Optimizing Open Inventor™ Applications

Release 2.0

Copyright © 1993, Silicon Graphics, Inc. All rights reserved.

Silicon Graphics, the Silicon Graphics logo, and IRIS are registered trademarks, and Inventor and Graphics Library are trademarks of Silicon Graphics.

Specifications are subject to change without notice.

Section 1 Introduction

This document describes a process for determining what is limiting the performance of your Open Inventor application, and provides suggestions on how to improve its performance.

The following sections are provided:

Section 2: Benchmarking Tips- Discovering where an application is spending its time is not a matter of random trial and error. This section describes a process for effectively investigating and improving performance.

Section 3: Optimizing Rendering- For each step in the rendering process, this section:

- Briefly describes that step.
- Presents a method for determining how much time your application is spending in that step.
- Describes techniques for reducing that time.

Section 4: Optimizing Everything Else- This section suggests ways of structuring your scene and application for maximum performance when doing tasks other than rendering (for example, picking or modifying the scene).

For more information on Inventor programming in general and on specific nodes, see *The Inventor Mentor*.

Section 2 Benchmarking T ips

Performance tuning is a necessary chore when developing any application, like fixing bugs or eliminating memory leaks. Proper organization and planning can make this chore much quicker and more pleasant.

The general approach to optimizing your application should be:

1. Set performance goals
2. Devise a method to measure your application's performance
3. Determine where your application is spending its time
4. Modify your application to reduce the bottleneck
5. Finally, measure your application against your goal again and repeat steps 2 and 3 until it meets your goal.

The rest of this section explains these steps in more detail, concentrating on optimizing the rendering performance of Inventor applications.

2.1 Measure Performance

It is very important to have an objective way of measuring your application's performance; watching your application run and trying to see if it *seems* faster is likely to cause you to waste your time on insignificant optimizations.

Adding code to your application that measures the number of polygons in your scene and how fast they are being rendered is fairly simple; see the source code for the `ivview` application in `/usr/share/src/Inventor/tools/ivview`, for example.

The `osview` program can also be useful. The *swapbuf* number in the *Graphics* section will also tell you how many frames per second your application is getting (assuming that your application is double-buffered, and that it is the only double-buffered application running). **Note:** don't confuse `osview` with its graphical counterpart, `gr_osview`.

Be sure to keep good records of your application's performance before you start optimization. Comparing before and after performance numbers will give you (and your boss) a feeling of accomplishment, and will ensure that you are not making things worse.

2.2 Determine Bottlenecks

Most applications spend most of their time executing a small part of the code. Optimizing a procedure that is taking up only 5% of the total time is probably not worthwhile; even if you manage to double the performance of the procedure, the application will speed up by only 2.5%. In fact, on some machines graphical operations can occur in parallel. For example, filling in polygons and transforming polygon vertices might occur at the same time. If the bottleneck is in the vertex transformation stage, increasing the pixel fill time may not increase performance at all! **Find the bottlenecks first, and then work on improving them**.

Finding bottlenecks is an experimental science. You should first come up with a theory on where the bottleneck might be, and then you should devise an experiment that will prove or disprove that theory. Create experiments that isolate one narrow part of your application's performance, and make sure you understand what you are measuring every time you run an experiment.

Section 3 and Section 4 describe the following bottlenecks, show how to determine the amount of time your application is spending in each of them, and give suggestions on improving them:

- Section 3.3: Window Clear Bottleneck
- Section 3.4: Traversal Bottleneck
- Section 3.7: Pixel Fill Bottleneck
- Section 3.8: Transforming Vertices Bottleneck
- Section 3.9: State Change Bottleneck
- Section 3.10: Culling Bottleneck
- Section 3.11: Level Of Detail Bottleneck

- Section 4.2: Memory Usage
- Section 4.4: Action Construction and Setup
- Section 4.5: Notification
- Section 4.6: Picking and Handling Events

If you have some reason to suspect one of these bottlenecks more than another, you can perform these sections in any order; just make sure you always know what you are measuring and keep good records of your experiments.

