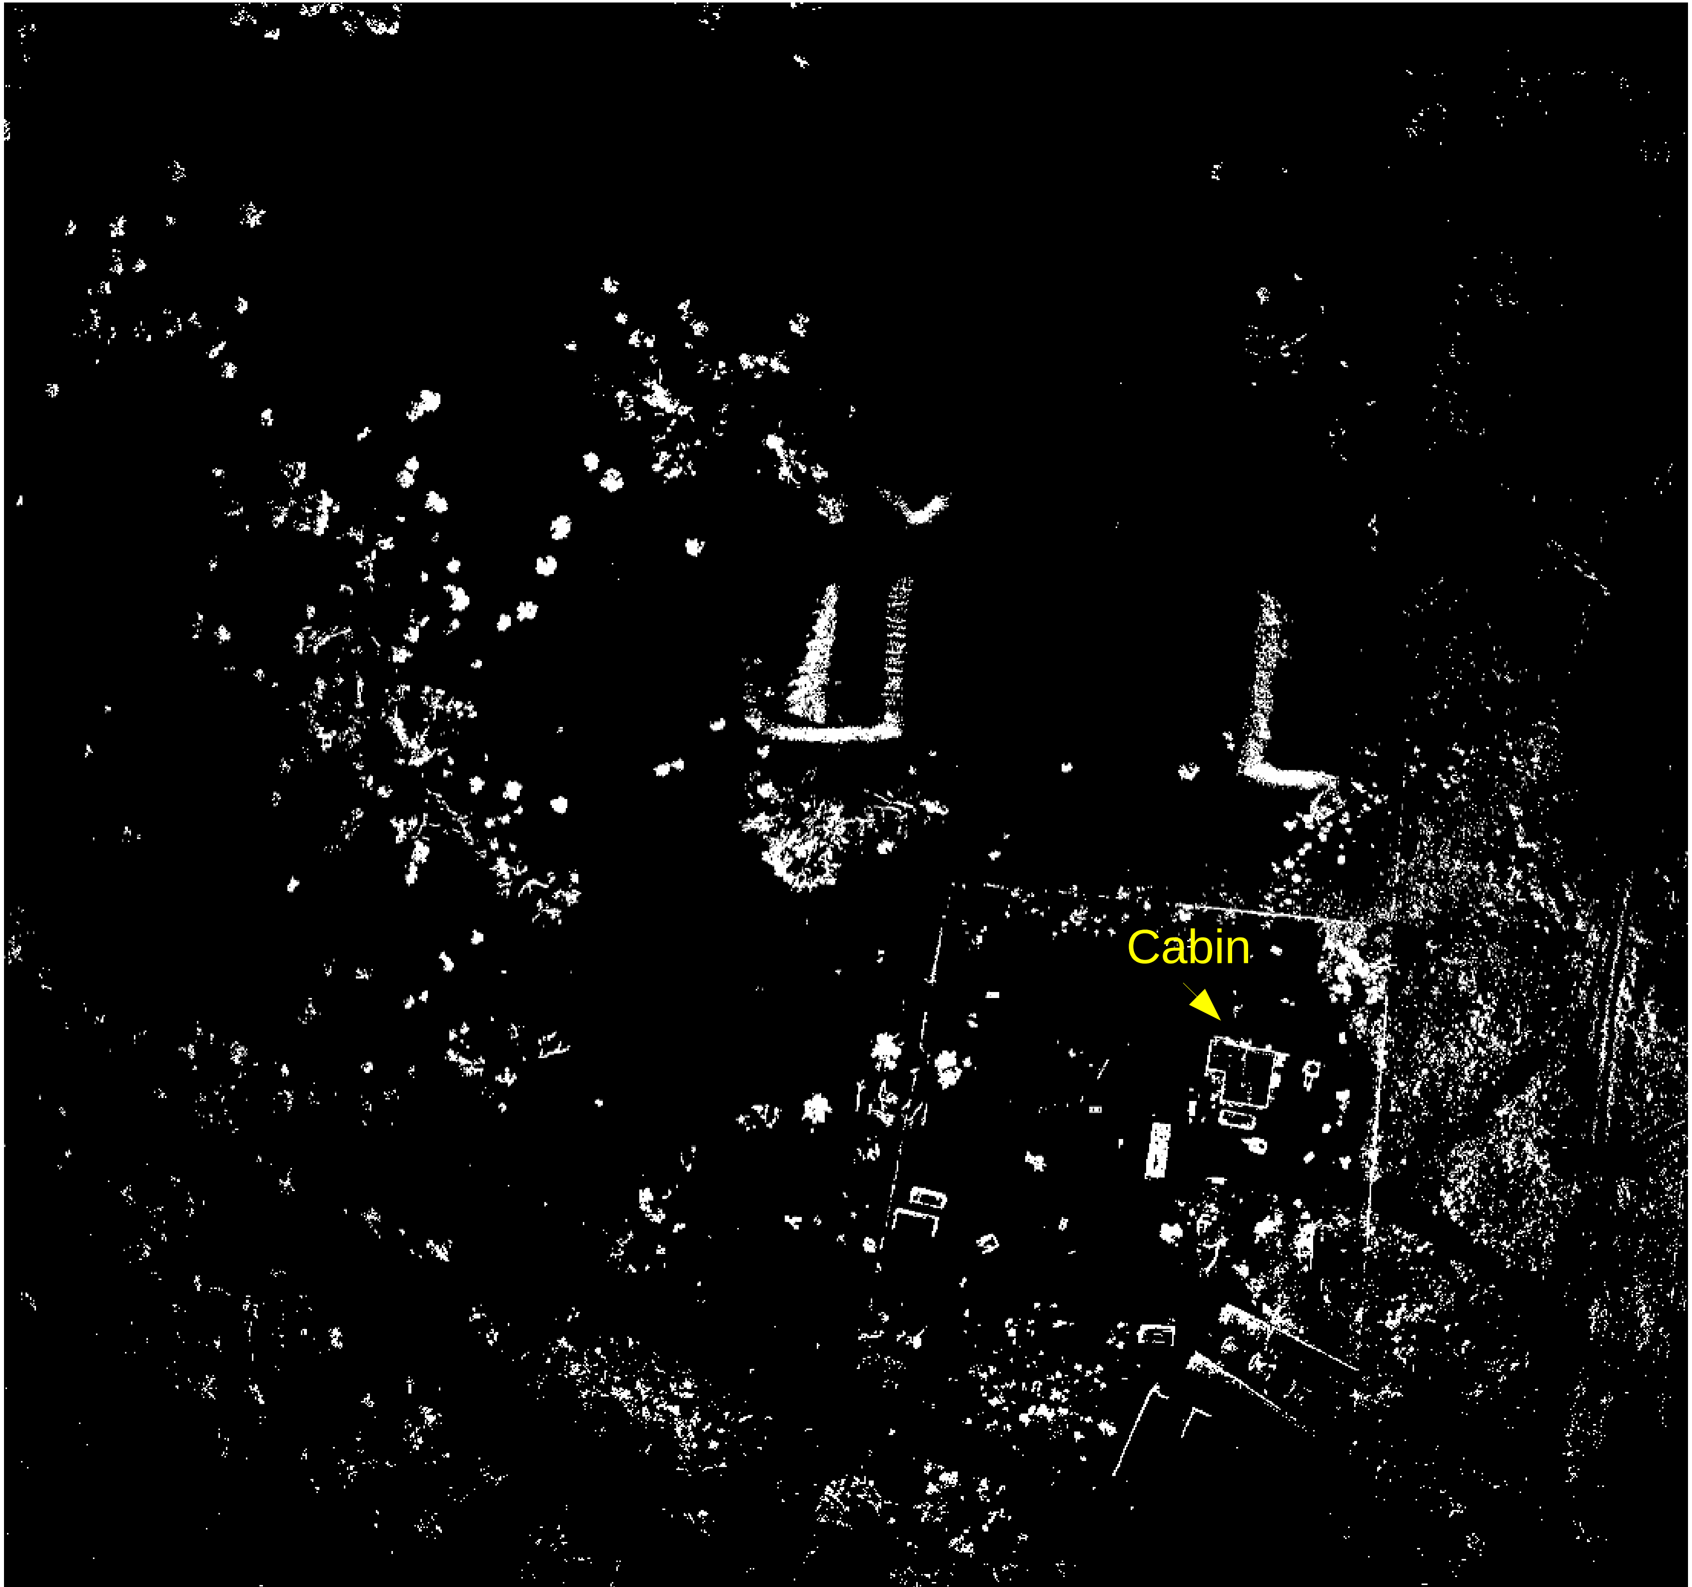
Cabin

Laser signature from the 2D feature map

- Take range readings every 2 degrees (ray tracing)
- Discard reading if it exceeds 50m
- Signature is an array of 180 range readings

