Real Time GPU Stereo Visual Simultaneous Localization and Mapping

Brent Tweddle

May 13, 2009

Mobile Robotics: DGC

- Darpa Grand Challenge vehicles represent the current state of the art in autonomous mobile robotics
- 3 Steps performed Online
  - Navigation (Localization and Mapping)
  - Path Planning
  - Control

Perception Video

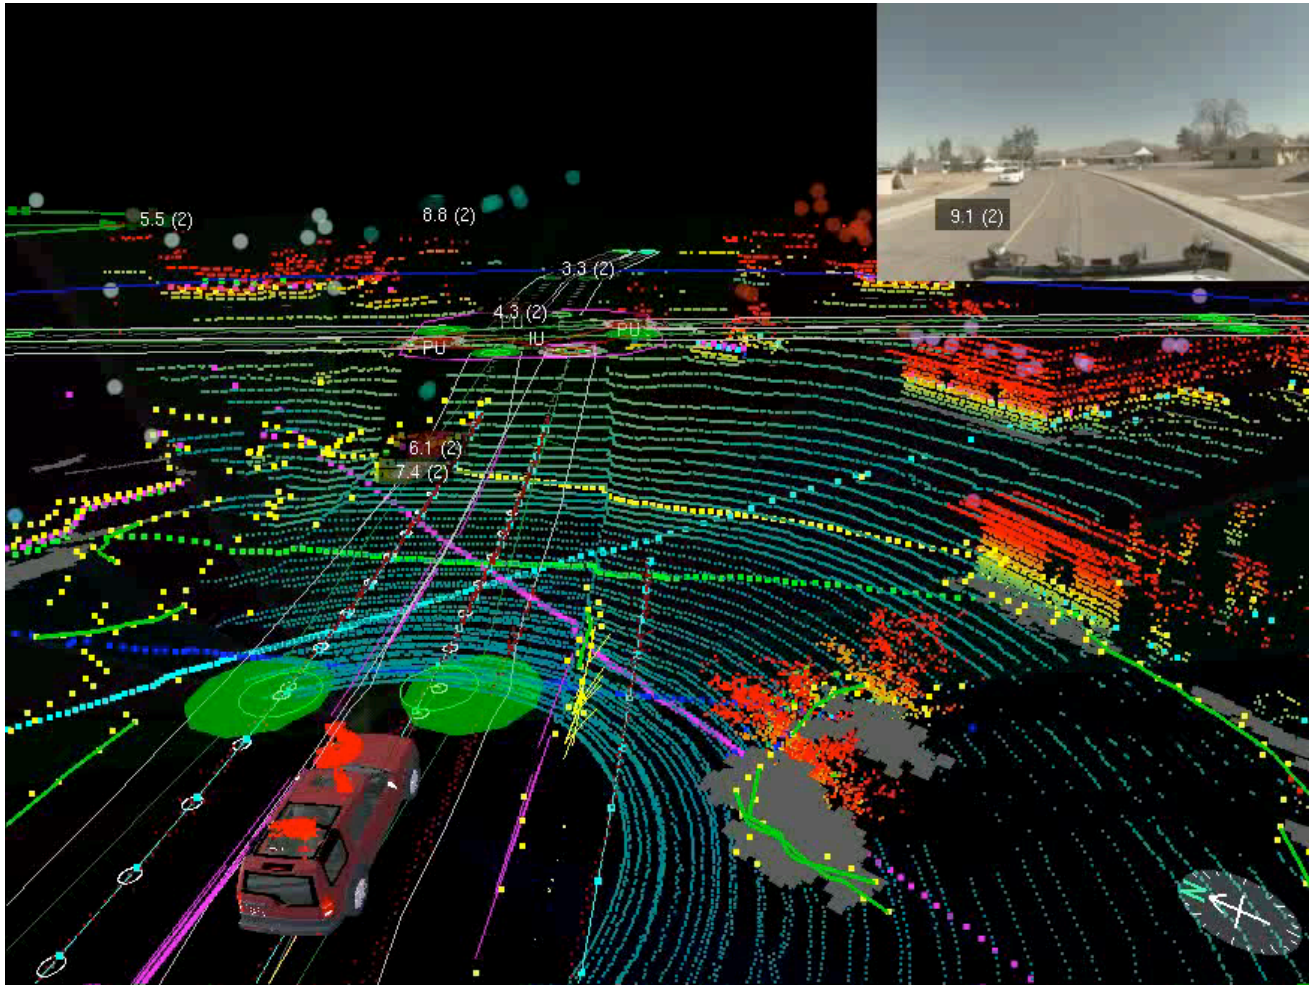

DGC Computing

- Sensors:
  - 15 radars
  - 12 Single axis lidars
  - 1 rotating lidar
  - 5 cameras
  - GPS
  - Inertial measurement system
- Computing
  - 10 Blade Cluster
  - Each Computer is Quad-Core 2.3 GHz Xeon
  - Total Power consumption: 4000W

This power consumption is impractical for a large number of applications, especially aerospace robotics.

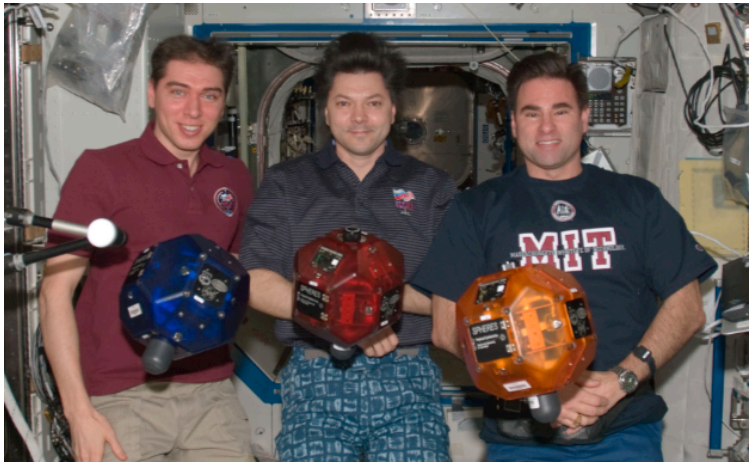

GPU's for Real Time Robotics

Processor	Theoretical Peak GFLOPS	Watts	Watts per GFLOPS
