Imperial College London

Semi-dense SLAM on an FPGA-SoC

Konstantinos Boikos, Christos-Savvas Bouganis
k.boikos14@imperial.ac.uk

FPL 2016, Lausanne
Sept 2016

HiPEDS

Imperial College London

What is SLAM?

- Simultaneous Localisation and Mapping
  Use a series of observations to:
  - Reconstruct an environment (map)
  - Track the observer's trajectory in it

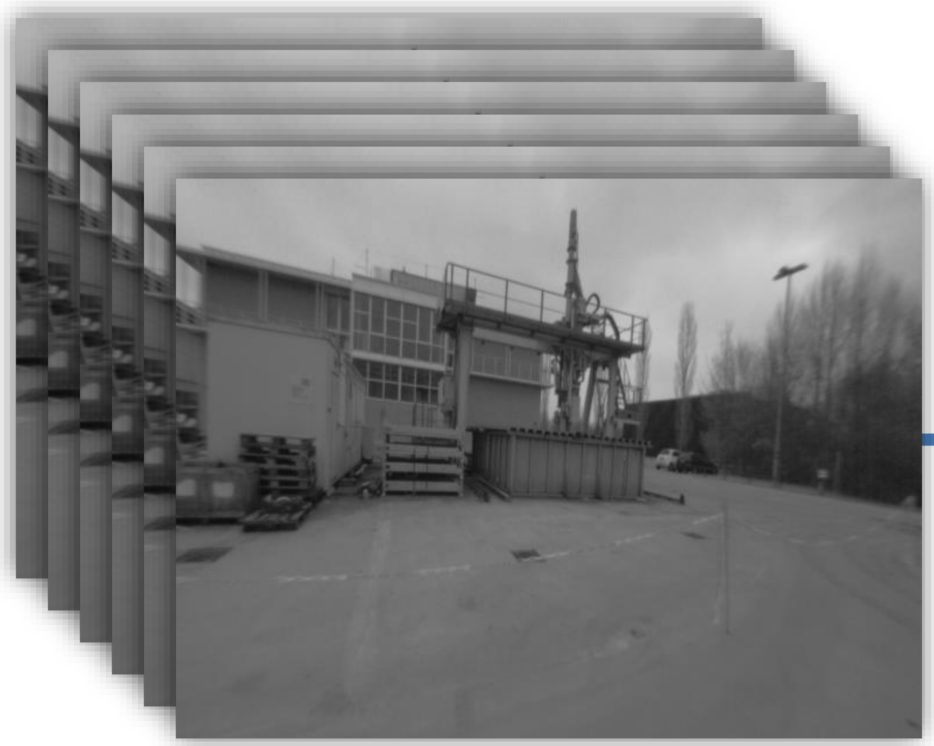

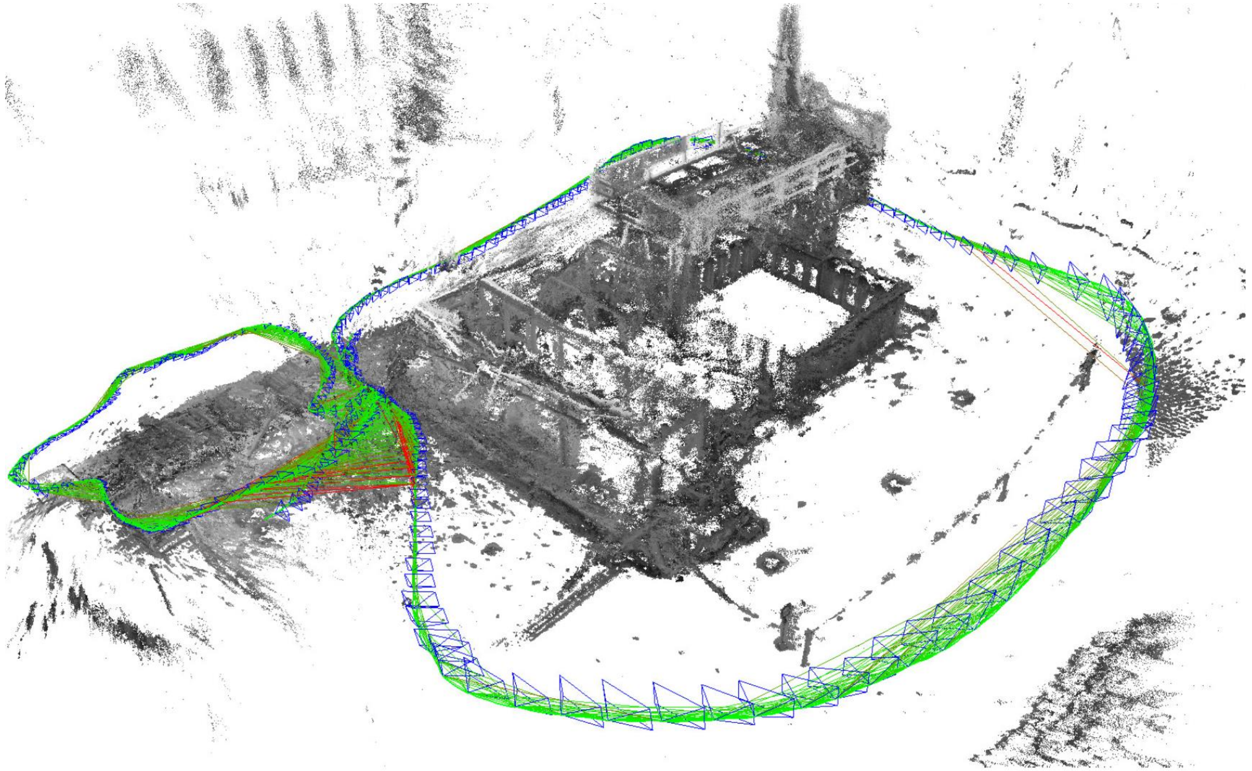

Why is it Important?

Advanced SLAM algorithms are fundamental in autonomous robotics and emerging applications

➢ Self-driving Cars

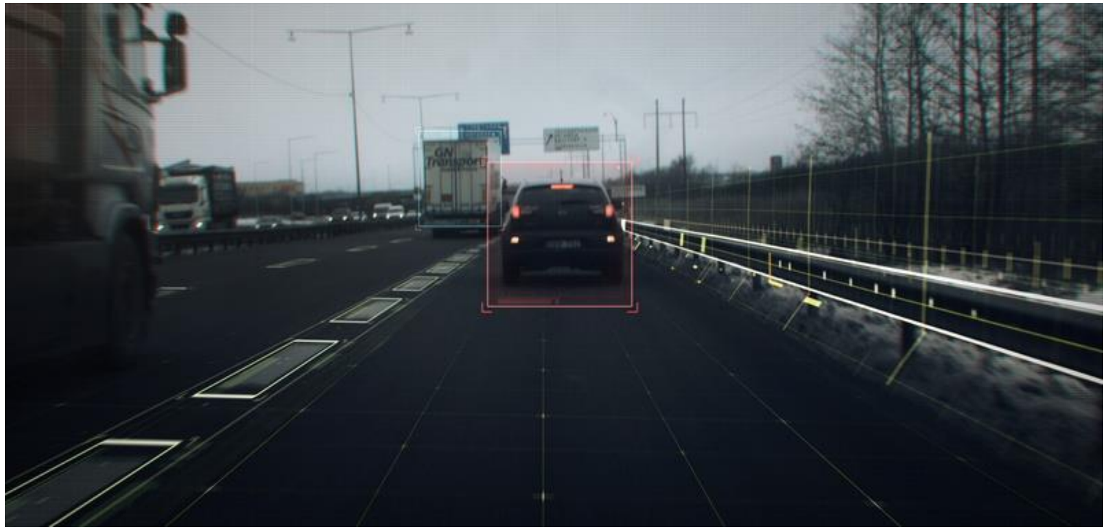

Why is it Important?

