

# GPUVision

## Image Processing on the GPU

Michael Lehr [michael.lehr@gmail.com](mailto:michael.lehr@gmail.com)  
 Ikjin Ahn [ikkjin@gmail.com](mailto:ikkjin@gmail.com)  
 Paul Turner [turnerpd@seas.upenn.edu](mailto:turnerpd@seas.upenn.edu)  
 May 5, 2005  
 CIS 700 – GPU Programming and Architecture  
 University of Pennsylvania

## Overview

- Purpose/Description
- Related Work
  - Basis
- Design
- Supporting Filters – Results – Performance
  - Convolutions
  - Edge/Feature Point Detection
  - Matrix/Vector Dense/Sparse mult/sum/max
    - Solver – Conjugate Gradient
  - Image Segmentation
  - Disparity Map
- Problems
- Future

Most Tests on 2 GHz Athlon, NVidia Quadro 3400 PCIe

## Purpose/Description

*“To create a windows based GPU Accelerated Image Processing Framework”*

- Users
  - Filter users
    - no CG knowledge
    - Template based graphics knowledge
  - Filter creator
    - No RenderTexture knowledge
    - Template based Filter Creation
      - Can create new filter structure (not CG) in under 5 min
        - » Focus on CG code

## Related Work

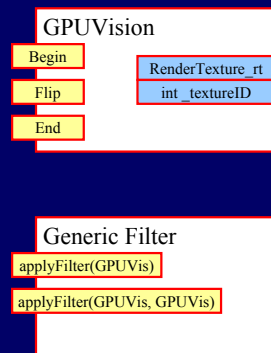
- OpenVidia
  - Linux based
  - Video processing computer vision
  - Some of the Algs are actually not quite right
  - Some CG code is portable
- Mac OSX ‘Tiger’
  - Core Image
    - Found after we came up with structure and Concept of Operations
    - Only on Mac
    - Very very similar
      - I guess we have a good structure =>
      - ‘Image Units’ versus ‘Filters’
- PyFx : Python IP
  - Very well abstracted

## Basis

- Image Segmentation
  - “Isoperimetric Graph Partition for Data Clustering and Image Segmentation” Leo Grady and Eric Schwartz, 2003
- Harris Corners
  - “A Combined Corner and Edge Detector” Chris Harris & Mike Stephens, 1988
- Color spaces
  - <http://www.couleur.org/index.php?page=transformations>
- Various other sources from computer vision (non GPU)


## GPUVision Design -Framework

- GPUVision class encapsulates
  - Up/down textures to GPU
  - Ping/Pong
  - Drawing to screen
- Generic Filter is basis for all filters
  - Filters hold CG code
  - Generally can be applied to one or two GPUVis
    - Different for each filter
    - Would like to make this more user-friendly
  - A GPUVision is applied to a Filter
  - GPUVision RenderTexture has results of filter
  - Filters can be chained

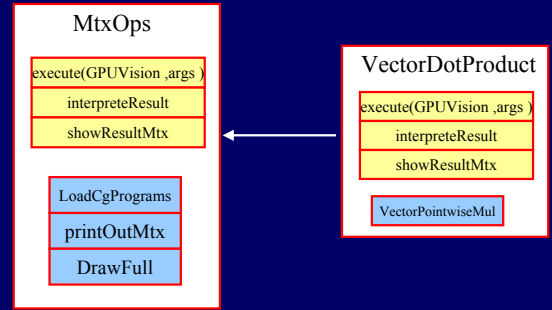


# GPUVision Design - Filters

- Can make more complex 'Filters' which are pre-defined sequence of filters
- Canny Edge uses 5 filters
  - Only one context switch
  - 5 'Flips'
  - Flip resets Read/Write and RenderTexture Target Buffer
- Can create any single filter or string filters together to make complex composite filters



# GPUVision Design - Matrix Ops



# GPUVision: Amt of Code

- Total
  - 150 files
  - 27 filters
  - 20 Matrix functions
  - 22 test class

# GPUVision vs OpenNvidia

|         |                   | Programmability |           |
|---------|-------------------|-----------------|-----------|
|         |                   | OpenGL style    | Abstract  |
| Purpose | General           | CG              | Brook     |
|         | Function specific | OpenVidia       | GPUVision |