Quad "Bloomfield" Xeon 3.2 GHz	25.6 GFLOPS	130 W	5.078
Core 2 Duo "Penryn" 2.53 GHz	20.2 GFLOPS	25 W	0.810
Cell Processor	152 GFLOPS	80 W	0.526
NVIDIA Tesla C870	518 GFLOPS	170 W	0.328
NVIDIA GeForce 9800 GT	504 GFLOPS	105 W	0.208
NVIDIA GeForce 8800M GTS	240 GFLOPS	35 W	0.145

- Assumptions:
  - Xeon issues 2 flops per cycle per core
  - Core2Duo issues 4 flops per cycle per core

http://icl.cs.utk.edu/hpcc/hpcc_desc.cgi?field=Theoretical%20peak

Inverse Depth Parametrization for Monocular SLAM: Loop Closing Sequence

AUTHORS:
Javier Civera, jcivera@unizar.es
Andrew J. Davison, ajd@doc.ic.ac.uk
J. M. M. Montiel, josemari@unizar.es

Proposed Algorithm

- Stereo Visual SLAM
  - Using stereo cameras create a map of your environment and locate yourself within it
- Grid Map
- Algorithm Flow:
  - Dense Stereo Correspondence (18.337)
  - Scan Matching (Thesis)
  - Particle Filter Grid Map (Thesis)

