
Subject: Re: Which Graphics Card?

Posted by [Randall Frank](#) on Sat, 09 Jun 2001 05:29:15 GMT

[View Forum Message](#) <> [Reply to Message](#)

Ok a topic worthy of a monthly rant...

A few notes from someone who has been using a GeForce 3 and a number of GeForce 2 variants (GTS, MX, Quadro and Ultra) for a bit now (as well as a fair helping of non-nVidia cards).

The GeForce 3 is the current "next great" gaming card, but there is no software out there that uses it yet. Thus, it is getting cheaper daily. I have seen the 64MB version for \$349 with a \$50 rebate already. It is a very "interesting" card, but rather immature (one might speculate that a competing card might show up before the apps do). On the average for a large distribution of scientific applications (IDL, EnSight, MeshTV, home brew code), the GeForce 3 was around 20-30% slower than a GeForce 2 Ultra or a GF2 Quadro 2 card. There is a lot of speculation about why this might be and it ranges from immature drivers to some bad RAM design decisions to questions about the hidden costs of the programmable vertex pipeline. But this is all speculation. The biggest issue has been the 3D texturing issue (not an IDL issue), but nVidia is supposed to issue a press release to straighten that out at some point. As hinted at in some of this thread, a quadro version of the GF3 would probably be the best option for an OG IDL user in the long run.

Advice is free, so here is a little bit:

Consider what you want to accelerate. If DG, stop and buy a Matrox card. If OG, decide if your application is dominated by drawing lines or polygons. If polygons, the best price/performance combo I have seen is the Kryo cards (~\$70 for an astonishing fill rate), but you need a very high speed CPU for that chip. The next best price/performer is the GeForce 2 Ultra. If you need lines, then you might consider a FireGL card, a great high quality CAD card, or wait for a "quadro" version of a GeForce 3. If you want to draw images, my advice is to time them. You might be very surprised at what generally wins the OG image drawing race.

IMHO the current GeForce 3 has little to offer the IDL user in its current state to warrant its higher cost, hopefully that will change...

Rick Towler wrote:

- >
- > You had to ask....
- >
- ...
- > nVidia segments their market by offering consumer and professional products
- > based on the same core technology (consumer lines are based on the
- > GF/GF2/GF3 and the professional lines are based on the Quadro). AFAICT, the
- > only real difference between the consumer cards and the pro cards is that
- > the pro cards sport a different BIOS, a few resistors and better OpenGL
- > drivers. How much better? I wish I knew since that might drive my next
- > purchasing decision. Maybe someone with a Quadro would be willing to do
- > some benchmarks?

Actually, the BIOS is the same (no comment on the resistors) and the chip spin is better (quadros are clocked higher). The big win is in line drawing performance. Overall polygon filling, I can get 20-25M tris/sec on a GF2 GTS and around 30M tris/sec on a Quadro version of the same chip, but both cards are seriously memory bandwidth bottlenecked. I don't have IDL numbers for the two cards as all my Quadro cards are in machines running Liunx.

- >
- > When purchasing a Geforce2/Quadro based card there are a few things to look
- > out for. All Geforce chips are limited to some extent by memory
- > bottlenecks. Card manufacturers have used this fact to segment the market
- > to the point of mass confusion. When shopping for a GF2, you will find 64
- > and 128 bit SDR based cards and 64/128 bit DDR based cards with RAM speeds
- > that vary from 7.5 to 4 ns. The best performance will be had with the
- > 128bit DDR based cards with the fastest RAM available. These products are
- > generally labeled "ultra" as in the WinFast Geforce2 Ultra.

The Ultra is acutally a little different chipset and is a bit more than just faster RAM. The Ultra is basically a Radeon killer...

- >
- > If you are looking at a Geforce3 things are a little simpler. This card
- > started hitting the streets last month and there are only a few variants
- > available. All seem to be shipping with 128 bit DDR running at 4ns. Your
- > only options look like TV out, DVI-I, and the amount of RAM (64 vs 128MB).
- > I highly doubt that you would ever make use of 32 let alone 64 MB of video
- > memory so don't waste your money on the 128 MB version.

Depends on the app. The 128MB version is great for volume rendering.
Now

if nVidia would just release nvfence extensions for texture maps, we would be really in business.

>
> A bonus option with the GF cards is that they offer full scene anti-aliasing
> (at the cost of raw speed). I have found this feature to be indispensable
> when rendering 3d scenes for animation and am now rendering final animations
> on my Geforce based workstations exclusively. (on side note, with newer
> driver revisions make sure your desktop is set to 32 bit and the default
> color depth for textures is set to "desktop color depth" or 32bpp otherwise
> IDL will bomb when opening an object graphics window when anti-aliasing is
> enabled)

Actually, if you like this option, please RUN out to the store and buy a GF3 or a 3DFX 5500. The FSAA support on a GF2 is really a hack and the quality is laughable compared to the 3DFX and GF3 implementations.

>
> I know I sound like a commercial here but stick with me....
>
> nVidia has also released the geforce2go chip. This is the first real step
> forward for portable 3d in years. For those looking for a portable only
> solution or for a laptop that can actually render high poly scenes this is
> your only choice. Don't think you are going to get this in that ultra slim
> vaio though.

They work pretty nicely in the Dell 8500s. A little hot, but this laptop and an 802.11b card means I can clobber my UT friends from the forest of my back yard...

>
...
> For the penguin's, nVidia is producing drivers for XFree86 4.x. Last time I
> checked they were lagging behind windoze platforms in performance but
> quality has been steadily improving.

The performance is pretty much a dead heat right now with raw, fenced triangle performance leaning toward Linux. Things like the GLX protocol do not work, but all the extensions are there and it works like a charm.

Does that help murky up the waters a bit?

--
rjf.

Randy Frank | ASCI Visualization
Lawrence Livermore National Laboratory | rjfrank@llnl.gov
B451 Room 2039 L-561 | Voice: (925) 423-9399

Livermore, CA 94550

| Fax: (925) 423-8704