2.3 Modify Your Application

Be careful when you apply a performance optimization. Make sure that your modification is actually an improvement: don't assume that all of the suggestions made in this (or any other) document will automatically apply to your application. For example, using Inventor render caching will usually increase performance. However, if the extra memory used by render caching causes your application to run out of memory and start swapping memory to disk, it might actually hurt performance.

Again, keep good records. Record what you did and how much it improved performance. **Try to minimize the number of things you change at any one time**; for example, if you make two "optimizations" and performance goes up by 10%, the speedup might be caused by a 5% improvement for each optimization, or might be caused by a 100% speedup caused by one optimization and a 90% slowdown caused by the other! It is very tempting to read a document like this, make lots of changes, and then see if the application gets faster. Not only does this waste time, but it could be counter productive.

2.4 Are We Done Yet?

One of the most frustrating things about optimizing an application's performance is that it can be difficult to determine when you are done. Once you have successfully eliminated one bottleneck, something else becomes the factor limiting performance. Before spending more time on optimization, you should ask yourself:

- **Is the application fast enough?** If you have a specific performance number that you must meet (it must render at least 15 frames per second) then this decision is easy. If not, come up with a specific performance number; for example, most users will find a frame rate of 10 frames per second acceptable for an interactive program (more is always better, of course). Try to get your application to run this fast for a typical scene.
- **Are your expectations reasonable?** If the absolute top speed for drawing polygons on your machine is 60,000 unlit, non-depth-buffered triangles per second and you are trying to get 10 frames per second while drawing 6,000 lit, depth-buffered triangles, you will be disappointed. Write short OpenGL benchmark programs, or feed test scene graphs to `ivview`, to help set your expectations.
- **Is your application running at 60 or 72 frames per second?** Double-buffered programs will never render faster than the refresh rate of your monitor, and Inventor sets its animation sensors to fire at a maximum of 60 frames per second (for example, the realTime global field is updated 60 times a second by default).
- **Do you need to experiment on different machines?** Different machines will have different bottlenecks; on a machine with very fast graphics, the bottleneck is more likely to be either in the application's code or in Inventor's code. On a machine with slow graphics and a fast CPU, the bottleneck is much more likely to be inside the OpenGL calls. If your application will be used on different types of machines, **make sure performance is acceptable on all of them.**

2.5 The Five Performance Commandments

1. Be scientific
2. Keep good records
3. Find bottlenecks
4. Change one thing at a time
5. Test all the types of machines your application supports

Section 3 Optimizing Rendering

The main goal of performance tuning is to make the application look and feel faster. However, just because the goal is to make the application render faster, don't assume that rendering is the bottleneck!

3.1 Is Rendering the Problem?

The first step is to modify your application so that it does everything it normally does except render, and then measure its performance. An easy way of getting your application to do everything except for render is to insert an `SoSwitch` node with its `whichChild` field set to `SO_SWITCH_NONE` (the default) above your scene. So, for example, modify your application's code from:

```
myViewer->setSceneGraph(root);
```

To:

```
SoSwitch *renderOff = new SoSwitch;
renderOff->ref();
renderOff->addChild(root);
myViewer->setSceneGraph(renderOff);
```

This experiment gives an upper limit on how much you can improve your application's performance by increasing rendering performance. If your application doesn't run much faster after this change, then rendering is not your bottleneck; see section Section 4 for information on optimizing the rest of your application.

3.2 Isolate Rendering

Once you have determined that your application is spending a significant amount of time rendering the scene, the next step is to isolate rendering from the rest of the things your application does. This makes it easier to determine where the bottleneck in rendering is occurring. The easiest way of doing this is to write your scene to a file and then use the `ivRender` program (source code in Appendix I; see the comments in Appendix I for ftp/www availability) to perform a series of rendering experiments.

The code for writing out your scene will look something like this:

```
SoOutput out;
if (!out.openFile("myScene.iv")) { ... error ... };
SoWriteAction wa(&out);
wa.apply(root);
```

If you have created your own node classes, either create DSO's for them (see Section 2.12 of the inventor_dev product's release notes for more information on creating

DSO's) or call their `initClass()` methods just after the call to `SoDB::init()` in the `ivRender` source and link their `.o` files into `ivRender`.