Stereo Correspondence Search

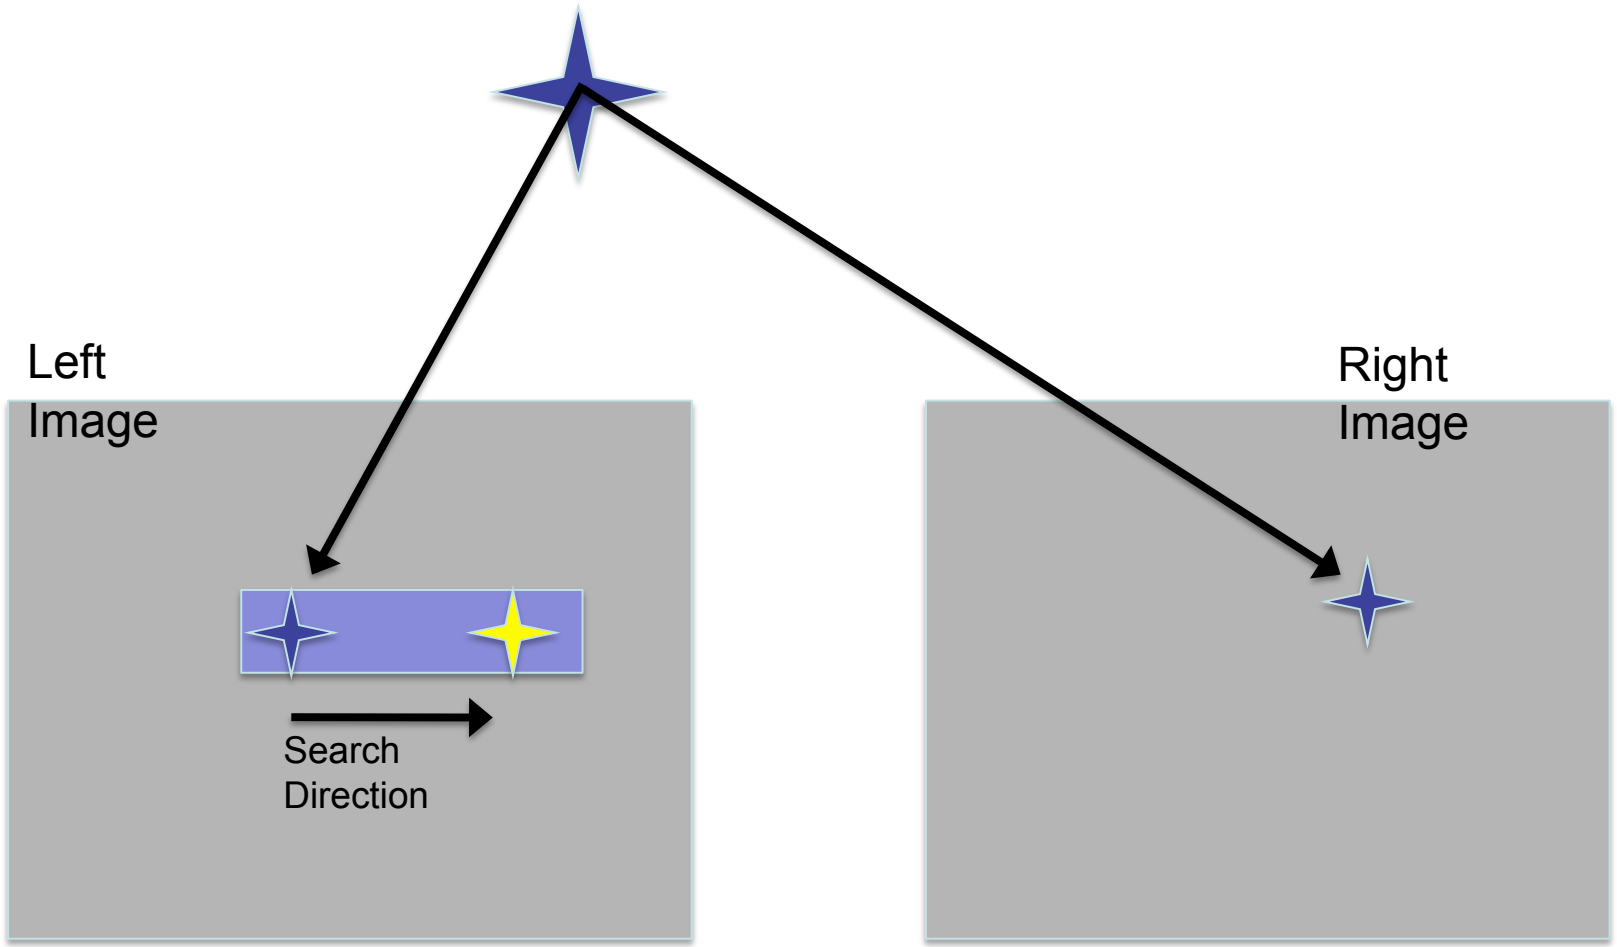

Dense Stereo Correspondence

(a) (b) (c) (d)

Fig. 6. Results with simulated stereo images. (a) Virtual left stereo image. (b) By S&S DP without noise. (c) By WTA MW5 lr with noise. (d) By S&S DP with noise.