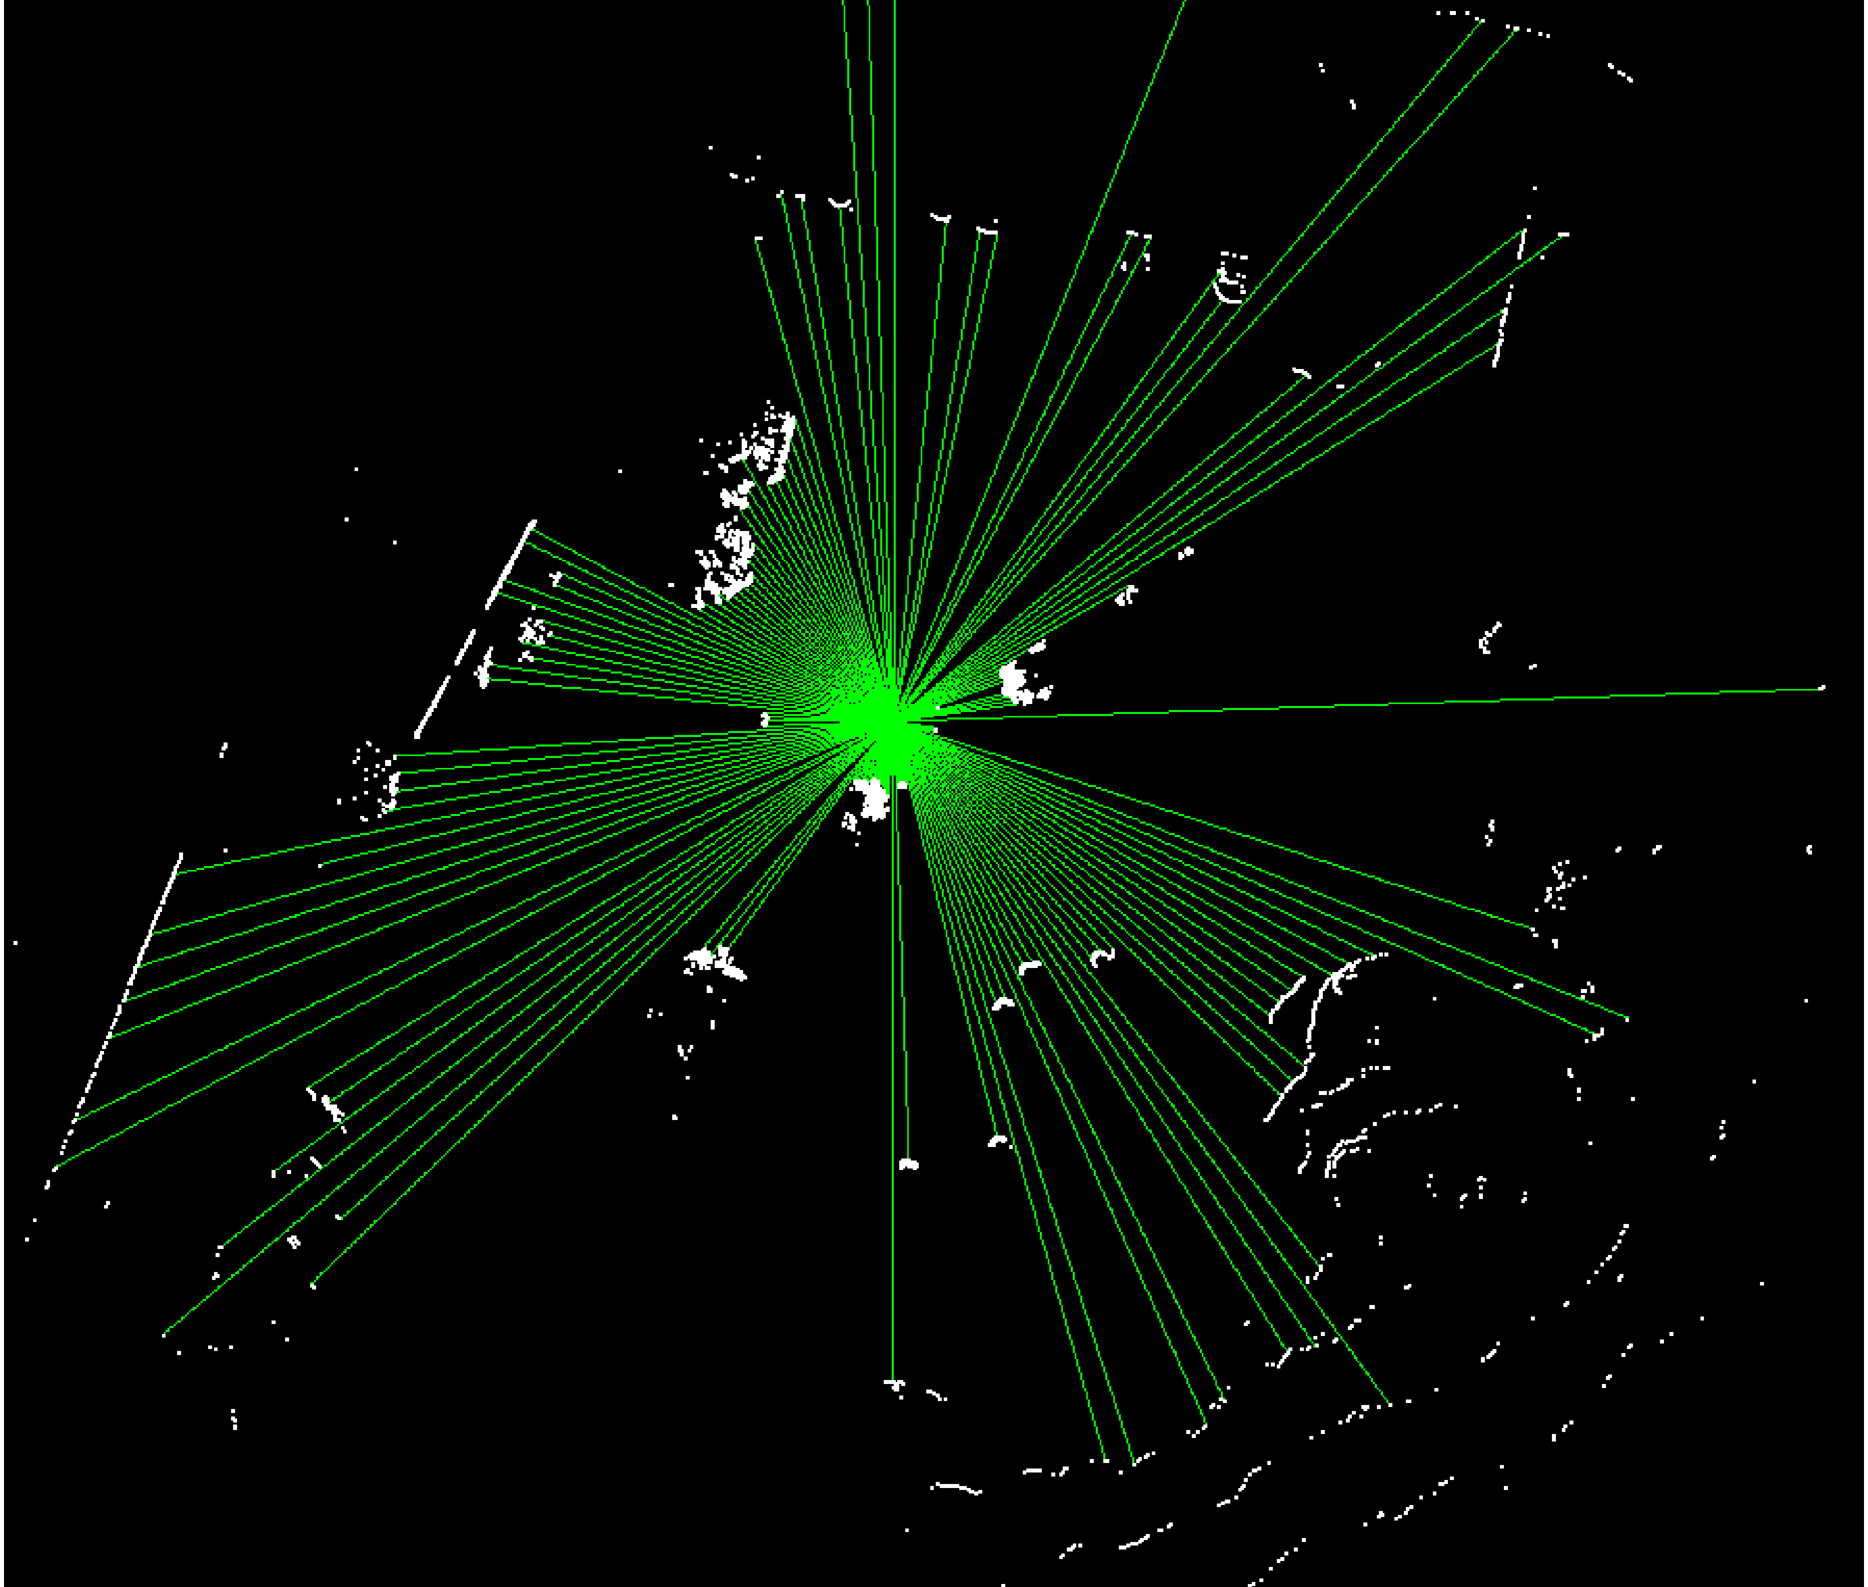

Scan matching

- A laser signature from the Velodyne is created in the same way as the Riegl (project 3D data to 2D image, grow features ...)
- Sum of absolute difference is used to score a match between two laser signatures
- To find the orientation, the signatures have to be scored for all possible rotation ($O(N^2)$, rather slow)

Localisation framework

- Non-probabilistic
- Global localisation is done by taking the best match from the cached signatures (one shot localisation, performed only once at the start)
- The vehicle is then tracked by taking the last estimated pose and performing a local search around it
- The user defines the radius and angle of the search area