# Real Code Example

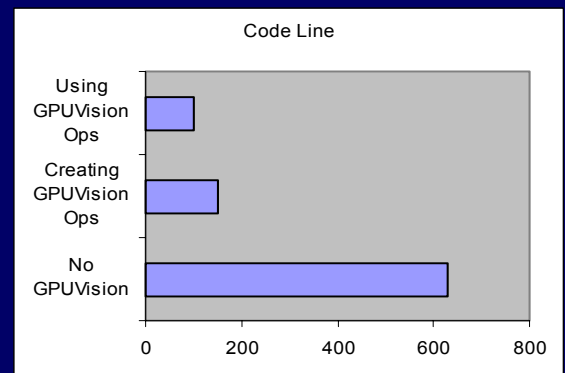
```
static void Display()
{
    glClear(GL_COLOR_BUFFER_BIT | GL_DEPTH_BUFFER_BIT);
    gpuVis = new GPUVision((unsigned int)0,n, m );
    simpMulSparseMtx = new IsoSimplifiedMulSparseMtx(gpuVis->GetContext());
    preProcessing();
    texIdx=gpuVis->UploadToTexture(data_x, n, m, GL_RGBA);
    texIDa=gpuVis->UploadToTexture(data_a, n, m, GL_RGBA);
    texIDremov=gpuVis->UploadToTexture(data_remov, 1, 1, GL_RGBA);

    static float theta = 0;
    simpMulSparseMtx->execute(gpuVis, texIDa, texIdx, texIDremov);
    printf("#####Result #####\n");
    simpMulSparseMtx->showResultMtx(gpuVis, 10, 10);

    glutSwapBuffers();
    glutPostRedisplay();


    system("pause");
    exit(0);
}
```

# GPUVision CodeCount



- In case of Sparse matrix multiplication.

# GPUVision vs OpenVidia



# Support Filters

- Add, Subtract, Multiply, Threshold
  - Can take in a number to add/sub/mult or can do pointwise add/sub/mult of two textures
  - Not for performance but for ease
    - Could easily combine these into custom filters
    - Good for proof of concept of filter chains and working with the GPU
- Very very fast.... but, not very useful by themselves..

# Kernel Based Filters

- Allow kernels of any shape in ConvolutionFilter
  - Generate the CG code on the fly
  - Can either pass in the array into the cg code (useful if it may change on the fly)
  - Or can write it into the code
    - Currently only do this for small kernels
  - Ex: 150 → 153 fps for putting in code
- Single Filters
  - Blur, Gaussian of Laplacian, Gabor, Packed Convolution, Convolution
- Composite Filters
  - Gaussian and Laplacian Pyramids, Sharpen, Laplacian


# Blur & Kernels in General

- 2 Vectors vs 1 square?
  - For small kernels (around 5 or less) squares are more efficient
  - Otherwise 2 Vectors
  - Code versitile enough to handle anything

```

_gauss5x5h = new ConvolutionFilter(_gauss5xKernel, 5,1,3,context);
_gauss5x5v = new ConvolutionFilter(_gauss5xKernel, 1,5,3,context);
_gauss3x3 = new ConvolutionFilter(_gauss3x3Kernel,3,3,3,context);
    
```


– params are  
Kernel(float\*), width, height, channels, CGContext



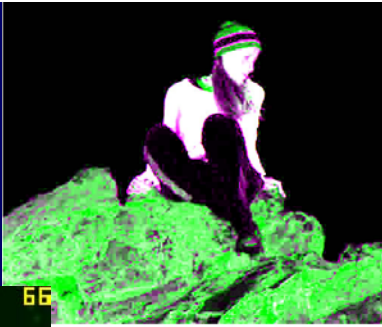
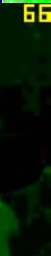
# Packing

