Energy-Efficient Application Specific Instruction-Set Processor for Feature Extraction in Smart Vision Systems

Lucas Ferreira[1], Steffen Malkowsky[1], Patrik Persson[2], Sven Karlsson[3], Kalle Åström[2], Liang Liu[1]

[1]Department of Electrical and Information Techonology (EIT), Lund University, Sweden
[2]Center for Mathematical Sciences, Lund University, Sweden
[3]High Performance Computing Department, DTU, Denmark

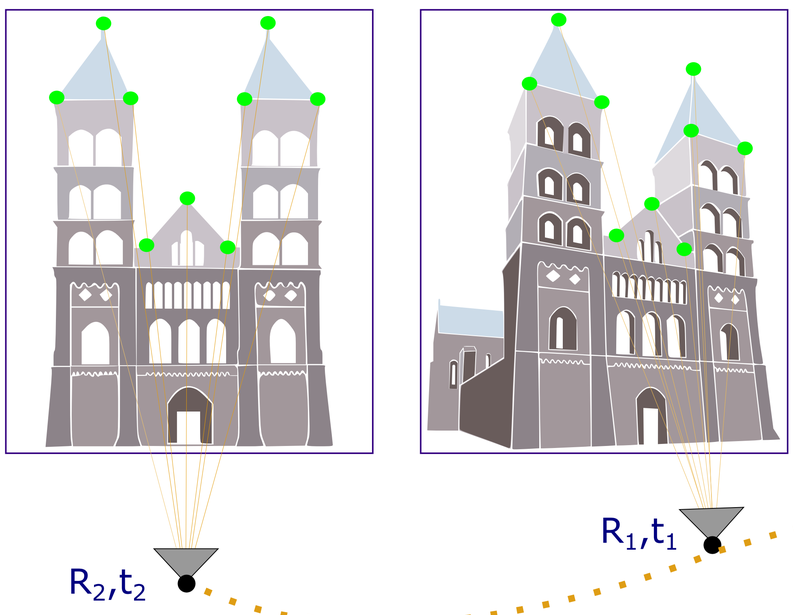

2 Vision-based SLAM: Given a set of point-correspondences in a image-pair, the relative camera pose, and 3D map points can be retrieved.

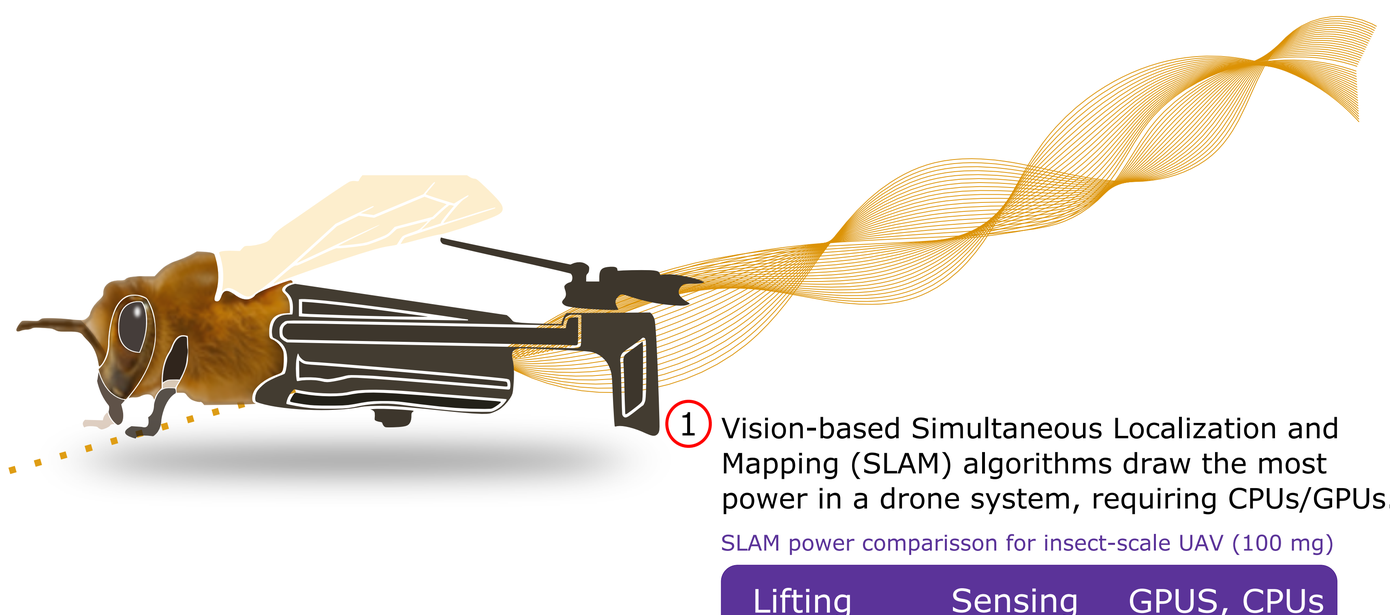

1 Vision-based Simultaneous Localization and Mapping (SLAM) algorithms draw the most power in a drone system, requiring CPUs/GPUs.