Advanced SLAM algorithms are fundamental in autonomous robotics and emerging applications

- Self-driving Cars
- Precision Agriculture

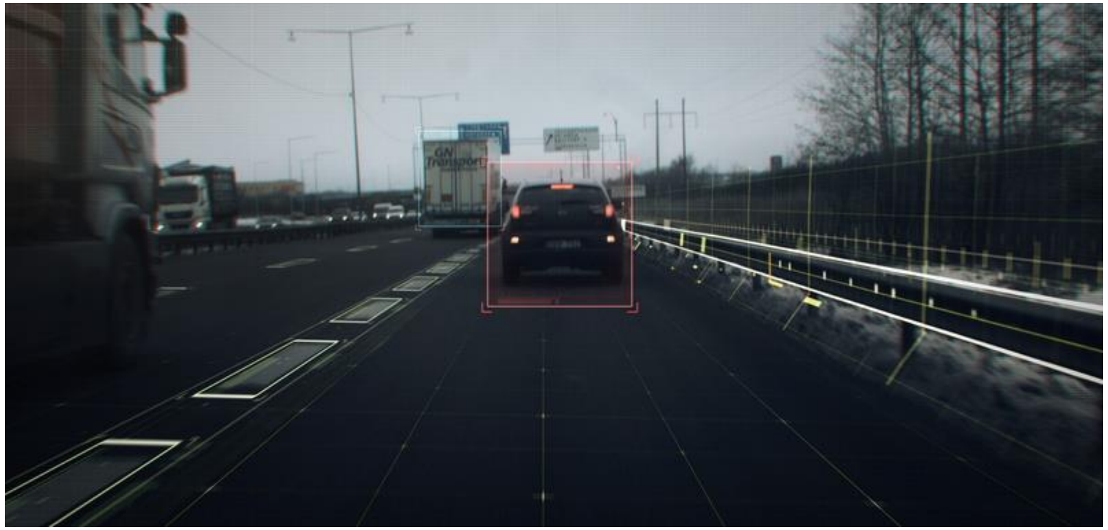

Why is it Important?

Advanced SLAM algorithms are fundamental in autonomous robotics and emerging applications

- ➢ Self-driving Cars
- ➢ Precision Agriculture
- ➢ Rapid environment exploration
- ➢ Much more effective search and rescue operations