- Convert RGB → XYZ → CIE LAB
- Take AB from LAB color channel and can use all for channels for neighboring pixel
  - Image is 1/2 the size
  - Can process 2 pixels values per packed pixel
  - For 5x5 goes from 25 texture lookups per pxl to 7.5 per pxl
- Pack 4 into 1 you reduce that to 2 1/4 tex lookups per pixel

$$(3*5)/2 = (\text{width*height})/\text{numPxls}$$



| Gaus blur | FPS | % Packed Faster |
|-----------|-----|-----------------|
| ur Pack   | 186 |                 |
| ur Normal | 101 | 184%            |
| ur Pack   | 117 |                 |
| ur Normal | 53  | 220%            |
| ur Pack   | 85  |                 |
| ur Normal | 36  | 236%            |
| ur Pack   | 67  |                 |
| ur Normal | 27  | 248%            |


66

- These are AB Channels
- Times for AB counts packing and unpacking images

19

## Gaussian and Laplacian Pyramids

| Gaussian             | GPU    | MipMap | CPU    |
|----------------------|--------|--------|--------|
| create RT            | .277   | --     | --     |
| create               | .00265 | --     | .01375 |
| ul+create            | .0086  | .14    | .01375 |
| ul+cr+dI             | .01485 | .15    | .01375 |
| Laplacian & Gaussian |        |        |        |
| create               | .0054  | --     | .094   |
| ul+create            | .0104  | --     | .094   |
| ul+cr+dI             | .01859 | --     | .094   |



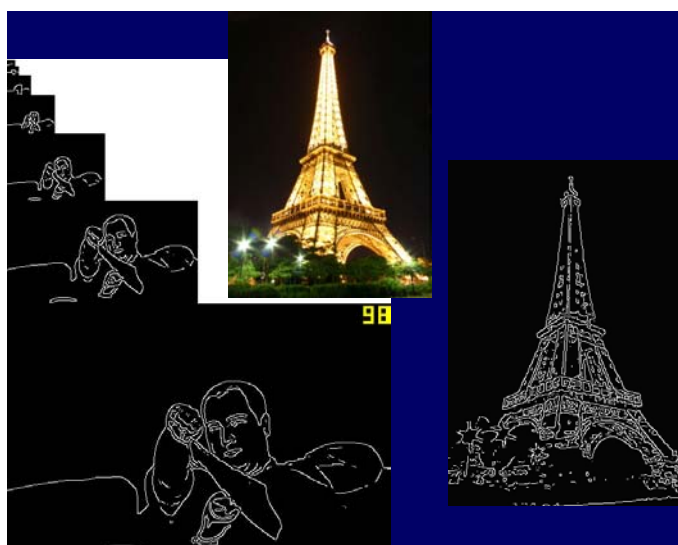
- Close to Matlab Gaus Pyramid Performance!
  - Demolished the Matlab Laplacian performance but...
    - Could not find efficient Laplacian Matlab code
- Once RTs are made, demolish regular mipmapping
- Am downloading ALL levels of the pyramid

Upenn - CIS 700 – Ahn, Lehr, Turner

## Canny

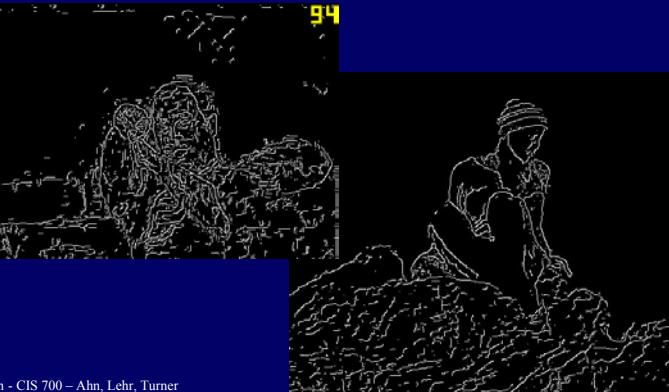
- Described earlier
  - Greyscale (Y from YUV colorspace)
  - Blur Image using 5x vertical and horizontal
  - Find X and Y magnitudes
    - Can find magnitude and orientation of edges
    - Threshold
  - NonMaxSupression
    - Convert multi-pixel line into single pixel
      - Only shrinks within an area – 5 pxl difference will create 2 lines

