Should I, or Should I Not, Collapse

Problem 1: (Graded for Correctness - 15 pts)

You have a triangle mesh represented by the graph shown below. You want to collapse the edge marked in bold. (Half edges on each side of the edge to collapse are shown.) Is this edge collapsable? In other words, is the mesh produced by collapsing this edge a valid manifold mesh? If yes, please draw the new mesh. If no, please explain why not.

A Triangle Mesh

Problem 2: (Graded for Correctness - 25 pts)

Consider the following mesh with five vertices and six triangle faces. (The mesh is represented in indexed triangle form.):

```
Vertices:
  1: (-1,  0,  0)
  2: ( 1,  0,  0)
  3: ( 0,  0, -1)
  4: ( 0,  1,  0)
  5: ( 0, -1,  0)

Faces:
  1: 1 2 4
  2: 2 3 4
  3: 1 3 4
  4: 5 2 1
  5: 5 3 2
  6: 5 1 3
```

    A. Please draw the mesh on the coordinate system below.

B. Assume that a graphics system expects triangles to have counter-clockwise windings. In other words, when viewing a triangle from the "front side", the vertices of the triangle should be positioned in counter-clockwise order. Another way of saying this is that the normal of the triangle, computed as edge i CROSS edge $i + 1$, should be oriented toward the outside the shape. **In the mesh definition above, one of the faces has an incorrect winding**. Which one is it? Please also give a correct ordering of the vertex indices for this face. (There are multiple possible correct orderings.)

C. Consider a new triangle mesh with only **two faces**. The positions of the vertices and the organization of vertices into triangles are given below. (No, this is not a typo!)

$P_0 = (.5, .5, .5)$
$P_1 = (.5, .5, .5)$
$P_2 = (.5, .5, .5)$

$T_1 = (P_0, P_1, P_2)$
$T_2 = (P_0, P_2, P_1)$

Is the mesh manifold? Please explain why or why not. Remember, the manifold property of a mesh is a statement about the mesh's topology, not about the location of the surface.

Problem 3: (Graded for Correctness - 20 pts)

A. Consider the scene below organized into a BVH. The structure of the BVH is shown in the top-right. Please give the ordered sequence of BVH nodes that ray R1 must visit when determining the **closest** intersection along the ray. **(We say a ray "visits" a node if during traversal it is determined that the node might contain the primitives resulting in the closest hit, and thus we must recurse into the node.)** You should assume that: (1) when a ray hits both child bounding boxes of the current node, the ray will visit the node with the closest bounding box first, (2) If a ray has already found an intersection that is **closer than** one of the child bounding boxes of the current node, it will skip visiting that node.

B. Consider the scene geometry (8 white triangles) and corresponding BVH shown below. The dotted boxes on the right indicate the bounding boxes of interior nodes N0-N4. Now imagine that triangle E moves so that it is now located at the position indicated by the shaded triangle. Assume that the BVH structure is not changed (the topology stays the same and so do the boundign boxes of interior nodes.) Precisely describe a ray-scene intersection test error that could occur as a result. In your answer describe or draw a ray that would trigger the error and describe what tracing this ray could not yield a correct hit result.
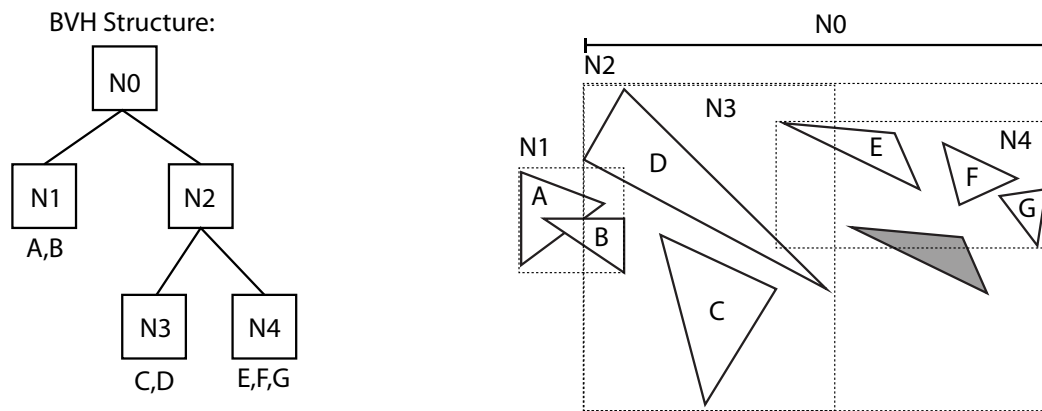
Problem 4: (Graded on Effort Only - 20 pts)

