

Evans & Sutherlands RealiMage3000



3D Labs Oxygen GMX



Nvidia Quadro2

Programmable Graphics Hardware

Performance & Paradigms



ATI Rage 3D Pro



ATI Radeon 9700Pro



Nvidia QuadroFX 3000



Nvidia GeForce 6800 Ultra

Hype

ATI X800 Pro

- 12 Pixel Pipes
- 6 Vertex Engines
- 160 M Transistors
- 475 Mhz Core
- 450 Mhz Memory
- 256 bit Memory Interface
- 256 Mb GDDR3 Memory



ATI X800 Pro

ATI Radeon 9800 Pro

- 8 Pixel Pipes
- 4 Vertex Engines
- 107 M Transistors
- 380 Mhz Core
- 340 (?) Mhz Memory
- 256 bit Memory Interface
- 128 Mb Memory



ATI Radeon 9700Pro

Nvidia Quadro FX 3000

- 12 Pixel Pipes
- 6 (?) Vertex Engines
- 130 M Transistors
- 450 Mhz Core
- 425 Mhz Memory Clock
- 256 bit Memory Interface
- 256 Mb Memory



Nvidia QuadroFX 3000

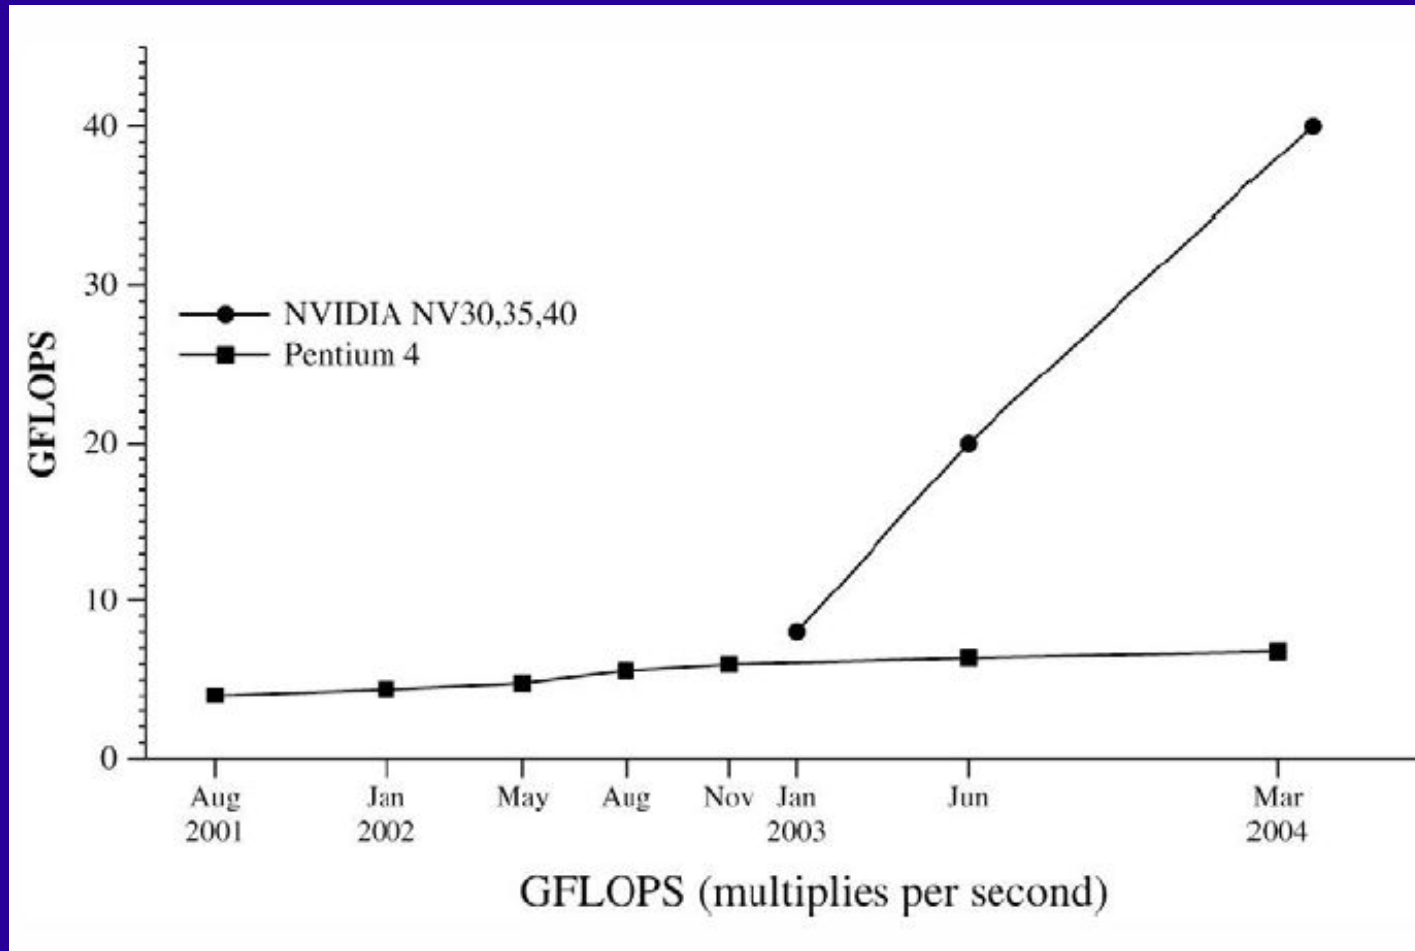
Nvidia 6800 Ultra

- 16 Fragment Pipes
- 6 Vertex Engines
- 222 M Transistors
- 400 Mhz Clock
- 550 Mhz Memory
- 256 bit Memory Interface
- 256 Mb Memory



Nvidia GeForce 6800 Ultra

Pentium vs. Nvidia




Ian Buck: ftp://download.nvidia.com/developer/presentations/2004/GPU_Jackpot/GeForce_6800_Performance.pdf