21



98

## Packed AB Canny

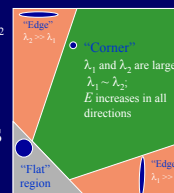


94

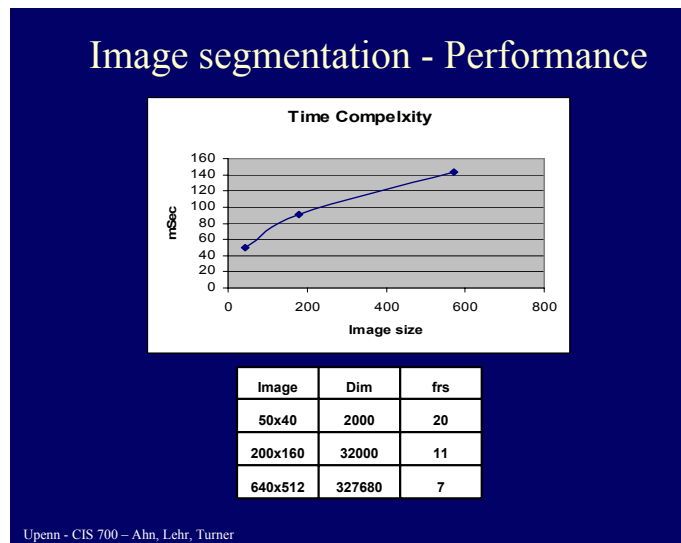
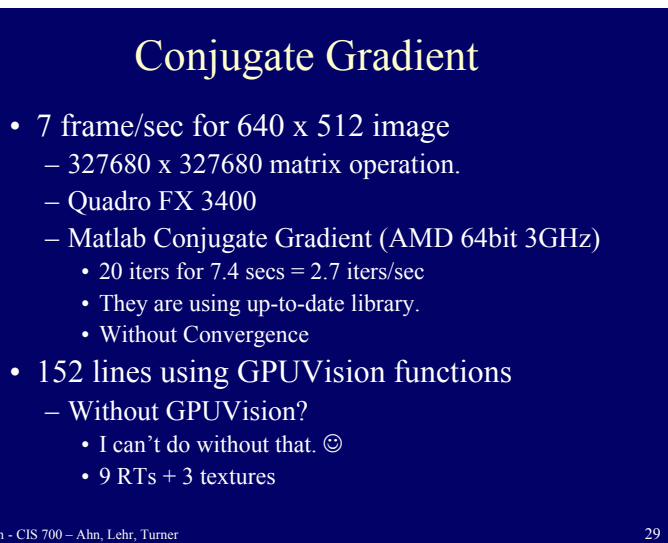
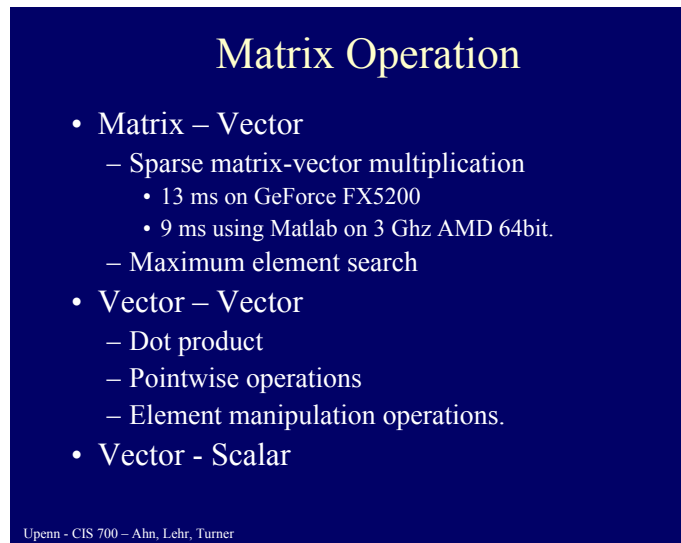
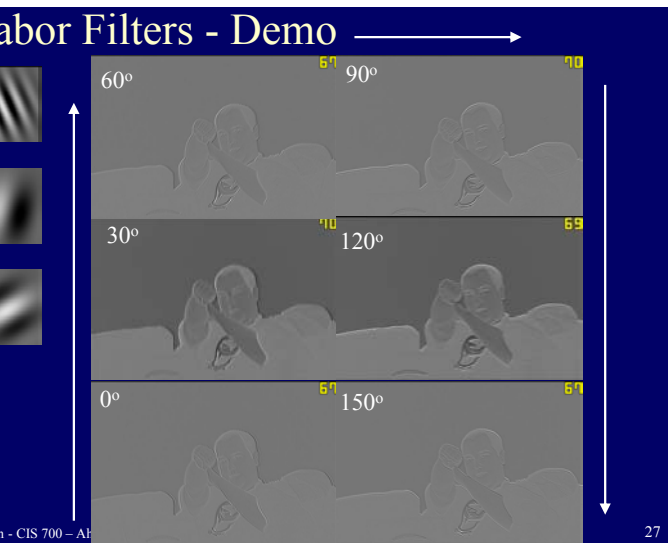