`ivRender`'s camera control is very simplistic: it does a `viewAll()` of the scene and just spins the scene around in front of the camera when benchmarking. If you have a sophisticated walk-through or fly-through application that uses level of detail and/or render culling, you should modify `ivRender` so that its camera motion, level of detail switching, and render culling are more appropriate for your application. For example, add something like the following to the beginning of your scene:

```
TransformSeparator {
     Rotor { rotation 0 1 0 .1 speed .1 }
     Translation { translation 100 0 0 }
     PerspectiveCamera { nearDistance .1 farDistance 600 }
}
```

... and modify `ivRender` so that it either updates the `realTime` global field in its rendering loop (so that the `SoRotor` node rotates) or so that it directly modifies the `SoRotor` node. Do the following to update `realTime`:

```
SoSFTime *realTime =
 (SoSFTime *)SoDB::getGlobalField("realTime");
realTime->setValue(SbTime::getTimeOfDay());
```

3.3 Window Clear Bottleneck

The first step in the rendering process is clearing the window. It is easy to forget about this step, but depending on the size of your application's window and the type of machine you are running on, clearing the window can take a surprisingly long time. If your application's main window is typically 800 by 800 pixels big, run `ivRender` like this:

```
ivRender -an -w 800,800 myScene.iv
```

The options given mean:

- -w 800,800 : Make `ivRender` use a window that is 800 pixels wide by 800 pixels high.
- -a : Do not actually apply an `SoGLRenderAction`. The purpose of this experiment is to see how fast the machine can clear the color and depth-buffers only, to give an upper limit on rendering performance.
- -n : Don't bother adding a camera, light, or transform, since we are only trying to measure screen clear speed.

For example, on an Indigo2 Extreme running IRIX 5.2, this experiment gives 48 frames per second for a 1,000 by 1,000 window. Put another way, if the application must run at 30 frames per second, then over 60% of the time you have to render a frame would be spent clearing the window.

Unfortunately, if clearing the window is taking too much time there is not a lot you can do about it. One possibility is to make your window's default size smaller (while still allowing users to resize the window if necessary).

3.4 Traversal Bottleneck

After the previous experiment, you know how much time your application is spending clearing the color and depth buffers. The next experiment is designed to find out how much time Inventor is spending **traversing** your scene. Traversal is the process of walking through the scene graph, deciding which render method needs to be called at each node. Run `ivRender` again on your scene, this time applying the `SoGLRenderAction` and using a reasonable camera:

```
ivRender myScene.iv
```

Run the experiment again, this time with **render caching** turned off:

ivRender -r myScene.iv

- -r : Do not build render caches at `SoSeparators`.

See *The Inventor Mentor*, Chapter 9, for an explanation of render caching. Basically, render caching eliminates Inventor's traversal of the scene. The difference between these two experiments is the amount of traversal overhead in your scene, when Inventor cannot build render caches. By default, Inventor will automatically build render caches at `SoSeparator`s where appropriate. However, if most of your scene is changing, or if your scene is not organized for efficient caching, Inventor may not be able to build render caches, and traversal might be the bottleneck in your application.

Compare the results of the cached experiment with your application's performance. If your application is about as fast as the fully cached experiment, your application is using caches efficiently and traversal during rendering is not your bottleneck; skip to Section 3.7 to look for other possible bottlenecks. Also skip to Section 3.7 if the cached and uncached times are fairly close to each other; this means that Inventor is able to render your scene efficiently even without the benefit of caching .

The next two sections discuss two ways of reducing traversal overhead: creating more efficient scene graphs and structuring your scenes so that Inventor is able to build render caches.

3.5 Reducing Traversal Overhead

Take a look at the number of nodes and triangles that are in your scene, as reported by `ivRender`. The number of nodes should be small compared to the number of triangles. Get rid of any redundant nodes, get rid of `SoLabel` nodes by using Open Inventor's naming feature (see Chapter 3 of *The Inventor Mentor* for more information). Try to arrange your scene so that property nodes affect a large part of the scene; for example, if all of the objects in your scene are solid, put one

SoShapeHints node at the top of the scene instead of inserting it before each of the objects.