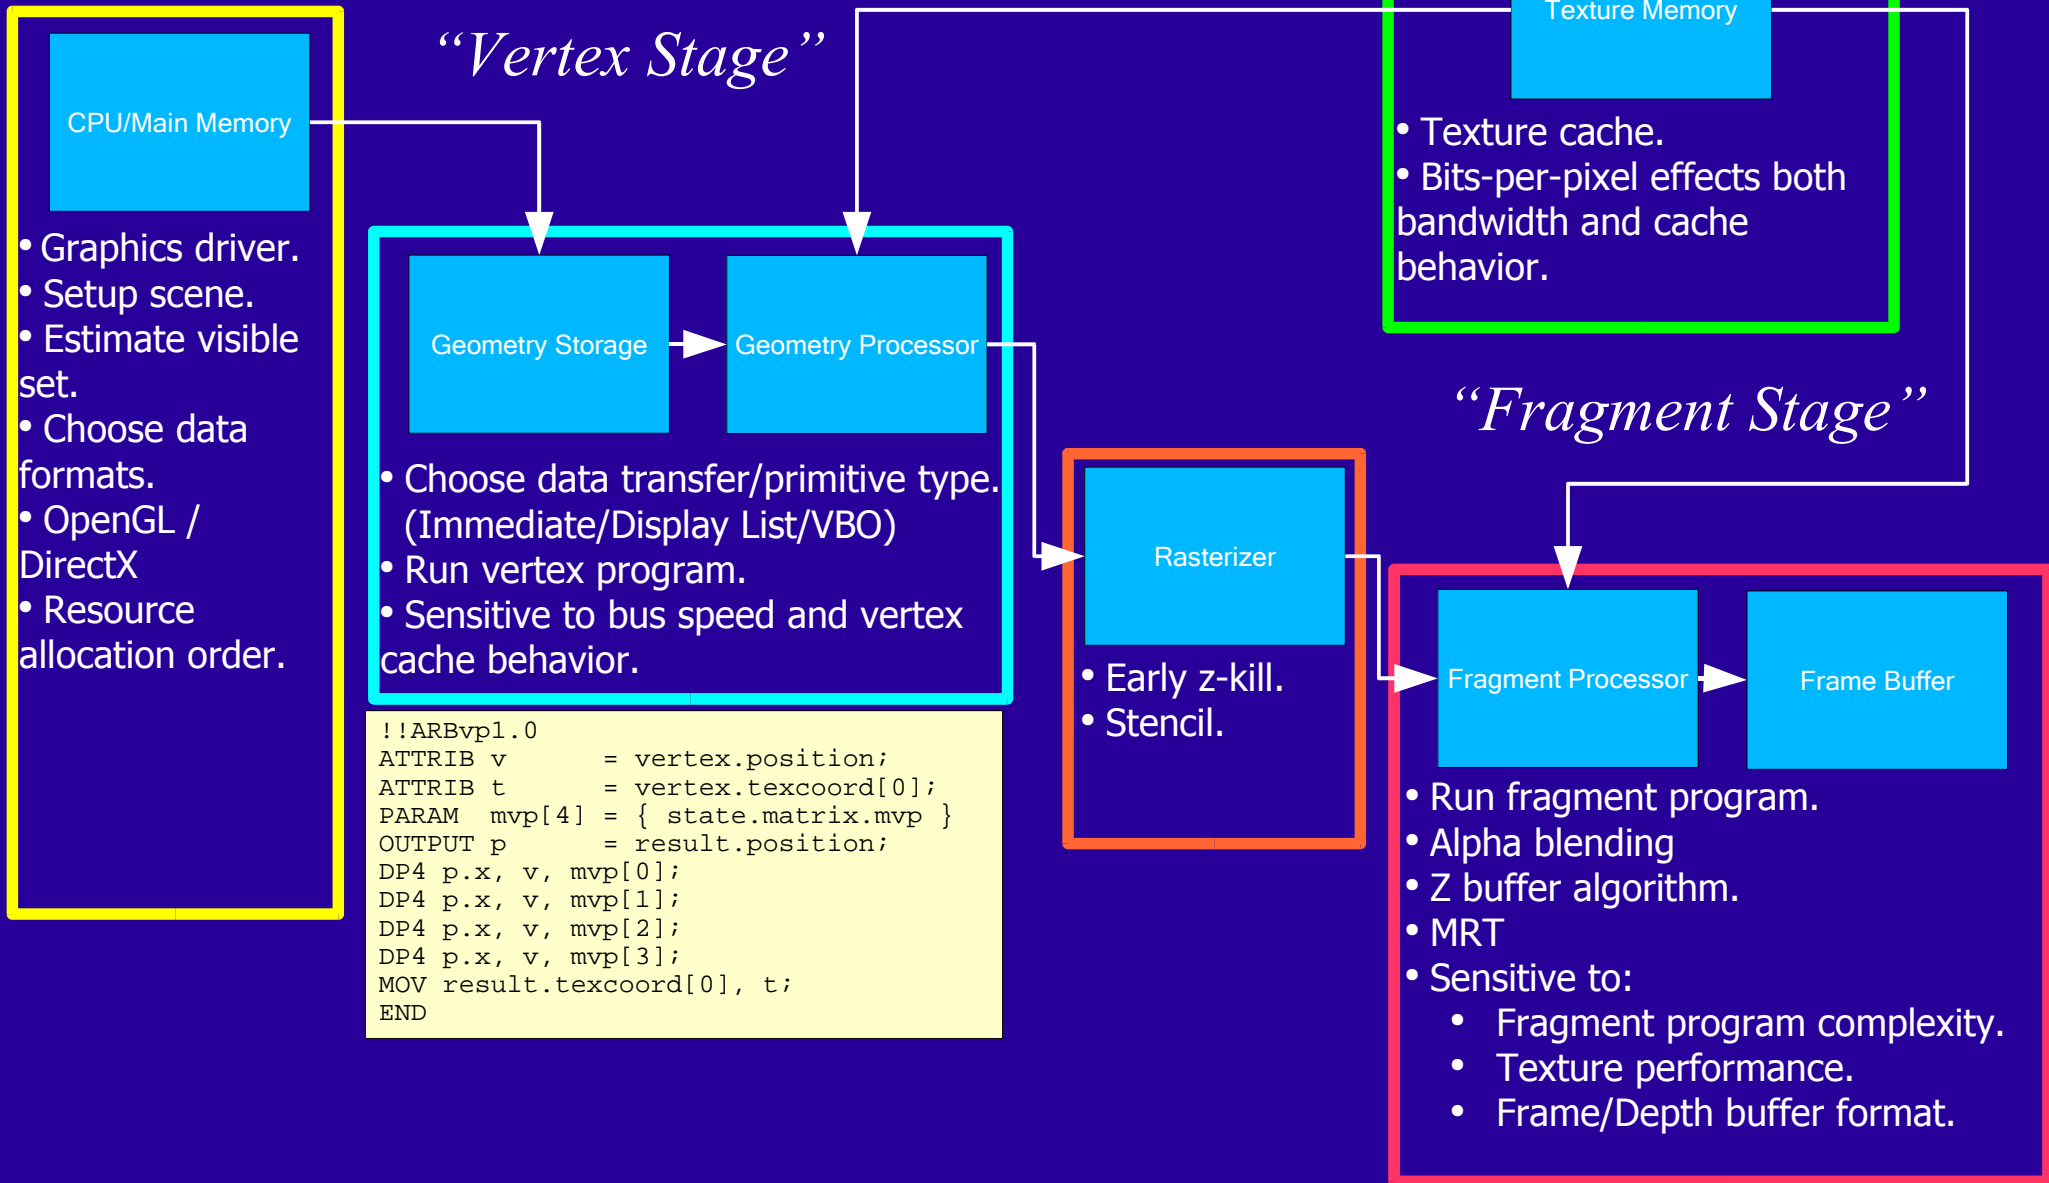
Highly Data Parallel

With 16 fragment