Second Generation Transform and Lighting (T&L)

Setting the Bar for 3D Graphics

With the GeForce2™ and Quadro2™ families of GPUs, second generation transform and lighting engines offer blazing fast graphics processing speeds. Animated characters running on a GPU are lifelike and detailed, with complex facial expressions and smooth movements. Developers and designers can create worlds lush with organic life and architectural details that we take for granted in the real world. With accelerated transform and lighting capabilities, NVIDIA GPUs create unsurpassed 3D experiences for everyone from the web surfer to the multimedia enthusiast.

What is a GPU?
A Graphics Processing Unit (GPU) offloads all transform and lighting calculations from your CPU, freeing it up for other functions such as physics and artificial intelligence. Not only will your graphics look better and run more smoothly with an NVIDIA GPU, but your computer's performance will improve as well.

Why T&L?

Geometrically complex worlds require exceptional processing speed. And transform and lighting are two very mathematically intense processes. Combined, transform and lighting radically enhance photo-realism to create worlds that come alive on your screen. NVIDIA GPUs use separate transform and lighting engines so that each can run at maximum efficiency. The transform engine converts 3D data from one frame of reference to the next. Every object that is redisplayed, and even some that are not, must be transformed every time the scene is redrawn. Lighting effects then provide high visual impact by enhancing the realism of the scene.

How does transform work?

Transform performance dictates how precisely software developers can "tessellate" the 3D objects they create, how many objects they can put in a scene and how sophisticated the 3D world itself can be. To tessellate an object means to divide it into smaller geometric objects, such as polygons. The images below are examples of a sphere tessellated by different degrees:

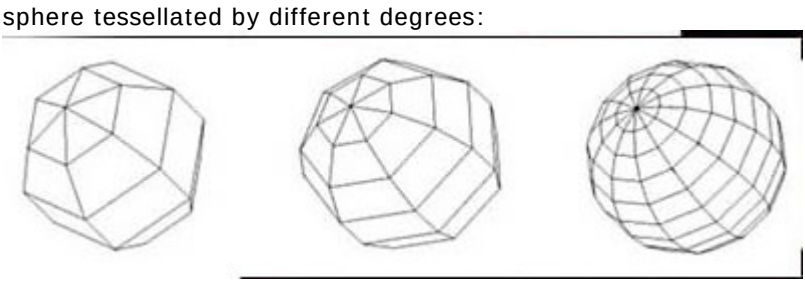

Each of the images above represents the same sphere, but the image on the far right is clearly the most realistic of the three. It has been carved up into five times as many polygons as the sphere on the far left, and therefore requires five times the transform performance as the sphere on the left. That may not seem very important for one sphere, but because hundreds to thousands of objects are often displayed in scenes, without a GPU those objects have to share the limited processing power of the CPU, forcing developers to budget processing tasks.

Now with an NVIDIA GPU transform calculations are offloaded from the CPU, allowing more detailed objects with higher polygon counts to be processed more quickly. With transformation a jungle scene can have lots of trees and bushes—rather than just a single tree—and each tree can consist of many leaves created by thousands of polygons. Since the GPU relieves the CPU of the burden of calculating the transforms, you will be able to view scenes rich with complex objects that look real and move like their real-life counterparts. Not only will the objects and characters be complex, but many more can exist.

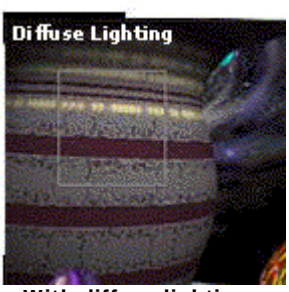

With diffuse lighting, we can't see where the light is concentrated on the object.
By adding specular lighting, we can see that the light is shining directly on the object.

How does lighting work?

The human eye is more sensitive to changes in brightness than it is to changes in color—which means that an image with lighting effects communicates more information to a viewer more efficiently. The discrete lighting engine on an NVIDIA GPU calculates distance vectors from lights to objects and from objects to a viewer's eyes within 3D scenes. Lighting calculations are an effective way to add both subtle and not-so-subtle changes in brightness to 3D objects in a manner that mimics real-world lighting conditions.

Specular highlights move on the object if the viewer or the object moves relative to the light source. For this reason they cannot be pre-computed or static. Specular lighting is also important for representing different materials for objects in a 3D scene. A silk shirt looks different than a cotton shirt—even if they are the same color—and will reflect light differently. Specular lighting combined with texture mapping creates more realistic objects because they have the visual properties of real materials. Only a GPU with a dedicated hardware lighting engine can support specular lighting without a severe loss in performance.

Diffuse vs. Specular Lighting

Lighting is divided into two main categories: diffuse and specular. Diffuse lighting assumes the light hitting an object scatters in all directions equally, so the brightness of the reflected light does not depend at all on the position of the viewer. For example, when sunlight hits a playground, the light is everywhere. Specular lighting is different because it depends on the position of the viewer as well as the direction of light and orientation of the object being lit. For example, the beam from a flashlight will bounce off a quarter differently than off a blade of grass. Specular lighting captures the mirror-like properties of an object so effects such as reflection and glare are achievable. (http://www.OrpheusComputing.com/)