Parallel Processing

Adam Gritt

Annotated References

Ordered as used in the presentation.

Sun, C. 2002. *Intel Pentium 4 2.8GHz Processor Review*. Available at: http://www.pcstats.com/articleview.cfm?articleid=1244&page=3
A. C. 2004. *Intel Pentium 4 3.2E GHz Prescott Processor Review*. Available at: http://www.pcstats.com/articleview.cfm?articleid=1535&page=7
D. M. 2005. *Intel Pentium D 840 Dual-Core LGA775 Processor Review.* Available at: http://www.pcstats.com/articleview.cfm?articleid=1838&page=13
Page, M. 2007. *Intel Core 2 Duo E6600 2.4GHz Processor Review.* Available at: http://www.pcstats.com/articleview.cfm?articleid=2097&page=7
Apong, J. 2009. *Intel Core 2 Duo E8400 3.0GHz 1333MHz FSB Processor Review.* Available at: http://www.pcstats.com/articleview.cfm?articleid=2394&page=6
Klar, J. 2010. *Intel Core i7 980X - 32nm Hexacore is here.* Available at: http://www.nordichardware.com/en/component/content/article/359-intel-core-i7-980x-32nm-hexacore-is-here.html?start=7

The above six references were used to provide a normalised comparison of desktop processors over the past decade. The figure used is the time taken to calculate one million digits of PI using the SuperPI application. Though the accuracy may be off by a small margin, it is likely to be in the order of 1-2%. The data was normalised and cross-referenced with the official release dates and number of cores in each processor. A graph was created to illustrate this data.

Intel, 2008. Intel Turbo Boost Technology in Intel Core Microarchitecture (Nehalem) Based Processors. Available at: http://download.intel.com/design/processor/applnots/320354.pdf?iid=tech_tb+paper

A whitepaper descibing Intel's "Turbo Boost" feature. The feature allows processor cores to run at above-standard speeds under certain conditions. The technology is an attempt to mitigate the effects of a divergence from single-threaded processor architectures.

de Campos, A. 2006. Animated Fractal Mountain. Available at: http://en.wikipedia.org/wiki/File:Animated_fractal_mountain.gif

An animated gif showing an iterative fractal algorithm. Such an technique is used as an example of a serial problem, would cannot be parallelised.

Masood, J. 2009. *NVIDIA CUDA*. Available at: http://www.hardwareinsight.com/nvidia-cuda/

This page contains an graph showing the rapid divergence in performance of typical CPUs and GPUs, when measured in FLOPs. In 2003, the performance difference is 2-3 fold, however by 2008 the difference is around 8 fold.

Super Micro Computer, Inc., 2010. *1U Superservers.* Available at: http://www.supermicro.com/products/system/1U/6016/SYS-6016T-NTF.cfm

A technical summary of a typical 1U server blade. Used to create a feature comparison between a typical server blade and NVIDIA's GPU counterpart.

NVIDIA, 2010. *Tesla S2050 / S2070 GPU Computing System.* Available at: http://www.nvidia.com/object/product_tesla_S2050_S2070_us.html

A detailed description and specification of nVidia's proprietary GPU Computing System. Used to to compare with the specifications of a typical 1U CPU-based server.

Huang, J. 2009. PhysX Destruction Demo from GPU Technology Conference (GTC) Keynote. Available at: http://www.youtube.com/watch?v=N1FOnpzUzZY

An example usage of GPGPU for real-time rigid body physics and shattering. The video perfectly illustrates an highly parallelisable problem. Binary space partitioning aside, each edge of each object in the scene must be compared in a collision test resulting in an extreme number of identical and parallel operations.

NVIDIA, 2010. *NVIDIA Collaborates With Weta to Accelerate Visual Effects for Avatar.* Available at: http://www.nvidia.com/object/wetadigital_avatar.html

Weta CTO Paul Ryan and NVIDIA Research senior architect Jacopo Pantaleoni discuss their collaboration for the production of Avatar. The article describes the problems which were solved using GPGPU technology; the extremely computationally heavy calculations of image-based lighting and occlusion. A 25 to 100 fold improvement in performance is claimed for this pre-computation step. It is stated that the output of these calculations is stored in an intermediary format, to be later used for a beauty-pass in RenderMan. Paul Ryan states Weta's intent to incorporate this technique into their production pipeline for their future project *TinTin*.

Bakkum, P. and Skadron, K. 2010. *Accelerating SQL database operations on a GPU with CUDA.* In *Proceedings of the 3rd Workshop on General-Purpose Computation on Graphics Processing Units* (Pittsburgh, Pennsylvania, March 14 - 14, 2010). GPGPU '10, vol. 425. ACM, New York, NY, 94-103. Available at: http://doi.acm.org/10.1145/1735688.1735706

Researchers from the University of Virginia collaborate to reimplement the SQLite virtual machine on the GPU. Some parallels can be drawn between this research, and the success of CUDA. The researchers reimplemented a widely understood language (SQL) using GPU acceleration. This compares with the introduction of CUDA, which reimplemented the C99 language for use in development of GPGPU applications. A 20 to 70 fold improvement in performance was achieved for SELECT queries, dependent upon the size of the result set.

NVIDIA, 2010. CUDA Community Showcase. Available at: http://www.nvidia.com/object/cuda_apps_flash_new.html

NVIDIA lists example projects that have utilised CUDA. Of particular interest are the categories listed, the vast majority of which are relevant to either CGI or games. That said, there are many less obvious application categories, including Finance and Oil & Gas exploration.

World Community Grid. Nutritious Rice for the World. Available at: http://www.worldcommunitygrid.org/research/rice/overview.do

World Community Grid is an organisation characterised by it's slogan "Technology Solving Problems"; a reference to the vast scope for which computers have use. In particular, the three fields of medical research, accelerated food production, and climate change. The most interesting example, in my opinion, is the use of GPGPU for protein folding of rice strains. It is hoped that the results of this research will yield clues as to how rice can be genetically modified in order to provide higher crop yields, increase disease and pest resistance, and improve the nutritional benefits of rice.

Pixar, 2006. GPU Shader Plugins for RenderMan. Available at: http://hradec.com/ebooks/CGI/RPS_13.5/prman_technical_rendering/AppNotes/GpuShaderPlugins.html

Use of RenderMan's plugin interface, allowing RSL (Renderman Shader Language) "hotspots" to be computed on the GPU. The plugin is meant to supplement, rather than replace, RSL shaders.

Patney, A. and Owens, J. D. 2008. Real-time Reyes-style adaptive surface subdivision. *ACM Trans. Graph.* 27, 5 (Dec. 2008), 1-8. Available at: http://doi.acm.org/10.1145/1409060.1409096

Two researchers from University of California, Davis present a GPU implementation of a REYES-style adaptive surface subdivision (a.k.a. 'dicing'). This represents an important step towards the goal of implementing a fully GPU-based REYES renderer. The level of parallelism achieved allowed the GPU-based implementation to perform in real-time, and claims "several times better" performance than Pixar's RenderMan.