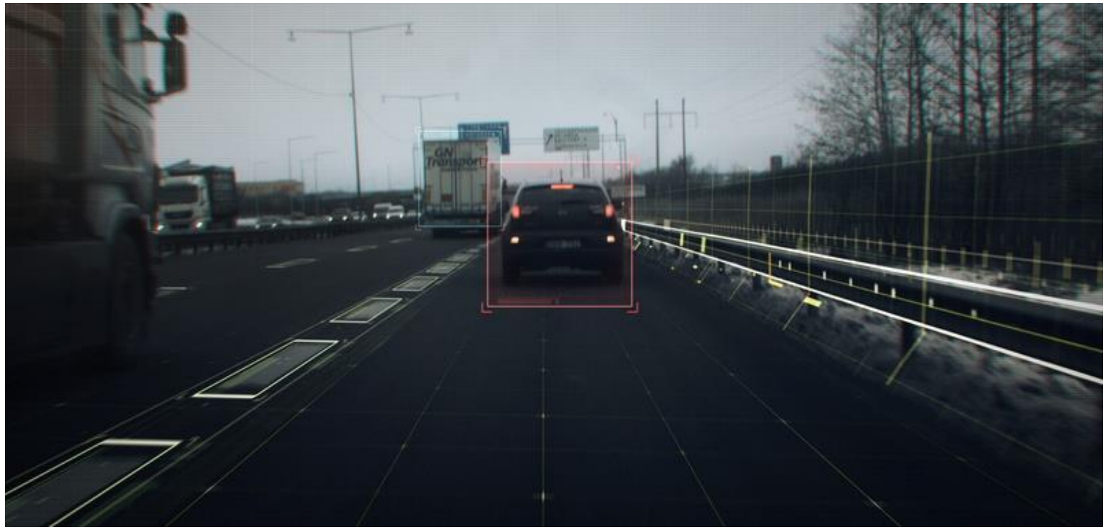

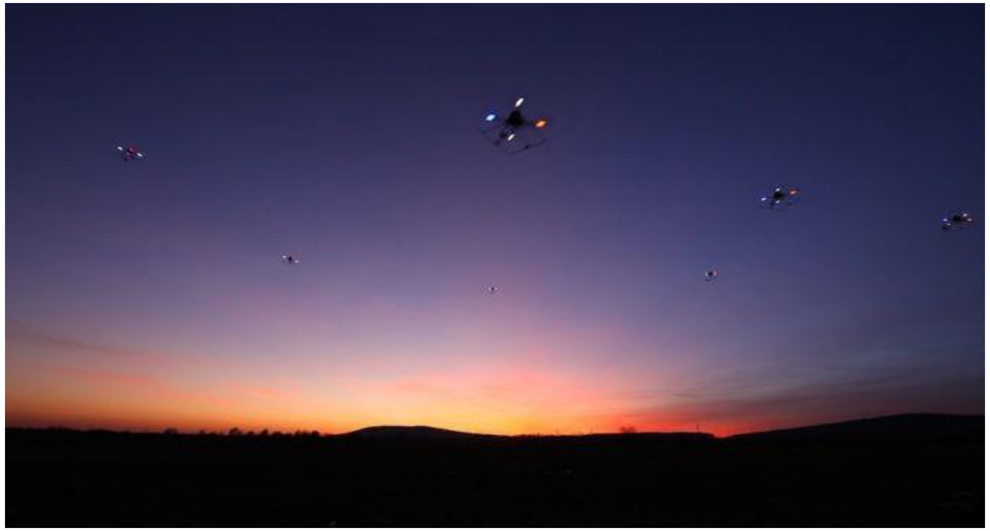

Challenges in Embedded SLAM

- High complexity and bandwidth requirements
- Low latency and high framerate crucial to keep track of a fast moving robot
- Emerging applications require an unprecedented richness and accuracy