- Large body of work exists on dense stereo:
  - Scharstein, Szeliski “A Taxonomy and Evaluation of Dense Tow-Frame Stereo Correspondence Algorithms”, IJCV 2002
  - Brown, Burschka, “Advances in Computational Stereo”, IEEE PAMI, 2003
- Optimized algorithms for CPU SIMD hardware (512x512: <0.1s)
  - Van der Mark, Gavrila, “Real-Time Dense Stereo for Intelligent Vehicles”, IEEE Trans. ITS, 2006
- Cuda Implementation by NVIDIA’s Joe Stam
  - Crude and no published timings

Images, Pixels, Threads & Blocks

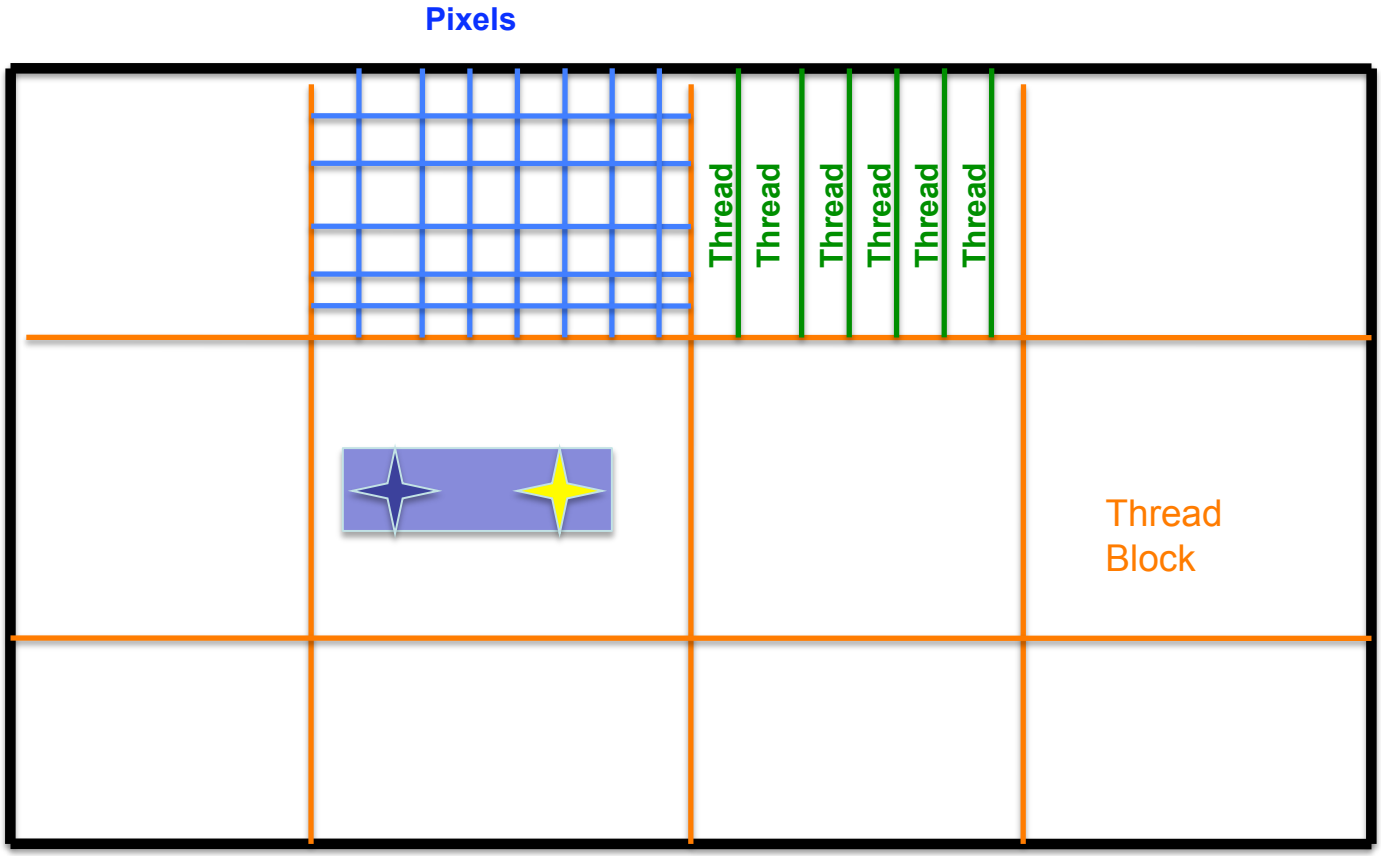

List of Improvements

- Left-Right Consistency Check
  - Perform correspondence on both sides and check that results match
    - Naïve implementation doubles FLOPS
