Unconstrained indoor localisation on a smartphone

Agata Brajdic, Robert Harle

The Challenge:
How to reliably determine user's location in indoor areas with a commodity smartphone and no environmental infrastructure?

The Good:

- Users now carry around various sensors embedded in their smartphones
- User's location can be estimated using only accelerometer and gyroscope ➔ pedestrian dead reckoning (PDR)
- Particle filters (PFs) with a building map can reduce accumulation error and allow use of less precise sensors (when applied to a foot-mounted PDR, PFs achieved accuracy below 0.78m 95% of the time [1])

How do Particle Filters work?

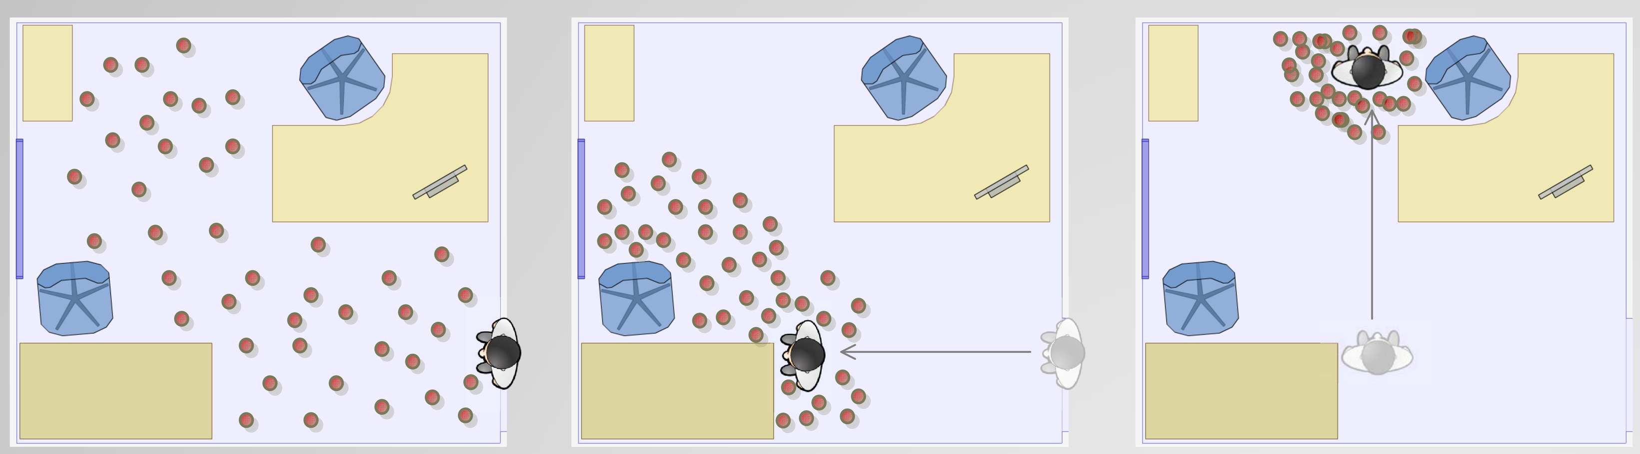

The Bad:

1. Greater number of particles are required in practice, incurring high computational cost:
   - smartphone sensors are often of lower grade so more particles are needed to reduce the drift
   - larger buildings impose more particles - e.g. 10^6 particles needed for reliable localisation in William Gates Building (8725 m^2)
2. Extracting user's steps from noisy sensor data is difficult
   - existing solutions are tailored for sensors attached firmly to user's body, but smartphones are often held loose

We extract steps solely from the accelerometer since it consumes least amount of energy among phone sensors and is present even in earlier generations of smartphones.

We may need to extract steps while the phone is held in hand, backpack, handbag, trousers back pocket, front pocket, or even while it is typed on.