Inventor is much more efficient at rendering multiple triangles if they are all part of one node. For example, you could create a multifaceted polygonal shape using a number of different coordinate and face set nodes, as shown in the lower half of Figure 1. A much better technique is to put all the coordinates for the polygonal shape into one SoCoordinate node, and the description of all the face sets into a second SoFaceSet node, as shown in the upper half of Figure 1.

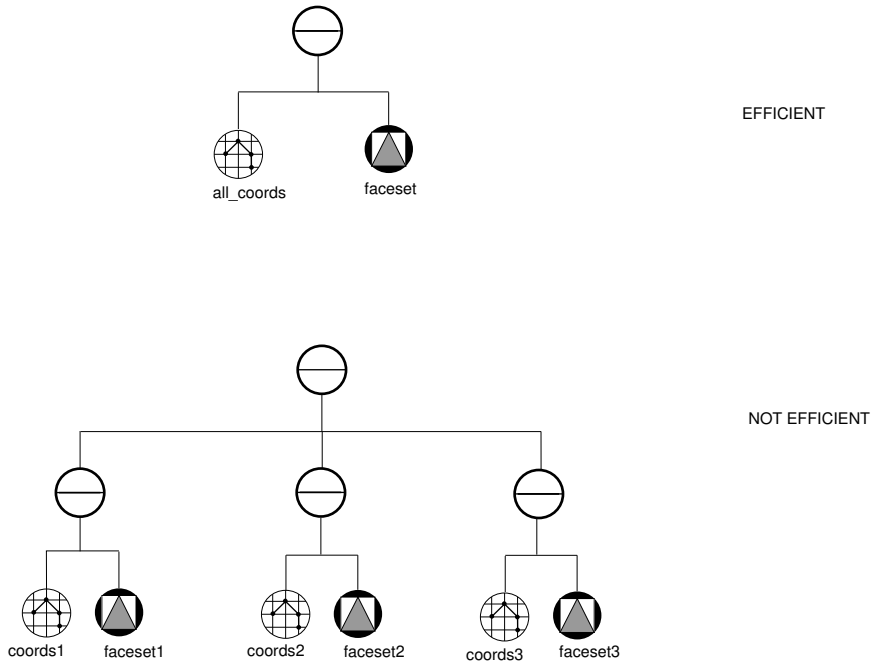

Figure 1 Condensing face sets into fewer nodes

Using fewer nodes to get the same picture reduces traversal overhead for scenes that cannot be cached. Arranging your scene so that each node makes more OpenGL calls also gives Inventor a better opportunity to produce efficient OpenGL calls.

For some applications, you should consider implementing your own nodes that implement the functionality of a subgraph of your scene. For example, a molecular modeling application might implement a BallAndStick node with fields specifying the atoms and bonds in a molecule, instead of using the more general SoSphere/SoCylinder/SoMaterial/SoTransform/SoGroup nodes. If the molecular modeling application changes the molecule frequently so Inventor cannot cache the scene, using a specialized node could make traversal orders of magnitude faster (e.g. a simple water molecule scene graph with three atoms and two bonds might consist of 20 nodes; replacing this with a single BallAndStick node would make traversal 20 times faster). The BallAndStick node could also perform application-specific

optimizations not done by Inventor, such as not drawing bonds between spheres whose radii were large enough that they intersected, sorting the spheres and cylinders by color, etc. See *The Inventor Toolmaker* for complete information on implementing your own nodes.

3.6 Organizing the Scene for Caching

You may be able to organize your scene so that Inventor can build and use render caches even if part of it is changing. You should be aware that the following things inhibit caching:

- Changing fields in the scene will destroy caches inside all `SoSeparator`s above the node that changed. Even fields that do not affect rendering, such as fields in the `SoLabel` or `SoPickStyle` nodes, will destroy caches if they are changed.
- The `SoLevelOfDetail` node will break caches above it whenever either the camera or any of the matrix nodes affecting it change. Make the children of the `SoLevelOfDetail` node `SoSeparator` nodes, so that they will be cached. See Section 3.11 for more information on efficient use of the `SoLevelOfDetail` node.
- Any shape using `SCREEN_SPACE` complexity will break caches above it whenever the camera or any of the matrix nodes affecting it change.
- The `SoText2` node will break caches above it whenever the camera changes (in order to correctly position and justify each line of text, it must perform a calculation based on the camera). Since most applications change the camera frequently, you should try to separate `SoText2` nodes from the other objects in your scene, to allow the other objects to be cached.
- Changing the override status of properties at the top of the scene, or changing global properties such as `SoDrawStyle` or `SoComplexity` that affect the rest of the scene, inhibits efficient caching. `SoSeparator` nodes will build multiple render caches (by default, a maximum of two) to handle cases in which a small set of global properties are changed back and forth, but you should avoid continuously changing a global property; for example, putting an engine on the value field of an `SoComplexity` node at the top of your scene would be bad for caching.
- Generating default normals for vertex-based shapes that cannot be cached is very expensive because it must be done every time the shape is rendered. Use `ivquicken` to generate normals for your objects if they will not be cached.