  - Implemented on the GPU by storing calculations in two 3D grids
    - Same number of FLOPS, but more memory is needed
  - Had to add additional kernel to avoid race conditions
- Threshold for minimization to avoid disparity noise in textureless regions
  - New < Best-250

Visual Results

- Visually appears much more accurate
- Runs in 25ms still less than
  - More than 16ms, but still less than most CPU implementations

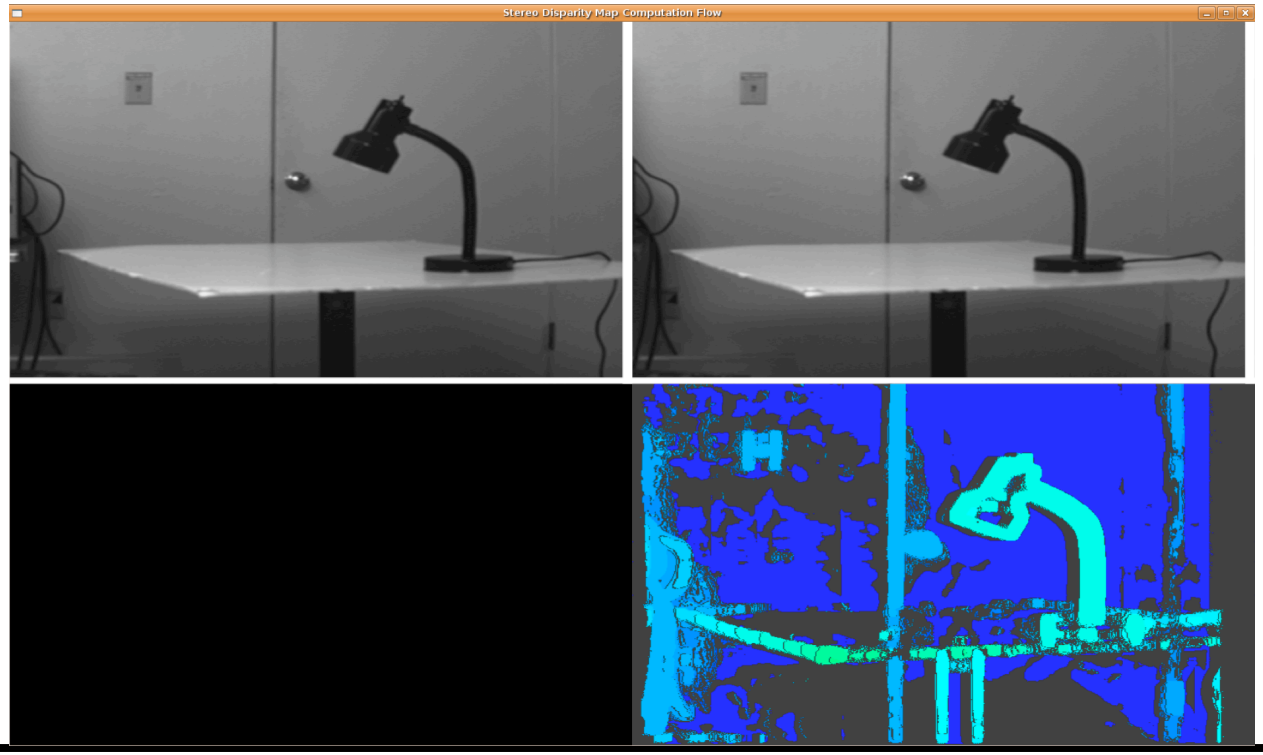

Performance Limitations

- Algorithm is not memory bandwidth limited
- However it is limited by:
  - Memory Latency
  - Multiprocessor Warp Occupancy
    - Compute 1.1, 20 registers, 80 bytes shared mem, 20 bytes constant mem
    - Register Limited

