CS 465 Final Exam (Fall 2007, Prof. James)

Thursday 6 December 2007—2.5 hours (150 minutes)

The exam contains *7 questions* worth a total of *100 points.*

Explain your reasoning in words and equations for full marks.

Problem 1: Image Filtering (15 pts)

Consider the following 4-by-4 image:

0	1	1	0
1	4	9	1
1	4	4	1
0	1	1	0

(a) Convolve the image with a 3-by-3 box blur filter. As in the assignment, renormalize the filter to account for missing values at the edges.

- *Answer:*

3/2	8/3	8/3	11/4
11/6	25/9	25/9	8/3
11/6	25/9	25/9	8/3
3/2	11/6	11/6	3/2

(b) Convolving an image with a box filter computes local mean values. Similarly, convolving an image with a median filter computes local median values–a nonlinear operation.

Recall that given an odd number of numbers, the median is the middle number. Given an even number of numbers, the median is the mean of the two middle numbers. Thus, the median of the numbers 2, 4, 7, 12 is (4+7)/2=11/2.

Analogous to image filtering using a 3-by-3 blur filter, **apply a 3-by-3 median filter to the 4-by-4 image**. Ignore missing values at the edges.

- *Answer:*

1	1	1	1
1	1	1	1
1	1	1	1
1	1	1	1

(c) The 3-by-3 box blur filter is a separable filter, but is the 3-by-3 median filter a separable filter?

- *Answer:* No. Recall that separable convolution filters are linear convolution filters that can have their kernel masks M represented by an outer product of a single vector with itself, e.g., $M = vv^T$. For example, the 3-by-3 box blur filter can be represented using $v = 1/3(1, 1, 1)'$. Since the median filter is a nonlinear filter it can not be represented this way, nor as any kind of linear convolution operation.

A number of people argued that since a horizontal (vertical) median operation followed by a vertical (horizontal) median operation led to different results, that this was not a separable filter. We did give marks for this, however this reasoning is does not constitute a proof and the linear/nonlinear distinction is more correct and simpler.

Problem 2: Ray-Quadrilateral Intersection (10 pts)

Consider a nonconvex planar quadrilateral whose four vertices are given by

$$\mathbf{p}_0 = \begin{pmatrix} 0 \\ 0 \\ -1 \end{pmatrix} \qquad \mathbf{p}_1 = \begin{pmatrix} 1 \\ 0 \\ -1 \end{pmatrix} \qquad \mathbf{p}_2 = \begin{pmatrix} 0 \\ 1 \\ -1 \end{pmatrix} \qquad \mathbf{p}_3 = \begin{pmatrix} -1 \\ -1 \\ -1 \end{pmatrix}$$

and a ray $\mathbf{r}(t) = \mathbf{e} + \mathbf{v}t,\ t > 0$, with

$$\mathbf{e} = \begin{pmatrix} 2 \\ 1 \\ 10 \end{pmatrix} \qquad \mathbf{v} = \begin{pmatrix} -0.15 \\ -0.01 \\ -1 \end{pmatrix}.$$

Does the ray hit the planar quadrilateral? You can draw a calibrated diagram (with labels) to support your claim. [Hint: Look at the particular geometry.]

- *Answer:* No, the ray misses.

  First observe that the polygon lies in the $z = -1$ plane. Intersecting the ray with that plane implies

  $$-1 = 10 - t^* \Rightarrow t^* = 11.$$

  Therefore

  $$\begin{aligned} x(t^*) &= 2 - 0.15t^* = 2 - 0.15 * 11 = 2 - 1.65 &= 0.35 \\ y(t^*) &= 1 - 0.01t^* = 1 - 0.01 * 11 = 1 - 0.11 &= 0.89 \end{aligned}$$

  and so the intersection point $(0.35, 0.89, -1)$ lies above the quadrilateral (see picture of $z = -1$ plane polygon and point). Unfortunately, a number of people interpreted the shape of the polygon incorrectly.