SLAM power comparisson for insect-scale UAV (100 mg)

Lifting	Sensing	GPUS, CPUs
100 mW	100 mW	100 W

Within SLAM, the feature extraction process, which identifies image corners in order to track them, is one of the bottlenecks given its complexity and real-time requirements.

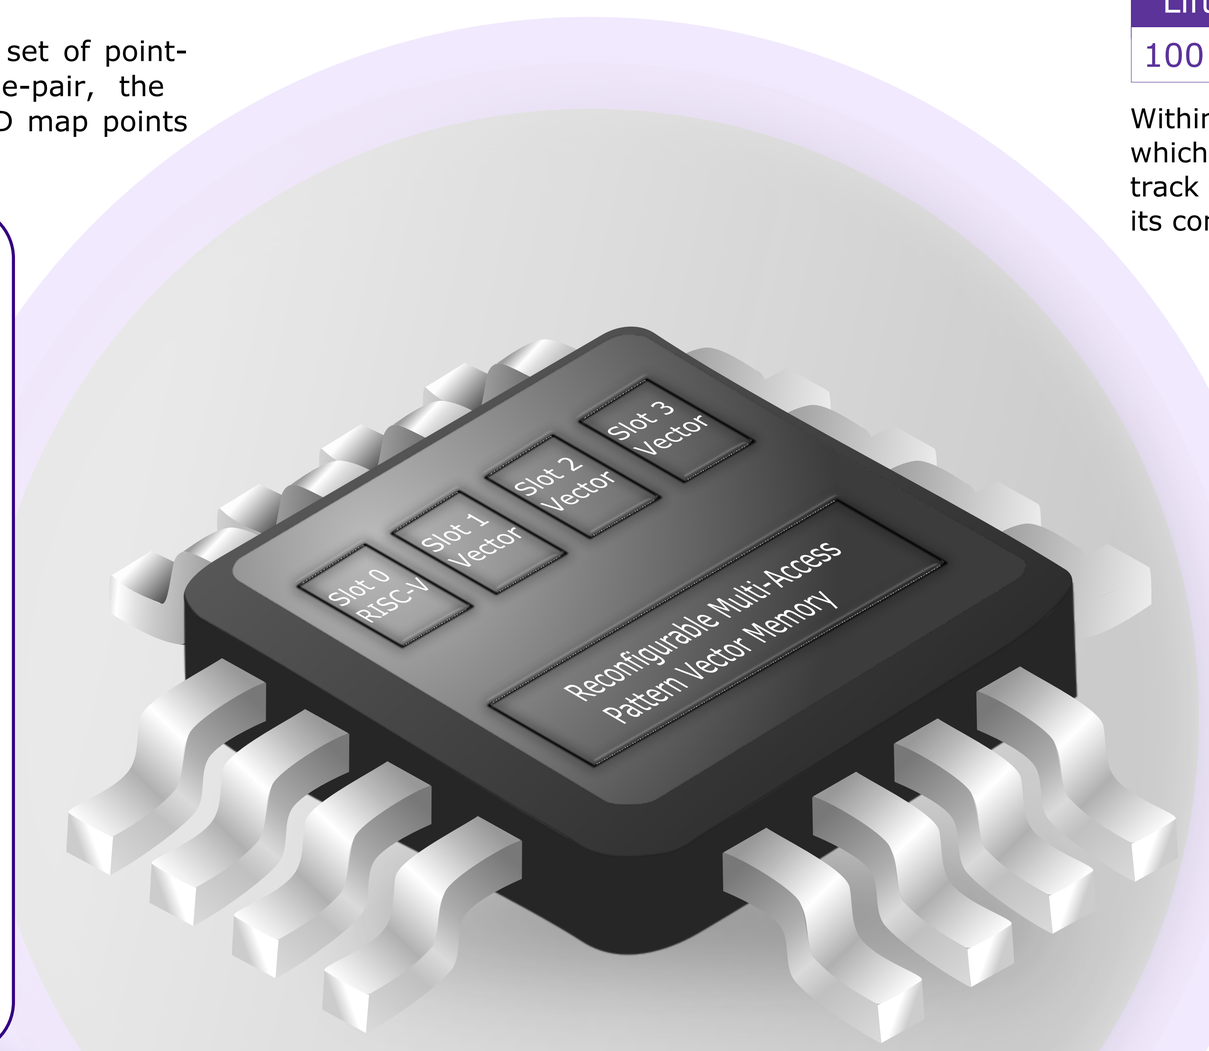

3 Slot 0 / RISC-V

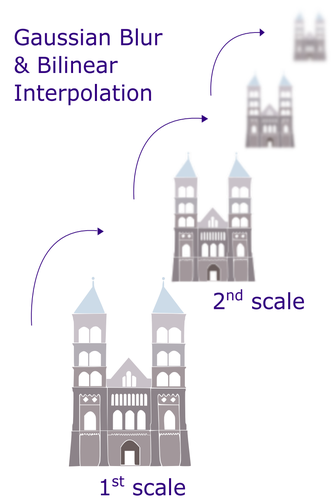

- Those operations are performed across the several scales for each frame
- Sequential and overlapping memory accesses oppose real-time processing and energy efficiency

4 Slot 1 / Vector

FAST / Non-Maximum-Suppresion (NMS)

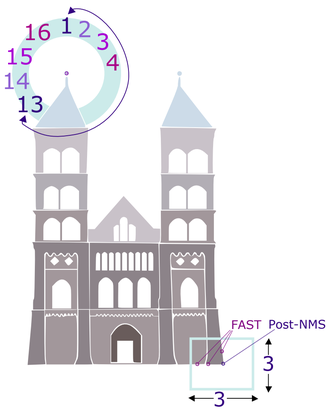

- Performed for all pixels across the scales
- Irregular memory access pattern to load the pixels' ring and NMS patch
- 4 x 12 sequential pixel comparisons for FAST evaluation

5 Slot 2 / Vector

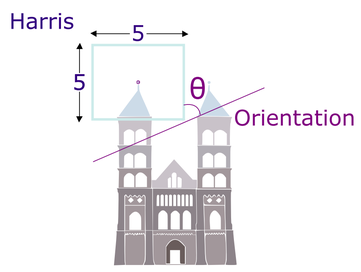

- Performed for the post-NMS keypoints
- Harris strings in a sequence of matrix operations
- ATAN2 involved to obtain the keypoint orientation

6 Slot 3 / Vector

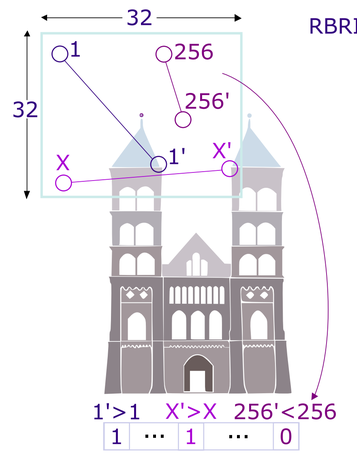

256-bit Descriptor made of binary pixel comparisons

- RBRIEF needs 256 pairs of random accesses per keypoint. For SLAM with VGA images, 1k keypoints per frame, 30 FPS are expected

7 Implementation Results

- C programable ASIP, implemented with Synopsys ASIP Designer, supporting real-time ORB feature extraction
- 4-slot VLIW processor. Scalar core consists of a GP RISC-V, extended with register move and control operations. Vector slots 1-3 implements intrinsics related to FAST & NMS, Harris and Orientation, and RBRIEF respectively
- Reconfigurable multi-pattern vector memory allows single-cycle access to recurrent loads ORB, except RBRIEF which takes 8 cycles
- Synthesized in 22-nm, with an area of 0.566 mm^2
- 1k keypoints can be extracted for VGA images at 140 FPS, for a single-scale (ASIP at 430 MHz)
- 90 mW

References

Reconfigurable Multi-Access Pattern Vector Memory for Real-Time ORB Feature Extraction, ISCAS 2021

Energy-Efficient Application-Specific Instruction-Set Processor for Feature Extraction in Smart Vision Systems, ASILOMAR 2021

Lucas Ferreira
PHD Student at Lund University
lucas.ferreira@eit.lth.se