For more information on Inventor's render caching, see Chapter 9 of *The Inventor Mentor*.

3.7 Pixel Fill Bottleneck

A common bottleneck on low-end machines is drawing the pixels in filled polygons. This is especially common for applications that have just a few large polygons, as opposed to applications that have lots of little polygons.

To determine whether or not your application is fill-limited, run `ivRender` again, this time instructing it to insert an `SoDrawStyle` node that will override the draw style in the scene so that everything is drawn as points. Do this by creating an "override.iv" file containing this:

```
#Inventor V2.0 ascii
DrawStyle { style POINTS }
```

and then running `ivRender` like this:

```
ivRender myScene.iv override.iv
```

`ivRender` reads the second file given, sets the override flag for all nodes in that file, and inserts them before the nodes in the first file.

The difference between these two experiments is the amount of time being spent filling in the polygons in your scene. If the difference is large, you can speed up pixel fill by:

- Rendering your scene, or parts of your scene, in wireframe or as points when possible. Viewers have "move wireframe" and "move points" modes built-in for exactly this case.
- If you are using texturing, see if turning off texturing or decreasing the `SoComplexity::textureQuality` field increases performance. A `textureQuality` of zero will turn off texturing; other values will cause OpenGL to use different filters for texturing, which can affect polygon fill performance. Put a `Complexity` node in the "override.iv" file and rerun `ivRender` to see if this makes a difference on your hardware with your scene.
- Some machines can fill flat-shaded polygons faster than Gouraud-shaded polygons. Inventor never changes the shade model from OpenGL's default, which is Gouraud shading. Try adding a call to `glShadeModel(GL_FLAT)` after the call to `glXMakeCurrent` in `ivRender` and re-running your scene. If there is a significant performance increase, you should determine which parts of your scene look good flat-shaded and insert `SoCallback` nodes to turn flat-shading on and off for those parts of your scene. Note that an `SoLightModel` node set to `BASE_COLOR` lighting does not turn on flat shading, since unlit primitives may still have different colors at each vertex and be Gouraud shaded.
- `SCREEN_DOOR` transparency (the default) is faster than the blended transparency on some machines (it is slower on some machines, too). Use the `setTransparencyType()` method on either `SoXtRenderArea` or `SoGLRenderAction` to change the transparency type.

3.8 Transforming Vertices Bottleneck

Modify the "override.iv" file created in the previous experiments, changing the `SoDrawStyle` from `POINTS` to `INVISIBLE`, and rerun `ivRender`. The difference between the `POINTS` and `INVISIBLE` experiments is the amount of time spent lighting, fogging, and transforming the vertices of the objects in your scene. If you

find that this time is a significant portion of the time it takes to render a frame, you can do the following to optimize per-vertex operations:

- Use fewer vertices in your objects, and use `SoComplexity` to turn down complexity for Inventor's primitive objects. If you are using a machine with hardware-accelerated texturing, texturing can be used to add visual complexity with very few vertices.
- Create less detailed versions of your objects and use `SoLevelOfDetail` nodes so that fewer vertices are drawn when objects are small. Use an empty `SoInfo` node as the lowest level of detail so that objects disappear when they get very small. A good rule of thumb for choosing levels of detail is that the switch between levels of detail should be fairly obvious if you are concentrating on the object; for most applications, the user concentrates on objects in the foreground and will not notice background objects "popping" between levels of detail. Beware that `SoLevelOfDetail` nodes will cause smaller caches to be built, which may slow down traversal. See Section 3.11 for more information on efficient use of level of detail.
- Make your vertices simpler. Try to use `OVERALL` rather than `PER_VERTEX` material binding (putting a `MaterialBinding { value OVERALL }` node in the "override.iv" file and re-running `ivRender` is a quick way of seeing if this might help). Turn off fog. Note that these suggestions are machine-specific; on machines with a lot of hardware for accelerated rendering, fogged vertices may be no slower than plain vertices. Be sure to do a quick `ivRender` test before spending time modifying your application or scene.
- Use fewer light sources, and use simpler lights (a `DirectionalLight` is simpler than a `PointLight`, which is simpler than a `SpotLight`). If possible, put lights inside `Separators` so that they affect only part of the scene, increasing performance for the rest of the scene.
- If you are using `SoFaceSets` or `SoIndexedFaceSets`, try using `ivquicken` to convert them into `SoIndexedTriangleStripSets`, which draw more triangles with fewer vertices. Note that `ivquicken` won't be able to create a mesh if your objects have sharp facets or `PER_FACE` material or normal bindings.
- Watch out for expensive primitives with lots of vertices, like `SoText3` and `SoSphere`. `ivRender` reports the number of triangles in your scene; make sure the number is reasonable for your desired performance.
- Organize your scene graph so that objects that are close to each other spatially are under the same `SoSeparator`, and turn on `renderCulling` so that Inventor won't send those objects' vertices when the objects are not in view. See *The Inventor Mentor*, Chapter 9, for more information on render culling.

3.9 State Change Bottleneck

You might have run the `DrawStyle INVISIBLE` experiment from the last section and been surprised that your scene still doesn't draw as quickly as you expect, even though nothing is being drawn. When `DrawStyle` is `INVISIBLE`, Inventor's shapes

do nothing, but Inventor's property nodes will still issue OpenGL calls, since the DrawStyle might be reset later during traversal and OpenGL's state must be correct when that happens. So, if rendering isn't as fast as you expect with a cached, INVISIBLE scene then executing the GL calls to change properties may be the culprit (other possibilities are level of detail switching and render culling, which the next two sections discuss).

To optimize state changes, first take a look at the structure of your scene. The following little shell command is a good way of getting an overview of your scene's structure:

```
ivcat myScene.iv | egrep '[{}]|USE'
```

Common causes of state change bottlenecks are SoTransform, SoMaterial, SoTexture2, and SoShapeHints nodes. If you see that your scene contains a lot of these nodes (or any other property node), create an override.iv file for ivRender that overrides that property. Of course, without these nodes, ivRender will render your scene incorrectly. Use the following techniques to get both faster and correct rendering if you have a state-change bottleneck:

- If overriding SoMaterial nodes increases performance, try the following:
  If your SoMaterial nodes have multiple values in them, note that having multiple values in just one of the fields is faster than having multiple values in several of the fields. For example, indexing into an SoMaterial node with 10 diffuseColors and 1 ambientColor will be faster than indexing into an SoMaterial node with 10 diffuseColors and 10 ambientColors.
  Changing between materials with different shininess values is much more expensive than changing any of the other material properties.
  If you are using shapes with a materialIndex field, try to sort their parts by material index to minimize material changes. For example, try to change: `IndexedFaceSet { materialIndex [ 0,1,0,1,0,1,0,1 ] ... }` to: `IndexedFaceSet { materialIndex [ 0,0,0,0,1,1,1,1 ] ... }` (this will only work for PER_PART or PER_FACE material bindings, of course).
- If overriding SoTransform nodes increases performance, see if you can use ivquicken to get rid of them. See the man page for ivquicken for more information.
- Use SoRotation, SoRotationXYZ, SoScale, or SoTranslation nodes instead of the general SoTransform node. Don't bother doing this if you would have to replace the SoTransform node with more than one of the simpler nodes to get the same transformation.
- If overriding SoTexture2 nodes increases performance and your machine supports textures in hardware, you may be running out of texture memory. Unlike Inventor 1, Open Inventor does not automatically use the same texture if two SoTexture2 nodes happen to have the same filename. Be sure to use instancing (DEF and USE in the file format) instead of creating several texture nodes with the same filename.

