


Abstract

The Stixel World is a medium-level, compact representation of road scenes that abstracts millions of disparity pixels into hundreds or thousands of stixels. The goal of this work is to implement and evaluate a complete multi-stixel estimation pipeline on an embedded, energy efficient, GPU-accelerated device. This work presents a full GPU-accelerated implementation of stixel estimation that produces reliable results at 26 frames per second (real-time) on the Tegra X1 for disparity images of 1024 × 440 pixels and stixel widths of 5 pixels, and achieves more than 400 frames per second on a high-end Titan X GPU card.

Problem: Compact depth representation

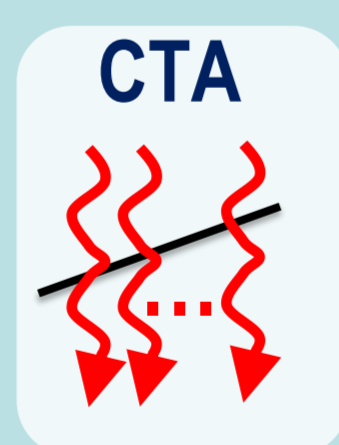



Stixels Overview

- Goal: Compact depth representation
- Extensions: tracking, grouping, semantics
- Dynamic Programming algorithm
- High computational complexity $O(w \times h^2)$


Parallelization

First level parallelism



- Independent / Task level: Typical CPU parallelization
- Each image column is processed by a CTA

Second level parallelism




- Extra parallelism level needed for efficient GPU use
- Sequentially perform h (image height) steps
- CTA threads collaborate sharing info each step
- Decreasing Parallelism: Each step uses one thread fewer

Results

	FPS	Speed Up	FPS / Watt
CPU Multi-thread	13.3	1	0.10
GPU Optimized	413	31	1.65
NVIDIA Drive PX ¹	26	1.95	2.6

CPU: Intel Core i7 980X
 GPU: NVIDIA Titan X
 Drive PX: NVIDIA Tegra X1

Image Size: 1024x440
¹ single-socket



Conclusions:

- Real-time performance for energy efficient GPU NVIDIA Tegra X1.
- NVIDIA Tegra X1 has better energetic efficiency than high-end GPUs.

References:

[1] D. Pfeiffer and U. Franke. Towards a global optimal multilayer stixel representation of dense 3D data. In British Machine Vision Conference, BMVC 2011, Dundee, UK, August 29 - September 2, 2011. Proceedings, pages 1–12, 2011.

Acknowledgements:

This research has been supported by MICINN-Spain under contract numbers TIN2014-53234-C2-1-R and TRA2014-57088-C2-1-R; also by the Generalitat de Catalunya projects 2014-SGR-1506, 2014-SGR-1562, and CERCA Programme; and by People Programme (Marie Curie Actions) FP7/2007-2013 REA grant agreement no. 600388 and ACCIO. We thank NVIDIA corporation for the donation of the systems used in this work.