POSTER P5332 | CONTACT NAME Ryan Beethe: rbeethe@mines.edu

EMBEDDED REAL-TIME OBSTACLE RECOGNITION

RYAN BEETHE COLORADO SCHOOL OF MINES

THE CHALLENGE

The robot will enter the Colorado Robot Challenge, a contest held annually in Great Sand Dunes National Park. Robots are required to chase a beacon signal across a series of obstacles without human assistance in a simulated Mars enviornment.

The goal of this project is to augment an existing chassis (Figure 1) with a stereo vision sensor and software. Using GPU technology, it will be capable of real-time obstacle recognition. Identified obstacles will serve as an input to the existing route-planning software.

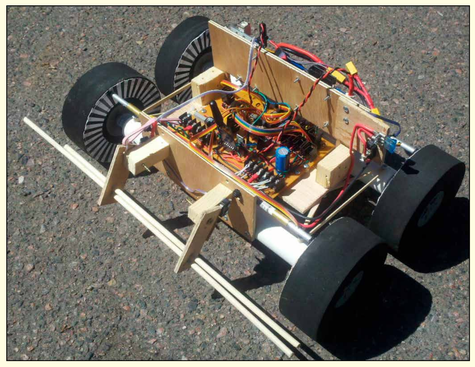

Figure 1: Last year's chassis

THE IMPLEMENTATION

The semi-global matching (SGM) algorithm is memory intensive. Path smoothness data must be kept for each pixel at each possible disparity. For 8-bit, monochrome, VGA-resolution images, with a modest maximum disparity value of 32 pixels, this memory requirement can theoretically be under 100MB. However, with additional GPU memory access optimizations, the memory requirement can easily be a few hundred MB. On the Jetson TK1, which shares 2GB of memory between CPU and GPU processes, including usual Ubuntu processes, a few hundred MB may not be available. Larger maximum disparities are necessary for identifying obstacles near to the stereo sensor, but are prohibitively memory intensive on this platform.

Also, shared memory is used extensively to minimize accesses to global memory. However, a typical thread block uses close to the full allotment of shared memory on a single CUDA core. This presents additional limitations to the maximum disparity that can be calculated with the current implementation.

THE STRATEGEY

Using the existing chassis and route planning software as a starting point, this project will combine the following elements:

- Stereo camera sensor
- Jetson TK1 embedded GPU to provide necessary computational power
- OpenCV library for basic image operations and stereo calibration
- Semi-global matching (SGM) algorithm to calculate stereo correspondence

Previous GPU implementations of SGM have not been open-source. This project aims to contribute a CUDA implementation of the SGM algorithm to the OpenCV library.

Figure 2: Stereo camera sensor

THE ALGORITHM

Stereo matching algorithms determine the disparity, or pixel shift, between corresponding features of stereo images. H. Hirschmuller's semi-global matching (SGM) algorithm [1] reduces noise in disparity images by maximizing smoothness along every path (left, right, and diagonal) through every pixel.

Figures 4 and 5 compare the output of the OpenCV CPU-implemented SGM to a standard block-matching technique.

Figure 4: SGM yields smooth disparity images

Figure 3: Sample stereo image (left image shown)

Figure 5: Standard block-matching technique

RESULTS

Currently, the CUDA implementation of the semi-global matching (SGM) is limited to path smoothing only in the horizontal directions, and smoothing parameters are not yet tuned.

Initial output from the algorithm (Figure 6) demonstrates early solutions to memory constraints and basic parallelizing schemes for the algorithm.

Figure 6: Current output of CUDA SGM, after horizontal path smoothing only. Smoothing in additional directions will reduce horizontal artifacts.

REFERENCES

[1] Heiko Hirschmuller. Stereo processing by semiglobal matching and mutual information. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 30(2):328–341, 2008.

[2] S. Birchfield and C. Tomasi. A pixel dissimilarity measure that is insensitive to image sampling. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 20(4):401–406, 1998.

FUTURE RESEARCH

The current research focus is to add a proof-of-concept object recognition application to the robot chassis in time for the competition (April) using existing CPU implementations for semi-global matching (SGM) (Figure 4).

After competition, additional research will seek to complete a functional CUDA version of the SGM algorithm. Additional performace optimizations will be necessary. Expected speed improvements have not yet been observed, and memory constraints will need to be readdressed to improve versatility.

When ready, this will be the first open-source GPU implementation of the SGM algorithm.

SPECIAL CONTRIBUTORS

Colorado Robot Challenge spacegrant.colorado.edu
Mines Self-Driving Car www.sdc-csm.com

CONTACT INFORMATION

Email rbeethe@mines.edu
Phone +1 (970) 556 2521