Evaluated step detection algorithms:
- Window Peak Detection (WPD)
- Mean Cross Counting (MCC)
- Normalised Autocorrelation (NASC)
- Hidden Markov Model (HMM)
- Short Term Fourier Transform (STFT)
- Continuous Wavelet Transform (CWT)
- Discrete Wavelet Transform (DWT)

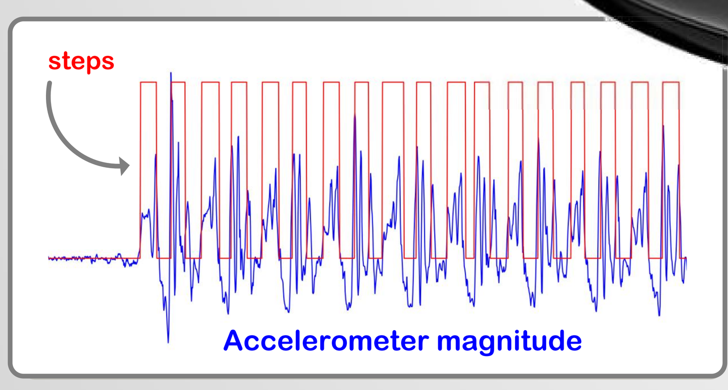

- nonrigid attachment violates many of the assumptions in previous work and makes existing algorithms prone to significant errors
- understanding behaviour of algorithms in different scenarios is essential for successful localisation

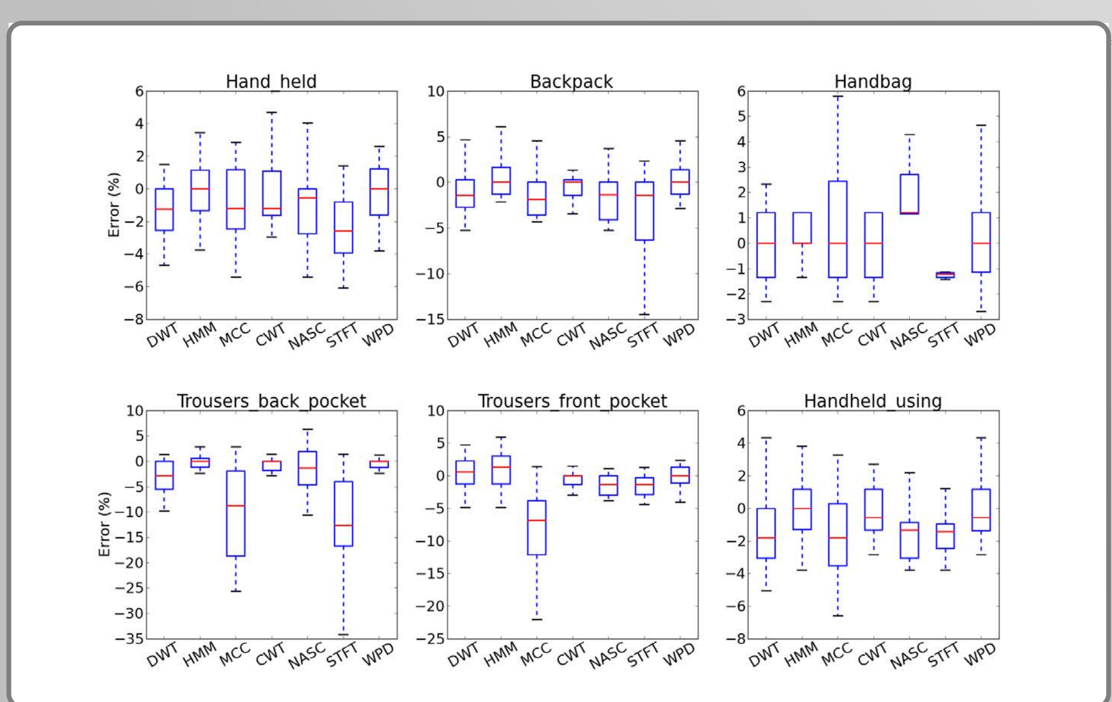

The key idea: combine modern smartphone sensing and processing capabilities with dynamic offloading and loading of particles to and from a remote node.