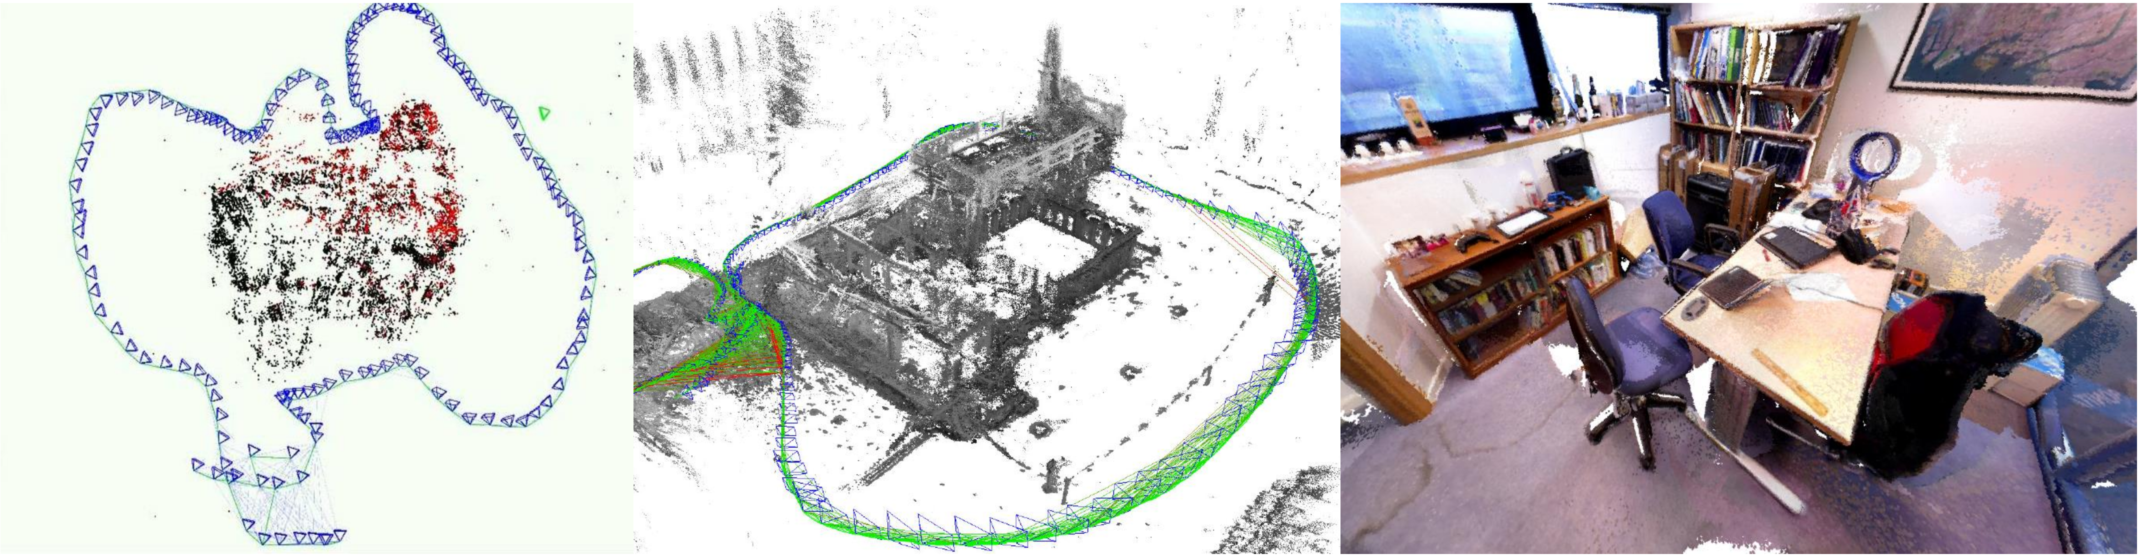

Sparse → Semi-Dense → Dense

Proposed Solution

- Based on LSD-SLAM, a state of the art semi-dense SLAM algorithm
- Design custom hardware to offload tracking (pose estimation) with a significantly higher performance-per-watt
- Dual-core ARM processor handles mapping and complex control flow

