VGP351 – Week 2

- Agenda:
  - Getting data to the GPU
  - Types of primitives
  - Transformations
    - Modeling
    - Viewing
    - Projection

13-April-2010

Graphics Pipeline

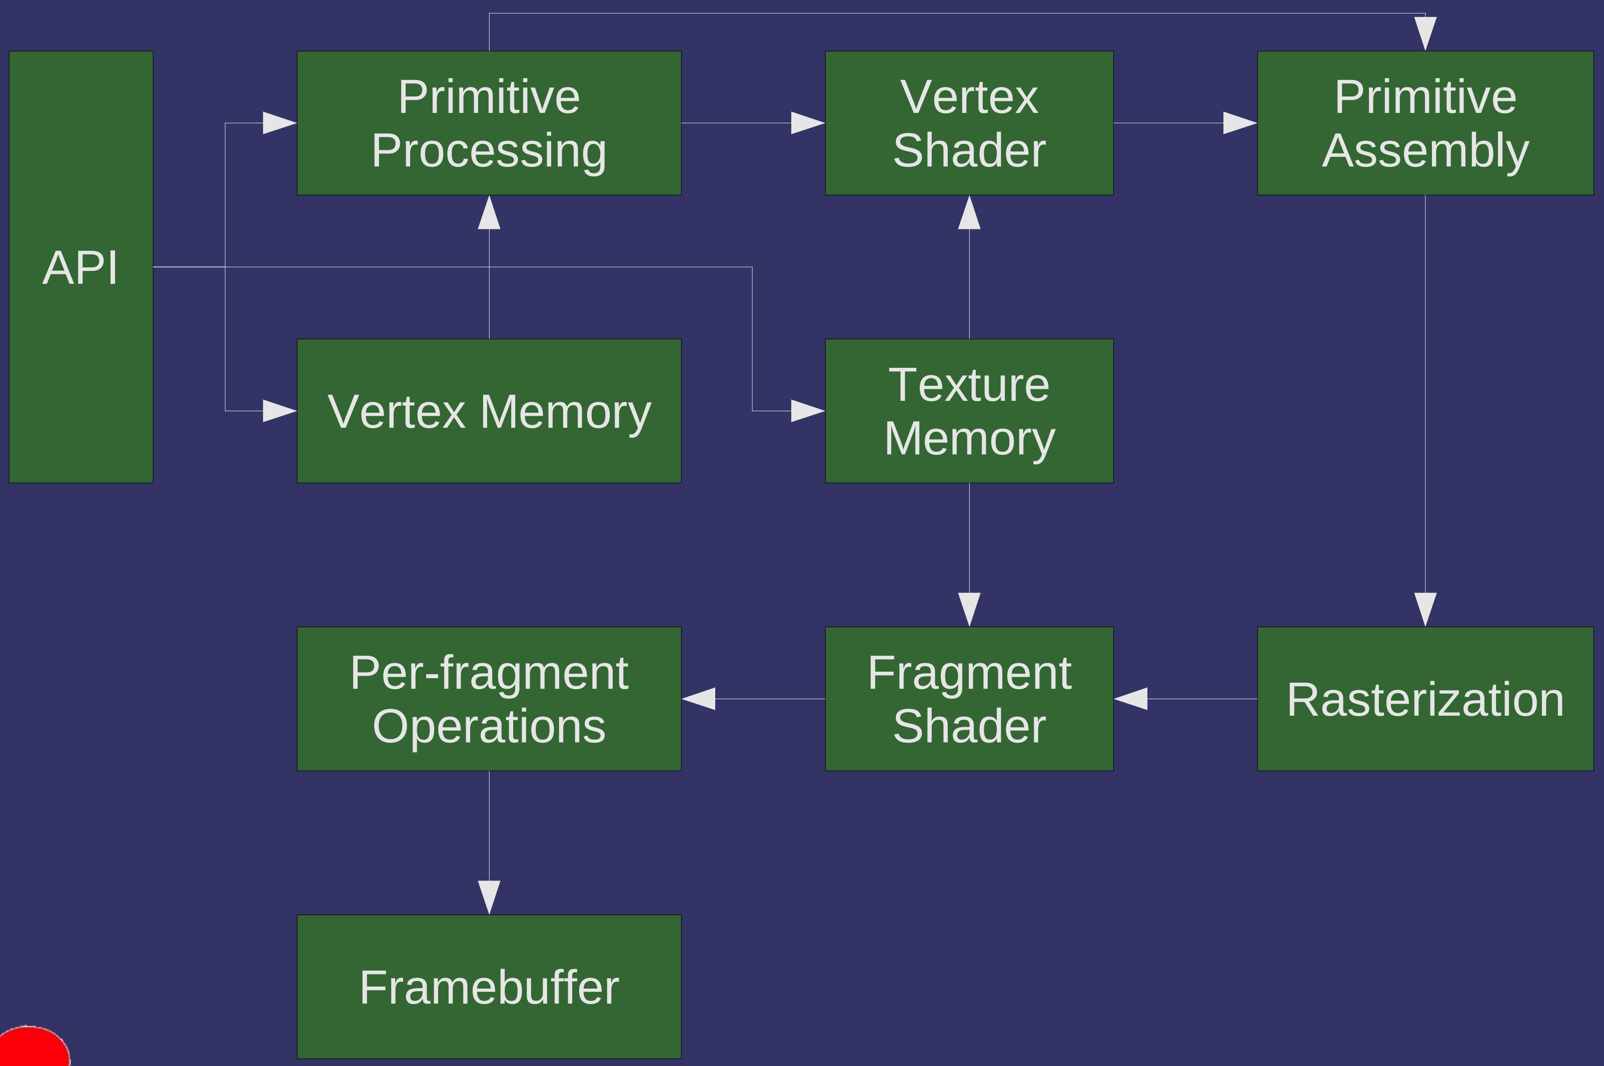

Graphics Pipeline

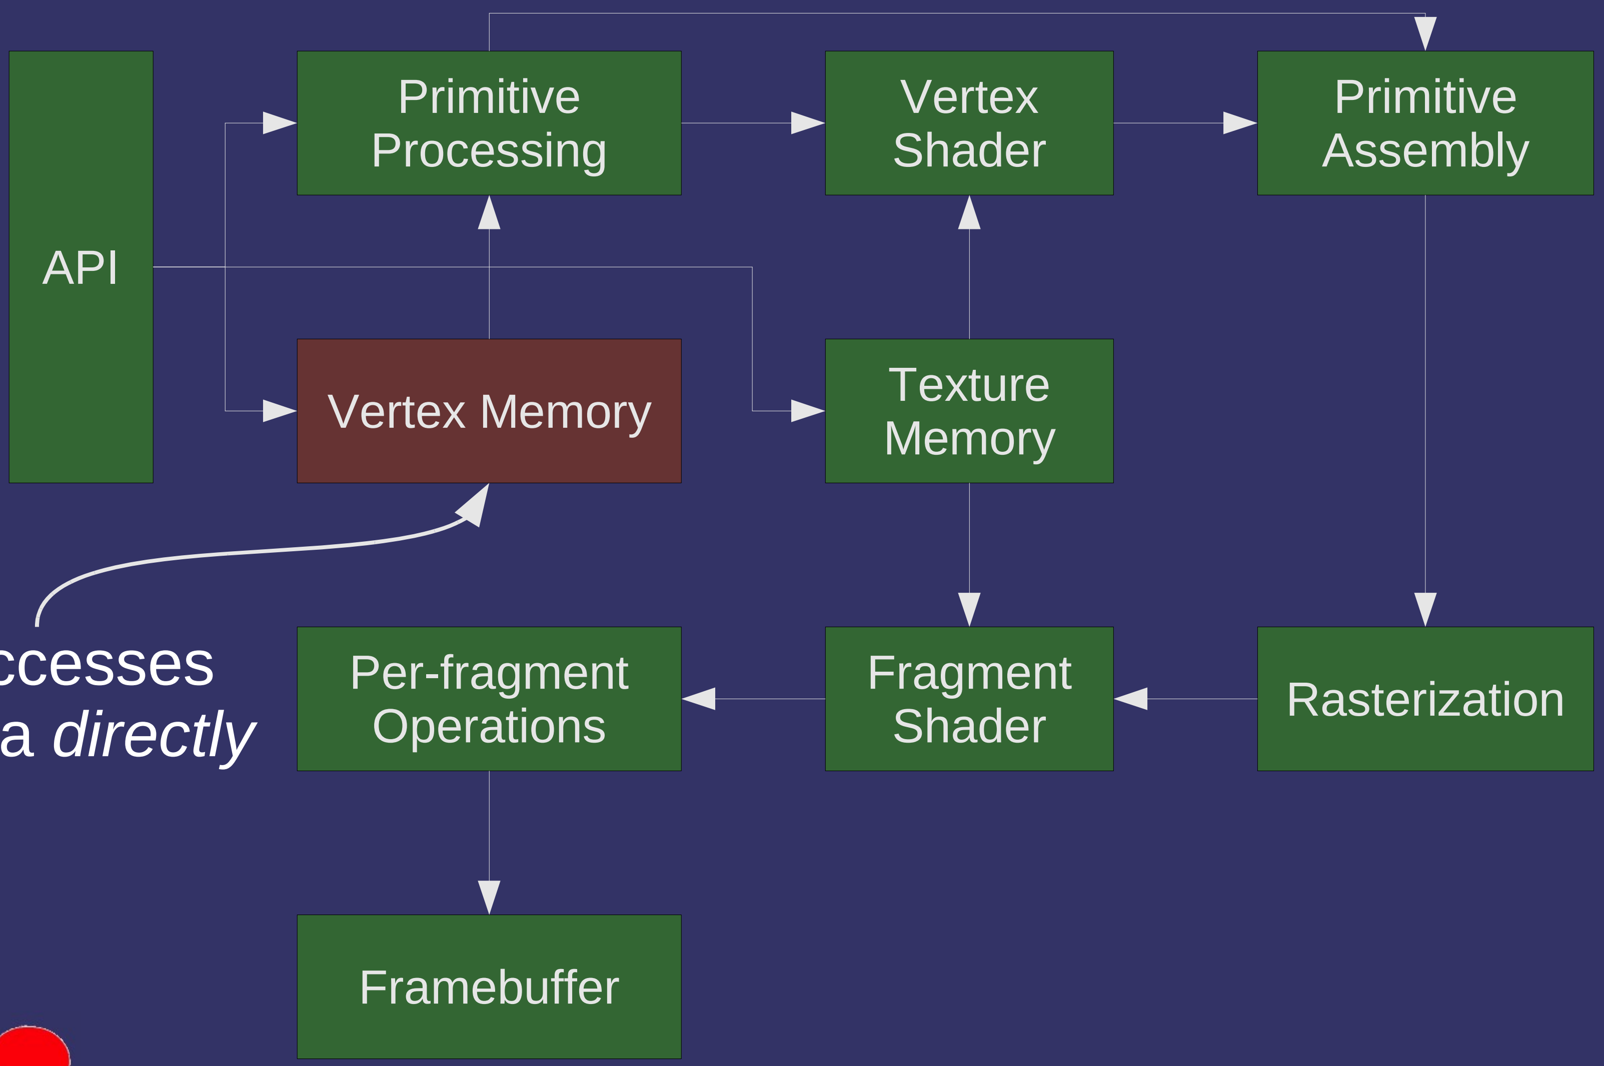

GPU accesses this data *directly*

13-April-2010

Memory Architecture

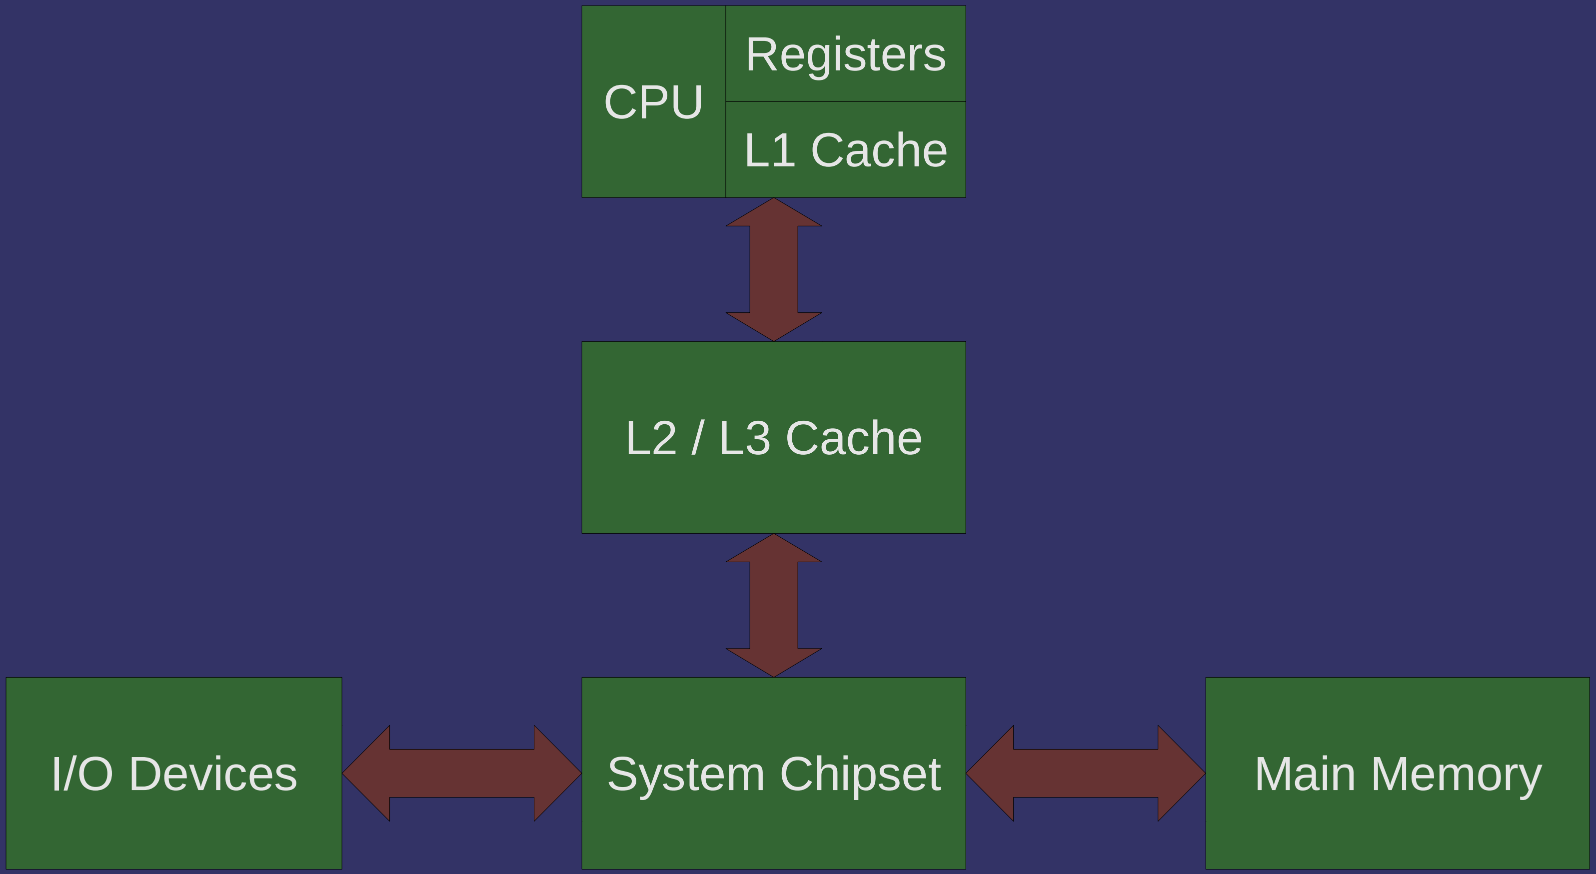

13-April-2010

Memory Architecture

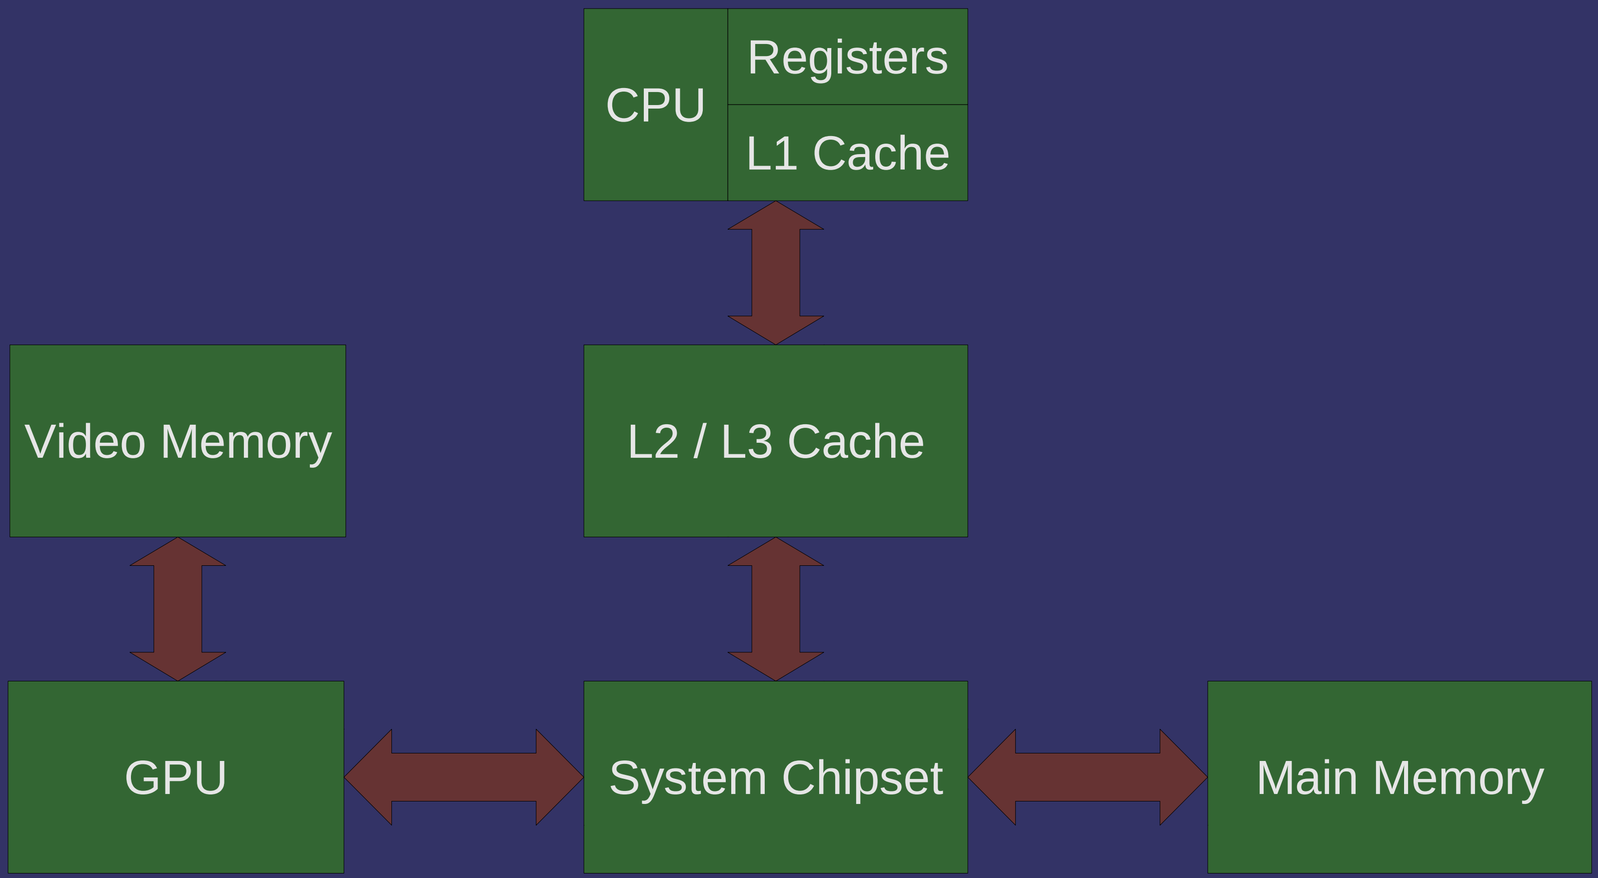

13-April-2010

Unified Memory Architecture

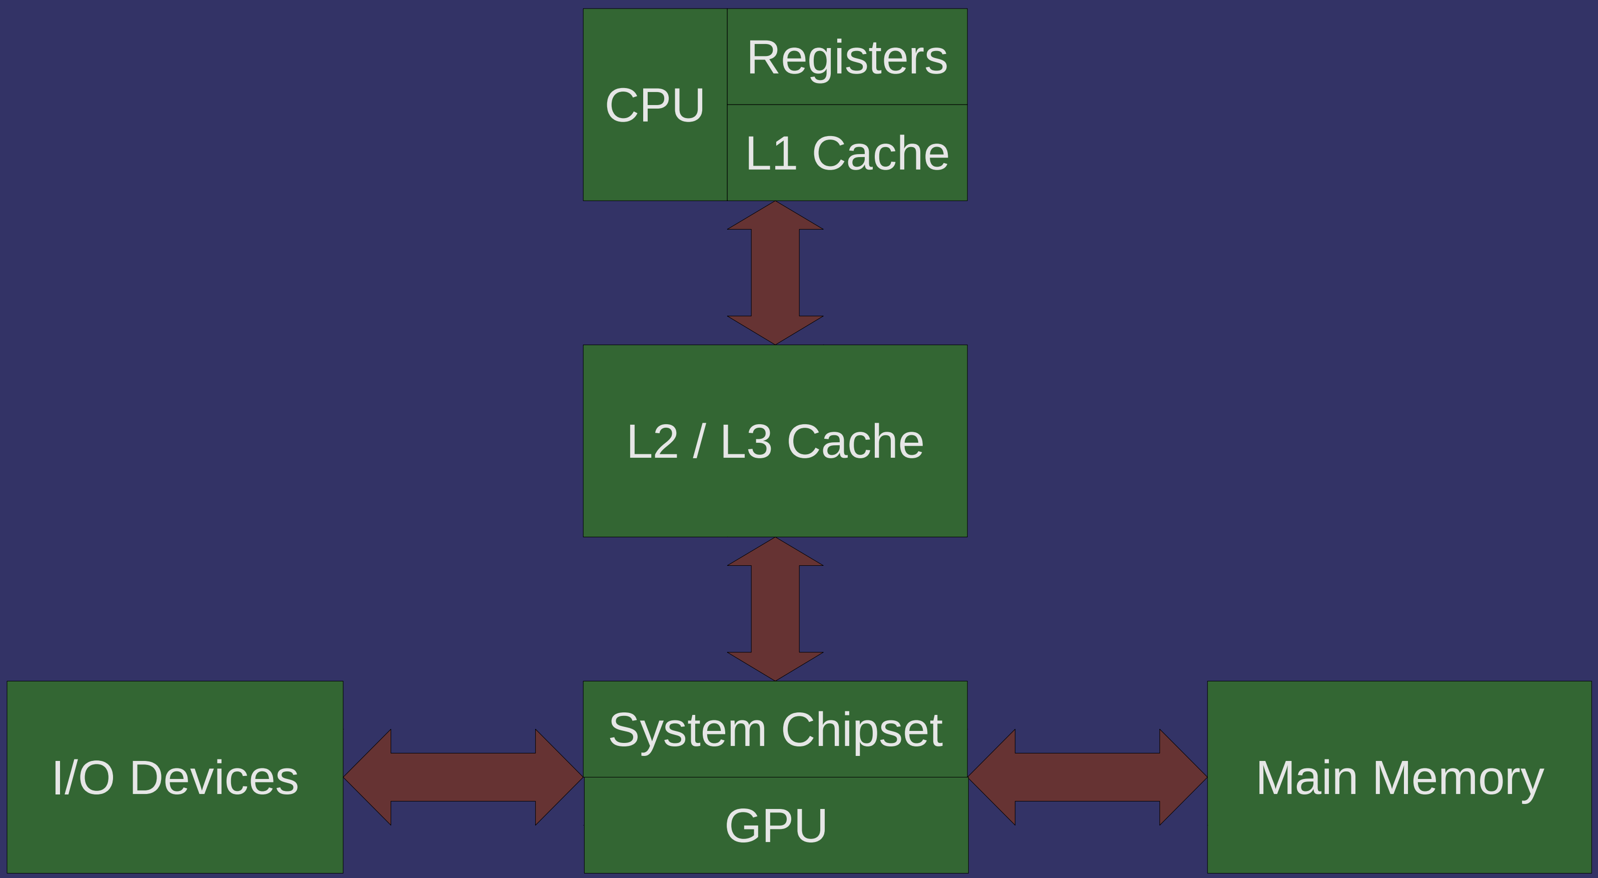

Memory Map

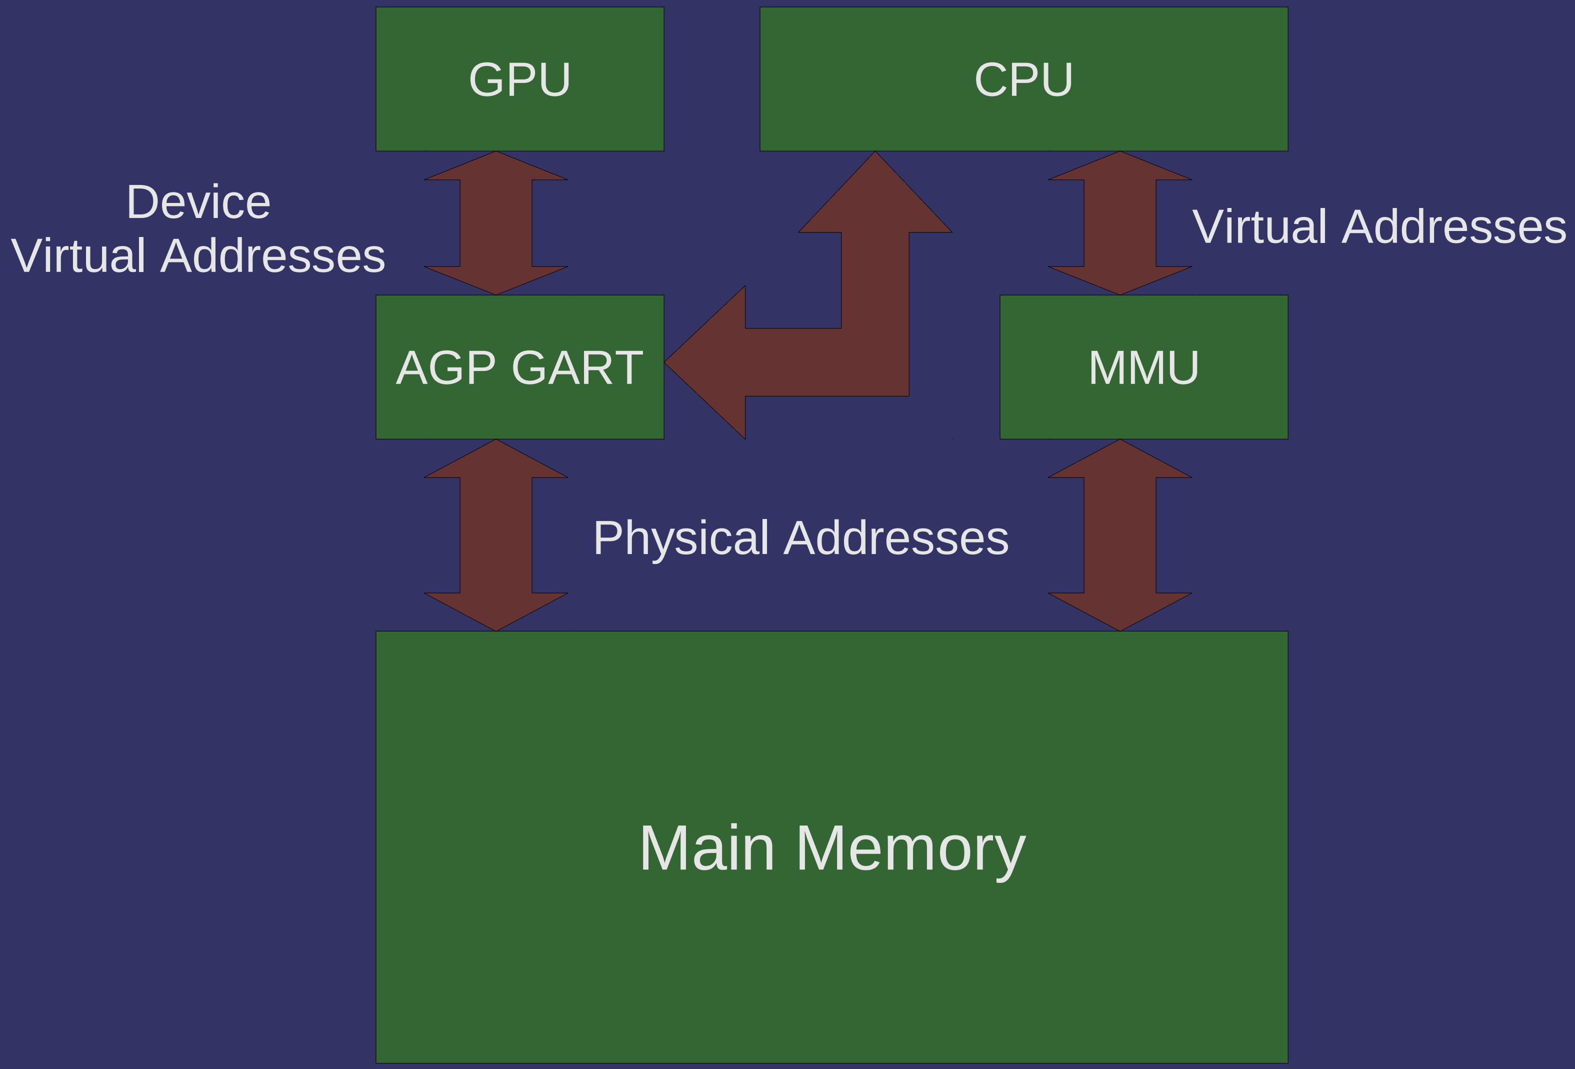

13-April-2010

Memory Map

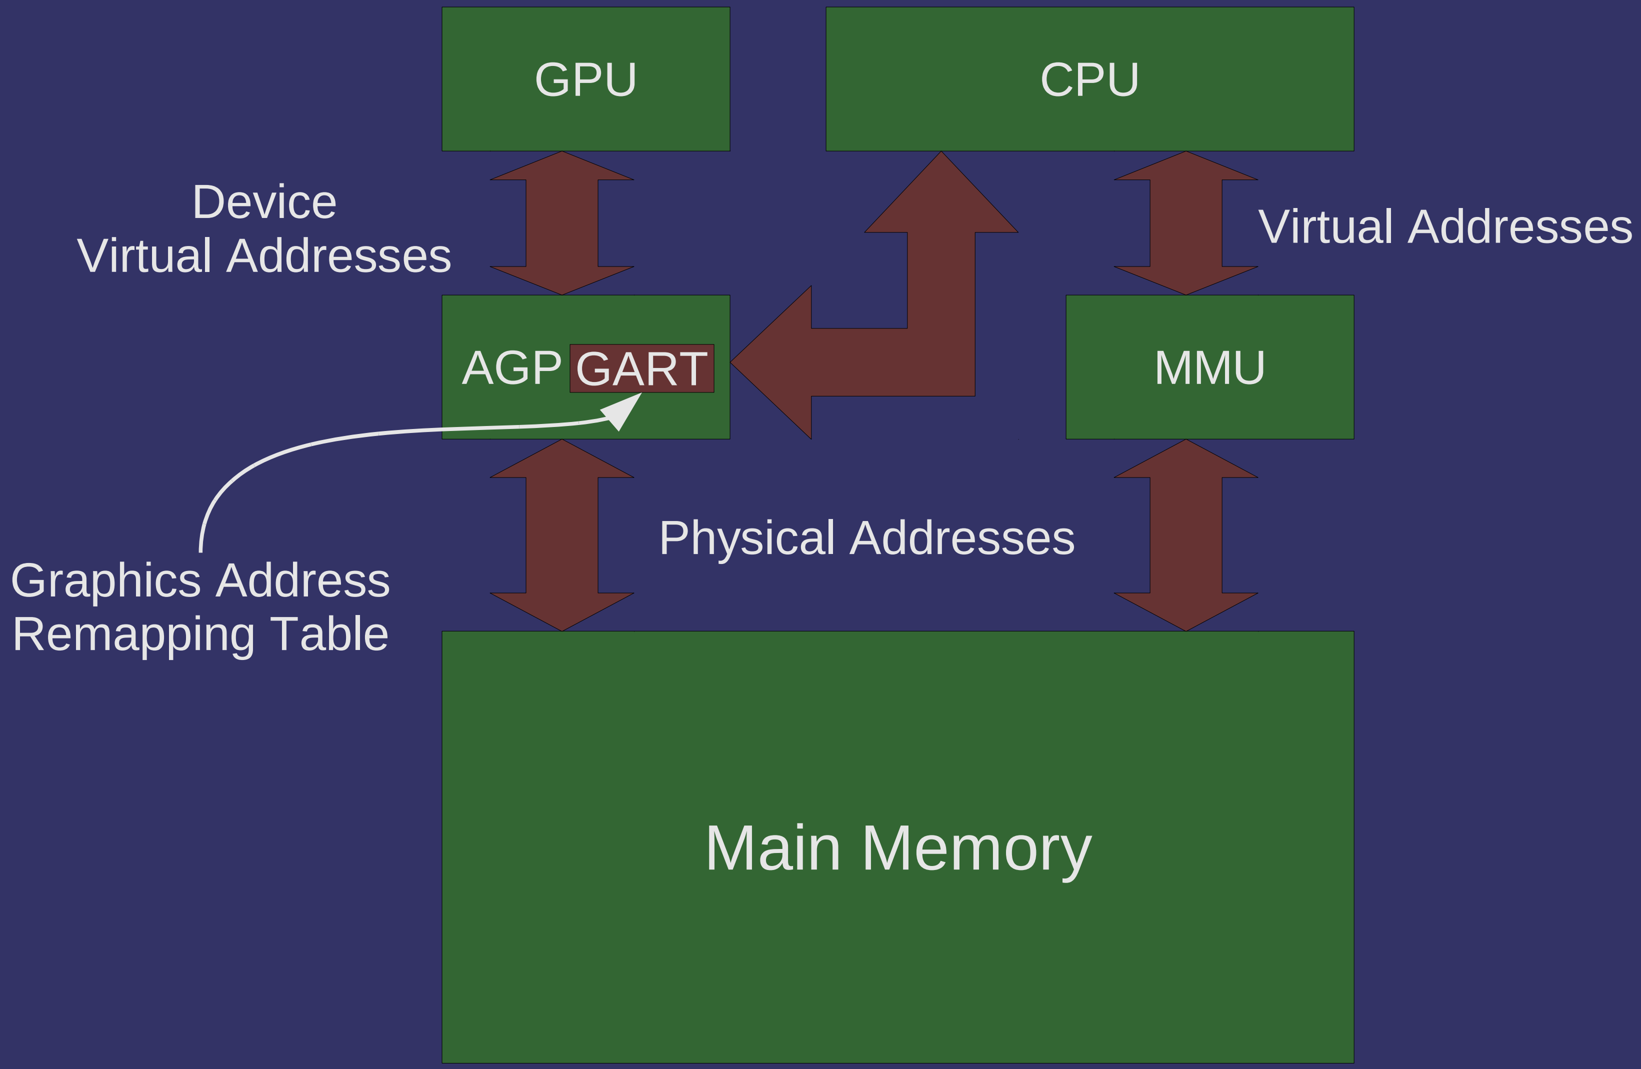
