Large Synoptic Survey Telescope

www.lsst.org

Accelerating LSST Source Catalog Simulations with Graphics Processing Units

Mario Juric[1,2], K. Cosic[3], D. Vinkovic[3], Z. Ivezic[4]

[1]Harvard Univ., [2]Hubble Fellow, [3]Univ. of Split, Croatia, [4]Univ. of Washington

The use of graphics processing units (GPUs) for computationally intensive tasks has gained a significant amount of traction in the field of high performance computing. For highly parallelizable applications, GPUs provide a factor of 20-100 increase in speed over conventional CPUs which could dramatically improve the performance of the LSST data management and simulation systems. Here, we present the results of porting a source catalog simulation code (galfast) to NVIDIA GPUs that resulted in typical speedups of 25-100x over the prior CPU-based code. Galfast is a generator of realistic mock catalogs of astronomical sources (e.g., stars in the Galaxy), originally written to interpret SDSS observations. Within LSST, galfast stellar catalogs are used as inputs to image simulations (ImSim), as well as by several Science Collaborations to estimate LSST yields of various astrophysical objects. It is available as a web service at http://hybrid.mwscience.net. The inner routines of galfast were modified to run on NVIDIA GPUs (~3000 lines of CUDA C code, four weeks of porting time). The effort resulted in typical *speedups of order 100x* (single Tesla S1070 GPU vs. a single contemporary CPU core). The time to generate simulated catalogs of Milky Way stars decreased from days to under an hour, becoming entirely bound by disk data storage speed. In an extreme example, a catalog computation that used to take 1hr was reduced to 0.37 seconds. While algorithmic changes are partially responsible for the speedup, we conservatively estimate that at least a factor of 25 speedup came from the use of GPUs. The success of this effort vividly demonstrates the ability of GPUs to bring significant speedup to LSST modeling and science. Work is ongoing to bring GPU-based acceleration to LSST ImSim code.

LSST: Observing and Measuring 10^{10} Stars and Galaxies

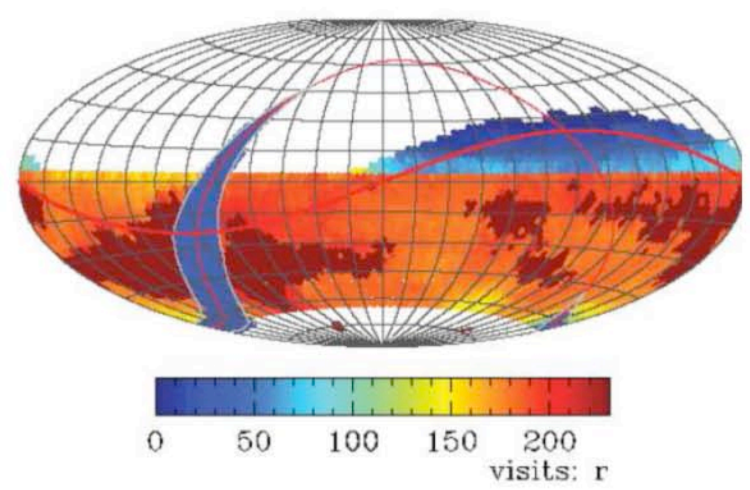

TABLE 1
THE LSST BASELINE DESIGN AND SURVEY PARAMETERS

Quantity	Baseline Design Specification
Optical Config.	3-mirror modified Paul-Baker
Mount Config.	Alt-azimuth
Final f-Ratio, aperture	f/1.234, 8.4 m
Field of view, étendue	9.6 deg^2, 318 $m^2 deg^2$
Plate Scale	50.9 μm/arcsec (0.2" pix)
Pixel count	3.2 Gigapix
Wavelength Coverage	320 – 1080 nm, *ugrizy*
Single visit depths[a] (5σ)	23.9, 25.0, 24.7, 24.0, 23.3, 22.1
Mean number of visits	70, 100, 230, 230, 200, 200
Final (coadded) depths[a]	26.3, 27.5, 27.7, 27.0, 26.2, 24.9

[a] The listed values for 5σ depths in the *ugrizy* bands, respectively, are AB magnitudes, and correspond to point sources and zenith observations.

