Cosmological calculations on the GPU

Deborah Bard, Matthew Bellis, Hasmik Yepremyan

SLAC NATIONAL ACCELERATOR LABORATORY

Northern Illinois University

Cosmological measurements often involve the calculation of non-trivial quantities over increasingly large datasets. The next generation of survey telescopes will yield information for billions of galaxies. The scale of the datasets, and the type of calculations involved, are ideal models for use of the GPU. We present two cosmological measurements, and describe the implementation and improvements found with the GPU.

Measuring the structure of the universe

The structure of matter in the universe can be characterised by the clustering of galaxies, and is measured using the angular correlation function. The two-point angular correlation function is based on galaxy counts, and measures the excess of pairs of galaxies as a function of separation, compared to to a random distribution. It provides an important test of galaxy formation models, and can differentiate between cosmological models describing the distribution and evolution of matter in the universe.

Galaxies are a good tracer of the matter distribution in the universe, but visible matter comprises only approx. 4% of the mass-energy in the universe. Dark matter cannot yet be detected directly, but the effect of its gravitational field can be measured indirectly using gravitational lensing.

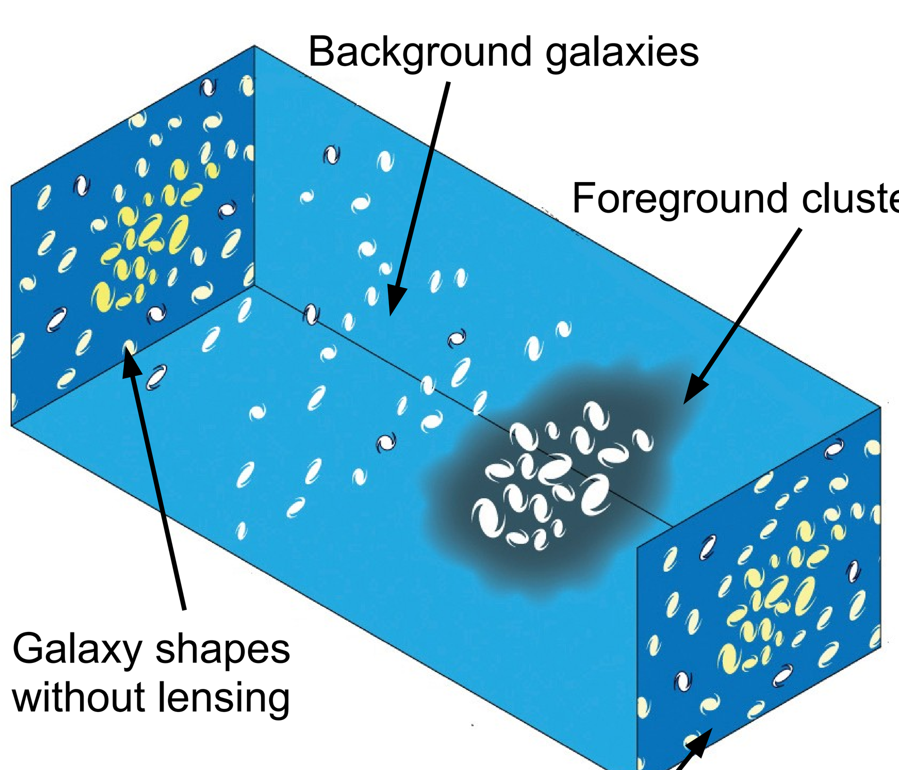

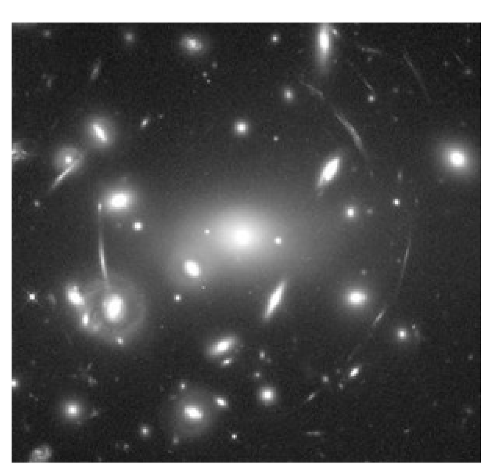

Hubble image of galaxy cluster Abell 2218, showing lensing

Gravitational lensing is one of the predictions of Einstein's theory of relativity. When light travels through the universe, its path is deflected by the gravitational potential of matter it passes by. Concentrations of dark matter, which exists around galaxy clusters, can be detected from the coherent distortions in the shapes of background galaxies. Background galaxies tend to have an ellipticity tangential to the mass concentration, visible in the figure.

Weak gravitational lensing is of the order of a percent and can only be detected through the study of large number of background galaxies, looking for systematic distortions. This kind of large-scale statistical calculation is an ideal case for processing on the GPU.