1 - CIS 700 – Ahn, Lehr, Turner

## Harris Corner Detector

- Find corners in a scene
- Solve following equation
 
$$M = \int_W g(g^T)w \, dx = \sum_{i,j} \begin{bmatrix} g_x(i,j)g_x(i,j) & g_x(i,j)g_y(i,j) \\ g_x(i,j)g_y(i,j) & g_y(i,j)g_y(i,j) \end{bmatrix}$$
- If falls in green area then corner
- Basically looking for x and y magnitude in a window to be large
- If  $\det(M) > \text{threshold}$
- Need to find the local maximum in a window so we don't get many points for the same corner!



Upenn - CIS 700 – Ahn, Lehr, Turner



## Code sample

```
while(r_result>threshold){
  vecDP->execute(rho_1, r_1, r_1);
  if(DEBUG) vecDP->showResultMtx(rho_1, 1, 1);
  vecPntDiv->execute(beta_1, rho_1, rho_2);
  vecSc1Mul->execute(temp, beta_1, p_1);
  vecPlus->execute(p, r_1, temp);

  mulMtx->execute(q, mtxId, p, dMaxPos);
  vecDP->execute(temp, p, q);
  vecPntDiv->execute(alpha, rho_1, temp);
  vecSc1Mul->execute(temp, alpha, p);
  vecPlus->execute(x, temp);
  vecSc1Mul->execute(temp, alpha, q);
  vecMinus->execute(r, r_1, temp);
  r->getResult(r_result);
  swapGPUVision(&r_1,&r);
  swapGPUVision(&rho_2,&rho_1);
  swapGPUVision(&p_1,&p);
}
```

## Image Segmentation



- Gestalt Theory
  - Can you see a triangle?
- Isoperimetric Graph Partitioning
  - Ideas from electronic circuits
  - Set a ground node (like a GND in circuit)
  - Calculate energy map (equivalent voltage)
  - Interactive




## Image Segmentation Examples



## More examples



## Examples





## Examples






## Disparity Map


- Calculate distance information
  - From stereo calibrated cameras
  - Vertically aligned



## Disparity Map



## Disparity Map



## Problems

- Using AB from LAB color space should result in *different* edges
  - No/little shadows
  - We got *bad* edges
- Flip/Begin was tricky
  - When switching between contexts the Read and Write buffer did not stay consistent.
  - Need to reset Read/Write and wglBindTexImageARB
  - These are less as costly as a context switch but not considerably less

## Future

- Program users
  - Uses a GUI and appropriate filters to create effect
  - Integrate into Photoshop (free SDK and implementation description)  
[http://download.developer.nvidia.com/developer/SDK/Individual\\_Samples/featured\\_samples.html](http://download.developer.nvidia.com/developer/SDK/Individual_Samples/featured_samples.html)
  - Video for real time computer vision (like OpenVidia)
- More optimized 4 packing and blurring

## Future

- Video Support
- Create an OpenSource Library for the community.
- Change GPUVision to allow holding of any number of textures and manages begin/end and flipping of all PingPong Units (move Ping/Pong to lower class)
  - Will hide more of the details from the users creating the Filters

# Future

- Debugging methods for GPUVision
  - We already have basic tools
    - Do you remember?

```
if(DEBUG) vecDF->showResultMtx(rho_1, 1, 1);
```
  - Print section
  - Check pixel value
  - Assert
- Render to multi-texture support
- **All these things for OpenGL novice**