pipes, each pipe is responsible for 1/16 of the screen through the entire pass.

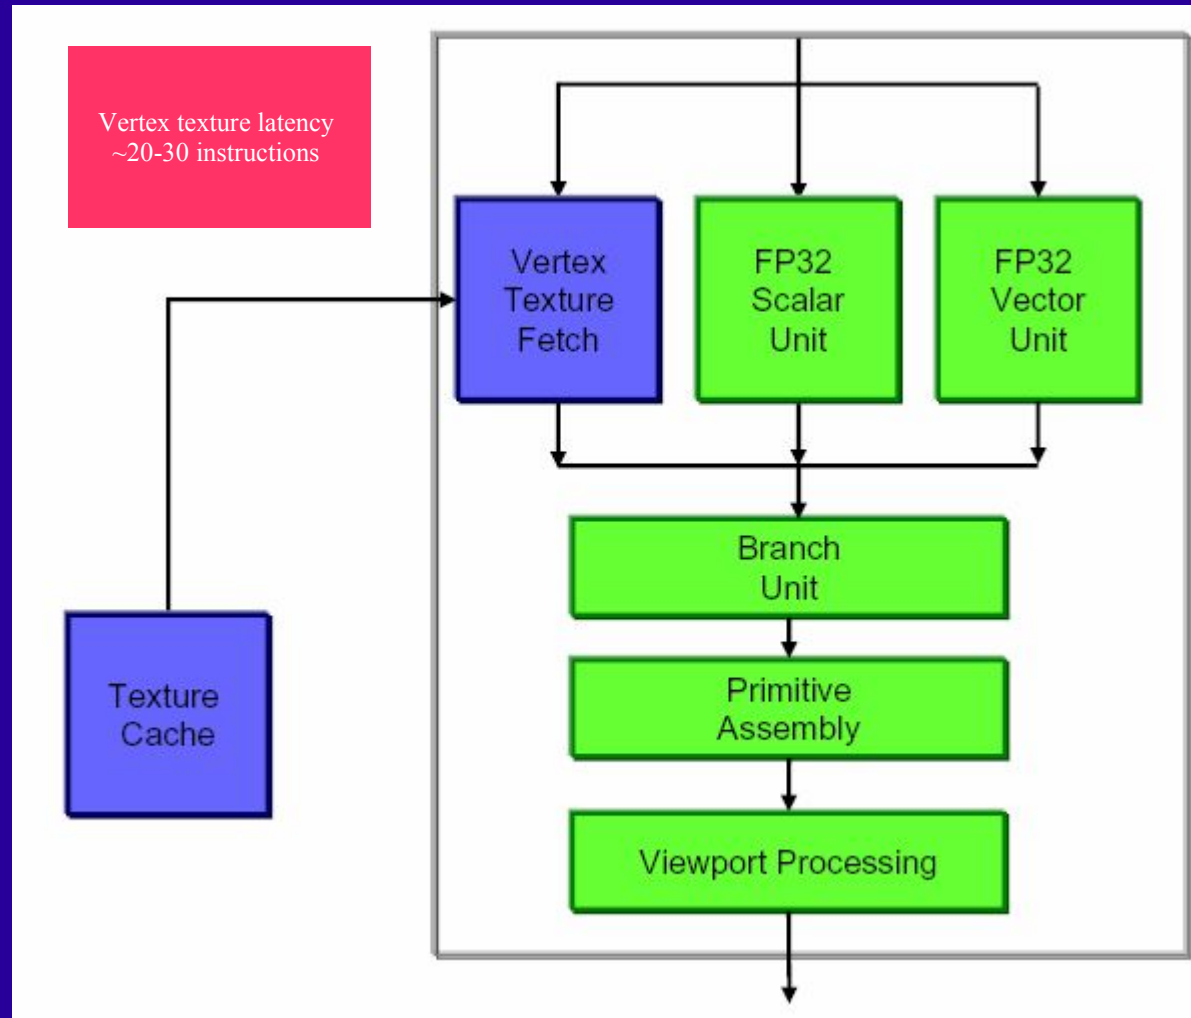
Each fragment is computed individually.