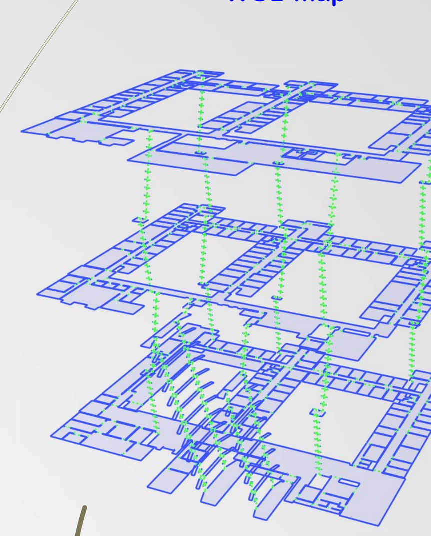

Requirements imposed on the remote computational unit:

- Must be able to sustain processing of required number of particles in real time.
- Must be able to simultaneously serve multiple users
- Must be able to reduce the initially large number of particles and offload them to the mobile device (KLD method)

Chosen architecture:
Graphics Processing Unit (GPU)

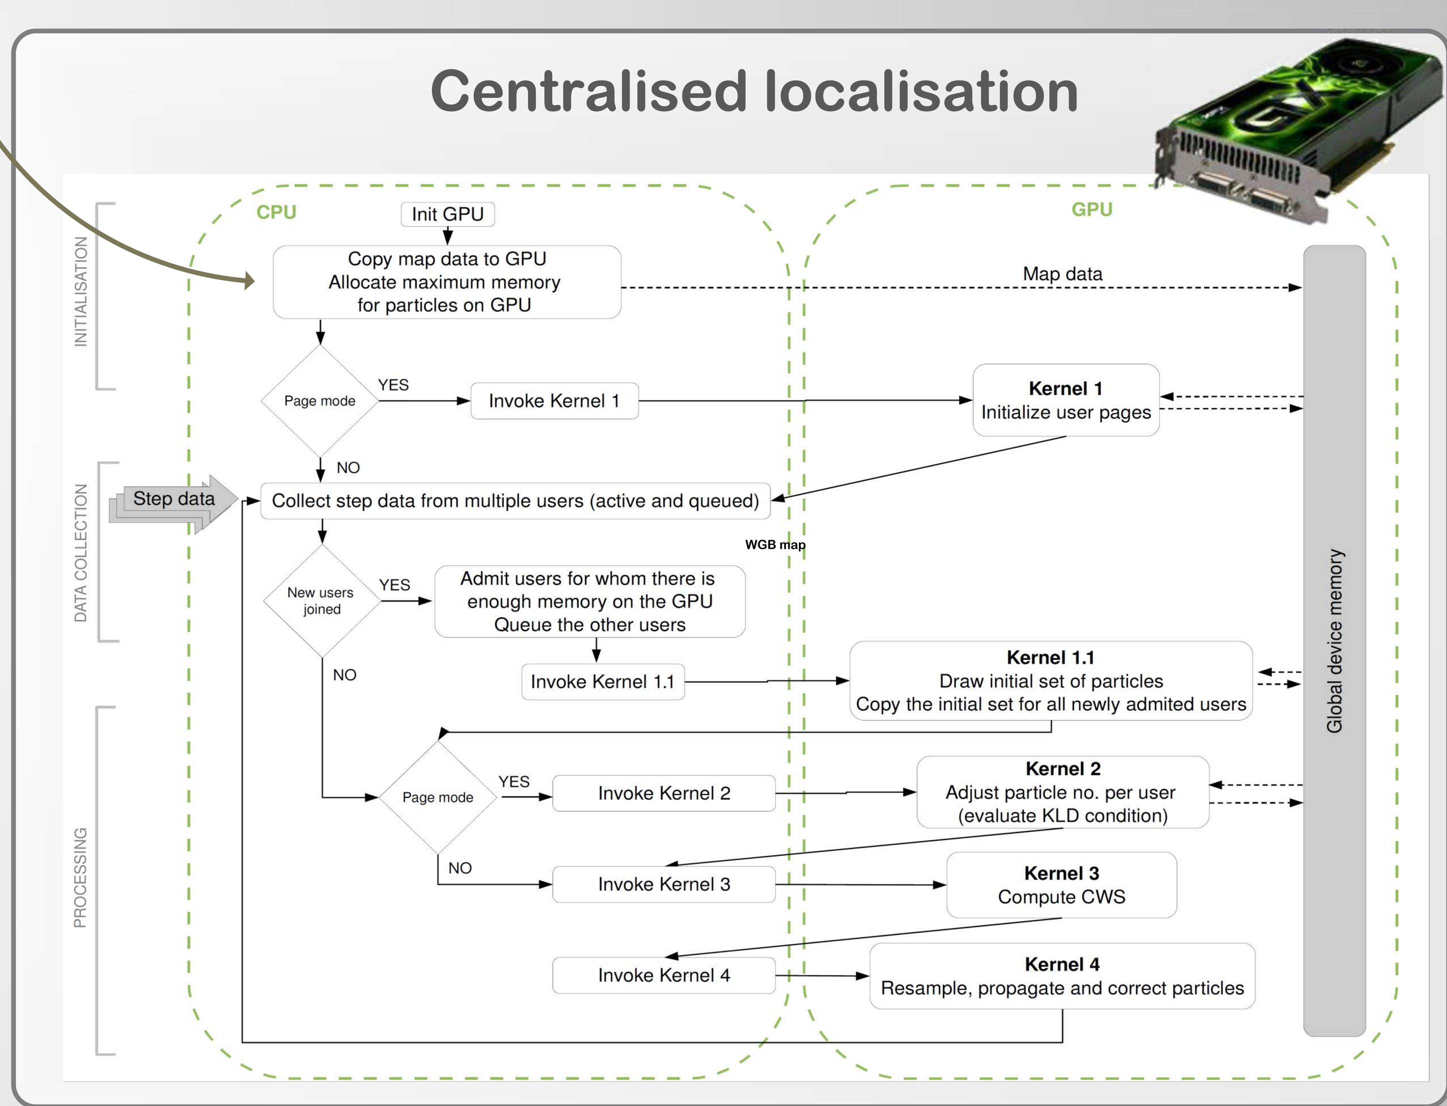

On the mobile device, the system alternates between a state in which the particle update is performed locally and a state in which the step data is streamed while the localisation is performed remotely.

Why should we use GPU for particle filtering?

GPUs are a massively parallel architecture which applies well to the structure of particle filter. Many thousands of threads then perform the same computation on different data in parallel.

