



Damascene: Highly Parallel Image Contour Detection

Bryan Catanzaro, Narayanan Sundaram, Bor-Yiing Su, Yunsup Lee, Mark Murphy, Kurt Keutzer

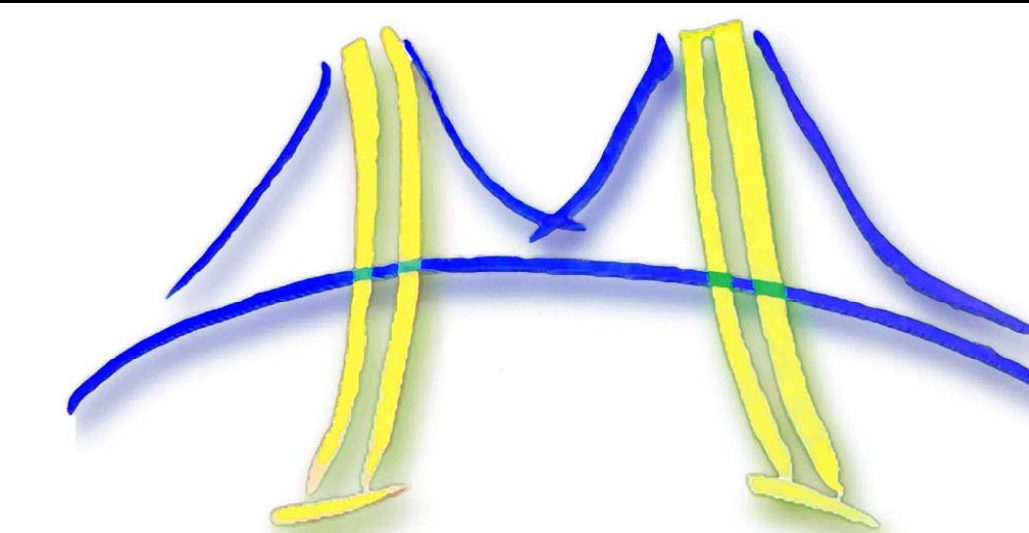


Image Contour Detection

- Image contour detection is fundamental to image segmentation and many other computer vision problems





Image Human Generated Contours Machine Generated Contours


gPb Algorithm: Current Leader

- global Probability of boundary [Maire, Arbelaez, Fowlkes, Malik, CVPR 2008]
- Currently, the most accurate image contour detector
- 5.8 mins per small image (481 by 321) limits its applicability
 - Too slow for interactive photo editing
 - Too slow even for Image Retrieval