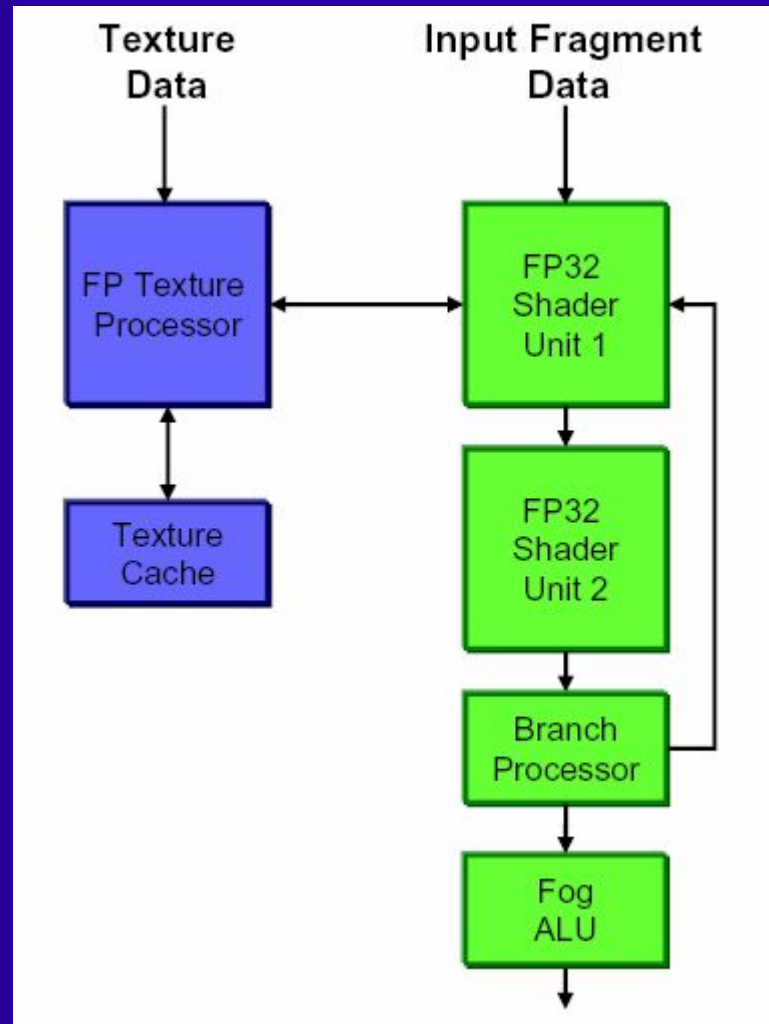
Pipelined Architecture



Vertex Pipeline



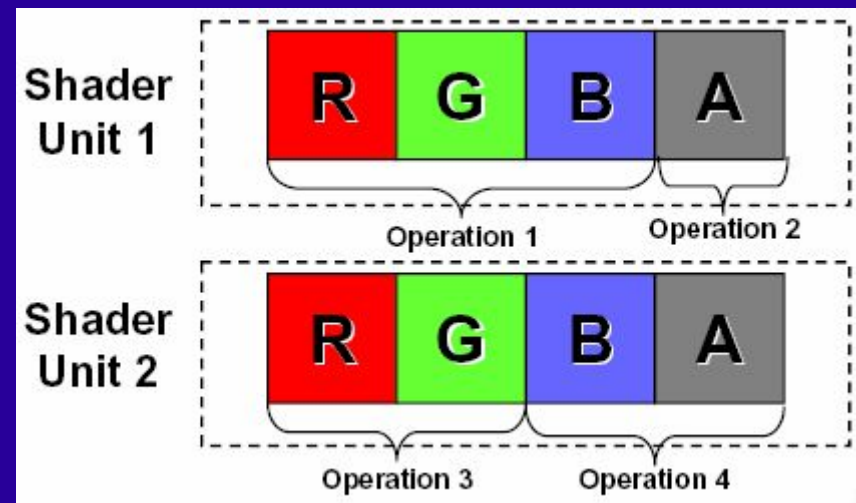
Fragment Pipeline



Instruction Issues

- Superscalar Execution.
Two 2-vector instructions.
One 3-vector instruction +
One scalar instruction
Two math units per pipe.
- Co-Issue.
Two instructions per cycle
per unit.

```
TEMP v, t;  
...  
SIN v.w, v.w;  
MUL v.xyz, v.xyz, t.xyz  
...
```






Measuring Performance

Microbenchmarks:



