

# PathFinder in CUDA


James Doyle '10, Prof. John MacCormick

Department of Mathematics and Computer Science, Dickinson College



## 1. Background

- PathFinder is an image segmentation algorithm that groups the hundreds of thousands of pixels in an image into a few thousand "superpixels"



Above: We find edge paths through an image of 180,000 pixels

Right: PathFinder's output image of 4,400 superpixels



- CUDA is a C-based software platform that permits the graphics processing unit (GPU) to perform computational tasks instead of the CPU
- The GPU excels at parallel computing

## 2. Goals

- Re-write sections of the PathFinder algorithm in CUDA to **utilize the GPU**
- Determine performance increase of Pathfinder in CUDA when compared to original implementation in Java

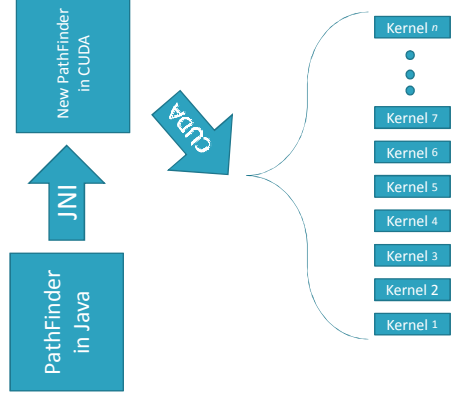
## 3. Results

- Overall, we achieved a **3-5 times speedup**
- In several areas, the CUDA algorithm performed its calculations over **100 times faster** than the Java implementation




## 4. How we did it

- Utilized the Java Native Interface (JNI) to call CUDA code from the Java implementation



- CUDA instructs GPU to execute hundreds of "kernels" concurrently

## The new GPU



## References

Fung, J., & Mann, S. (2008). Using graphics devices in reverse: GPU-based Image Processing and Computer Vision. 2008 IEEE International Conference on Multimedia and Expo, 9-12.

Halfhill, T. (2008). Parallel Processing With CUDA. *Microprocessor Report*, 1.

NVIDIA Corporation (2009). NVIDIA CUDA – Programming Guide. Retrieved October 10, 2009, from [http://developer.download.nvidia.com/compute/cuda/2\\_3/toolkit/docs/NVIDIA\\_CUDA\\_Programming\\_Guide\\_2.3.pdf](http://developer.download.nvidia.com/compute/cuda/2_3/toolkit/docs/NVIDIA_CUDA_Programming_Guide_2.3.pdf).