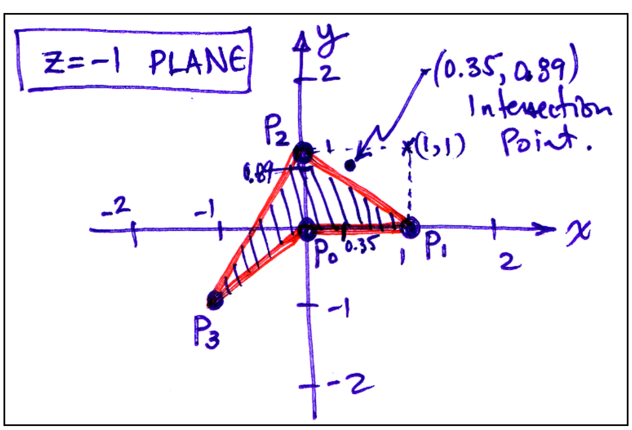

Problem 3: Curve Rasterization (20 pts)

You have rasterized lines such as $y = mx + b,\ \ m \in [0, 1]$ using a digital differential analyzer (DDA) that exploits cheap additions and comparisons to update quantities. For example, recall the Pixel-walk line rasterizer on an integer grid ($\Delta x = \Delta y = 1$):

```
x = ceil(x0)
y = round(m*x + b)
d = m*x + b - y
while x < floor(x1)
    if d > 0.5
        y += 1; d -= 1;
    else
        x += 1; d += m;
    if -0.5 < d ≤ 0.5
        output(x, y)
```

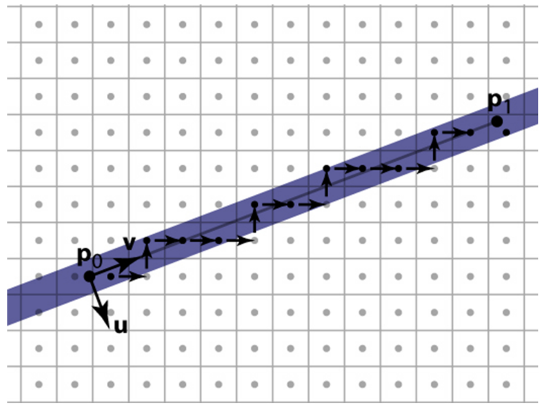

In this question you will explore how quadratic splines can be rasterized using a DDA-type algorithm. Consider rasterizing a section of the quadratic curve

$$y = ax^2 + mx + b.$$

For simplicity, let us assume "shallow curves" with $\frac{dy}{dx} \in [0, 1]$, analogous to the $m \in [0, 1]$ condition for lines.

(a) First, derive *addition-only updates* of the implicit difference value

$$d = d(x, y) = ax^2 + mx + b - y$$

for both

(i) unit increments of y: $\quad d$ += $d(x, y+1) - d(x, y)$;

- *Answer (3 pts):* The increment is the same as before,

$$(d(x, y+1) - d(x, y)) = \big(ax^2 + mx + b - (y+1)\big) - \big(ax^2 + mx + b - y\big) = -1.$$

(ii) unit increments of x: $\quad d$ += $d(x+1, y) - d(x, y)$.
(Hint: You may need to additively update more than just x, y, and d quantities.)

- *Answer (7 pts):* The increment is

$$(d(x+1, y) - d(x, y)) = \big(a(x+1)^2 + m(x+1) + b\big) - d = 2ax + a + m = D_x$$

where we have defined the increment to be $D_x \equiv 2ax + a + m$, which changes with x and therefore can also be incremented to avoid dreaded multiplication in the DDA's inner loop. Therefore consider the effect of an x increment on D_x (which would be done *after* we update d):

$$D_{x+1} = 2a(x+1) + a + m = (2ax + a + m) + 2a = D_x + 2a \quad \Rightarrow \quad \texttt{Dx += 2a}$$

(b) Second, rewrite the aforementioned Pixel-walk rasterizer for the quadratic case using addition-only updates in the inner `while` loop.
(Hint: You should only need to modify the parts indicated by gray boxes.)

- *Answer:* The key phrase here is "using addition-only updates" and not multiplication. The modified algorithm (including recurrences and optimizations discussed in answer to (a)) is

```
x  = ceil(x0);
Y  = a*x*x + m*x + b; // just compute once
y  = round(Y);
d  = Y - y;
a2 = 2*a;// Needed to update Dx (after d is updated!)
Dx = a2*x + a + m;
while(x < floor(x1)) {
  if(d > 0.5) {// Increment y:
    y += 1;
    d -= 1;
  }
  else {// Increment x:
        // Note: Only ONE ADD MORE than line rasterization!
    x  += 1;
    d  += Dx
    Dx += a2;// post-d update of Dx (see def'n)
  }
  if(-0.5 < d <= 0.5)
    output(x,y);
}
```

  Some people updated `Dx` before they updated `d`, which is incorrect–except for those who defined D_x in terms of $(x+1)$. Note that we could've broken off the constant $(a+m)$ from D_x, but it doesn't buy us anything, since (1) both $(2ax+a+m)$ and $2ax$ have the same recurrence, and (2) if we break it off we need to do an extra `add` in the inner loop which is less efficient. Finally, to avoid evaluating $2a$ via multiplication in the inner loop, we precompute the constant `a2 =` $2a$, so that `Dx += a2`); some people used two "`+= a`" operations to increment `Dx`, which is not incorrect but obviously less efficient.