The implicit form of an ellipsoid centered at point (x_0, y_0, z_0) can be given by:

$$\frac{(x - x_0)^2}{A^2} + \frac{(y - y_0)^2}{B^2} + \frac{(z - z_0)^2}{C^2} = 1$$

Recall that the equation for all points on an array with origin $\mathbf{o}$ and direction $\mathbf{d}$ is given by $\mathbf{o} + t\mathbf{d}$. Please give an espression for the values of t that correspond to the ray-ellisoid intersection points (Hint: set up a quadratic equation and solve for t using the quadratic formula.) After giving an expression for t, please describe how many unique solutions for t there can be. Why is this the case?

Problem 5: (Graded on Effort Only - 20 pts)

After a raining winter week, you find yourself craving some fresh air and sunlight. Therefore, you've decided to go cherry picking with your CS248A friends. However, instead of the traditional method of picking cherries by hand, you'll be shooting light rays at a box of cherries.

Each cherry is represented by two spheres: a large sphere for the fruit, and a smaller internal sphere for the cherry pit. Assume the light rays can travel through the flesh of all cherries and continue to intersect with primitives behind them, but stops if it hits the box or any cherry pit. Each time the ray hits any part of a cherry, you get to keep it!

You are given two subroutines below to use. Both functions will return true when an intersection occurs (and false if not), and populate the ray's t-values with information for up to two points of impact. In cases where there's only one point of intersection, like when the ray just grazes the shape, you can consider both t-values to be identical. If there is no intersection, then it's assumed that the t-values are assigned a significantly large number.

```
// intersect ray with the box
bool rayBoxIsect(Box b, Ray r, float* t1, float* t2);

// intersect ray with a sphere centered at the origin
bool raySphereIsect(Ray r, float radius, float* t1, float* t2);
```

Question continues onto the next page...

You can assume that you have the following information for each cherry.

```
Cherry {
  vec3 cherry_center;
  float cherry_radius;
  CherryPit cherry_pit;
}

CherryPit {
  vec3 pit_center;
  float pit_radius;
}
```

Warning!: Notice that `raySphereIsect()` assumes that the sphere is centered at the origin.

Given that you get to keep all cherries that your ray hits, please give an algorithm for finding the number of cherries you get to take home given an input ray r, Box b, and list of `Cherrys` `cherry_list`. Remember that the ray stops if/when it hits the Box or a `CherryPit`. Psuedocode is fine!

```
// TODO: Implement the function below...

int getNumCherries(Ray r, Box b, std::vector<Cherry> cherry_list) {
```

```
}
```

Refitting a BVH

PRACTICE PROBLEM 1:

One of the nice properties of a bounding volume hierarchy (BVH) is that it can be efficiently updated ("refit") when a single primitive contained in the hierarchy is moved. Refitting modifies the BVH so that the BVH property is maintained: *the bounding box of a node is a bound containing the primitives in all child nodes*. Refitting does not modify the structure of the BVH – **it only updates bounding boxes for some of the tree's nodes based on the new bounds of leaf nodes.**

Below you're given a definition of a BVHNode and interfaces for getting the bounding box of a primitive. Please implement the function refitBVH(node) which should update bounding boxes of the *necessary nodes* in the BVH so that the BVH property holds. **You should assume that only primitives in the specified leaf node (node) have moved.** To keep things simple, all BVH nodes hold a pointer to their parent in the BVH. **NOTE THAT EVEN THOUGH WE GAVE YOU VALID C STRUCTS IN THE PROBLEM FOR CLARITY, YOU ONLY NEED TO SKETCH OUT AN APPROACH TO A SOLUTION. VALID C CODE IS NOT REQUIRED. A DESCRIPTION OF AN APPROACH IS FINE.**

```
struct BBox {
   void clear();               // resets bbox to empty
   void union(const BBox& b);  // enlarges bbox to include volume in b
};

struct Primitive {
  BBox getBBox() const;     // returns primitive's world-space axis-aligned bbox
};

struct BVHNode {
   BVHNode* left, *right, *parent  // parent is NULL if root node, left/right NULL if leaf
   BBox bbox;          // node's bbox, you need to update this!
   int numPrims;       // 0 if node is interior node, non-zero otherwise
   Primitive* prims;  // prims[i]->getBBox() returns bbox of i'th prim in node
};

// Refit the entire BVH assuming that the contents of the leaf node 'node' have changed.
void refitBVH(BVHNode* node) {



}
```