GPU
CPU
Device
Virtual Addresses
Virtual Addresses
AGP GART
MMU
Physical Addresses
Graphics Address
Remapping Table
Main Memory

Vertex Memory

- Practically, the GPU can only access:
  - Memory physically on the graphics card
  - Memory mapped in the GART
- To get GART or card memory, we have to allocate it using the driver
  - Only the driver knows what *kind* of memory to use
  - ...but we have to give it some hints

Vertex Memory

- In OpenGL this memory is called *buffer object*
  - It is used somewhat like a file:
    - Bulk I/O via accessor routines
    - Direct mapping and access via a pointer

Buffer Objects

- Generate "names" for the buffer objects:

  `glGenBuffers(GLsizei num, GLuint *names);`

- "Bind" a buffer for use:

  `glBindBuffer(GLenum target, GLuint name);`

  - `target` selects which buffer we're talking about
    - `GL_ARRAY_BUFFER` is used for vertex data
    - `GL_ELEMENT_ARRAY_BUFFER` is used for vertex indices
      - More on that *later*...
    - There are other targets we'll cover later in the term
  - Binding creates the object, but it still has no storage

Buffer Objects

- Storage is created and *optionally* initialized with:

```
void glBufferData(GLenum target,
    GLsizeiptr size, const GLvoid *data,
    GLenum usage);
```

  - `usage` tells the GL how the app will utilize the buffer

- Storage is updated with:

```
void glBufferSubData(GLenum target,
    GLintptr offset, GLsizeiptr size,
    const GLvoid *data);
```

Buffer Object Usage Hints

- Usage conveys information along two axes:
  - Data “frequency”:
    - Stream – data is specified once and used a few times
    - Static – data is specified once and used many times
    - Dynamic – data is specified and used many times
  - Data “usage”:
    - Draw – data used as source for drawing
    - Read – data copied from GL and read back to client
    - Copy – data copied from GL and used as source for drawing
  - Combine these to create the enums (e.g., `GL_STATIC_DRAW`)