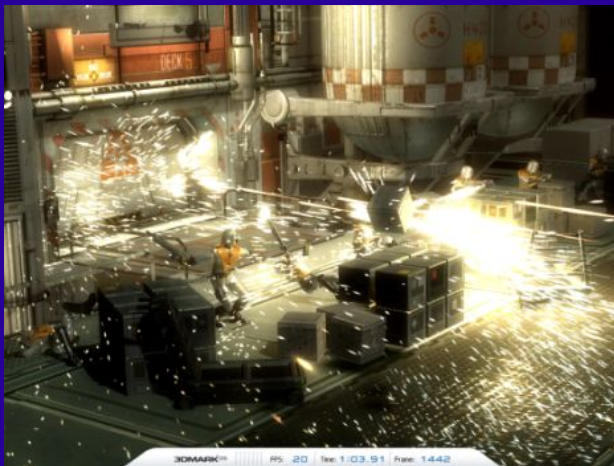
- Instruction Issue
- Texture Cache
- Readback
- Fill




Example output from GPUBench.
(Instruction Issue and Cache Hit Fetch from Quadro FX 3000, Readback from 6800 Ultra PCI-e)

Application benchmarks:

- Useful between application fields?
- Great for your own applications.






3DMark'05 "Return to Proxycon"



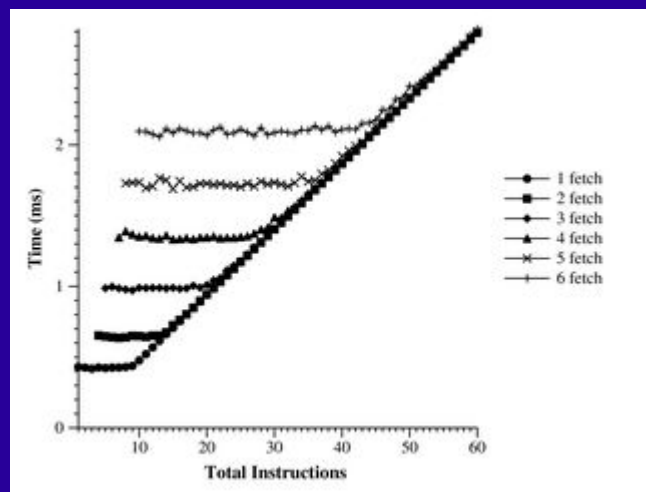


Benchmark screenshots of scientific data (points vs. spheres) and synthetic particle datasets.