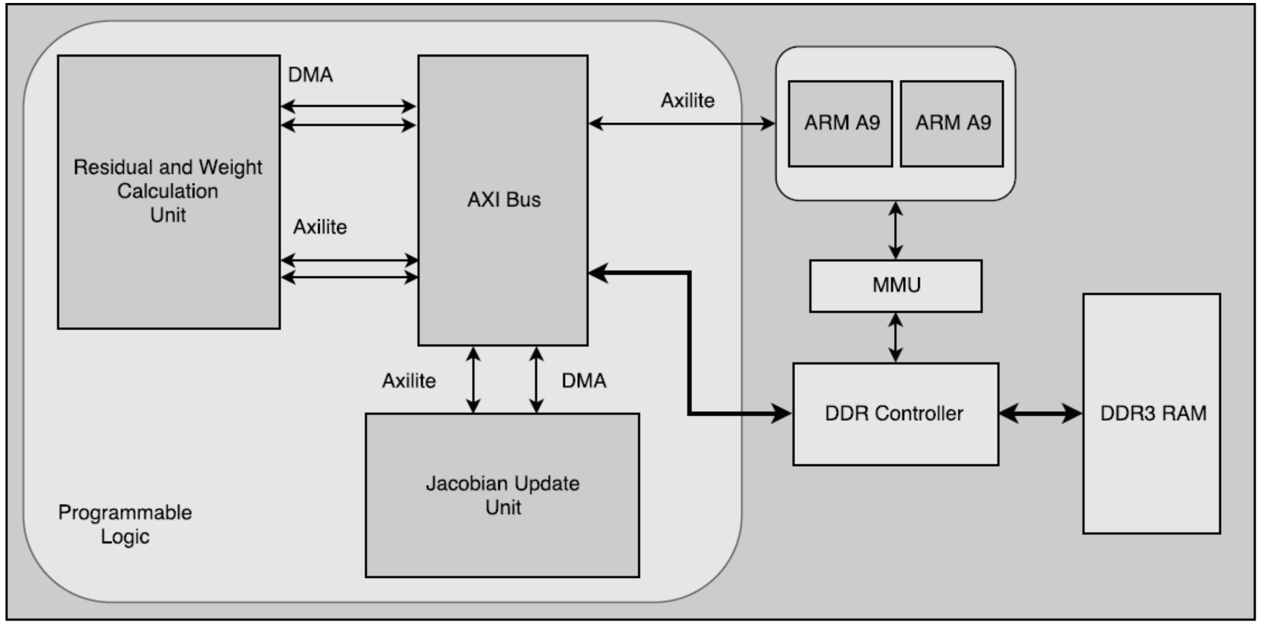

Proposed Solution – Pt. 2

- Prefetching the current frame in a local cache
- Interleaved all other input data in the DDR memory for optimised read access
- Reads and writes in batches to take advantage of AXI4 burst transfers
- Accelerator units and ARM cores operate in a common memory space