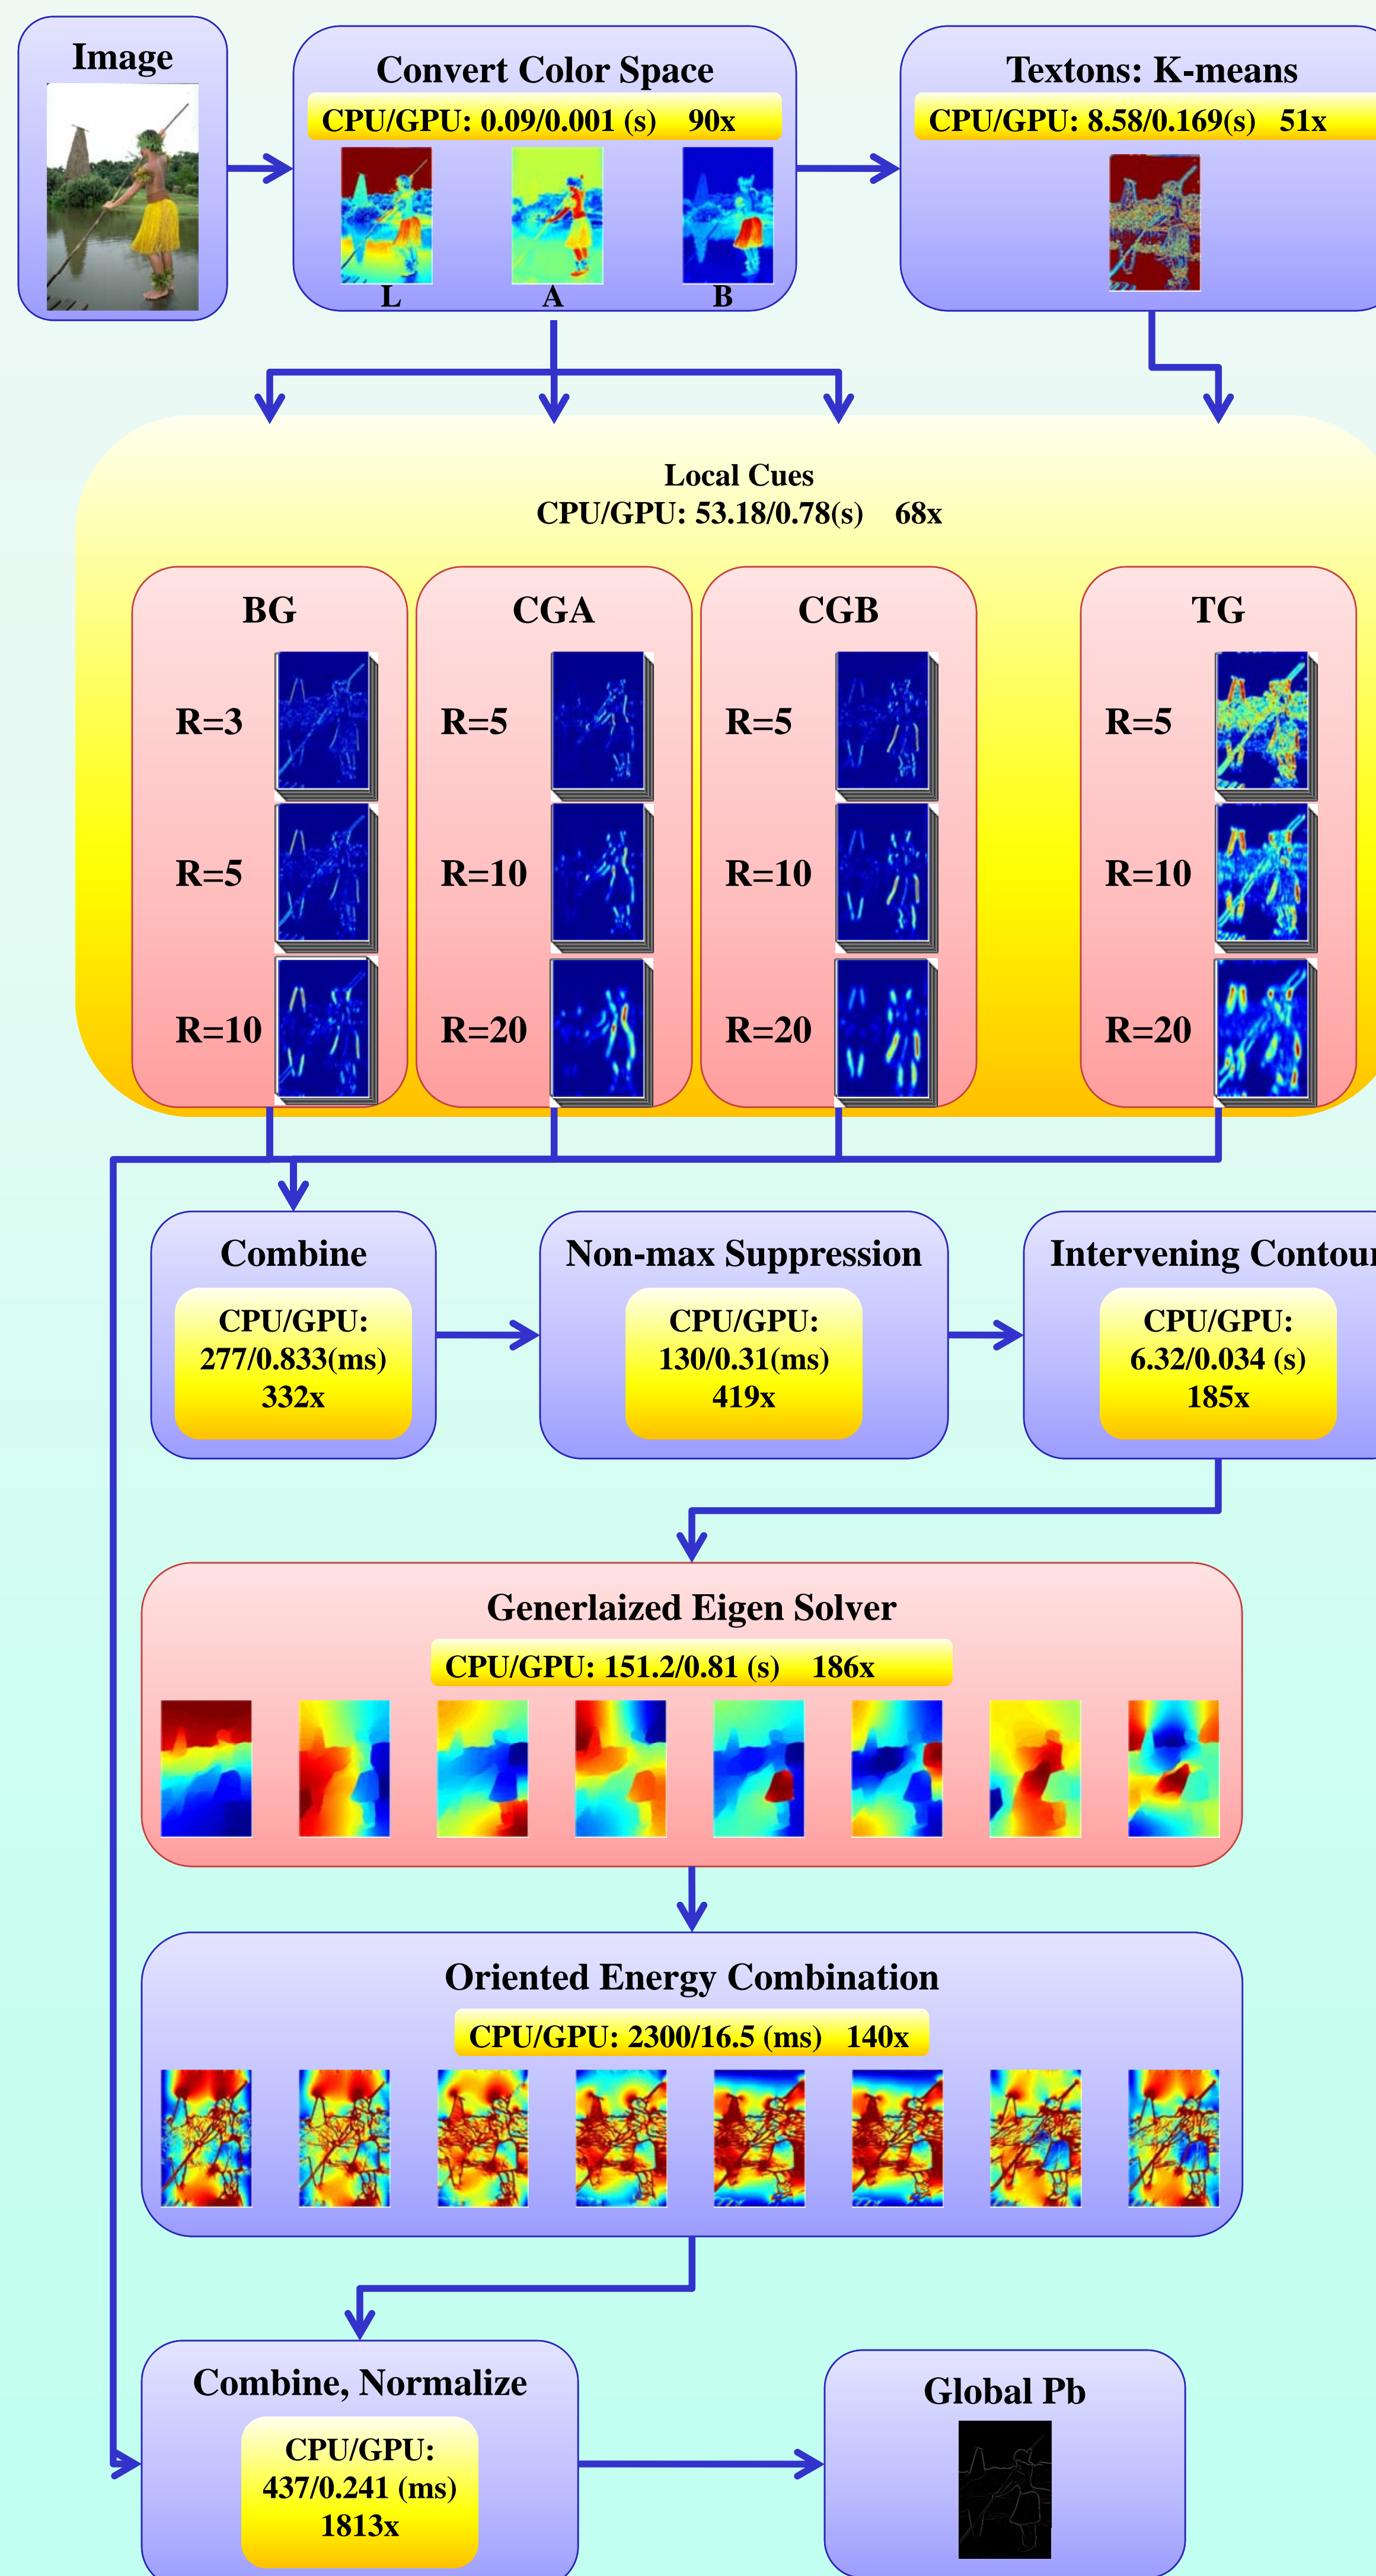


Platform: Nvidia GTX200 Series

Specifications	GTX280
Processors	30 @ 1.3 GHz
Physical SIMD Width	8
SP GFLOPS	933
Memory Bandwidth	141.7 GB/s
Register File	1.875 MB
Local Store	480 kB
Memory	1 GB