Buffer Object Usage Hints

Buffer Object Usage Hints

`GL_STATIC_DRAW`

Buffer Object Usage Hints

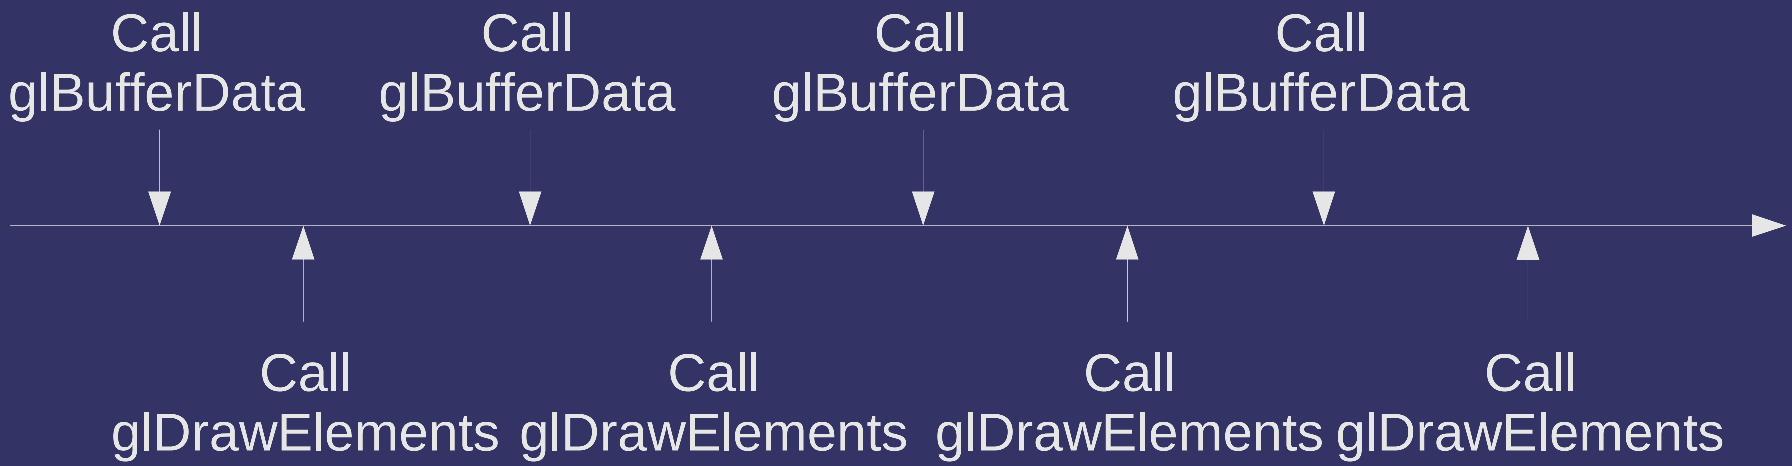

Buffer Object Usage Hints

GL_STREAM_DRAW

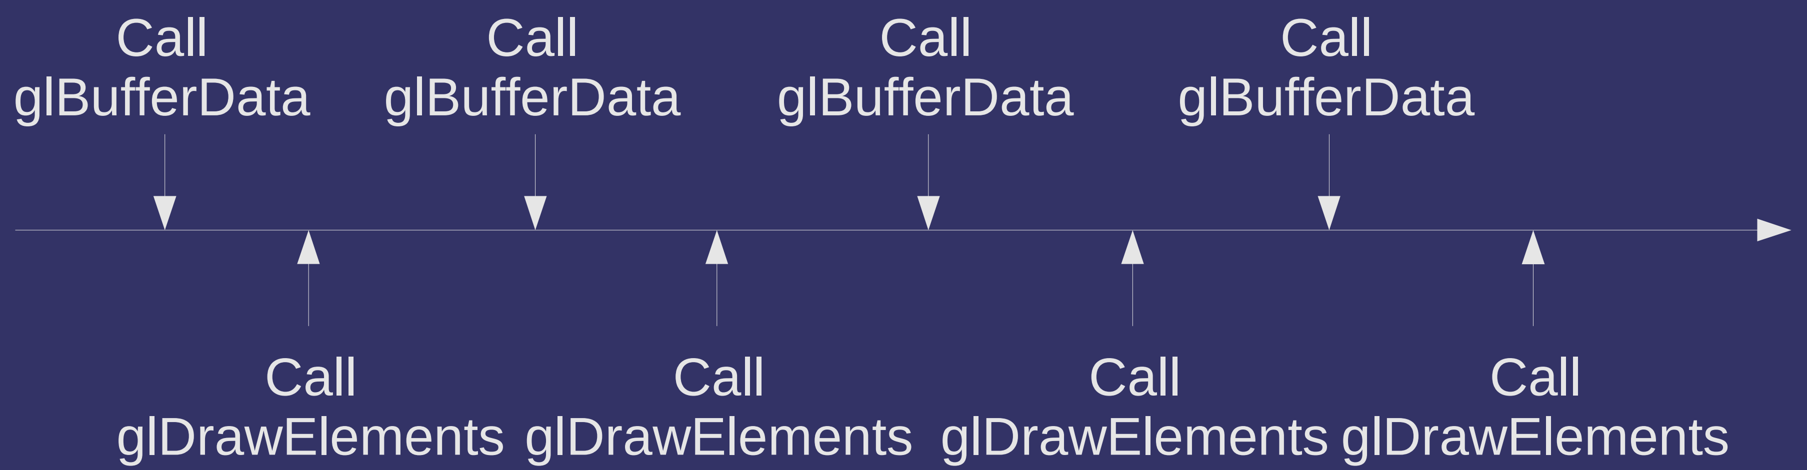

Buffer Object Usage Hints

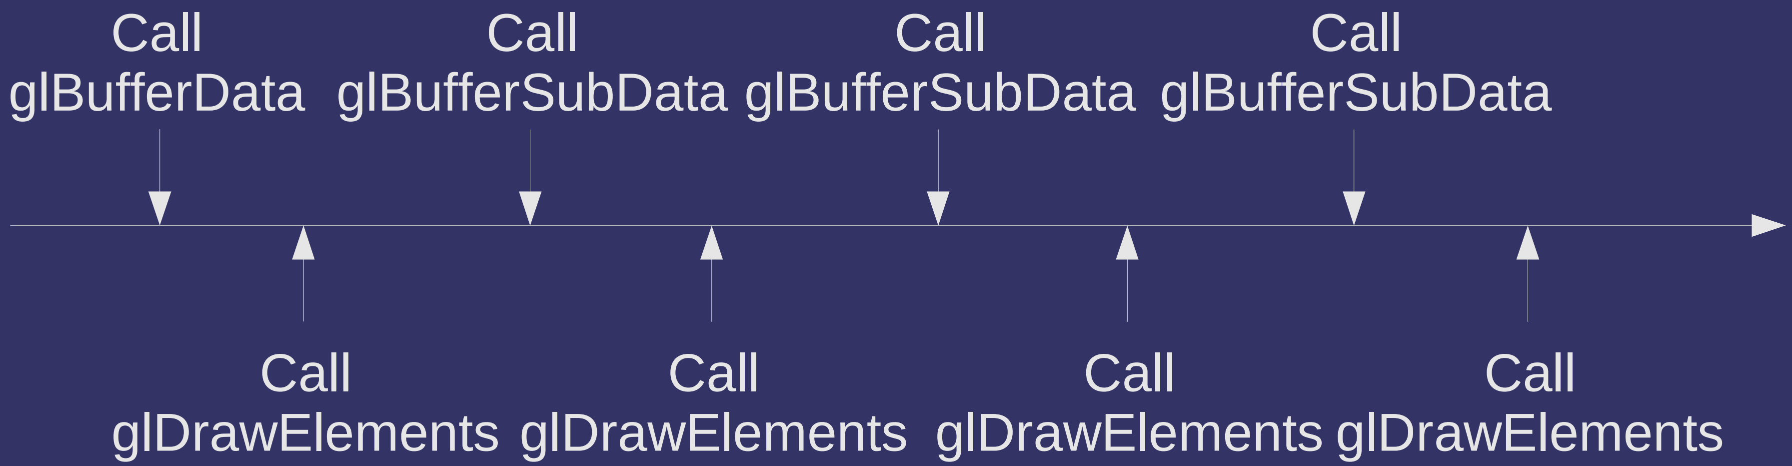

Buffer Object Usage Hints

GL_DYNAMIC_DRAW

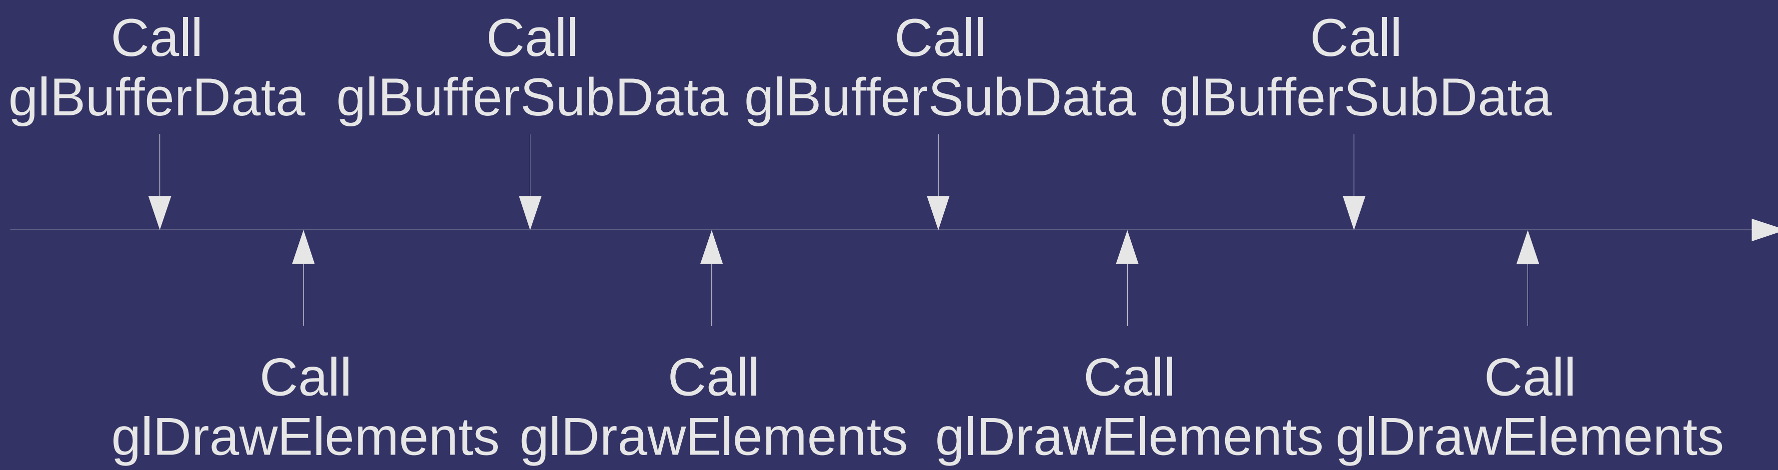

Buffer Objects

- Memory backing the buffer can be mapped into CPU space:

```
GLvoid *glMapBuffer(GLenum target,
                    GLenum access);
```

  - `access` tells the driver how the application will access the mapped buffer:
    - `GL_READ_ONLY`
    - `GL_WRITE_ONLY`
    - `GL_READ_WRITE`
- Unmap the buffer with:

```
GLboolean glUnmapBuffer(GLenum target);
```

Now what?

- The vertex data is in a buffer object...how do we tell the GPU know where to get it?

Vertex Attribute Pointer

- Set the location and format of a vertex attribute with:

```
void glVertexAttribPointer(GLuint index,
    GLint size, GLenum type,
    GLboolean normalized, GLsizei stride,
    const GLvoid *pointer);
```

Vertex Attribute Pointer

- Set the location and format of a vertex attribute with:

```
void glVertexAttribPointer(GLuint index,
    GLint size, GLenum type,
    GLboolean normalized, GLsizei stride,
    const GLvoid *pointer);
```

In the API, attributes are numbered

Vertex Attribute Pointer

- Set the location and format of a vertex attribute with:

```
void glVertexAttribPointer(GLuint index,
    GLint size, GLenum type,
    GLboolean normalized, GLsizei stride,
    const GLvoid *pointer);
```

Number of components in each element

Type of data (e.g., `GL_FLOAT`)

Vertex Attribute Pointer

- Set the location and format of a vertex attribute with:

```
void glVertexAttribPointer(GLuint index,
    GLint size, GLenum type,
    GLboolean normalized, GLsizei stride,
    const GLvoid *pointer);
```

For integer data, specifies whether it is normalized or not

Vertex Attribute Pointer

- Set the location and format of a vertex attribute with:

```
void glVertexAttribPointer(GLuint index,
    GLint size, GLenum type,
    GLboolean normalized, GLsizei stride,
    const GLvoid *pointer);
```

Number of bytes from
the start of one element
to the start of the next

Vertex Attribute Pointer

- Set the location and format of a vertex attribute with:

```
void glVertexAttribPointer(GLuint index,
    GLint size, GLenum type,
    GLboolean normalized, GLsizei stride,
    const GLvoid *pointer);
```

Offset, in bytes, from the start of the buffer where the data starts

Enable Attribute

- Attributes that will be used must also be enabled:

```
void glEnableVertexAttribArray(GLuint index);
```

- Attributes can later be disabled:

```
void glDisableVertexAttribArray(GLuint index);
```

Setting Attribute Numbers

- GLSL uses names for attributes:

```
attribute vec4 color;
```

- The API uses numbers:

```
void glVertexAttribPointer(GLuint index,
    GLint size, GLenum type,
    GLboolean normalized, GLsizei stride,
    const GLvoid *pointer);
```

- How do we connect the two?

Setting Attribute Numbers

- Bind the attribute name to the index we want:

```
void glBindAttribLocation(GLuint programObj,
    GLuint index, const GLchar *name);
```

  - Can only call *before* linking the program
  - Changes to attribute locations do not take effect until the program is linked (or linked again)

Drawing

- Draw a series of vertices:

```
void glDrawArrays(GLenum mode, GLint first,
    GLsizei count);
```

Drawing

- Draw a series of vertices:

```
void glDrawArrays(GLenum mode, GLint first,
    GLsizei count);
```

Sets the primitive type

Drawing

- Draw a series of vertices:

```
void glDrawArrays(GLenum mode, GLint first,
    GLsizei count);
```