<*Intermission!*>

Question: Why doesn't $\mathbf{p} = \mathbf{N}\mathbf{p}$ iff $\mathbf{N}$ is the identity matrix?

Answer: Because $\mathbf{p}$ could be the zero vector. <*Ha ha ha. Boo...*>

<*Exeunt*>

Problem 4: Constructive Solid Geometry (CSG) (15 pts)

(a) CSG allows you to build objects from simple point sets using set operations (op): intersection ($\cap$), union ($\cup$), and difference ($-$). **Give a CSG expression for the model M (shown in solid gray)** in terms of the circular point sets A, B, and C.

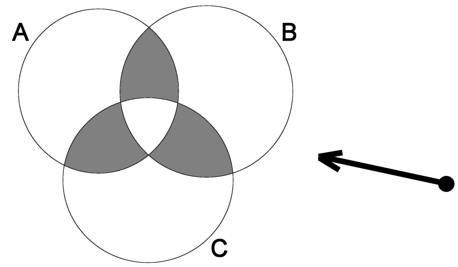

- *Answer:* For convenience, denote the missing convex center piece by

$$D = A \cap B \cap C.$$

The model M is the the union of three pieces, with each piece the intersection of two circles with D subtracted,

$$M = ((A \cap B) - D) \cup ((B \cap C) - D) \cup ((C \cap A) - D),$$

which can be written more succinctly as

$$\begin{aligned} M &= (A \cap B) \cup (B \cap C) \cup (C \cap A) - D \\ &= (A \cap B) \cup (B \cap C) \cup (C \cap A) - (A \cap B \cap C). \end{aligned}$$

(b) Consider intersecting a ray with your CSG model, M. Assume that your implementation supports intersecting a ray ($\mathbf{r}(t) = \mathbf{e} + \mathbf{v}t,\ t > 0$) with any solid circle, X, to generate a t interval, $\mathcal{I}_X = [t_X^{min}, t_X^{max}]$ (or possibly the empty set). Denote the ray-intersection operator by $\bullet\!\!\rightarrow$, so that

$$\bullet\!\!\rightarrow A = \mathcal{I}_A,$$

and use the fact that $\bullet\!\!\rightarrow$ passes through CSG operands,

$$\bullet\!\!\rightarrow (A \text{ op } B) = (\bullet\!\!\rightarrow A) \text{ op } (\bullet\!\!\rightarrow B).$$

Given the CSG model M from part (a), derive an expression for $\bullet\!\!\rightarrow M$, the set of intersected t intervals.

- *Answer:* For starters, consider the t interval $\mathcal{I}_D$ produced by the ray on D (note that this is indeed a single interval or $\emptyset$ since D is convex as it is the intersection of three circles):

$$\mathcal{I}_D = \bullet\!\!\rightarrow D = \bullet\!\!\rightarrow (A \cap B \cap C) = (\bullet\!\!\rightarrow A) \cap (\bullet\!\!\rightarrow B) \cap (\bullet\!\!\rightarrow C) = \mathcal{I}_A \cap \mathcal{I}_B \cap \mathcal{I}_C.$$

Similarly, M can be represented as

$$\begin{aligned} \bullet\!\!\rightarrow M &= \bullet\!\!\rightarrow \{(A \cap B) \cup (B \cap C) \cup (C \cap A) - (A \cap B \cap C)\} && (1) \\ &= ((\bullet\!\!\rightarrow A) \cap (\bullet\!\!\rightarrow B)) \cup ((\bullet\!\!\rightarrow B) \cap (\bullet\!\!\rightarrow C)) \cup ((\bullet\!\!\rightarrow C) \cap (\bullet\!\!\rightarrow A)) - (\bullet\!\!\rightarrow D) && (2) \\ &= (\mathcal{I}_A \cap \mathcal{I}_B) \cup (\mathcal{I}_B \cap \mathcal{I}_C) \cup (\mathcal{I}_C \cap \mathcal{I}_A) - \mathcal{I}_D. && (3) \end{aligned}$$

Finally one could write this in terms of t intervals explicitly, but this doesn't tell us much more.