The LSST will cover about 30,000 deg^2 with $\delta < +34.5$, imaged multiple times in six bands (ugrizy), with zenith 5σ depth for point sources of r=24.7. About 90% of the observing time will be devoted to a wide-fast-deep survey mode which will observe a 20,000 deg^2 region and yield a coadded map up to r=27.7. These data will result in databases including 10 billion galaxies and a similar number of stars.

An empirically constrained simulation of this dataset, the "source catalog", should be used as the basis for LSST planning: image simulations, Science Book estimates, data management, as well as operations simulations. However, the generation of this dataset proved difficult: for example, using existing tools the generation of stellar source catalogs would take on order of ~weeks per realization.

By using Graphics Processing Units (GPUs) as math accelerators, we were able to speed up this process by typical factors of >100. This enables the efficient generation of multiple source catalogs, and the full exploration of allowable model parameter space. Ultimately, the same code will be used to fit the dynamical models of the Milky Way to actual survey data, thus providing a direct scientific as well as operational benefit.

galfast: The Fast, Scalable, Synthetic Stellar Catalog Generator

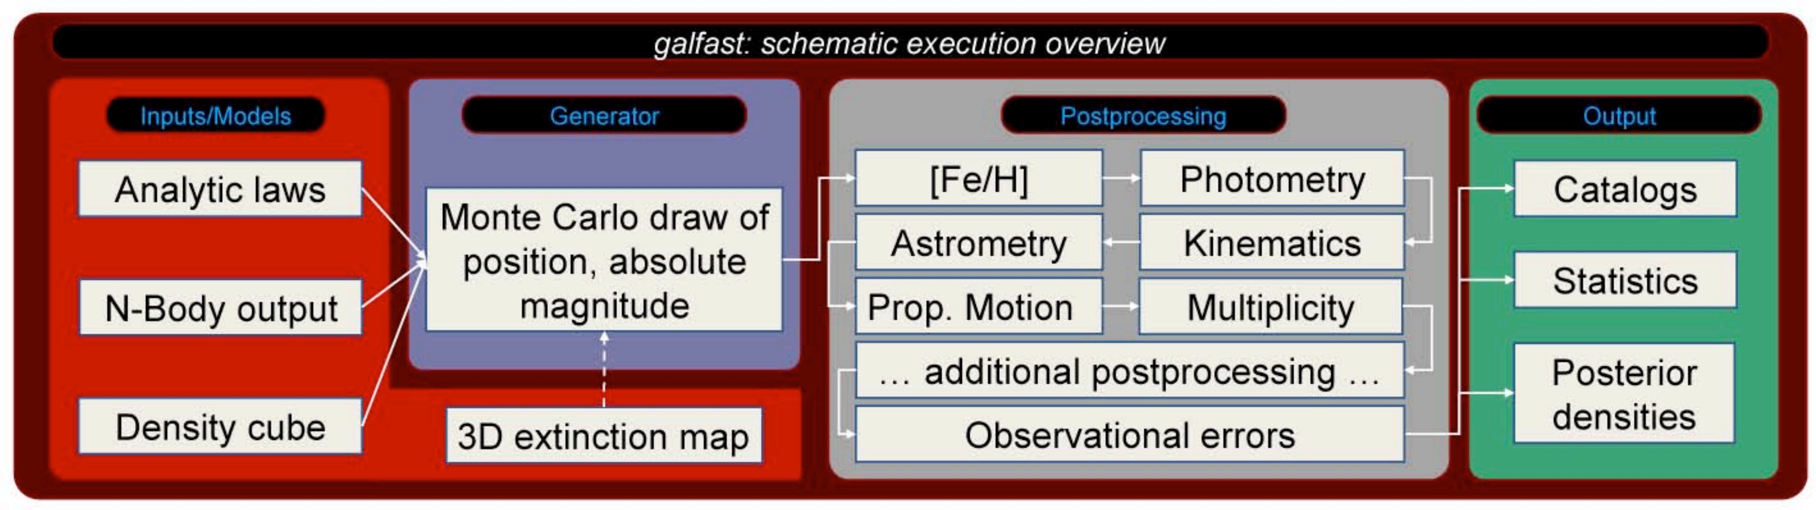

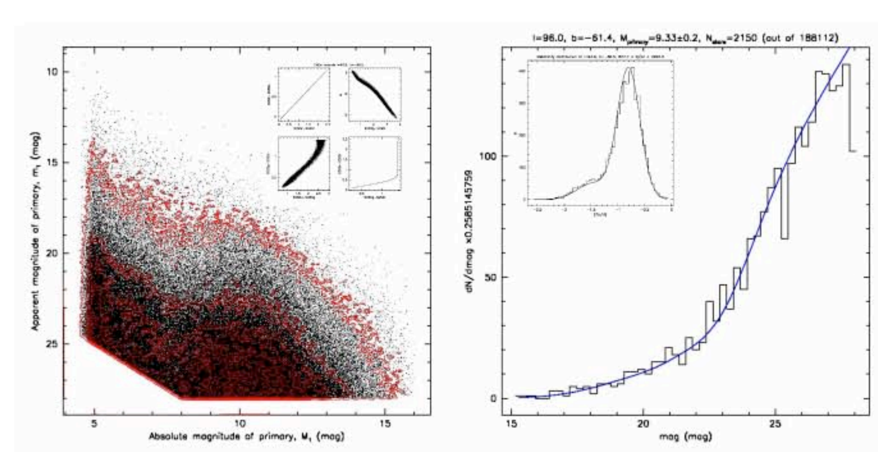

galfast (top) is a new, fast, and scalable synthetic star catalog generator designed for simulating and fitting point source catalogs from large surveys. While similar to the widely used TRILEGAL (Girardi 2005) or the Besançon model (Robin et al. 2003), in *galfast* we prefer to use empirical calibrations (rather than models) for the necessary inputs (usually derived from SDSS; Juric et al. 2008, Ivezic et al. 2008, Bond et al. arXiv:0909:0013, Covey et al. 2007, Bochanski et al.). This makes it suitable as a basis for LSST source catalog simulations.

Left: Sample color-magnitude and color-color diagrams (left), and density and [Fe/H] profiles (right), of a *galfast*-generated catalog.

