Subject: Re: Vector output of idlgrpolygon models
Posted by D D on Wed, 09 Nov 2011 19:28:12 GMT
View Forum Message <> Reply to Message

On Nov 9, 5:27 pm, Karl <Karl.W.Schu...@gmail.com> wrote:
> I was going to point out that Select could be called for each pixel, but you figured that out :-).
>
> But I also didn't point that out because I still don't think it is a perfect solution.
>
> Did you actually send the resulting polygon list to vector output and inspect the results? I think that rendering the resulting polygon list to the display would look pretty good, since you're getting help from the depth buffer. But that won't be the case on a vector device.
>
> If two of your surfaces (toroids) intersect, there are going to be some faces that intersect with each other.
>
>
> EYE
> |
> V
>
> \ /
> \ /
> \ /
> ^ \/
> | /\
> Z / \
> X - - >
>
> Your algorithm would (correctly) select both surfaces and put them both in the list to send to the vector device. That device is going to render then both in the order you supply them and you'll see one surface or the other near the intersection. The part of the second surface to draw that is behind the first drawn surface will draw on top of the first surface, which is incorrect.
>
> But perhaps your data does not do this and you may be fine.
>
> The GL2PS approach is very interesting. It uses the same technique that the IDL Clipboard uses (OpenGL Feedback buffer) and sorts things in a similar way if you pick the simple sort option. I had also mentioned using BSP trees and I note that GL2PS offers a BSP tree sort option. This would address the problem I illustrated above with the intersecting surfaces. These two surfaces would each be split at the intersection line, resulting in four surfaces. The two "back" pieces would not be in the remaining list and would not draw, leaving the two "front" pieces to represent the correct rendering of the intersection. Might still be worth a look if you need this sort of precision.

I've sent a couple of preliminary attempts to vector output and things
do seem to work quite well (the file size decreased by an order of
magnitude, and then some!) however I need to tweak my settings a bit I

think as currently things are a bit over excited and throw away a few more polygons than needed. This should be solvable at the expense of more cpu time.

One puzzle i'm having is to why the execution time of the function I gave previously varies by an order of magnitude between machines. It's possible my setup is at fault: one machine is a fedora box with software rendering whilst the other is a virtual machine running lubuntu on top of vista using hardware graphics. The virtual machine is much slower for this (although much more responsive when interacting with the widget plot).

GL2PS does indeed look interesting, i'd got 80% of the way to compiling Paraview with GL2PS support enabled when I ran out of disk space :S
I think a BSP tree type approach would be interesting, and certainly a good project for me to try sometime but for now I'll stick with the Select method as I can't afford the time (the long document I'm writing is my thesis so trying to keep distractions to a minimum but this IDL challenge keeps nagging at me!)

Thanks,
David