- Try organizing your scene to minimize state changes. For example, if you have several objects that use the same texture, group them together so that the texture is traversed only once. Note that organizing your scene for efficient state changes might interfere with attempts to organize your scene for efficient render culling, where you want objects that are spatially near each other to be grouped together. You will have to find the optimal balance for your application.
- Use render culling to avoid traversing objects with state changes that are outside the view. See Chapter 9 of *The Inventor Mentor* for an explanation of culling.
- Create simpler levels of detail with fewer state changes that can be used for objects that appear small.

3.10 Culling Bottleneck

If your application is using render culling, it might be spending most of its time deciding whether or not objects should be culled. To find out whether this is the case, use `prof`, `pixie`, or the CaseVision/WorkShop Performance Analyzer tools to look for a lot of CPU time being spent in the `SoSeparator::cullTest()` or `SoBoundingBox::transform()` routines. See the man pages for `pixie`, `prof`, or `cvspeed` for information on using these tools.

If a large percentage of the rendering time is being spent doing cull tests, try to re-organize your scene so that more triangles are culled for each culling `SoSeparator`. For example, if you have a city scene with thousands of buildings, it may be better to perform one cull test for each city block rather than the thousands of cull tests needed to decide whether or not each individual building is visible. Doing this will also allow Inventor to build larger render caches, which will increase traversal speed.

Also, remember that render culling breaks render caches when the camera or transformation matrices change, so double-check to make sure that no `SoSeparators` above an `SoSeparator` doing render culling have their `renderCaching` fields set to `ON`.

3.11 Level Of Detail Bottleneck

Like render culling, if your application is using `SoLevelOfDetail` nodes it might be spending a significant amount of time deciding which level of detail should be drawn. One way of testing to see if this is the case is to temporarily replace all of the `SoLevelOfDetail` nodes in your scene with `SoSwitch` nodes set to traverse the highest level of detail. Then run `ivRender` again (with `DrawStyle` still set to `INVISIBLE` so that the high-complexity objects aren't actually drawn) and compare the results. If the `SoSwitch` node scene is much faster, try doing the following:

- Use a level of detail node with a simpler, faster level of detail test. See Appendix 2 for the source to `LODD`, which has a very simple, very fast test.
  Note: There is a bug in Inventor 2.0 which causes `SoLevelOfDetail` nodes to be slow when used with render culling. This bug will be fixed in Inventor 2.0.1; use the `LODD` node as a workaround until then.

- Try to group objects together so that one level of detail test determines the level of detail for several objects. For example, if you have a group of 10 buildings that are near each other, use one level of detail node instead of 10 level of detail nodes. Doing this will also make it easier for Inventor to build larger render caches, which will increase performance by increasing traversal speed.
- Remember that level of detail nodes break render caches when the camera or transformation matrices change, so make sure that no `SoSeparators` above an `SoLevelOfDetail` have their `renderCaching` fields set to `ON`.

3.12 Making it Feel Faster

Sometimes it is worthwhile to sacrifice features temporarily to make your application **seem** faster to the user. Inventor has several features that make this easier:

- Use the `SoGLRenderAction::setAbortCallback()` method to interrupt rendering before the entire scene has been drawn. For this to be most effective, you must organize your scene so that the most important objects are drawn first, and you should only abort when it is important that rendering happen quickly, even if the rendering is not complete, such as when the user is interactively manipulating the scene.
- Use one of the “Move ...” draw styles if you are using a viewer, so that a simpler version of the scene is drawn when the user is interacting with the viewer.
- Use the start and finish callbacks of manipulators and components to temporarily modify the scene to make it simpler while the user is interacting with it.

Section 4 Optimizing Everything Else

Once you have determined that rendering is not your bottleneck, or if you have already optimized rendering as much as possible and a significant amount of time is still being spent doing something other than rendering, this section suggests ways of finding other bottlenecks, and suggests Inventor-specific things to look for.

4.1 Useful Tools

The standard performance analysis tools (`prof`, `pixie`, or the CaseVision/WorkShop Performance Analyzer) make performance analysis of the non-graphics part of your application easy. See the man pages for `pixie`, `prof`, or `cvspeed`