Number of vertices to draw

Selects which vertex in the buffer to start drawing with

Primitive Types

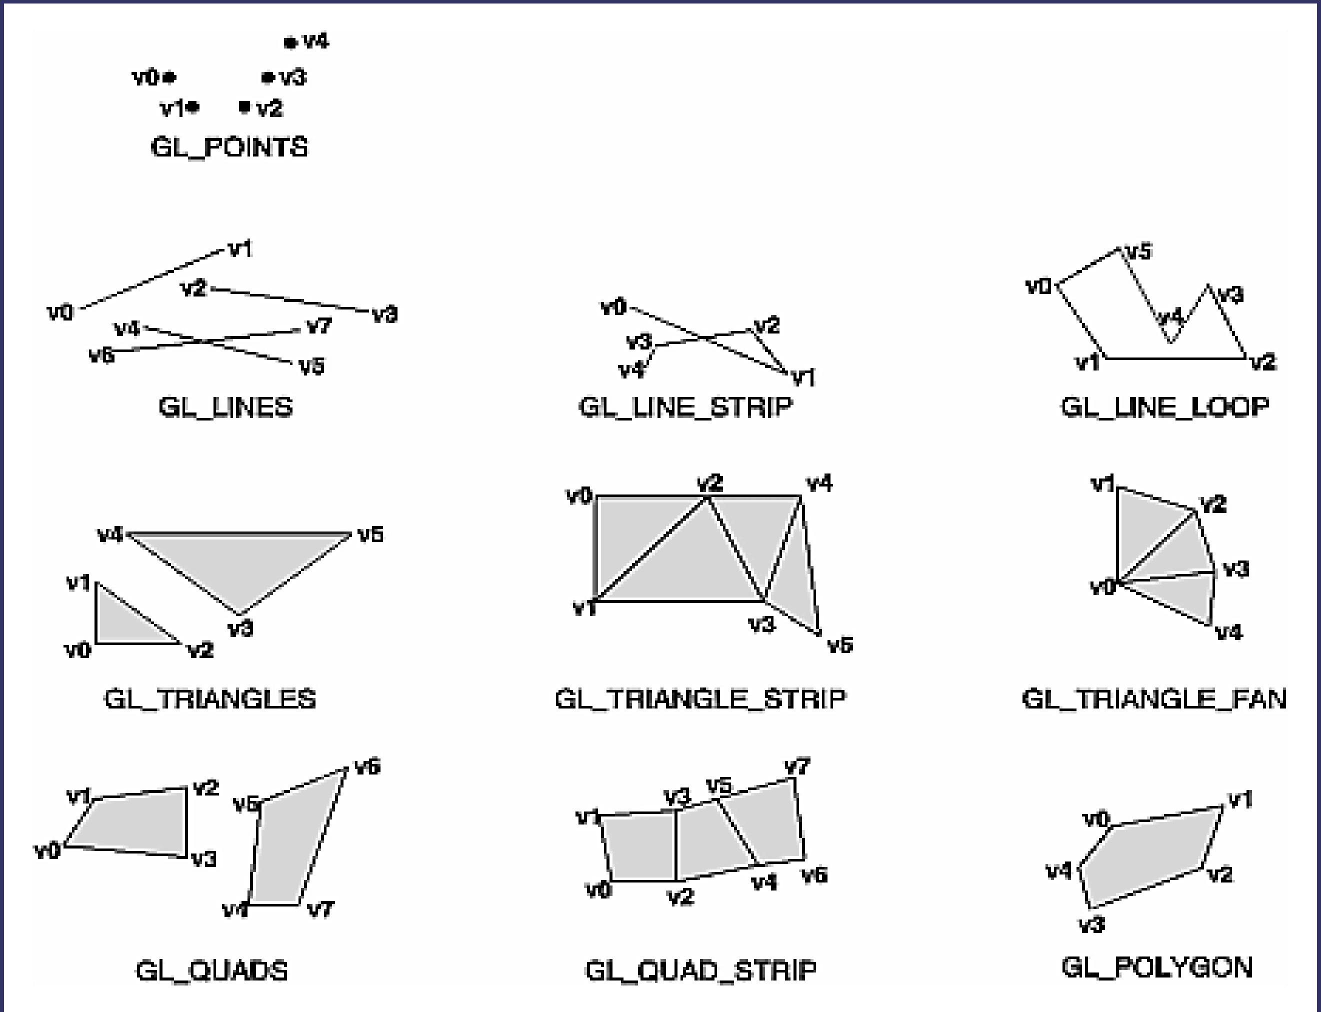

Image borrowed from “OpenGL Programming Guide”.

13-April-2010

Primitive Types

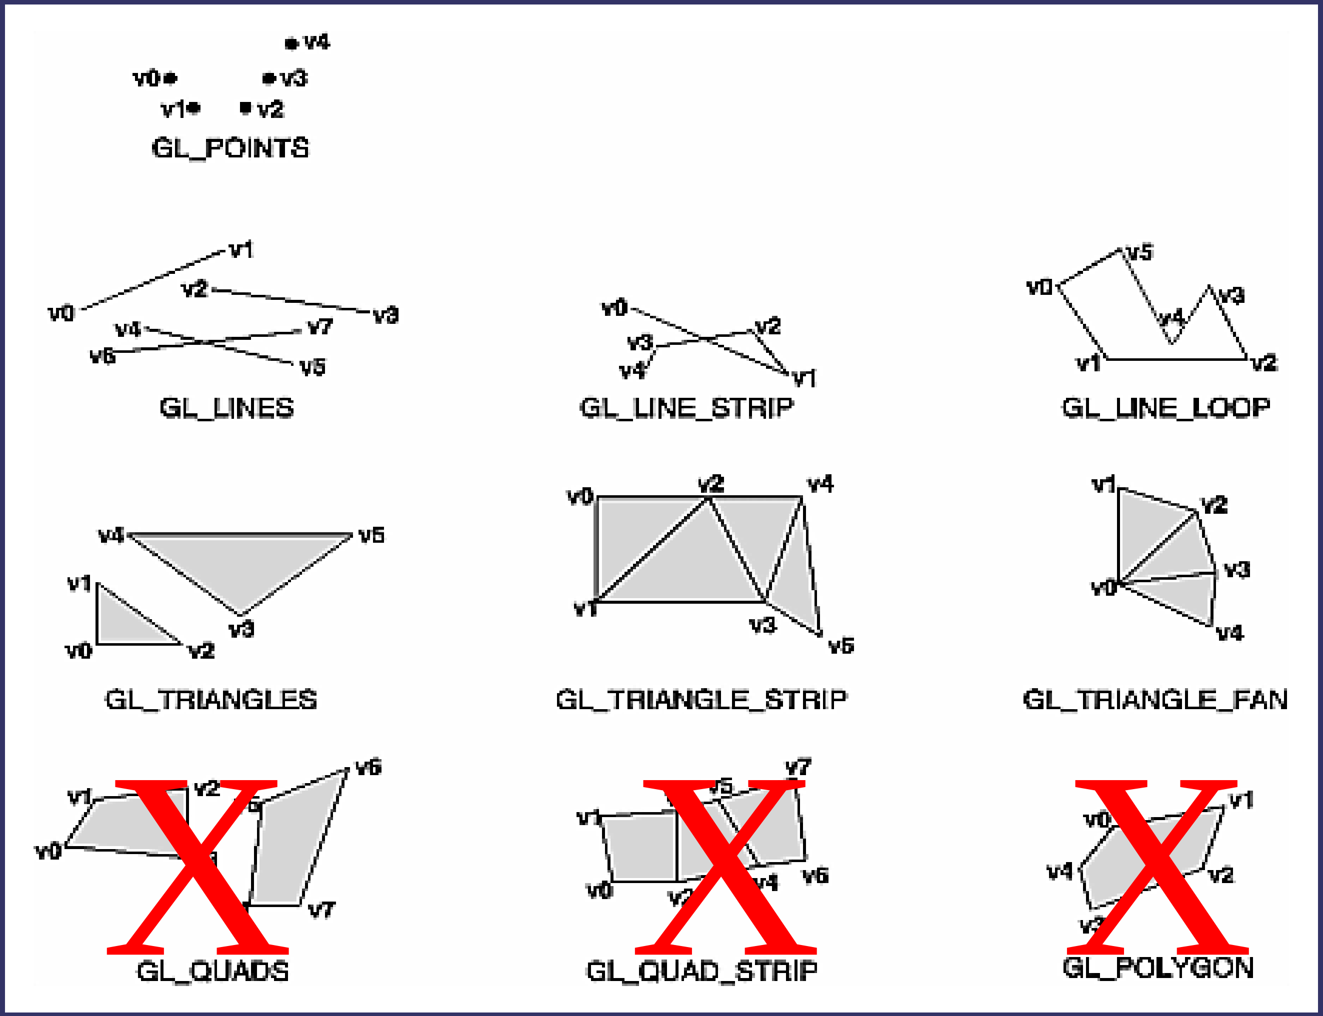

Image borrowed from “OpenGL Programming Guide”.

13-April-2010

References

- More information about I/O MMUs in general:
  http://en.wikipedia.org/wiki/IOMMU
- Nvidia whitepaper about using VBOs:
  http://developer.nvidia.com/object/using_VBOs.html

Linear Algebra Primer

- Three important data types:
  - Scalar values
  - Row / column vectors
    - 1x4 and 4x1 are the most common sizes
  - Square matrices
    - 4x4 is the most common size...to match the 1x4 & 4x1 vectors

Notation

- Try to use the same notation as the textbook:
  - Angle: θ (lower-case Greek)
  - Scalar: s (lower-case, italic, serif)
  - Vector or point: $\mathbf{v}$ (lower-case, bold, serif)
    - Sometimes $\hat{\mathbf{u}}$ is used to differentiate vectors from points
  - Matrix: $\mathbf{M}$ (upper-case, bold, serif)
  - Plane: π: $\mathbf{n} \cdot x + d$ (π: a vector and a scalar)
  - Triangle: $\triangle\mathbf{abc}$ ($\triangle$ 3 points)
  - Line segment: $\mathbf{ab}$ (2 points)
  - Geometric entity: A (upper-case, italic, serif)

Row Vectors

- These are special matrices that have multiple columns but only one row
  - Example: $\begin{bmatrix} 5.0 & 3.14 & 37 \end{bmatrix}$
- Addition and subtraction is component-wise:
  - Example: $\begin{bmatrix} 1 & 2 & 3 \end{bmatrix} + \begin{bmatrix} 9 & 10 & 11 \end{bmatrix} = \begin{bmatrix} 10 & 12 & 14 \end{bmatrix}$
  - Both vectors must be the same size
- Operations with scalars also component-wise:
  - Example: $3.2 \times \begin{bmatrix} 1 & 2 & 3 \end{bmatrix} = \begin{bmatrix} 3.2 & 6.4 & 9.6 \end{bmatrix}$
- Notice that vector multiplication is missing...

Column Vectors

- These are special matrices that have multiple rows but only one column
  - Example: $\begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}$
- Work just like row vectors
- Notationally convert a row to a column with a T in the exponent
  - Example: $\mathbf{v}^T$
  - We'll talk more about this notation later...

Vector Operations

- There are a few operations specific to vectors that are really important to graphics:
  - Dot product
  - Vector magnitude / normalization
  - Cross product

Dot Product

- Component-wise multiply, then sum components
  - Example:

$$\begin{bmatrix}2.3 & 1.2\end{bmatrix} \cdot \begin{bmatrix}1.7 & 6.5\end{bmatrix} = (2.3 * 1.7) + (1.2 * 6.5) = 11.71$$

  - Noted as $\mathbf{u} \cdot \mathbf{v}$ or $\langle \mathbf{u}, \mathbf{v} \rangle$
  - Also known as the *inner product* or *scalar product*

Vector Magnitude

- Noted by vertical bars around the vector
  - Like absolute value...which is the scalar magnitude
  - Can also be thought of as the length of the vector
- Square-root of dot-product of vector with itself
  - Like absolute value
  - Example: $\left\|\begin{bmatrix}\frac{\sqrt{2}}{2} & \frac{\sqrt{2}}{2}\end{bmatrix}\right\| = \sqrt{\begin{bmatrix}\frac{\sqrt{2}}{2} & \frac{\sqrt{2}}{2}\end{bmatrix} \cdot \begin{bmatrix}\frac{\sqrt{2}}{2} & \frac{\sqrt{2}}{2}\end{bmatrix}} =$

$$\sqrt{\left(\frac{\sqrt{2}}{2}\right)^2 + \left(\frac{\sqrt{2}}{2}\right)^2} = \sqrt{\frac{2}{4} + \frac{2}{4}} = 1$$

Normal

- *Normal* is an overloaded term in graphics and linear algebra
  - Sometimes it means a vector has unit length
    - $|\mathbf{u}| = 1.0$
    - Can say the vector is "normalized"
  - Sometimes it means a vector is perpendicular to a surface or another vector
    - This mean the angle between the vectors is 90°
    - Can say that the vectors are "normal to each other"
    - Can say that the vectors are "orthogonal"
  - Can combine for even more fun!
    - "Use normalized surface normals in the calculation."

Normalize

- Can normalize a vector by dividing it by its magnitude
  - Example: $\dfrac{\mathbf{u}}{|\mathbf{u}|}$
  - Vector has the same direction, but the magnitude will be 1.0
  - Also works with scalars

Dot Product

- Why is the dot product so interesting?