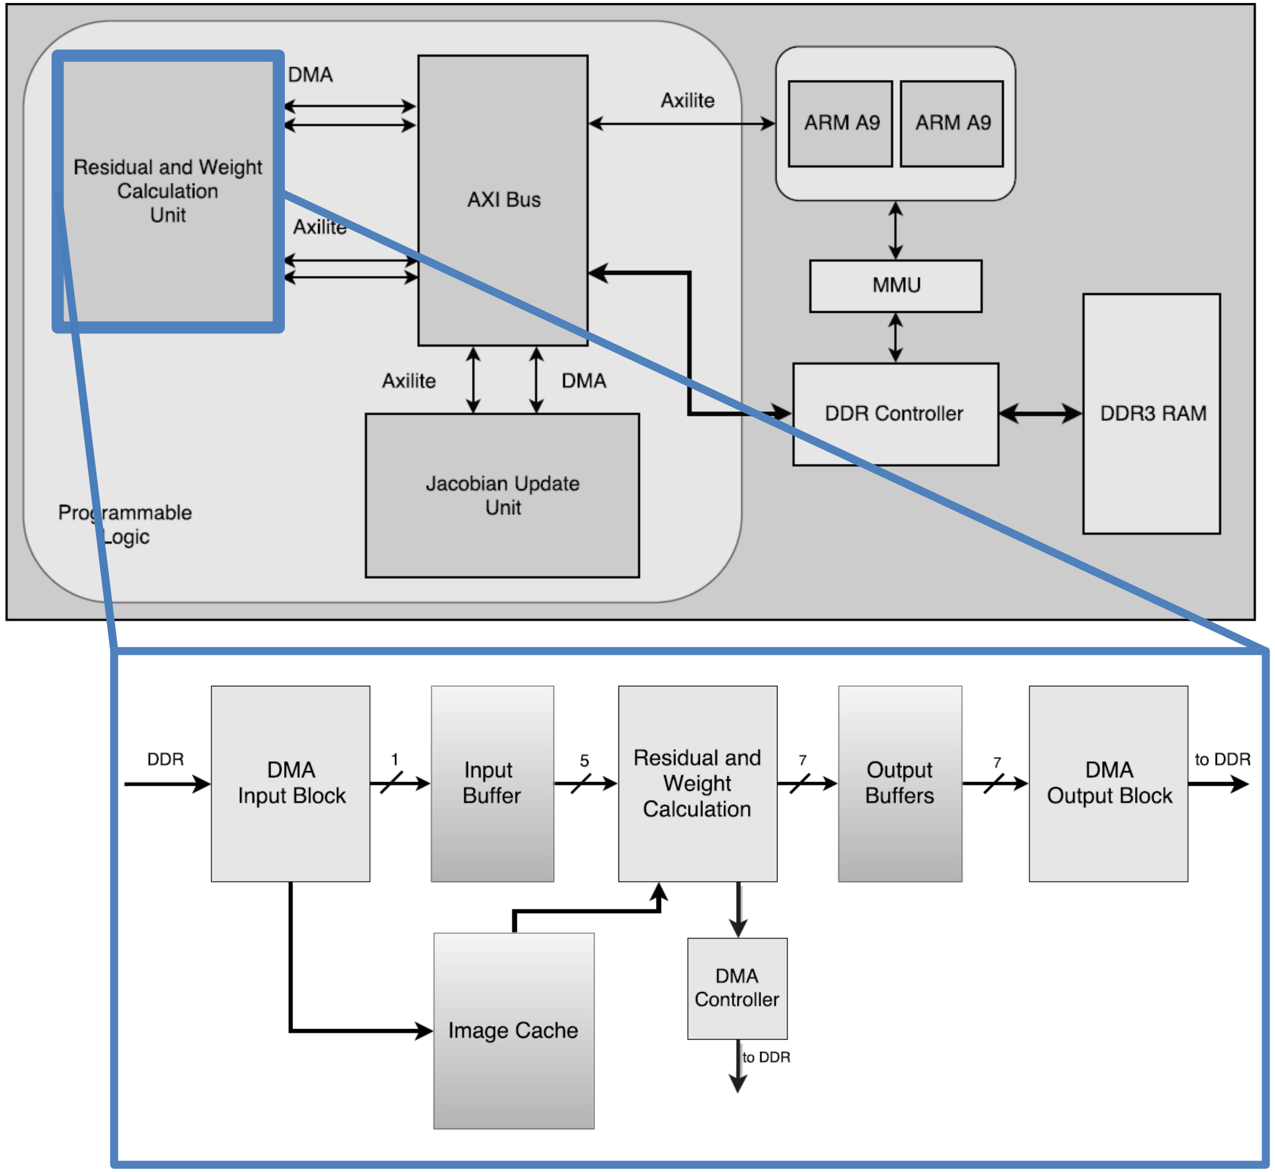

Imperial College London

Evaluation

- Full SLAM system running on a Zynq-7020
- Overall, achieved a 2X speedup and more than 4.3X the energy efficiency
- Speedup is consistent for simpler (blue) to more complicated scenes with a larger number of textures (gray)
- Graph highlights sensitivity to the amount of information we want to recover