GPU Computing: Massively Multi-core Scientific Computing

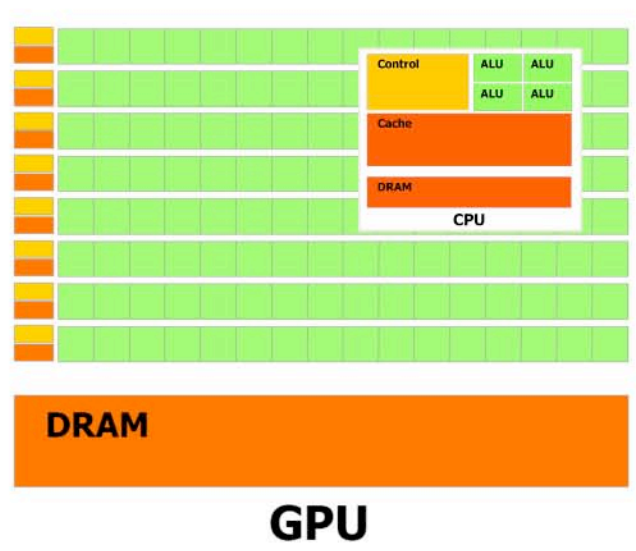

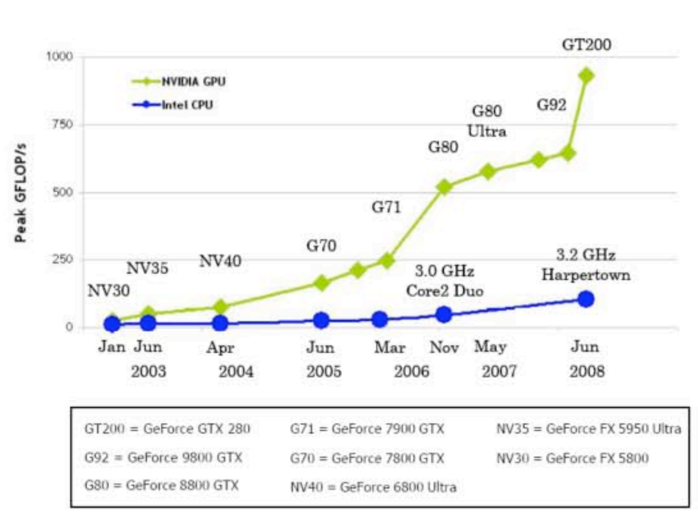

- **GPUs: Massively Multi-Core Processors**
  - Hundreds of simple cores (ALUs)
  - Minimal program control logic – all execute the same program on different data
  - Explicit memory hierarchy (global, local, texture, …)
  - Zero-overhead thread context switching
- **10x more computational power**
  - 80% of GPU transistors are devoted to math
  - Fast basic arithmetic (single precision; DP ~8x slower)
  - Hardware implementation of common functions (sin, cos, exp, ln, …)
- **Economies of scale**
  - **~$400/GPU (!!)**
- **Challenge**
  - Code for a thousand core shared memory machine
  - Take into account the memory hierarchy
  - Lower level issues: caches, latencies, bank conflicts, etc...
- **Programming interface**
  - "C for CUDA": C with extensions for GPU