Dot Product

- Why is the dot product so interesting?
  - The dot product of two vectors is related to the cosine of the angle between those vectors
  - Formally: $\mathbf{u} \cdot \mathbf{v} = |\mathbf{u}|\ |\mathbf{v}| \cos\theta$
- We often want to know the angle between two vectors
  - This is the basis of all lighting calculations in 3D graphics!
  - $(\mathbf{u} \cdot \mathbf{v}) / (|\mathbf{u}|\ |\mathbf{v}|) = \cos\theta$

Cross Product

- From Wikipedia:

  [T]he cross product is a binary operation on two vectors in a three-dimensional Euclidean space that results in another vector which is perpendicular to the plane containing the two input vectors.

  - Noted as an $\times$ between two vectors
  - Calculated as:

$$\mathbf{a} \times \mathbf{b} = \begin{bmatrix} \mathbf{a_y b_z - a_z b_y} & \mathbf{a_z b_x - a_x b_z} & \mathbf{a_x b_y - a_y b_x} \end{bmatrix}$$

  - Not associative
  - Anti-commutative: If $\mathbf{u} \times \mathbf{v} = \mathbf{w}$, then $\mathbf{v} \times \mathbf{u} = -\mathbf{w}$

[1] From http://en.wikipedia.org/wiki/Cross_product

Cross Product

- Why is the cross product so interesting?
  - Cross product of two vectors results in a new vector that is normal both
  - The cross product of two vectors is related to the sine of the angle between the vectors
    - Formally: $\mathbf{u} \times \mathbf{v} = |\mathbf{u}|\ |\mathbf{v}| \sin\theta\ \mathbf{n}$

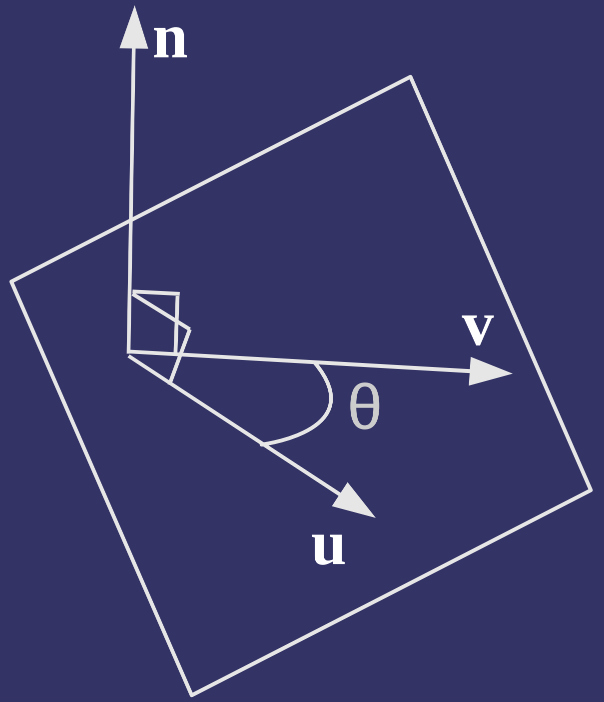

Matrices

- Like vectors, but have multiple rows and columns
  - Example: $\begin{bmatrix} 1.0 & 0.0 & 0.0 & 0.0 \\ 0.0 & 1.0 & 0.0 & 0.0 \\ 0.0 & 0.0 & 1.0 & 0.0 \\ 0.0 & 0.0 & 0.0 & 1.0 \end{bmatrix}$
- Add and subtract like you would expect
  - Like vectors, both matrices must be the same size...in both dimensions

Matrix Multiplication

- Special rules make matrix multiplication different from scalar multiplication
  - **NOT** commutative! e.g., $\mathbf{M} \times \mathbf{N} \neq \mathbf{N} \times \mathbf{M}$
  - Associative e.g., $\mathbf{A}(\mathbf{BC}) = (\mathbf{AB})\mathbf{C}$
  - Column count of first matrix must match row count of second matrix
    - If $\mathbf{M}$ is 4-by-3 matrix and $\mathbf{N}$ is a 3-by-1 matrix, we can do $\mathbf{M} \times \mathbf{N}$ but not $\mathbf{N} \times \mathbf{M}$
  - If the source matrices are n-by-m and m-by-p, the resulting matrix will be n-by-p

Matrix Multiplication

- To calculate an element of the matrix, **C**, resulting from **AB**:

$$\mathbf{C}_{ij} = \sum_{r=1}^{n} \mathbf{A}_{ir} \mathbf{B}_{rj}$$
$$= \mathbf{A}_{i,0}\mathbf{B}_{0,j} + \mathbf{A}_{i,1}\mathbf{B}_{1,j} + \mathbf{A}_{i,2}\mathbf{B}_{2,j} + \ldots + \mathbf{A}_{i,n}\mathbf{B}_{n,j}$$

- What does this look like?

Matrix Multiplication

- To calculate an element of the matrix, $\mathbf{C}$, resulting from $\mathbf{AB}$:

$$\begin{aligned}\mathbf{C}_{ij} &= \sum_{r=1}^{n} \mathbf{A}_{ir}\mathbf{B}_{rj} \\ &= \mathbf{A}_{i,0}\mathbf{B}_{0,j} + \mathbf{A}_{i,1}\mathbf{B}_{1,j} + \mathbf{A}_{i,2}\mathbf{B}_{2,j} + \ldots + \mathbf{A}_{i,n}\mathbf{B}_{n,j}\end{aligned}$$

- What does this look like?
  - The dot product of a row of $\mathbf{A}$ with a column of $\mathbf{B}$!
  - This is why the column count of $\mathbf{A}$ must match the row count of $\mathbf{B}$...otherwise the dot product wouldn't work

Multiplicative Identity

- There is a multiplicative identity for matrices

$$\mathbf{I} = \begin{bmatrix} 1 & 0 & \cdots & 0 \\ 0 & 1 & & \vdots \\ \vdots & & \ddots & 0 \\ 0 & 0 & \cdots & 1 \end{bmatrix}$$

  - Just like any other multiplicative identity, $\mathbf{AI} = \mathbf{A}$
  - If you pretend that a scalar is a 1×1 matrix, this should make sense

Transpose

- Rows become columns and columns become rows
  - Noted with a T in the exponent position (e.g., $\mathbf{M}^T$)
  - Example:

$$\begin{bmatrix} 2 & 3 \\ 4 & 5 \\ 6 & 7 \end{bmatrix}^T = \begin{bmatrix} 2 & 4 & 6 \\ 3 & 5 & 7 \end{bmatrix}$$

Matrix Multiplication

- Can rewrite the dot product (inner product) of two row vectors as:

$$s = \mathbf{u}\,\mathbf{v}^{\mathrm{T}}$$

- Can write the *outer product* of two row vectors as:

$$\mathbf{M} = \mathbf{u}^{\mathrm{T}}\,\mathbf{v}$$

  - Notation is $\mathbf{u} \otimes \mathbf{v}$

$$\mathbf{u} \otimes \mathbf{v} = \begin{bmatrix} \mathbf{u}_1\mathbf{v}_1 & \mathbf{u}_1\mathbf{v}_2 & \mathbf{u}_1\mathbf{v}_3 & \dots & \mathbf{u}_1\mathbf{v}_n \\ \mathbf{u}_2\mathbf{v}_1 & \mathbf{u}_2\mathbf{v}_2 & \mathbf{u}_2\mathbf{v}_3 & \dots & \mathbf{u}_2\mathbf{v}_n \\ \cdots & \cdots & \ddots & & \cdots \\ \mathbf{u}_m\mathbf{v}_1 & \mathbf{u}_m\mathbf{v}_2 & \mathbf{u}_m\mathbf{v}_3 & \dots & \mathbf{u}_m\mathbf{v}_n \end{bmatrix}$$

Matrix Multiplication

- Not commutative

$$\mathbf{M} \times \mathbf{N} \neq \mathbf{N} \times \mathbf{M}$$

- But...

$$\mathbf{M} \times \mathbf{N} = \left(\mathbf{N}^T \times \mathbf{M}^T\right)^T$$

- How is this useful?

Matrix Multiplication

- Not commutative

$$\mathbf{M} \times \mathbf{N} \neq \mathbf{N} \times \mathbf{M}$$

- But...

$$\mathbf{M} \times \mathbf{N} = \left(\mathbf{N}^T \times \mathbf{M}^T\right)^T$$

- How is this useful?
  - Assume v is a vector we want to transform by a matrix M, but we only have M^T in our program...

$$\mathbf{M} \times \mathbf{v} = \left(\mathbf{v}^T \times \mathbf{M}^T\right)^T$$

  - A vector and its transpose are represented the same way (`vec4` in GLSL), so we don't have to do the transpose of the matrix

References

http://en.wikipedia.org/wiki/Matrix_multiplication

http://en.wikipedia.org/wiki/Dot_product

http://en.wikipedia.org/wiki/Cross_product

http://en.wikipedia.org/wiki/Outer_product

Rotation

- Rotation around the Z-axis
  - If θ is 0°, this is the identity matrix

$$\begin{bmatrix} \cos\theta & -\sin\theta & 0 & 0 \\ \sin\theta & \cos\theta & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

- Rotation around the Y-axis

$$\begin{bmatrix} \cos\theta & 0 & \sin\theta & 0 \\ 0 & 1 & 0 & 0 \\ -\sin\theta & 0 & \cos\theta & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

Rotation

- From the previous equations, we can rotate using 4 multiplies and 2 adds, but a matrix multiply requires 16 multiplies and 12 adds
  - $x' = x \cos\theta + y \sin\theta$
  - $y' = -x \sin\theta + y \cos\theta$
  - $z' = z$
- Why use the matrix method?

Rotation

- A series of rotations can be implemented as:

$$\mathbf{v}' = \mathbf{M}_1 \mathbf{v}$$
$$\mathbf{v}'' = \mathbf{M}_2 \mathbf{v}'$$
$$\mathbf{v}''' = \mathbf{M}_3 \mathbf{v}''$$

- Which is the same as:

$$\mathbf{M}_3 (\mathbf{M}_2 (\mathbf{M}_1 \mathbf{v}))$$

- What can we do with this?

Rotation

- A series of rotations can be implemented as:

$$\mathbf{v}' = \mathbf{M}_1 \mathbf{v}$$
$$\mathbf{v}'' = \mathbf{M}_2 \mathbf{v}'$$
$$\mathbf{v}''' = \mathbf{M}_3 \mathbf{v}''$$

- Which is the same as:

$$\mathbf{M}_3 \left(\mathbf{M}_2 \left(\mathbf{M}_1 \mathbf{v} \right) \right)$$

- What can we do with this?

$$\left(\mathbf{M}_3 \mathbf{M}_2 \mathbf{M}_1 \right) \mathbf{v}$$

  - Matrix multiplication is associative!

Rotation

- A series of rotations can be implemented as:

$$\mathbf{v}' = \mathbf{M}_1 \mathbf{v}$$
$$\mathbf{v}'' = \mathbf{M}_2 \mathbf{v}'$$
$$\mathbf{v}''' = \mathbf{M}_3 \mathbf{v}''$$

Notice that the matrices are composed in the reverse order of how they are applied to the vector!