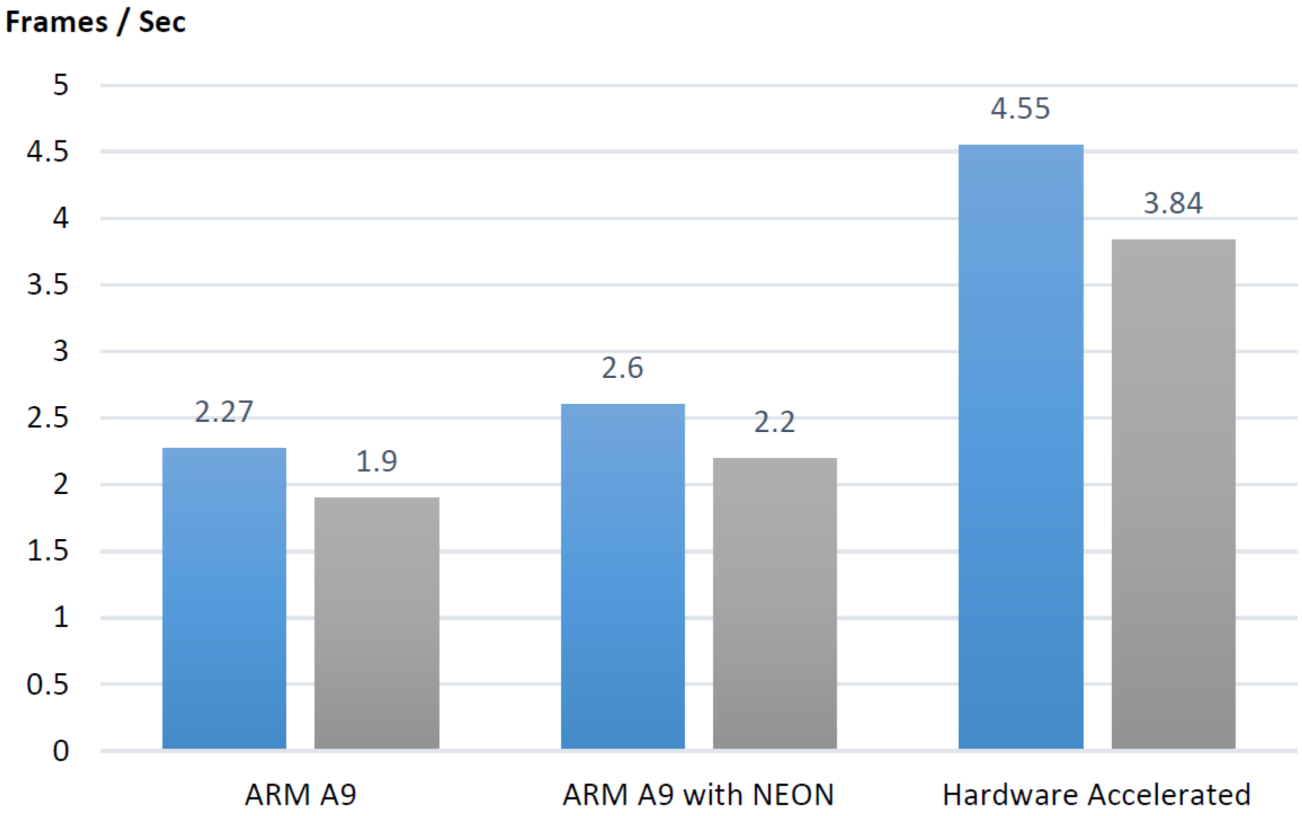

Poster Session

- If you are interested to learn more and discuss my work, look for my poster in the afternoon session that starts at **3.15pm**.

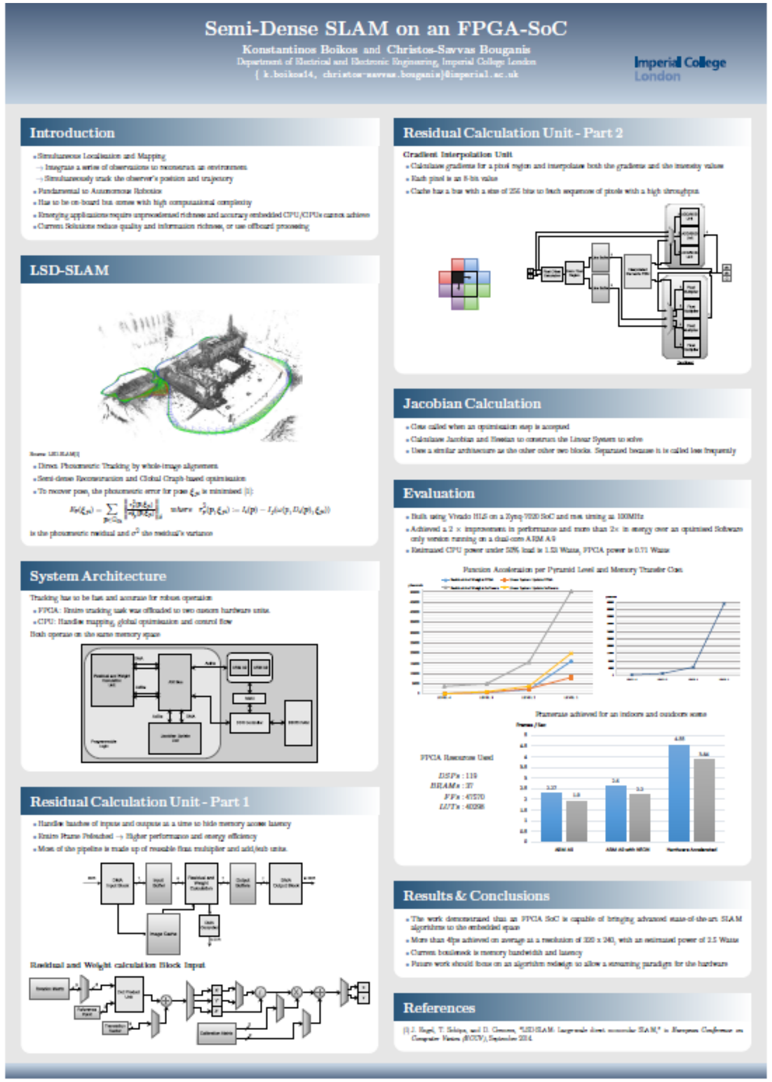

Thank you!