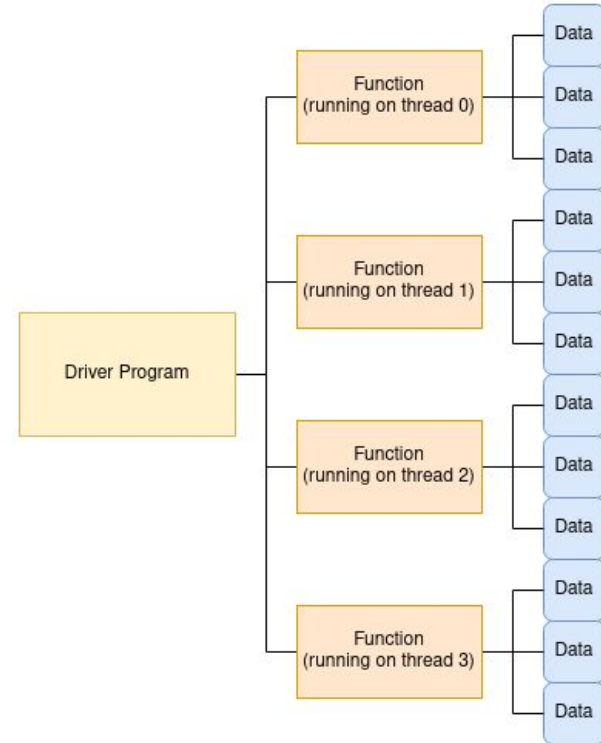


GPU Programming and the Graphics Pipeline


CSE 113, Devon McKee

Single-Program-Multiple-Data (SPMD) Parallelism

```
File: thread.cpp
1  #include <iostream>
2  #include <thread>
3
4  #define NUM_THREADS 4
5  #define SIZE 1024
6
7  using std::thread;
8
9  void worker(int *data, int start, int offset) {
10     for (int i = start; i < start + offset; i++)
11         data[i] = i;
12 }
13
14 int main() {
15     int data[SIZE] = {0};
16     thread threads[4];
17     int chunk_size = SIZE / NUM_THREADS;
18     for (int i = 0; i < NUM_THREADS; i++)
19         threads[i] = thread(worker, data, chunk_size * i, chunk_size);
20     for (int i = 0; i < NUM_THREADS; i++)
21         threads[i].join();
22     return 0;
23 }
```




Limits of CPU SPMD Parallelism



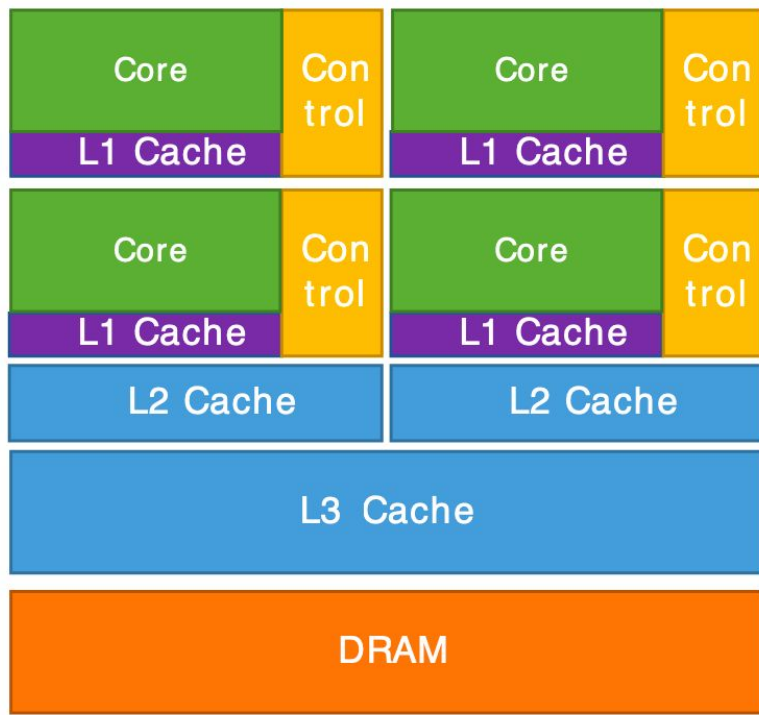
Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
New plot and data collected for 2010-2015 by K. Rupp