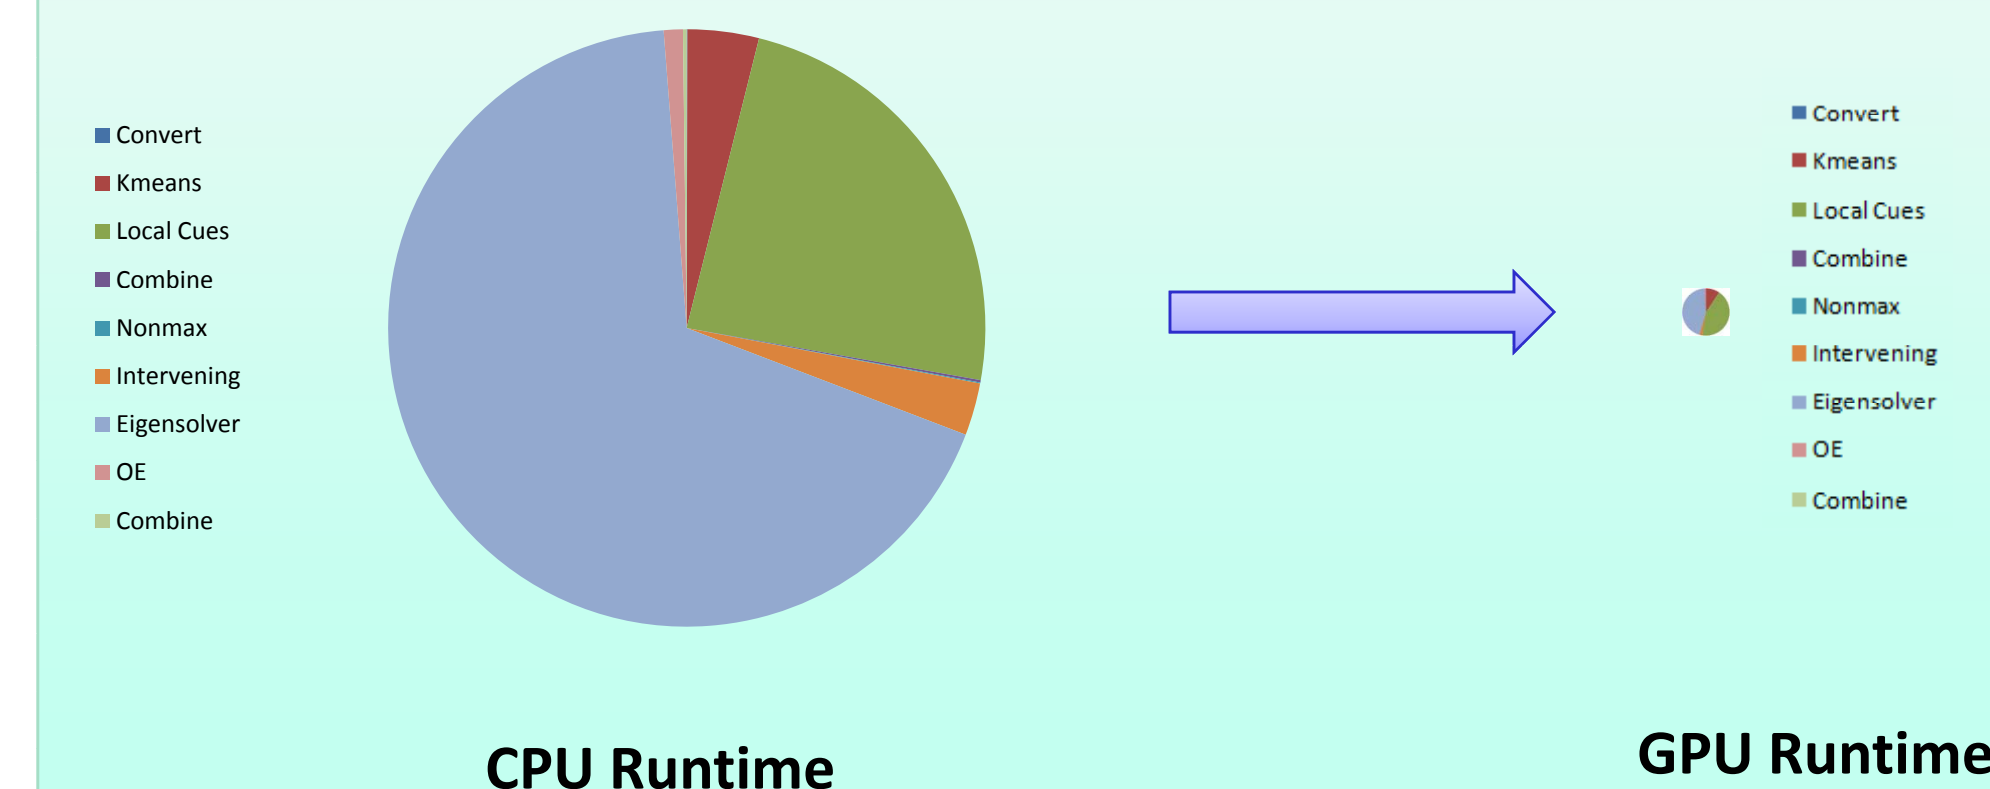


Program Flow



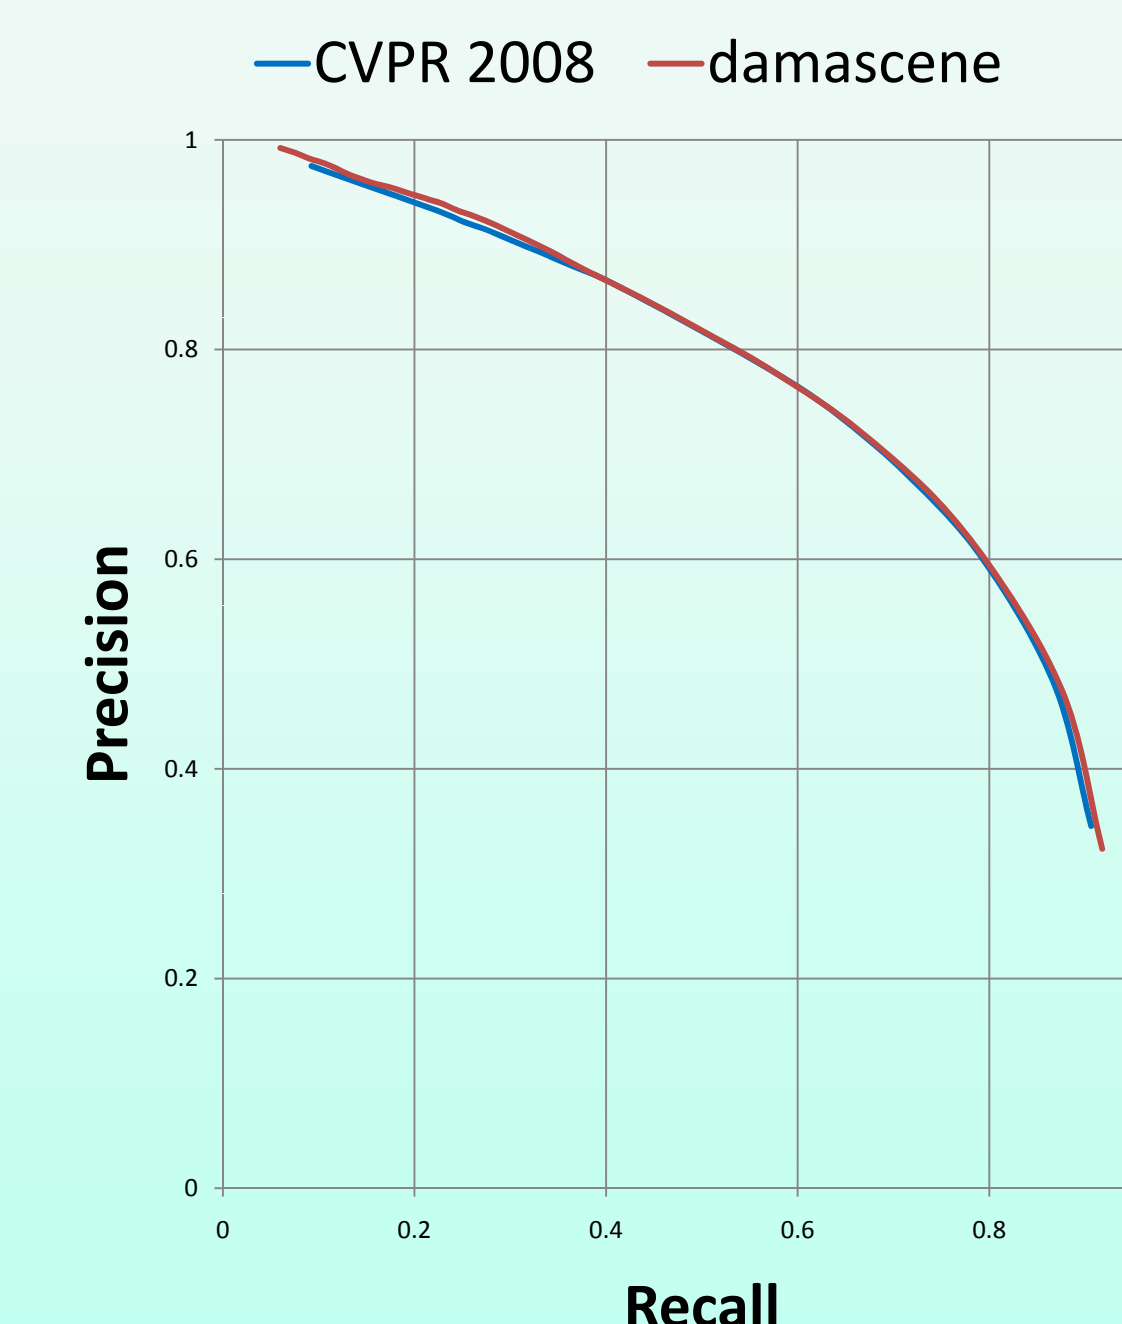
Overall Performance

- Image Size: 481 by 321, 154401 pixels in total
- CPU/GPU Runtime: 236.7s/2.081s
- Speedup: 114x



Precision-Recall Graph

- We achieve slightly better accuracy on the Berkeley Segmentation Dataset
 - Comparing to human segmented "ground truth"
- F-measure 0.70 for both



Conclusion

- Damascene provides highest quality image contour detection at user acceptable rates
- It demonstrates the transformational speedup potential of manycore architectures
- Damascene was enabled by the collaborative environment at the Berkeley UPCRC
- Future work will generalize Damascene into a case study for application and programming frameworks (stay tuned)