Problem 5: FOX News Moon Hoax (10 pts)

FOX Television Network recently[a] advanced the claim that NASA's Apollo Moon missions were a hoax, largely due to alleged "anomalies" in Apollo photos. Here is an excerpt:

FOX narrator: *"In this picture the Sun is directly behind the astronaut, his figure should be a silhouette, yet even the smallest characteristics of his suit are recognizable."* A photographic expert then says: *"He seems like **he's** standing in the spotlight. I can't explain that."*

Can *you* defend NASA's image? What factors are contributing to the brightness of the astronaut's figure?

[a] FOX Television Network, "Conspiracy Theory: Did We Land on The Moon?" February 15, 2001. (see http://users.erols.com/igoddard/moon01.htm)

- *Answer:* For context, please see the higher resolution version of this photo taken during the Apollo XVI mission on 1972-04-21 with 70mm film (NASA Photo ID: AS16-114-18423):

  Title: Astronaut Charles Duke photographed collecting lunar samples at Station 1.

  Description: Astronaut Charles M. Duke Jr., lunar module pilot of the Apollo 16 lunar landing mission, is photographed collecting lunar samples at Station no. 1 during the first Apollo 16 extravehicular activity at the Descartes landing site. This picture, looking eastward, was taken by Astronaut John W. Young, commander. Duke is standing at the rim of Plum crater, which is 40 meters in diameter and 10 meters deep. The parked Lunar Roving Vehicle can be seen in the left background.

  Clearly the sun is shining down on Duke, and the question is asking how can the front of him (which is clearly facing away from the sun, and thus shadowed) be so bright?

  First, notice that there is essentially just one luminous source (the sun), since the Duke's ground shadow is pitch black. In the high-res image you can see that it is totally black (except for the lens flare artifacts). This essentially rules out the Earth (or stars) as a major light source contributor in *this* particular photo. However, if the Earth were in the sky, since it is much larger than the moon, it could be a major light source: people estimate that it is approximately 68.4X brighter than the moon is in our sky (See `http://web.archive.org/web/20010802062000/http://www.forteantimes.com/artic/97/moon.html`). Clearly there is no atmosphere to fill in the shadow like on Earth.

  The very dark shadow is consistent with the claim that most of the light on the astronaut is coming from ground reflection. The reflection in clearly visible in the astronaut's visor. You can also see a little bit of the Apollo lunar module in the visor–that's what took the picture, and no, there wasn't any flash. Consequently, most of the illumination of the front of the astronaut is due to indirect (global) illumination. You can also see darkening at the feet of astronaut (aside from just moon dirt stains) which is consistent with ground reflection.

To make matters more extreme, the indirect illumination is extra bright at nonnormal reflection angles because of the unusual roughness of the moon's regolith. For example, recall the Oyen-Nayar reflectance model for rough surfaces in an earlier homework when discussing the moon's near-constant shading when viewed from Earth.

The astronaut's space suit is also a very reflective white color to provide shielding, so a fair amount of light will be reflected toward the camera.

Finally, some people may have mentioned that the camera person may have had a light, but there was not such a light on the lunar module, and you can't see a reflection of any unusually bright light in the astronaut's visor–mostly just the moon's surface and the lunar module.

<*Intermission!*>

A font walks into a bar.

The bartender says, “Welcome! We don’t see many of your *type* around these parts.” <*Ha ha ha. Boo...*>

<*Exeunt*>

Problem 6: Triangle Meshes (20 pts)