Ray-Primitive Intersection Practice (Two Boxes)

PRACTICE PROBLEM 2:

You wish to intersect a ray with a shape defined by two axis aligned boxes B_1 and B_2. The shape is defined by the volume given by $B_1 - B_2$. In other words, B_2 cuts volume out of B_1. Depending on the size and shapes of the two boxes, you can get some interesting shapes. Two examples are given below. B_2 **need not intersect with** B_1**, and when this is the case, the resulting shape is just** B_1**. (**$B_1 - B_2 = B_1$**).**

Assume you are given code for computing the closest point of intersection with a box:

```
// Returns true if ray r hits box b, fills in t with distance to closest hit.
// IMPORTANT: If r originates inside the box, or if ray R originates on the
// surface of the box and points towards the inside of the box, then the closest
// hit will be the point where r leaves the box, not the origin of r.
bool rayBoxIsect(Ray r, Box b, float* t);
```

Please use `rayBoxIsect` as a subroutine in an algorithm for computing the closest intersection between a ray R **and the shape defined by** B_1 **and** B_2. To make things easier, you can assume that the ray R originates "outside" both boxes, and you can ignore the case where the ray grazes any box.

You can assume that `Ray.o` and `Ray.d` give the origin and direction of a ray. Pseudocode is fine (it need not compile), but be precise enough for the grader to understand how to perform key steps. **Hint: first give pseudocode for how to find the point where the ray r enters each box AND the point where the ray leaves each box. Then what do you with this information to solve the problem?**

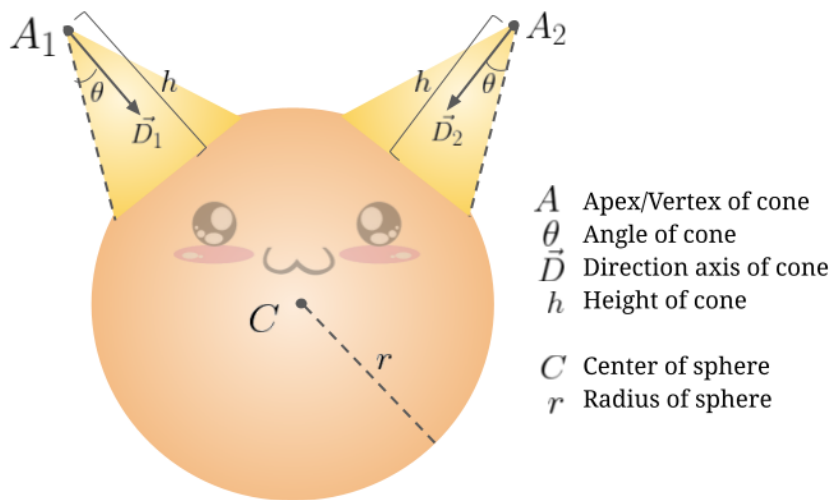**Please provide your answer on the next page.**

```
// returns true if ray hits shape
// fills in t to give the position of the closest hit
bool hitsShape(Ray r, Box b1, Box b2, float* t) {

    float b1_t1 = INF;  // distance to first hit with B1 (if applicable)
    float b1_t2 = INF;  // distance to second hit with B1 (if applicable)
    float b2_t1 = INF;  // distance to first hit with B2 (if applicable)
    float b2_t2 = INF;  // distance to second hit with B2 (if applicable)
```

PRACTICE PROBLEM 3:

One way of modeling a cat head is with a union of a sphere centered at C with radius r and two cones of height h).

A Apex/Vertex of cone
θ Angle of cone
$\vec{D}$ Direction axis of cone
h Height of cone

C Center of sphere
r Radius of sphere

Assume you have access to the following:

(1) an `InfiniteCone` struct (a cone with a height of infinity) that stores A, D and θ
(2) a `Sphere` struct that stores sphere center point C and radius r
(3) a `rayInfiniteConeIsect()` function
(4) and a `raySphereIsect()` function

Questions are on the next page...