CPU vs. GPU: Cores

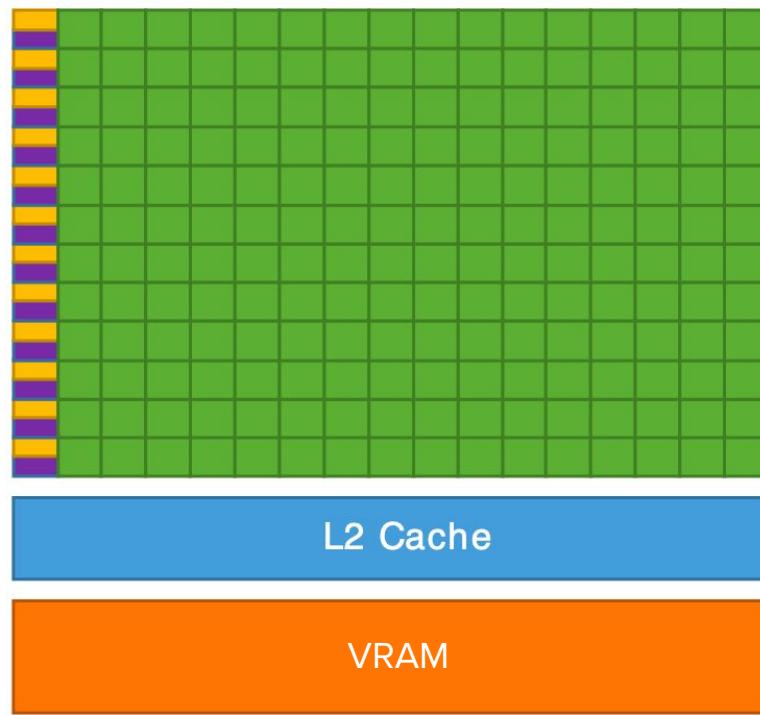
NVIDIA RTX 4090: 16384 cores!