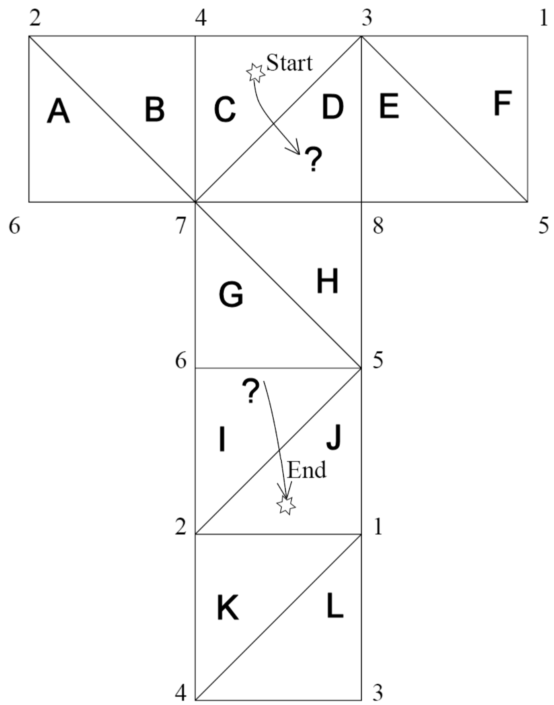

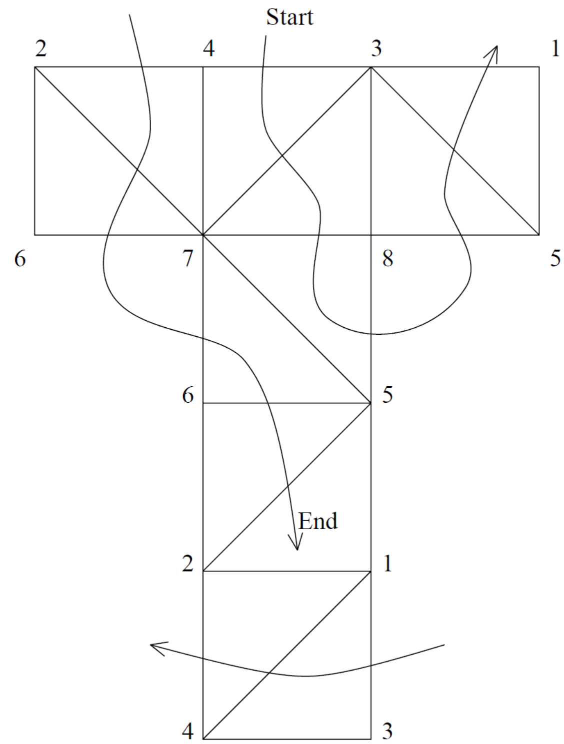

Consider the 12-triangle mesh of a cube, with indexed vertices, shown here unfolded and flattened–analogous to a cube map texture. Look at it for a moment to reconstruct it in your mind.

(a) **Find a triangle strip** containing all triangles that starts with triangles C, D,. . . and ends with . . . I, J. State your strip as a sequence of 12 triangles:

C, D, ?, ?, ?, ?, ?, ?, ?, ?, I, J.

Remember adjacent strip triangles must share an edge. (If you get really stuck, state another strip, else you can state two strips.)

- *(a) Answer:* See image. The sequence is

  C, D, H, E, F, L, K, B, A, G, I, J.

  FYI, the vertex sequence associated with the tri-strip is

  4, 3, 7, 8, 5, 3, 1, 4, 2, 7, 6, 5, 2, 1.

(b) **Find three nonoverlapping triangle fans** that span the surface. State each fan as the vertex followed by a sequence of edge-sharing triangles, e.g.,

3: C, D, E, F.

- *(b) Answer:* There are many possibilities. Here is one:

```
7: G, H, D, C, B, A
2: K, J, I
3: E, F, L
```

Problem 7: Quaternion Splines (10 pts)

Spherical linear interpolation, or SLERP, is widely used to interpolate key-framed rotations represented as quaternions. Recall that a unit quaternion can be represented as a four-vector $\mathbf{q} = (\cos\theta, \hat{\mathbf{v}}\sin\theta)$, and that the SLERP of two unit quaternions, $\mathbf{q}_0$ and $\mathbf{q}_1$, is

$$\mathsf{SLERP}(\mathbf{q}_0, \mathbf{q}_1; \alpha) = \mathbf{q}_0 \frac{\sin(\Omega(1-\alpha))}{\sin\Omega} + \mathbf{q}_1 \frac{\sin(\Omega\alpha)}{\sin\Omega}, \qquad \alpha \in [0,1],$$

where $\cos\Omega = \mathbf{q}_0 \cdot \mathbf{q}_1$.

One criticism of SLERP is that interpolating a temporal sequence of quaternions, $\mathbf{q}_0, \mathbf{q}_1, \mathbf{q}_2, \ldots$, sampled at equally spaced times $t = 0, 1, 2, \ldots$, can produce nonsmooth rotations at quaternion nodes.