A. Give an algorithm for finding the closest intersection of a ray with the cat above. Your solution can be described in words, but make sure it's clear what your algorithm is. Be precise how you use `rayInfiniteConeIsect()` to determine ray-cone intersection with a cone of height h.

```
struct InfiniteCone {
  Vec3D A;        // apex position
  Vec3D D;        // direction of axis
  float theta;    // cone angle
};
struct Sphere {
  vec3D C;        // center of sphere
  float r;        // radius
};

// in the functions below t1 is the ''closer hit'': t1 <= t2
void rayInfiniteConeIsect(InfiniteCone b, Ray r, float* t1, float* t2);
void raySphereIsect(Sphere c, Ray r, float* t1, float* t2);
```

B. Ah, you thought you got off easy without having to solve a ray-primitive intersection problem! Consider a simpler problem where you need to intersect a ray with a cone that has an apex at the origin, is oriented along the z axis (D=(0,0,1)), and has half angle θ. An implicit form of this infinite cone is $x^2 + y^2 = (z \tan \theta)^2$. (Convince yourself this is true!)

Please give an equation for the t value of the intersection point between the infinite cone and a ray with origin $\mathbf{o}$ and ray direction $\mathbf{d}$. You do not need to apply the quadratic formula to solve for t (just give us an equality involving t and components of $\mathbf{o}$ and $\mathbf{d}$).

C. Now imagine you want to compute the intersection of a ray with an infinite cone has apex at point A and is oriented in the direction D (just like in part A). How would you modify the ray's origin and direction to reduce this problem to intersection with the cone from the previous problem that has apex at the origin and is oriented along $(0, 0, 1)$? A description (in words) that involves rotations and translations is fine. But in what direction do you translate, and how do you define the rotation?

Ray-Grid Intersection

PRACTICE PROBLEM 4:

We are interested in casting rays onto a uniform 3D grid shown below. The grid is 4×4×4 and each grid cell is a cube with side length 1. The origin of the grid (**O**) is located at the origin of the coordinate system: **O** = $(0, 0, 0)$.

Recall that a ray can be represented as $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$. Assume that the origin of the ray is ALWAYS outside of the grid and recall that a plane can be represented as $\mathbf{N}^T \mathbf{x} = c$.

Hint: A plane parallel to the XY plane has $\mathbf{N} = (0, 0, 1)$, *a plane parallel to the XZ plane has* $\mathbf{N} = (0, 1, 0)$ *and a plane parallel to the YZ plane has* $\mathbf{N} = (1, 0, 0)$.

Assume that you have access to the following:
(1) Struct `Plane` that stores normal `N` and offset `c`
(2) Struct `Ray` that stores origin `o` and direction `d`
(3) Function `bool rayPlaneIsect(Plane plane, Ray r, float* t)` that takes as input a plane and a ray and set `t` to be the intersection. A boolean is returned to indicate if an intersection is found.

THE QUESTIONS BEGIN ON THE NEXT PAGE

A. Give an algorithm for finding the index of the first cell getting hit by the ray. The index of a cell should be represented as a 3D vector (a, b, c) corresponding to the indices in x, y, and z axis, as shown in the figure above. **The complexity of your algorithm SHOULD NOT scale as the number of cells in the grid increases.** Your solution can be described in words or rough pseudocode, but make sure it's clear what your algorithm is. Be precise about how you initialize the given structs and how you use `rayPlaneIsect()` to determine the ray-grid intersection. Hint: figure out where the ray first hits the volume of the grid, then turn that into an index.

```
struct Plane {
  vec3 N;        // unit normal
  float c;       // offset
};
struct Ray {
  vec3 o;        // ray origin
  vec3 d;        // ray direction
};

bool rayPlaneIsect(Plane plane, Ray r, float* t);
```

B. Now we want to remove some cells from the grid. Given a collection of cell indices that are removed `removedCellIndices`, give an algorithm for finding the index of the first cell hit by the ray. **Now, the complexity of your algorithm CAN scale linearly with the number of grid cells.** Your solution can be described in words.

```
vec3[M] removedCellIndices = [vec3(0,1,0), vec3(1,2,1), ...]
```