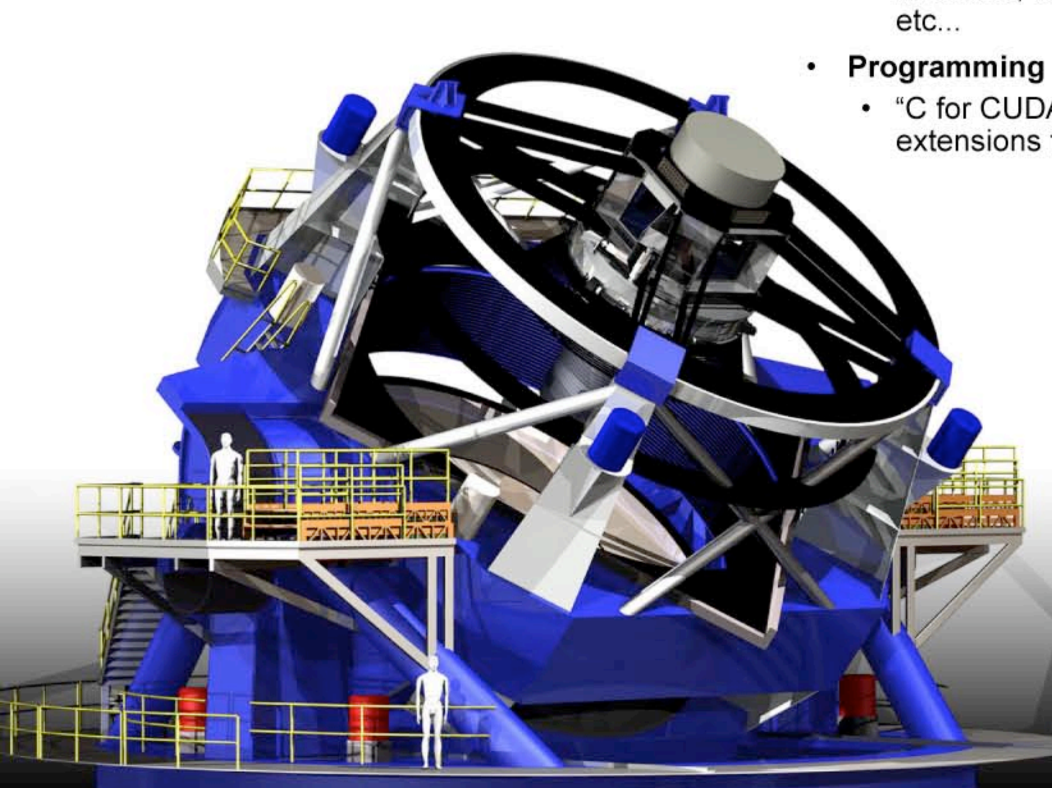

GPU Implementation of *galfast*

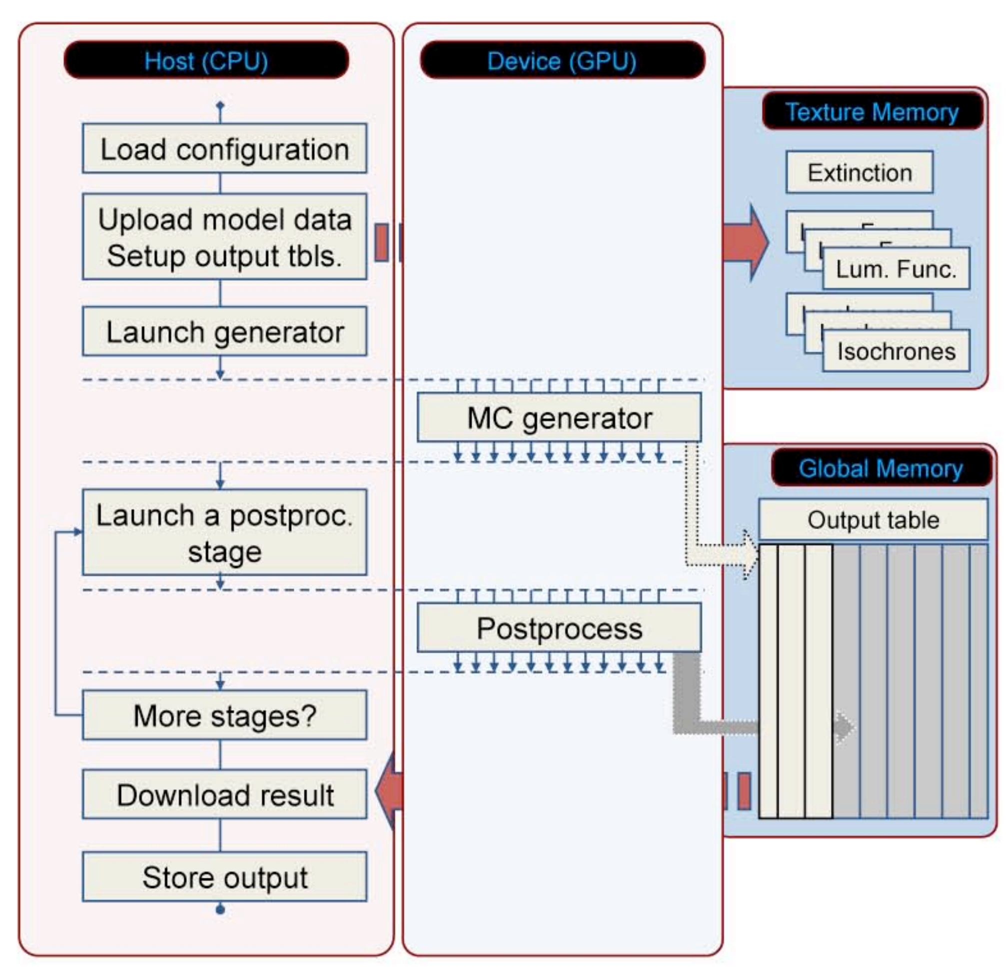

galfast samples stars from the input $\rho(X,Y,Z,M)$ distribution, and calls a series of postprocessing modules that assign further observables (e.g. ugrizy magnitudes). This design allows for straightforward parallelization to N threads, where each thread samples 1/N[th] of the input (X,Y,Z,M) space. Note: (X,Y,Z) are the spatial coordinates, and M is the absolute magnitude.

The most computationally intensive parts of galfast (MC sampler and postprocessing modules, ~3000 lines of C++ code) were ported to NVIDIA GPUs using C for CUDA (NVIDIA, 2009). The effort took less than a month.

The above graphic illustrates the organization of the GPU accelerated port. Besides parallelization, we also make use of the GPU texturing hardware to accelerate sampling from 3D extinction maps, luminosity functions, and photometric isochrones. The results are accumulated to GPU-card memory (4GB) until the very end when they're transferred to the host and stored.