(a) Consider the case of interpolating the quaternion pair $(\mathbf{q}_0, \mathbf{q}_1)$ for $t \in [0,1)$, followed by $(\mathbf{q}_1, \mathbf{q}_2)$ for $t \in [1,2]$. State the function for this piecewise-SLERP quaternion curve,

$$\mathbf{Q}(t) = \begin{cases} ?, & t \in [0,1), \\ ?, & t \in [1,2]. \end{cases}$$

- *Answer:* Here you need to specify the piecewise definition of $\mathbf{Q}(t)$ explicitly as a function of t–not some undefined $\alpha(t)$ values. Defining $\cos\Omega_{ij} \equiv \mathbf{q}_i \cdot \mathbf{q}_j$, we have

$$\mathbf{Q}(t) = \begin{cases} SLERP(\mathbf{q}_0, \mathbf{q}_1, t), & t \in [0,1), \\ SLERP(\mathbf{q}_1, \mathbf{q}_2, t-1), & t \in [1,2]. \end{cases} \tag{4}$$

$$\mathbf{Q}(t) = \begin{cases} \{\mathbf{q}_0 \sin(\Omega_{01}(1-t)) + \mathbf{q}_1 \sin(\Omega_{01} t)\} / \sin\Omega_{01}, & t \in [0,1), \\ \{\mathbf{q}_1 \sin(\Omega_{12}(2-t)) + \mathbf{q}_2 \sin(\Omega_{12}(t-1))\} / \sin\Omega_{12}, & t \in [1,2]. \end{cases} \tag{5}$$

  Note that there are *two* different Ω values here. Also, notice that $\mathbf{Q}(t)$ is continuous as $t = 1$ since

$$\lim_{t\to 1^-} \mathbf{Q}(t) = \lim_{t \to 1^+} \mathbf{Q}(t) = \mathbf{q}_1.$$

(b) Show that this quaternion spline has a derivative discontinuity at $t = 1$.

- *Answer:* In this question you needed to establish that

$$\lim_{t\to 1^-} \mathbf{Q}'(t) \neq \lim_{t\to 1^+} \mathbf{Q}'(t).$$

  First you needed to take the derivative of $\mathbf{Q}(t)$. You can start by taking the partial derivative of $SLERP(\mathbf{q}, \mathbf{r}, \alpha)$ with respect to α, since

$$\mathbf{Q}'(t) = \frac{d\mathbf{Q}(t)}{dt} = \begin{cases} \frac{\partial}{\partial\alpha} SLERP(\mathbf{q}_0, \mathbf{q}_1, \alpha)\Big|_{\alpha=t}, & t \in [0,1), \\ \frac{\partial}{\partial\alpha} SLERP(\mathbf{q}_1, \mathbf{q}_2, \alpha)\Big|_{\alpha=t-1}, & t \in [1,2]. \end{cases}$$

  Differentiating, and using $\frac{d}{dx}\sin(ax) = a\cos(ax)$, we find (note that this doesn't require any knowledge of quaternions or SLERP whatsoever)

$$\begin{aligned} \frac{\partial SLERP(\mathbf{q}, \mathbf{r}, \alpha)}{\partial\alpha} &= -\mathbf{q}\frac{\Omega\cos(\Omega(1-\alpha))}{\sin\Omega} + \mathbf{r}\frac{\Omega\cos(\Omega\alpha)}{\sin\Omega} \\ &= \frac{\Omega}{\sin\Omega}\left\{-\mathbf{q}\cos(\Omega(1-\alpha)) + \mathbf{r}\cos(\Omega\alpha)\right\}, \end{aligned}$$

with $\cos\Omega = \mathbf{q}\cdot\mathbf{r}$. It follows by the chain rule (to support $\alpha = t$ or $t-1$) that the left and right limits of dQ/dt at $t=1$ are

$$\lim_{t\to 1^-} \frac{d\mathbf{Q}(t)}{dt} = \frac{\Omega_{01}}{\sin\Omega_{01}}\left\{-\mathbf{q}_0 + \mathbf{q}_1\cos\Omega_{01}\right\}, \quad (6)$$

$$\lim_{t\to 1^+} \frac{d\mathbf{Q}(t)}{dt} = \frac{\Omega_{12}}{\sin\Omega_{12}}\left\{-\mathbf{q}_1\cos\Omega_{12} + \mathbf{q}_2\right\}, \quad (7)$$

and since these are not equal, in general, there is a derivative discontinuity.

<*Intermission!?!*>

A freshly shaded "exam pixel" walks onto the frame buffer.

"Whew! I'm sure glad I got by that Early-Z test." <*Ha ha ha. Boo...*>

<*Exeunt (for good)*>