C. Now consider the case where the grid is translated and rotated such that the grid origin **O** is located at position **p** and the three sides next to **O** now have unit directional vector **u**, **v** and **w**. An illustration is provided below:

We want to find the index of the first grid cell hit by a ray. Assume that you can access your solution to part B via function `getCellIndex(Ray r)`. Please implement the function:
`getCellIndex(Ray r, vec3 p, vec3 u, vec3 v, vec3 w)`.
Be precise about how to construct arguments to `getCellIndex()`. **Specifically, please be precise about how you construct and use transformation matrices.**

```
vec3 getCellIndex(Ray r, vec3 p, vec3 u, vec3 v, vec3 w) {

// implement your solution here...
```

Solution:

Create a transform:

PRACTICE PROBLEM 5:

Consider a scene contain 15 rectangular-boxes as shown below. The scene is formed by starting with a box of length=1.0, with it's center-bottom at the origin, and extending down the Y axis (the width and depth of the box do not matter in this problem). The next box begins at the end of the first box (at point (0,1,0)), but it has a length that is 0.9 times the first box, and rotated with an angle θ relative to the first box. Note that the rotation is about the Z axis, in the *clockwise direction* as viewed when looking down the -Z axis. This pattern of starting the next rectangle at the end of the previous, rotating by θ, and shrinking length by a factor of 0.9 repeats for a total of 15 boxes.

You are given a function void ray_rect_isect(Ray r, float rect_length) that computes the first hit of a ray r, with a box aligned with the Y axis of length rect_length (see code on next page) For example, the call ray_rect_isect(r, 1.0) would compute the intersection of a ray with the first box. Upon return r.min_t would reflect the distance to the closest hit. *If the original r.min_t of the ray passed into the function is smaller than the t computed from the intersection, then r.min_t is unchanged.*

You are also given a function rotatez that takes a 3-vector v and rotates it θ degrees about the Z axis in the **counter-clockwise direction.**. Using only these routines, on the next page, please implement the function isect_funky_scene(Ray& r), which will fill in r.min_t to contain the closest point on the ray. You should make no assumptions about the origin or direction of the ray in 3D, except that it does not originate from inside any of the boxes.

Hint: be careful about the direction of your rotations. The geometry rotations in the figure are clockwise, and the function rotatez() specifies a rotation in the counter-clockwise direction as was typical in class.

```cpp
// students are free to assume that useful constructors, or common ops like
// addition, multiplication on vectors is available..
struct vec3 {
  float x, y, z;
};

struct Ray {
  vec3 o;
  vec3 d;
  float min_t;  // assume this is initialized to INFINITY
};

// rotate counter-clockwise about Z axis (counter clockwise defined when looking down -z)
vec3 rotatez(float theta, vec3 v);

// fills in r.min_t if intersection with rectangular box is closer than current r.min_t
void ray_rect_isect(Ray r, float rect_length);

// assume r.min_t is INFINITY at the start of the call
// result: fill in r.min_t as a result of intersecting r with the 15 segments
void isect_funky_scene(Ray& r) {




}
```

PRACTICE PROBLEM 6:

Consider the 2D shape given below, which is made up of the intersection of volumes contained by two circles and a rectangle.


Assume you are given two functions for ray-shape intersection as given below. Both functions return true if there is an intersection (false otherwise), and fill in ray t values for up to two hits. If there is a single intersection point, such as when the ray grazes the shape, just assume that both t values are the same. If there is no intersection, assume the t's are set to a very large number.

```
bool rayBoxIsect(Box b, Ray r, float* t1, float* t2);
bool rayCircleIsect(Circle c, Ray r, float* t1, float* t2);
```

A. Assuming we call the circles c_1 and c_2, and the box b, give an algorithm for finding the closest intersection of a ray with the shape above.

B. Now consider a new kind of shape, which is formed by *subtracting the region inside one circle from the region inside one box. (a few cases are given to help your understanding.)*

Give an algorithm for finding the closest intersection of a ray with this new type of shape. *Note, please make sure you handle the case where the shape is actually two disjoint pieces, and the case where there is no intersection! However, to make things simpler, you can assume the ray origin is outside the area of the box or the circle.*