Fragment Stage Performance

Cache Hit Fetch



Streaming Access



Nvidia 6800 Ultra

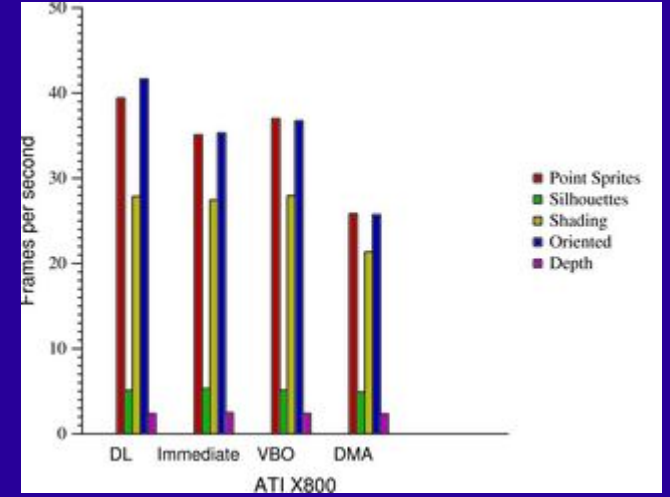

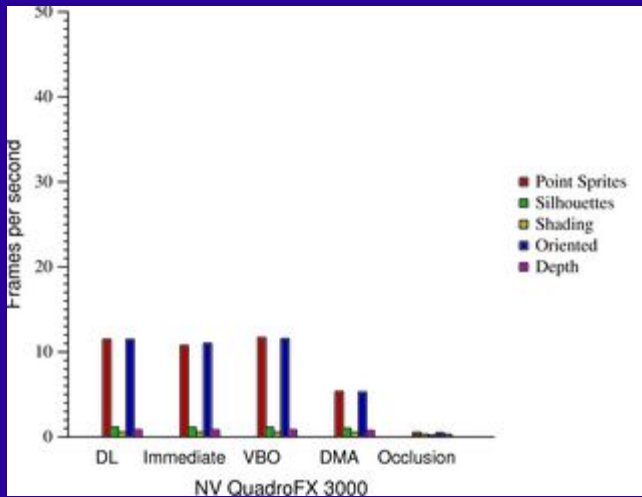
Nvidia QuadroFX 3000

ATI X800

marketing mascots:



Shader Complexity



Nvidia 6800 Ultra



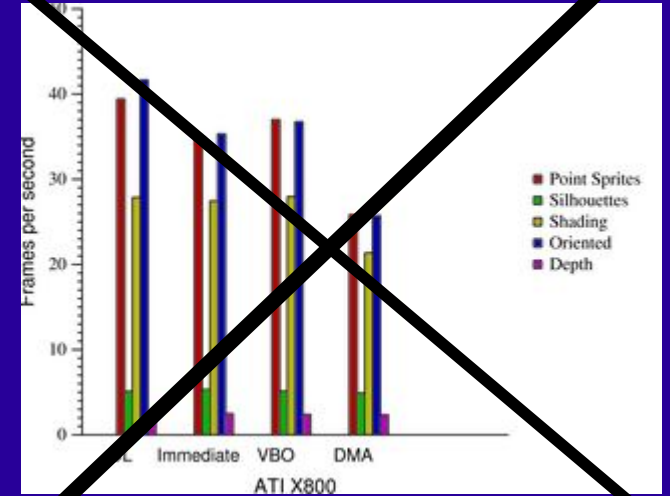
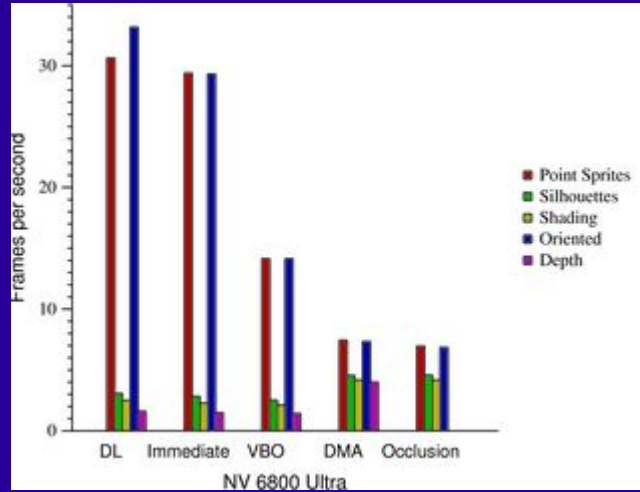
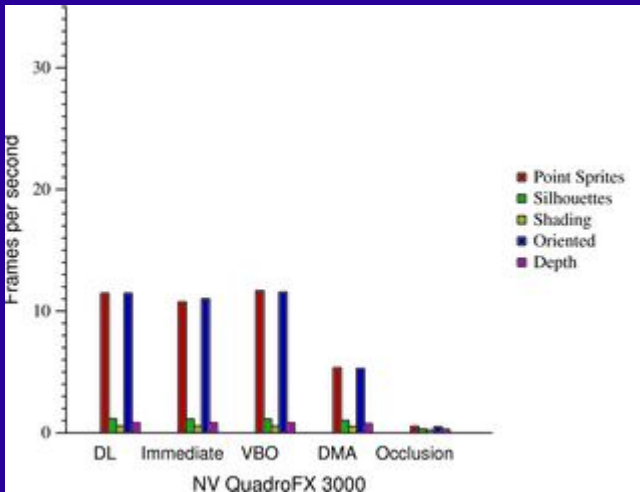
Nvidia QuadroFX 3000