- Which is the same as:

$$\mathbf{M}_3 \left(\mathbf{M}_2 \left(\mathbf{M}_1 \mathbf{v} \right) \right)$$

- What can we do with this?

$$\left(\mathbf{M}_3 \mathbf{M}_2 \mathbf{M}_1 \right) \mathbf{v}$$

  - Matrix multiplication is associative!

Arbitrary Rotation

- Given a vector, $\mathbf{v}$, and an angle, θ, we can create an arbitrary rotation matrix:

$$\tilde{\mathbf{V}} = \begin{bmatrix} 0 & -\mathbf{v}_z & \mathbf{v}_y & 0 \\ \mathbf{v}_z & 0 & -\mathbf{v}_x & 0 \\ -\mathbf{v}_y & \mathbf{v}_x & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

$$\mathbf{R} = \left(\mathbf{I}\cos\theta\right) - \left(\left(1-\cos\theta\right)\left(\mathbf{v}\otimes\mathbf{v}\right)\right) + \left(\tilde{\mathbf{V}}\sin\theta\right)$$

Translation

- Points are stored as $\mathbf{p} = [\ x\ y\ z\ 1\ ]$
- Remember the definition of matrix multiplication:

$$
\begin{aligned}
\mathbf{p}_x' &= \mathbf{p}_x \mathbf{M}_{11} + \mathbf{p}_y \mathbf{M}_{12} + \mathbf{p}_z \mathbf{M}_{13} + \mathbf{p}_w \mathbf{M}_{14} \\
\mathbf{p}_y' &= \mathbf{p}_x \mathbf{M}_{21} + \mathbf{p}_y \mathbf{M}_{22} + \mathbf{p}_z \mathbf{M}_{23} + \mathbf{p}_w \mathbf{M}_{24} \\
\mathbf{p}_z' &= \mathbf{p}_x \mathbf{M}_{31} + \mathbf{p}_y \mathbf{M}_{32} + \mathbf{p}_z \mathbf{M}_{33} + \mathbf{p}_w \mathbf{M}_{34} \\
\mathbf{p}_w' &= \mathbf{p}_x \mathbf{M}_{41} + \mathbf{p}_y \mathbf{M}_{42} + \mathbf{p}_z \mathbf{M}_{43} + \mathbf{p}_w \mathbf{M}_{44}
\end{aligned}
$$

- Since $\mathbf{p}_w$ is always 1, the 4th column of the matrix acts as a translation

Scaling

- To scale a vector, multiply each component by a scale factor

$$\mathbf{M} = \begin{bmatrix} \mathbf{s}_x & 0 & 0 & 0 \\ 0 & \mathbf{s}_y & 0 & 0 \\ 0 & 0 & \mathbf{s}_z & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

Coordinate Spaces

- Coordinates are always relative to some “space”
  - Object space: Local coordinate system of the object
  - World space: Global coordinate system relative to the 3D “world”
  - Eye / camera space: Coordinate system relative to the viewer
- When we translate objects relative to other objects, we may talk about other spaces
  - If the hand of a 3D model is rotated relative to the arm of the model, we may talk about “hand-space” or “arm-space”

Orthonormal Basis

- It's a mouthful...what does it mean?
- A vector space where all of the components are *orthogonal* to each other, and each is *normal*
  - Normal meaning unit length
  - Orthogonal meaning at right angles
    - The *other* meaning of normal
- Every pure rotation matrix (i.e., no scaling) is an orthonormal basis
  - As is the identity matrix

Viewing

- Q: Given a world position for a camera, a world position to point the camera at, and an “up” direction, how can we construct a transformation using just rotations and translations?

Viewing

- Q: Given a world position for a camera, a world position to point the camera at, and an “up” direction, how can we construct a transformation using just rotations and translations?
- A: We can't. We need 3 vectors to construct an orthonormal basis
  - [Hughes 99] presents a method to construct from just one vector

Viewing

- Given:
  - $\mathbf{e}$: Position of the eye (or camera) in world-space
  - $\mathbf{v}$: The point being viewed
  - $\mathbf{u}$: the “up” direction
- Calculate the unit vector from the viewpoint to the eye:

$$\mathbf{f}' = \mathbf{v} - \mathbf{e}$$
$$\mathbf{f} = \frac{\mathbf{f}'}{|\mathbf{f}'|}$$

  - This is the Z axis

Viewing

- Calculate a vector orthogonal to the Z-axis and the up vector:

$$\mathbf{s} = \mathbf{f} \times \mathbf{u}$$

  – This is the X-axis

Viewing

- Calculate a vector orthogonal to the Z-axis and the up vector:

$$\mathbf{s} = \mathbf{f} \times \mathbf{u}$$

  - This is the X-axis
- Calculate a vector orthogonal to the X-axis and the Z-axis:

$$\mathbf{t} = \mathbf{s} \times \mathbf{f}$$

  - This is the Y-axis
  - Why can't we just use $\mathbf{u}$?

Viewing

- Drop these vectors into a matrix:

$$\mathbf{M}_v = \begin{bmatrix} \mathbf{s}_0 & \mathbf{s}_1 & \mathbf{s}_2 & 0 \\ \mathbf{t}_0 & \mathbf{t}_1 & \mathbf{t}_2 & 0 \\ -\mathbf{f}_0 & -\mathbf{f}_1 & -\mathbf{f}_2 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \times \begin{bmatrix} 1 & 0 & 0 & -\mathbf{e}_0 \\ 0 & 1 & 0 & -\mathbf{e}_1 \\ 0 & 0 & 1 & -\mathbf{e}_2 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

  - The translation moves the eye to the origin

References

General information about rotation matrices and orthonormal bases:

http://en.wikipedia.org/wiki/Rotation_matrix

http://www.wikipedia.org/Orthonormal_basis

Really good explanation of arbitrary rotation matrices:

http://www.euclideanspace.com/maths/geometry/rotations/conversions/angleToMatrix/index.htm

Hughes, J. F., and Möller, T. Building an Orthonormal Basis from a Unit Vector. *Journal of Graphics Tools* 4, 4 (1999), 33-35. http://www.cs.brown.edu/research/pubs/authors/john_f._hughes.html

Projection

- Once objects are transformed to camera-space, they're still 3D
  - The screen is still 2D
  - Camera parameters (e.g., field of view) need to be applied
- Three steps remain:
  - Projection from camera space to normalized device coordinates (NDC)
  - Perspective divide
  - Conversion from NDC to screen coordinates
    - Remaps the ± 1 cube to (0,0)-($width$, $height$)

Projection

- Perspective:
  - Simulates visual foreshortening caused by the way light projects onto the back of the eye
  - Represents the view volume with a frustum (a pyramid with the top cut off)
  - The real work is done by dividing X and Y by Z
- Orthographic:
  - Represents the view volume with a cube
  - Also called *parallel projection* because lines that are parallel before the projection remain parallel after

Perspective Projection

- A few parameters control the view volume:
  - Near: Distance from the camera to the near viewing plane. Objects in front of this plane will be clipped
  - Far: Distance from the camera to the far viewing plane. Objects behind this plane will be clipped
  - θ: Field-of-view in the Y direction
  - Aspect ratio: Ratio of the width of the view to the height of the view

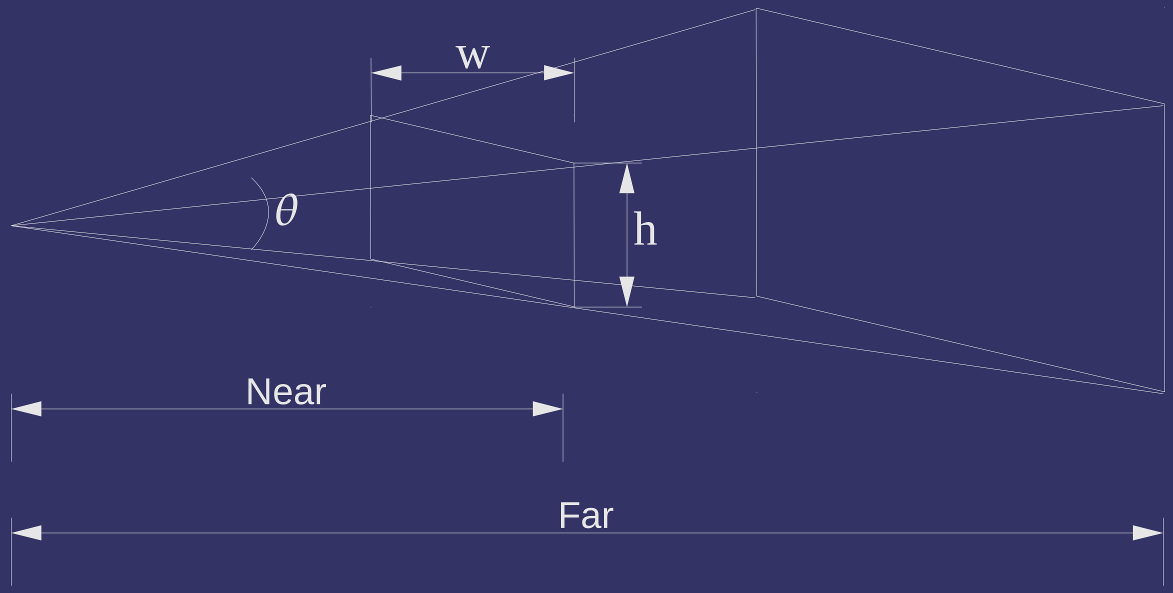

Perspective Projection

$$f = \cot\left(\frac{\theta}{2}\right)$$

$$\mathbf{M}_p = \begin{bmatrix} \frac{f}{aspect} & 0 & 0 & 0 \\ 0 & f & 0 & 0 \\ 0 & 0 & -\frac{far+near}{far-near} & -\frac{2 \times far \times near}{far-near} \\ 0 & 0 & -1 & 0 \end{bmatrix}$$

- Limited form of projection matrix that assumes symmetry in X and Y directions
- *near* and *far* are distances
  - We're actually looking down the negative Z axis in camera space

Putting it all together

- Typically have a *modeling* transform, a *viewing* transform, and a *projection*
  - Combine these into a single “model-view-projection” matrix: $\mathbf{M}_{mvp} = \mathbf{M}_p \times \mathbf{M}_v \times \mathbf{M}_m$
  - Transform a vertex by this single matrix:

```
uniform mat4 mvp;
void main(void)
{
        gl_Position = mvp * gl_Vertex;
}
```

References

http://en.wikipedia.org/wiki/3D_projection

- Especially the third step: perspective transform

http://en.wikipedia.org/wiki/Orthographic_projection_%28geometry%29

http://en.wikipedia.org/wiki/Isometric_projection

Next week...

- Quiz #1
  - Will cover material from last week and this week
- Hidden surface removal / occlusion
  - Backface culling
  - Painters algorithm
  - Z-buffer
  - Frustum culling
- Assignment #2, part 1

Legal Statement

This work represents the view of the authors and does not necessarily represent the view of Intel or the Art Institute of Portland.

OpenGL is a trademark of Silicon Graphics, Inc. in the United States, other countries, or both.

Khronos and OpenGL ES are trademarks of the Khronos Group.

Other company, product, and service names may be trademarks or service marks of others.