CPU vs. GPU: Architecture

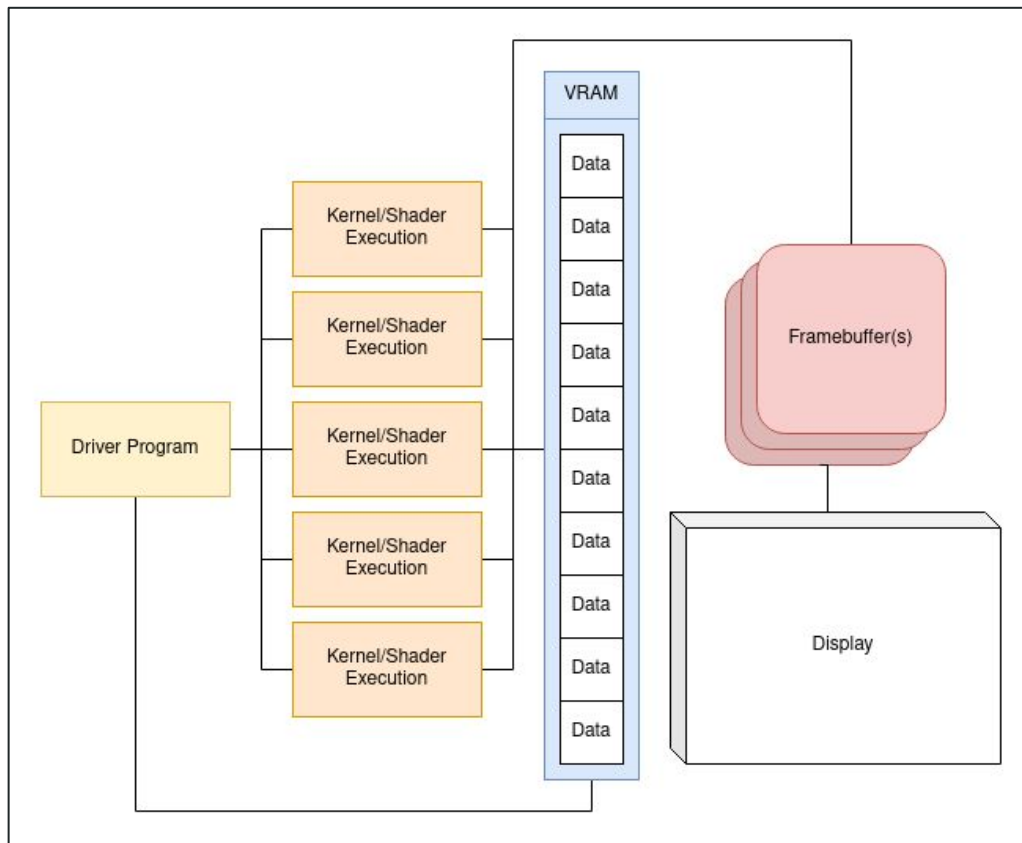


CPU




GPU

GPU Programming




GPU Programming: Graphics vs. Compute

Typical Graphics Program Flow



Typical Compute Program Flow



CUDA: Vector Add

```
#include <stdio.h>
#include <assert.h>

__global__ void vectorAdd(int* a, int* b, int* c, int n) {
    int idx = blockIdx.x * blockDim.x + threadIdx.x;
    c[idx] = a[idx] + b[idx];
}

int main() {
    int n = 1<<24;
    int val = 1024;

    int *h_a, *h_b, *h_c; // Pointers to host memory
    int *d_a, *d_b, *d_c; // Pointers to device memory

    // Allocating host memory
    h_a = (int*)malloc(sizeof(int) * n);