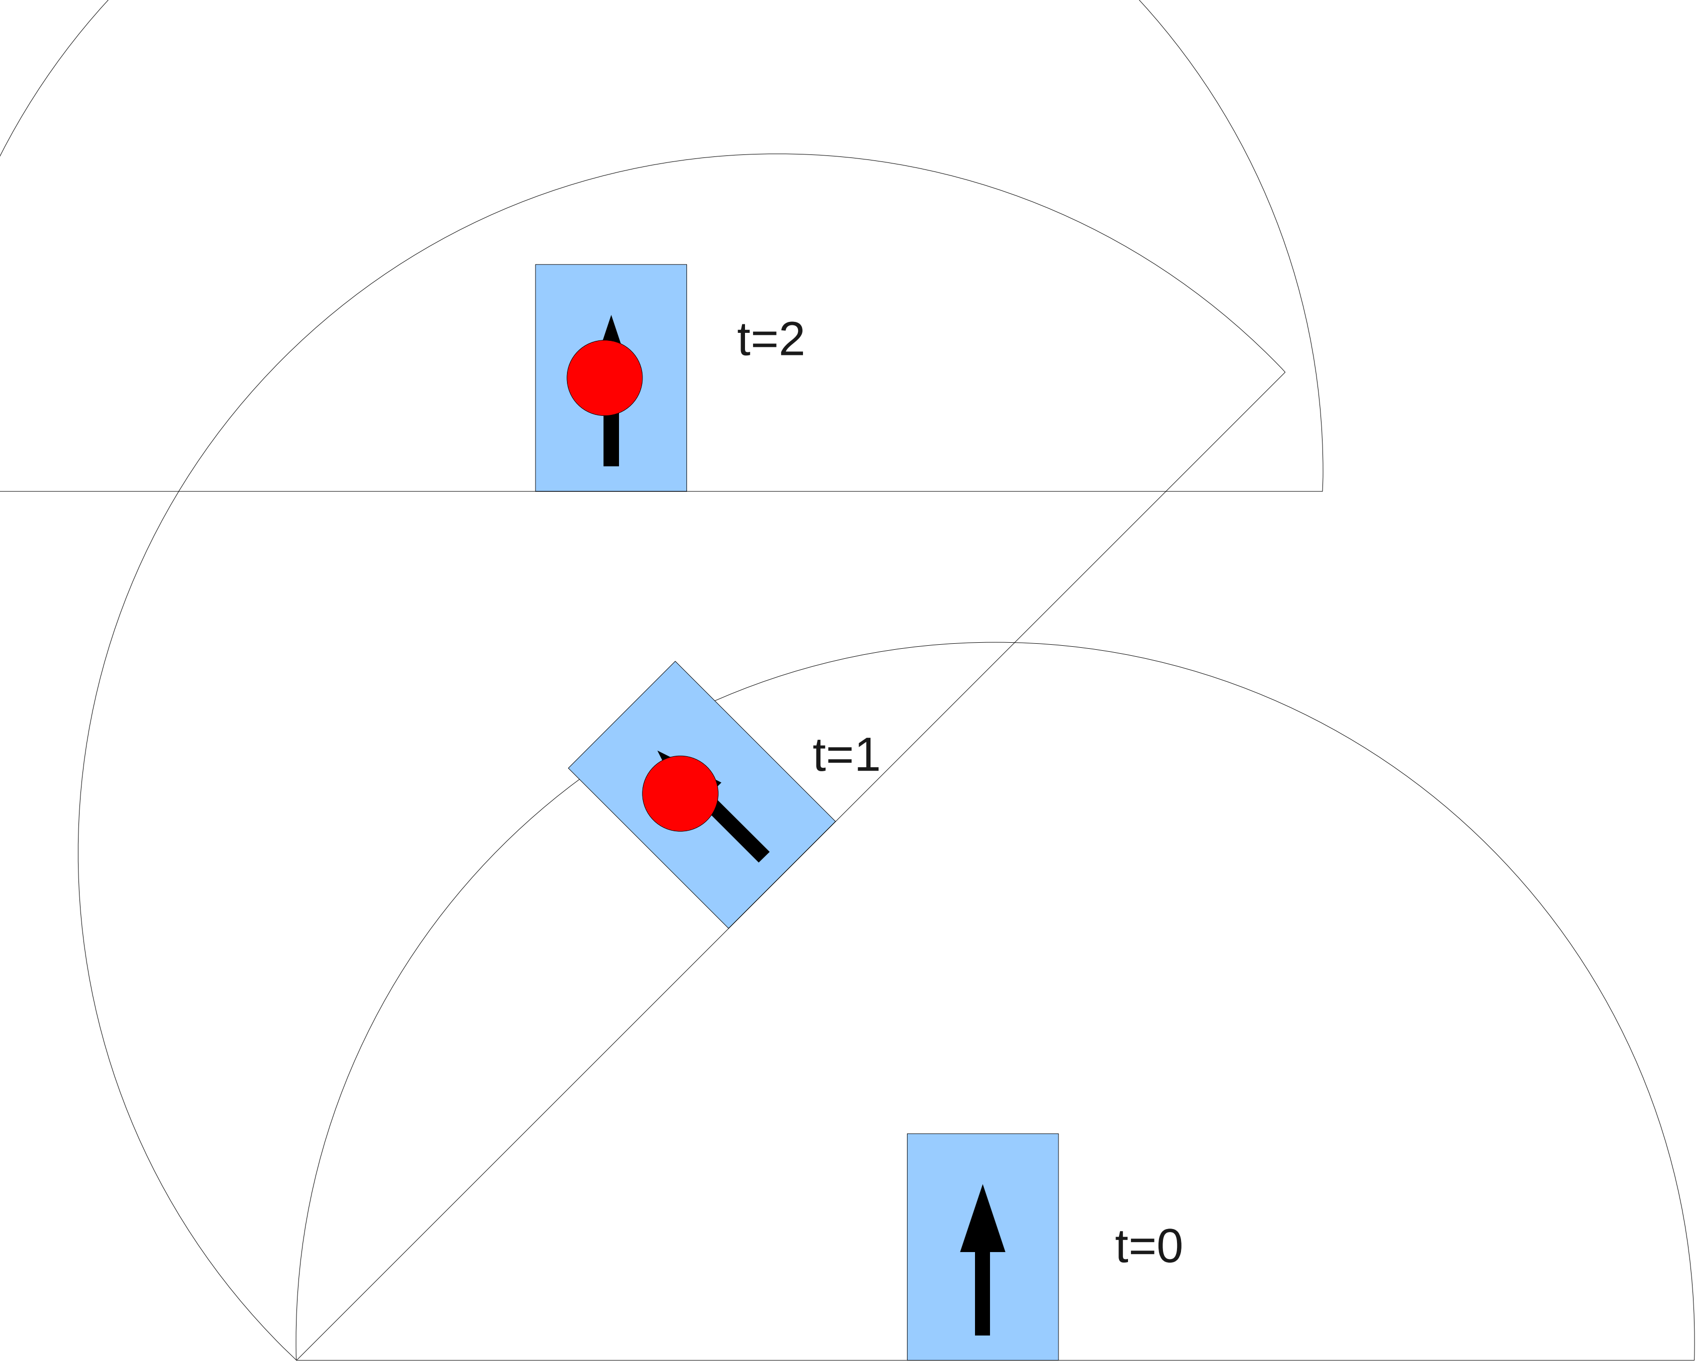
t=2
t=1
t=0

Localisation issues

- Problem: The vehicle sometimes ventures into areas not adequately scanned by the Riegl
- Solution: The global map is augmented with the previous Velodyne data if the best scan matching score is below a threshold. The global map is augmented temporarily for the current reading only.
- Ray tracing needs to be performed (can't use the cached laser signatures), this is slow

Implementation

- Using an Intel i7 (4 cores) with 6GB of RAM
- Scan matching process is parallelised (default to 8 threads)
- (show demo)

Results

- Localisation errors about 14 cm and 1 degree
- Manually registered Velodyne scans (registered to the Riegl 3D model) were used as ground truth
- Runs at a usable speed (could be much faster if odometry information was available)

Conclusions

- Accurate laser localisation was achieved using very simple algorithms
- Can still localise without odometry if we make certain assumptions, though at the cost of an increase in search space
- Non-probabilistic localisation does work, but was used to demonstrate a proof of concept only
- Exploiting multi-core is an easy way to gain near linear speed up of parallelisable tasks

Some random pics from Pomonal ...


EXIDE
EXIDE
ENDURANCE


SPECIFIED
ROOFING
Telephone
9311 2004
KEEP CLEAR WHEN SLEWING
411591
US431