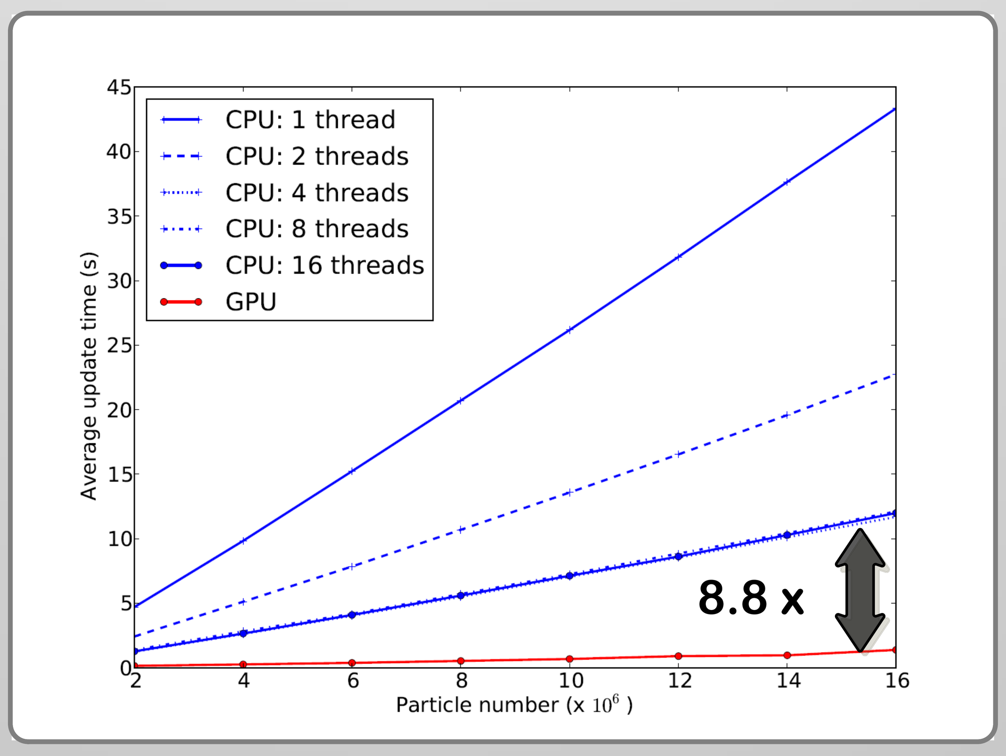

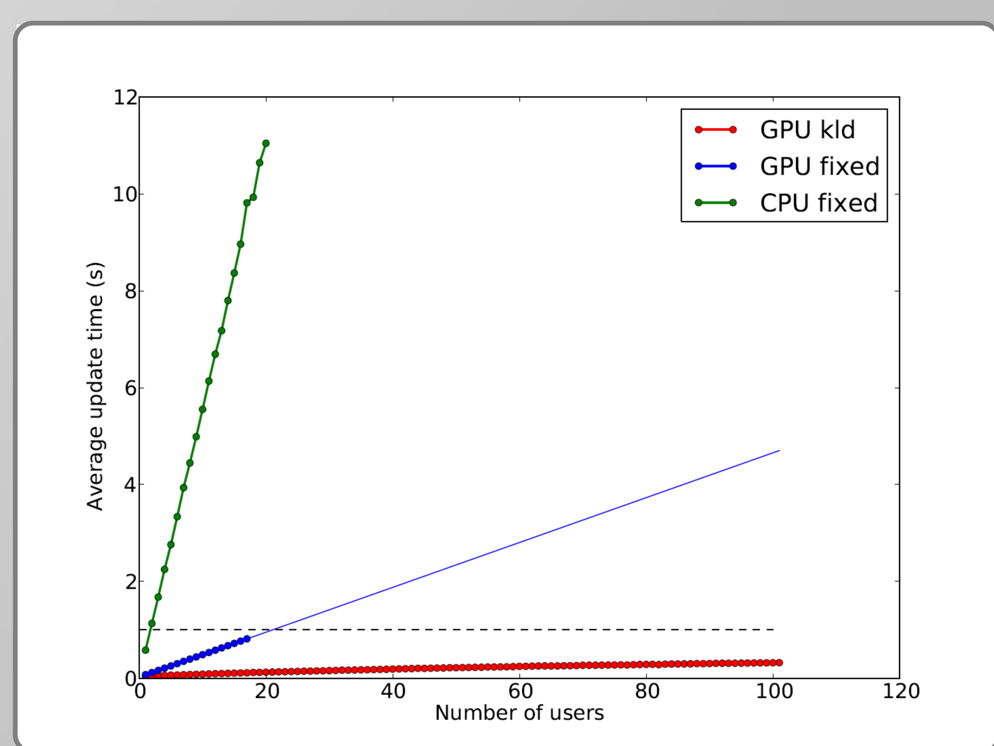

References:

[1] Oliver Woodman and Robert Harle. Pedestrian localisation for indoor environments. In UbiComp 2008, volume 344, pages 114–123. ACM, 2008.
[2] Dieter Fox. Adapting the sample size in particle filters through kld-sampling. I. J. Robotic Res., 22(12):985–1004, 2003.
[3] N. Gordon, D. Salmond, and A. F. Smith. Novel approach to nonlinear/non-gaussian bayesian state estimation. IEE Proc.-F, Radar Signal Processing, 140:107–113, 1993.
[4] M. Addlesee, R. Curwen, S. Hodges, J. Newman, P. Steggles, A. Ward, and A. Hopper. Implementing a sentient computing system. Computer, 34(8):50–56, August 2001.