# Scaling the Edge

Jon Crowcroft

# A tour through some recent innovations

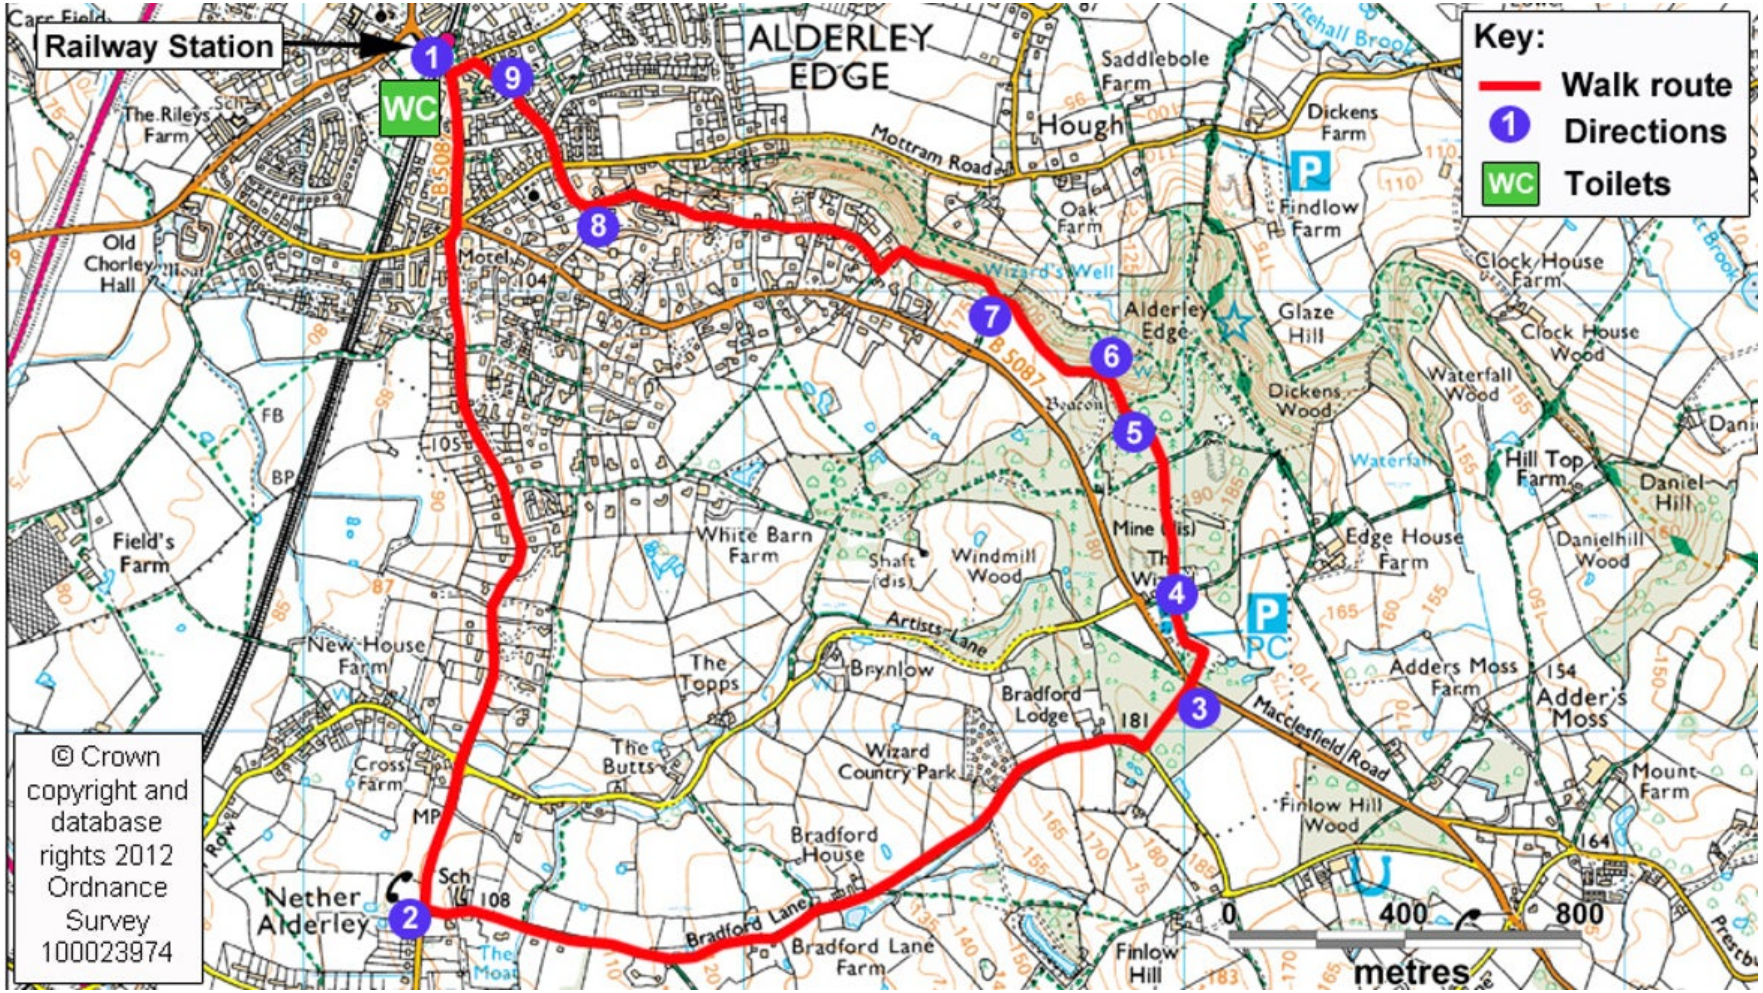

# Extreme Federation…

- Edge processing data in networked systems becoming mainstream:
  - It reduces load on the uplinks,
  - it saves energy &
  - potentially provides better privacy for personal data.

  - and possibly higher availability (on aggregate)

  - but must cope with asymmetric access link speeds
  - and highly heterogeneous individual node availability

  - would like to retain some aspects of business models
  - i.e. analytics/ML/AI – so…

# A variety of edge techniques

- simple aggregation,
- compressive sensing, &
- edge-machine learning
  - where models are locally acquired, and
  - model parameters are distributed,
  - so nodes can further refine their models.

# Challenges #1

- Firstly to scale federated learning to billions of nodes needs some way to scale
- even just sharing model parameters e.g. https://arxiv.org/abs/1907.08059
- including sampling of model parameters
  - thinning,probabilistic update &
  - self organising hierarchies of aggregation (model parameter servers). e.g. https://arxiv.org/abs/1709.07772
  - For some Machine Learning algorithms, there may be updates from the federated model back to nodes
  - to adjust their learning (e.g. regret) as well.
  - indeed, what even is initial placement system?
  - it sure isn’t kubernetes
  - Could be

https://www2.eecs.berkeley.edu/Pubs/TechRpts/2018/EECS-2018-119.pdf

# Challenges #2

- Some schemes may require synchronisation of learning steps.
- All these need to scale out, &