Benchmarks and Results

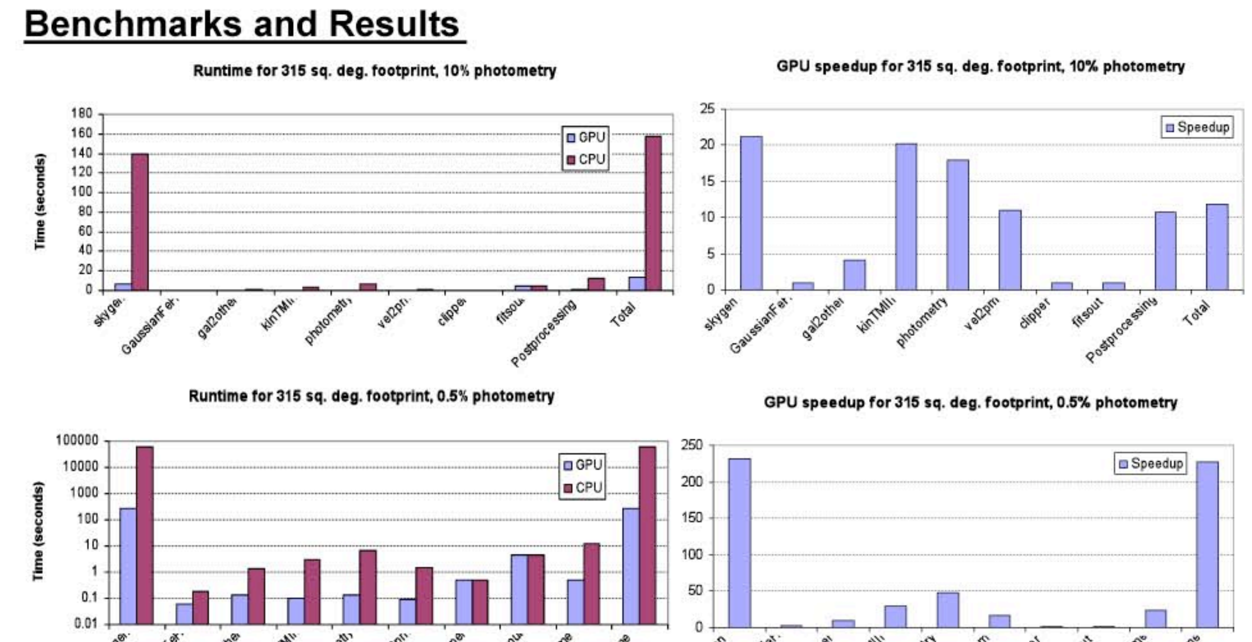

Above: The graphs to the left show the execution time of various *galfast* stages running on a single Tesla S1070 GPU (blue) vs. a single core of an Intel Xeon E5405 2.0GHz CPU (purple). The graphs to the right show the ratio of the two (the GPU speedup). In all cases the code was generating a flux-limited catalog in a 20 deg diameter pencil beam towards the North Galactic pole; for the top row the photometric accuracy was set to 0.1 mag, while it was 0.005 for the bottom (as needed for generation of realistic LSST source catalogs).

The GPU accelerated *galfast* outperforms the CPU version by a factor of ~20x for the low accuracy case, and a factor of >200x in the (relevant!) high accuracy scenario. The difference can be attributed to initial kernel startup costs associated with the GPU and CUDA runtime, that dominate runtimes of short kernels. The overall speedup is due to a) parallel computation on 240 cores b) hardware implementation of texture lookups, and c) fast arithmetic and transcendental function implementation on the GPU.

Conservatively extrapolating from the above, with a single Tesla S1070 GPU we will be able to generate and store a complete LSST stellar source catalog in ~5 hours.

References

Covey et al., 2007, AJ, 134, 2398

Bond et al., subm., arXiv:0909.0013

Girardi et al., 2005, A&A, 436, 895

Ivezic et al., 2008, ApJ, 684, 287

Ivezic et al., 2008, astro-ph/0805.2366 0

Juric et al., 2008, ApJ, 673, 864

Robin et al., 2003, A&A, 409, 523

http://nvidia.com/cuda

http://mwscience.net/galfast

LSST is a public-private partnership. Design and development activity is supported in part by the National Science Foundation. Additional funding comes from private gifts, grants to universities, and in-kind support at Department of Energy laboratories and other LSSTC Institutional Members.