    h_b = (int*)malloc(sizeof(int) * n);
    h_c = (int*)malloc(sizeof(int) * n);

    // Allocating device memory
    cudaMalloc(&d_a, sizeof(int) * n);
    cudaMalloc(&d_b, sizeof(int) * n);
    cudaMalloc(&d_c, sizeof(int) * n);

    // Setting up host memory
    for (int i = 0; i < n; i++) {
        h_a[i] = i;
        h_b[i] = 1024 - i;
    }

    // Copying memory from host to device
    cudaMemcpy(d_a, h_a, sizeof(int) * n, cudaMemcpyHostToDevice);
    cudaMemcpy(d_b, h_b, sizeof(int) * n, cudaMemcpyHostToDevice);

    // Executing kernel
    int blockSize = 1024; // Number of threads in each thread block
    int gridSize = (int)ceil((float)n / blockSize); // Number of thread blocks in grid
    vectorAdd<<<gridSize, blockSize>>>(d_a, d_b, d_c, n);

    // Copying results from device to host
    cudaMemcpy(h_c, d_c, sizeof(int) * n, cudaMemcpyDeviceToHost);

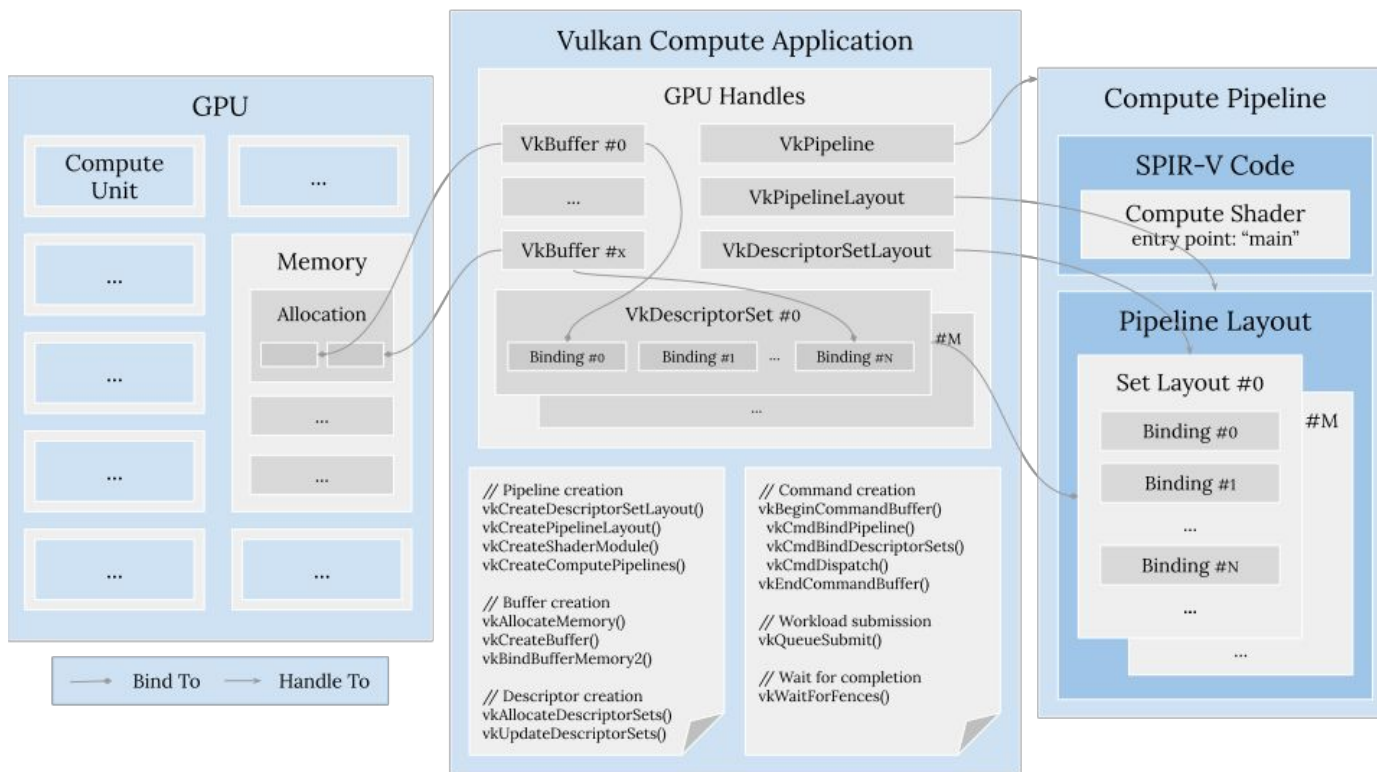
    // Checking results
    for (int i = 0; i < n; i++)
        assert(h_c[i] == val);

    // Deallocating device memory
    cudaFree(d_a);
    cudaFree(d_b);
    cudaFree(d_c);

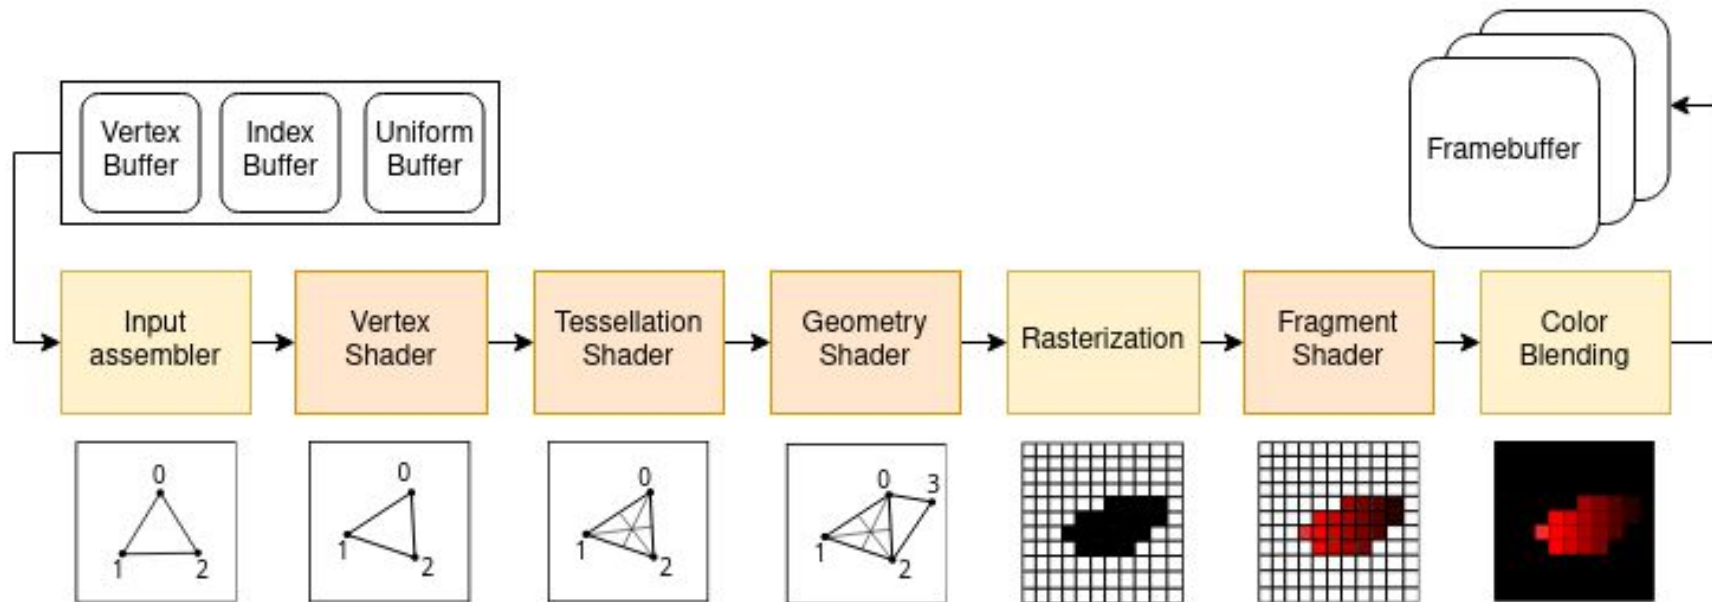
    // Deallocating host memory
    free(h_a);
    free(h_b);
    free(h_c);

    return 0;
}
```