ATI X800



Shader Complexity



Shaders didn't produce correct images on X800.
Occlusion program didn't run.

More debugging necessary.

Nvidia 6800 Ultra



Nvidia QuadroFX 3000



ATI X800



Estimating Performance


```
Pass 1: 13 cycles 492.31 MP/s  
Pass 2: 20 cycles 320.00 MP/s  
Pass 3: 19 cycles 336.84 MP/s  
Pass 4: 3 cycles 2.13 GP/s
```

Average: ~383 MP/s

262144 pixels / 383 MP/s

= 0.0006 seconds * 4 passes

Nice FPS!???



```
$ NVShaderPerf -a NV40 particle_pca_frag_bi2.txt
```

```
-----  
Running performance on file particle_pca_frag_bi2.txt
```

```
----- NV40 -----
```

```
Target: GeForce 6800 Ultra (NV40) :: Unified Compiler: v71.80
```

```
Cycles: 19.00 :: R Regs Used: 4 :: R Regs Max Index (0 based): 3
```

```
Pixel throughput (assuming 1 cycle texture lookup) 336.84 MP/s
```

```
Pass 1: 13 cycles 492.31 MP/s
Pass 2: 20 cycles 320.00 MP/s
Pass 3: 19 cycles 336.84 MP/s
Pass 4: 3 cycles 2.13 GP/s
```


Average: ~383 MP/s

262144 pixels / 383 MP/s

= 0.0006 seconds * 4 passes

Nice FPS!???

- First pass induces considerable overdraw.
- One cycle texture lookup?



```
$ NVShaderPerf -a NV40 particle_pca_frag_bi2.txt
```

```
-----  
Running performance on file particle_pca_frag_bi2.txt
```


```
----- NV40 -----
```

```
Target: GeForce 6800 Ultra (NV40) :: Unified Compiler: v71.80
```

```
Cycles: 19.00 :: R Regs Used: 4 :: R Regs Max Index (0 based): 3
```


```
Pixel throughput (assuming 1 cycle texture lookup) 336.84 MP/s
```

Application Paradigms



Mesh based geometry & effects.


- “Conventional”
- Geometry is projected & rasterized.
- Effects & Shading applied in fragment programs.
- Vertex & texture cache.



Rasterized geometry.

- Volume rendering
- Deferred shading.
- Image is produced upon rasterized geometry.


Application Paradigms




Fragment Based Applications

- Move input directly to the fragment shader on screen filling triangle/quads.
- Ray tracer
- GP-GPU.

Good References.



One pass volume rendering



PS3 Ray Tracer

Questions?

http://developer.nvidia.com/object/gpu_jackpot_2004_presentations.html