From: Joey Kendall <joey.kendall@gmail.com>
Subject: **Catching up and an interesting new programming idea**
Date: May 19, 2009 11:47:22 AM MDT
To: chuck.allison@gmail.com

Hey Professor, how goes your time on the faculty senate? I'm sure it keeps you plenty busy.

I realized the other day that I never updated you after I last talked to you and after my interview. The interview went really well (or so it seemed to, more on this later) they really didn't ask me nearly as much programming questions as I was expecting. After the interview went they told me that they were just getting ready to ship a project and so they might not respond for a week or two and that if they hadn't by that time I should send them an email to check the status. So I did just that, but I never heard back from them. Ever. Seems that with that kind of follow through that may not have been a company I wanted to work for anyways.

I also wanted to drop you a little line about something I just read about the other day. It's called CUDA, (I'll be brief in case you've heard about it already). It is basically an API to get a hold of the GPU (certain nVidia cards) and use it for general purpose processing. One of the great things about it is that because of the architecture of GPUs it allows for great parallelism. In one article I read they took a program that takes a height map and calculates the corresponding normal map and converted it to use the CUDA SDK.

They ran the program on 1 thread then threading it for 2 and 4 threads (their test machine was a mobile core 2 duo) processing an image 2048 x 2048 achieving times 1419 ms, 749 ms, and 593 ms respectively. Then running the CUDA enabled app under 2 different data sizes they achieved times of 109 ms and 94 ms on their test machine (a 8600M GT). Running it on a higher end graphics card (an 8800 GTX) and tweaking the data sizes again they achieved a time of 31 ms. If you want to check out more of the nitty gritty this is the article I mentioned http://www.tomshardware.com/reviews/nvidia-cuda-gpu,1954.html they cover more about architecture and some history about GPGPU usage.

I thought it might make an interesting subject to bring up if you're ever talking about new things that could be added to the curriculum (because I'm sure there aren't already a billion ideas vying for a spot in the line up *wink*).