Vulkan




The Graphics Pipeline




Fragment (Pixel) Shader

```
1  precision mediump float;
2
3  uniform vec2 u_resolution;
4
5  void main() {
6      vec2 st = gl_FragCoord.xy/u_resolution.xy;
7      vec3 color = vec3(st.x, st.y, st.x + st.y / 2.0);
8      gl_FragColor = vec4(color, 1.0);
9  }
```



Fragment (Pixel) Shader


```
1  precision mediump float;
2
3  void main() {
4      if (mod(floor(gl_FragCoord.x / 40.0), 2.0) == 0.0)
5          if (mod(floor(gl_FragCoord.y / 40.0), 2.0) == 0.0)
6              gl_FragColor = vec4(0, 0, 0, 1);
7          else
8              gl_FragColor = vec4(1, 1, 1, 1);
9      else
10         if (mod(floor(gl_FragCoord.y / 40.0), 2.0) == 0.0)
11             gl_FragColor = vec4(1, 1, 1, 1);
12         else
13             gl_FragColor = vec4(0, 0, 0, 1);
14     }
```




Fragment (Pixel) Shader

File: phong.glsl

```
1 precision mediump float;
2 varying vec3 normalInterp; // Surface normal
3 varying vec3 vertPos; // Vertex position
4 uniform int mode; // Rendering mode
5 uniform float Ka; // Ambient reflection coefficient
6 uniform float Kd; // Diffuse reflection coefficient
7 uniform float Ks; // Specular reflection coefficient
8 uniform float shininessVal; // Shininess
9 // Material color
10 uniform vec3 ambientColor;
11 uniform vec3 diffuseColor;
12 uniform vec3 specularColor;
13 uniform vec3 lightPos; // Light position
14
15 void main() {
16     vec3 N = normalize(normalInterp);
17     vec3 L = normalize(lightPos - vertPos);
18
19     // Lambert's cosine law
20     float lambertian = max(dot(N, L), 0.0);
21     float specular = 0.0;
22     if(lambertian > 0.0) {
23         vec3 R = reflect(-L, N); // Reflected light vector
24         vec3 V = normalize(-vertPos); // Vector to viewer
25         // Compute the specular term
26         float specAngle = max(dot(R, V), 0.0);
27         specular = pow(specAngle, shininessVal);
28     }
29     gl_FragColor = vec4(Ka * ambientColor +
30         Kd * lambertian * diffuseColor +
31         Ks * specular * specularColor, 1.0);
32 }
```



Viewer


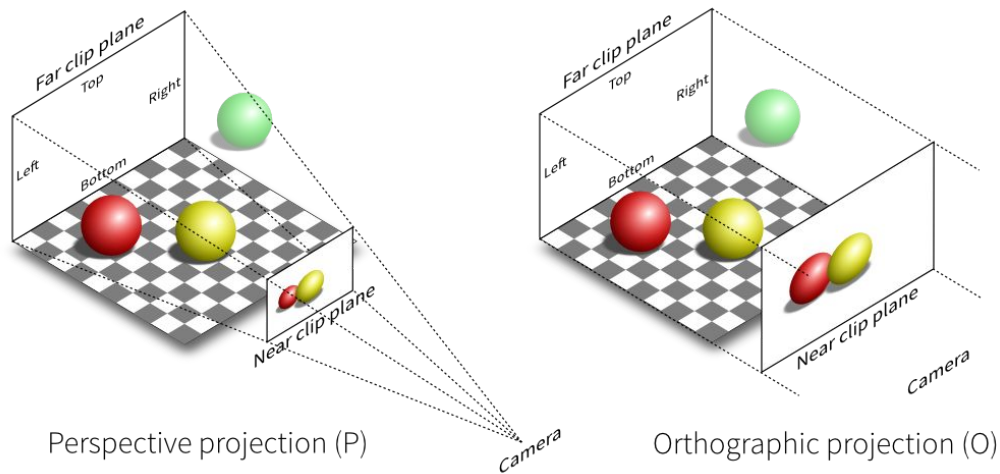



Vertex Shader

```
File: vertex.glsl
1  #version 410 core
2  in vec2 point;
3  in vec3 color;
4  out vec4 vColor;
5  uniform float radAng = 0;
6  uniform float time = 0;
7
8  vec2 Rotate2D(vec2 v) {
9      float c = cos(radAng), s = sin(radAng);
10     return vec2(c*v.x-s*v.y, s*v.x+c*v.y);
11 }
12 void main() {
13     vec2 point_r = Rotate2D(point);
14     point_r.x *= (sin(time));
15     point_r.y *= (sin(time));
16     vColor = vec4(sin(color.x + point.x + time),
17                 cos(color.y + point.y + time - 5),
18                 (cos(color.x + time)+sin(color.y + time))/2,
19                 1);
20 };
21 gl_Position = vec4(point_r, 0, 1);
22 }
```



3D Graphics



Combining Graphics and Compute



GPU Resources

- Shaders:
 - <https://thebookofshaders.com/>
- Graphics/OpenGL:
 - <https://learnopengl.com/>
 - <http://www.opengl-tutorial.org/>
- CUDA:
 - <https://developer.nvidia.com/blog/even-easier-introduction-cuda/>
 - <https://cuda-tutorial.readthedocs.io>
- GPU Architecture
 - <https://bit.ly/appendix-c>