- techniques from data centers may, surprisingly be applicable, even though
  - we are often in a much less rich networking environment,
  - even without full connectivity or symmetric bandwidth or reachability.

# Challenges #3

- Federation alone is not a complete solution to privacy, &
- some further techniques may be needed to reduce the loss of confidentiality –
  - e.g. differential privacy is useful, but also
  - more fundamental approaches such as secure multi-party computation, in extreme cases.

# Challenges #4

- Secondly, there is the problem of bad actors injecting false data
  - pollution,
- Then there is the omnipresent presence of possible DDoS attacks.
- Scale federated trust?
- Ownership of derived models?
- Decentralised trust (transparency) – what tools?
- Hybrids – like
  - GAEN, https://developers.google.com/android/exposure-notifications/exposure-notifications-api
  - original Skype Supernode architecture,
  - Chainspace shards https://arxiv.org/abs/1708.03778

# Challenges #5

- Thirdly, a federated model may present some challenges to model explain-ability or interpret-ability.
- What if we have ensembles of (many) different models?
- Interesting trade-offs between these requirements  & privacy.
- e.g. Identity and Personhood in Digital Democracy:
  - https://arxiv.org/abs/2011.02412

# Recent Examples of Threat & Opportunity

- Google outage - single point of failure for all services
- versus
- Google Apple Exposure Notification - decentralised (hybrid)

# Conclusions

10 thousand data centres with a million cores

10 billion edges – add some structure?