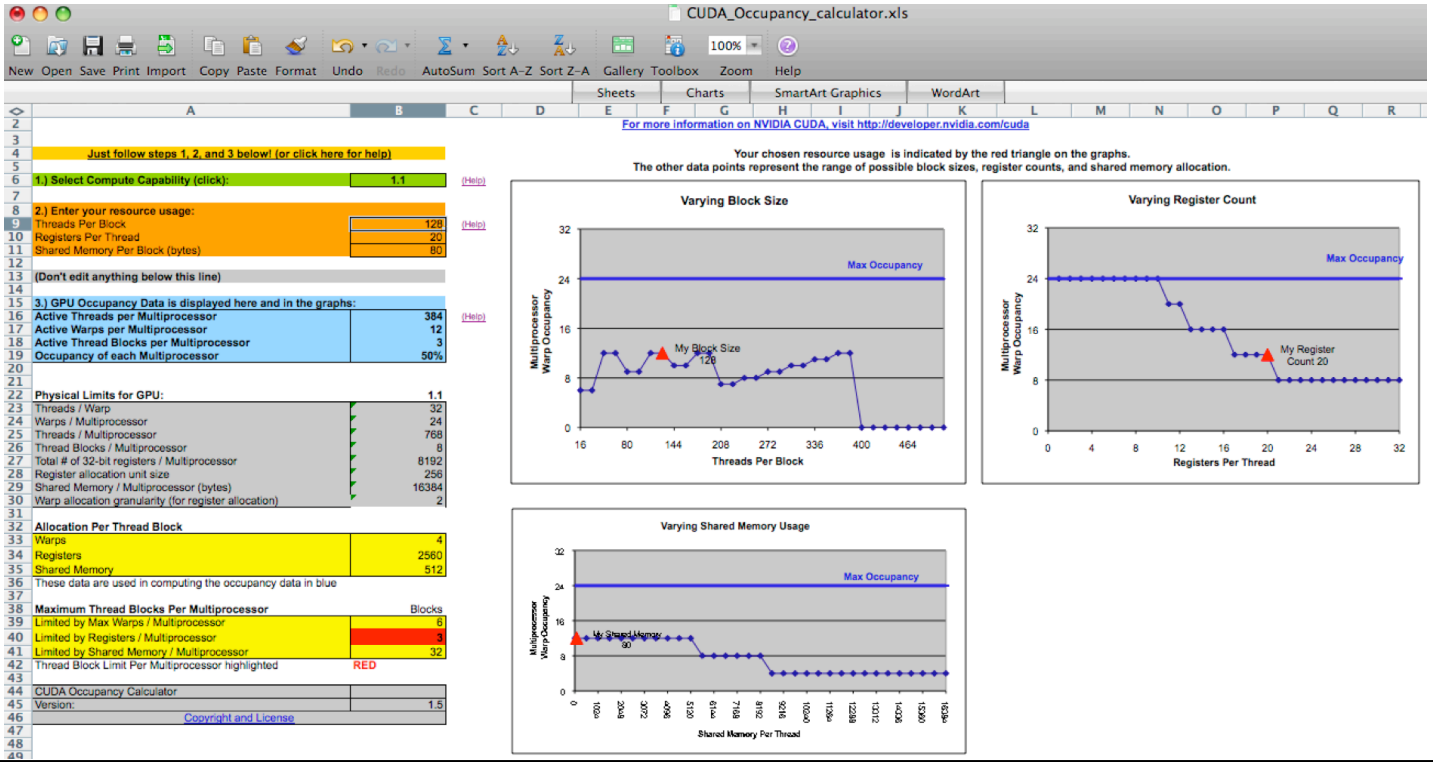

Conclusion

- GPU's are a valid method to use for robotic navigation
- Showed implementations of first step of navigation algorithm (accurate stereo vision)
- Analyzed performance limitations of implementation and suggested future recommendations