The authors thank Pat Burchat of Stanford University for her support and the NSF and DOE for funding. We also thank Stanford's Office of Science Outreach and Foothill College for their contributions.

Two-point angular correlation function

The angular distance between two galaxies can be calculated using :

$$\theta = \frac{180^\circ}{\pi} \arctan\left(\frac{\sqrt{\cos^2\delta_2 \sin^2(\alpha_2-\alpha_1) + [\cos\delta_1 \sin\delta_2 - \sin\delta_1 \cos\delta_2 \cos(\alpha_2-\alpha_1)]^2}}{\sin\delta_1 \sin\delta_2 + \cos\delta_1 \cos\delta_2 \cos(\alpha_2-\alpha_1)} \right)$$

where α and δ are the galaxy coordinates. The angular correlation function is calculated using the estimator [1]:

$$w(\theta) = \frac{DD - 2DR + RR}{RR}$$

where DD, DR and RR and pair counts of data-data, data-random and random-random objects, binned in angular separation θ. We use publicly available data from the 2MASS all-sky redshift survey (2MRS) [2]. For the entire angular correlation function over the full dataset, we need to make over 3 billion non-trivial calculations.

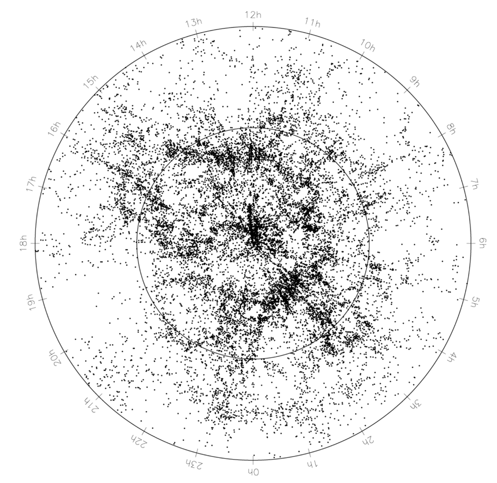

Projection of galaxy positions in 2MRS, looking from the north celestial pole.

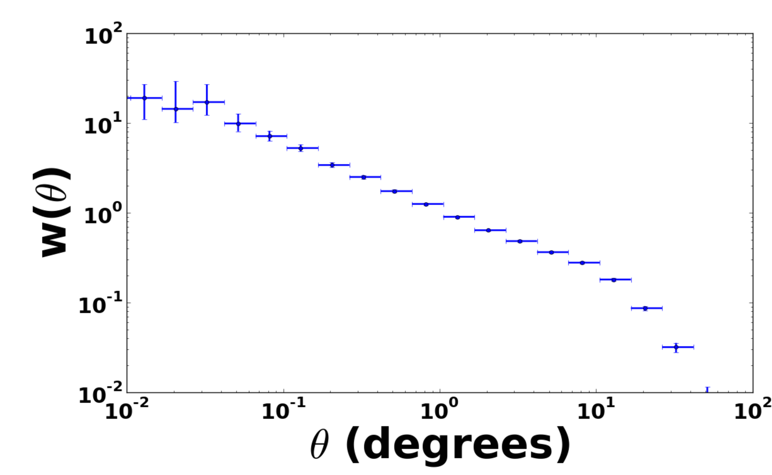

Angular correlation function

GPU Implementation

The matrix of calculations is broken into chunks, which are then passed to the GPU in global memory. The histogramming is also performed on the GPU, with each thread given its own histogram array to increment. This array of histograms is then passed back to the CPU, where they are combined for the final answer.

	1000 galaxies	10,000 galaxies	100,000 galaxies
Compiled C on Intel Xeon 2.53 GHz CPU	0.96 s	85.7 s	8570 s
CUDA on Nvidia Tesla GPU	6.5 s	7.4 s	74.0 s
Improvement	7 x slower	8 x faster	116 x faster

In addition, we estimate the error bars using a bootstrapping technique. We sample with replacement from our galaxies to create 1000 new datasets and also generate simulations of flat distributions, for the random datasets. We then calculate the angular correlation function for *each* of the 1000 datasets and use these to establish 68% confidence intervals for the error bars. Using CUDA, we have reduced the total calculation of uncertainties from 360 CPU hours to 5.5 hours on the GPU.

Shear Peak Statistics

Peaks in the lensing signal can be interpreted as the signature of massive dark matter halos. As such, it is a very useful cosmological probe and can be used to test predictions from theory. An appropriate measure for the lensing signal is the aperture mass (M_{ap}) - the weighted sum of the tangential components of shear (γ_t) of the galaxies surrounding a position in the sky (θ_0). The filter function (Q) used for the weighting can have a generic (e.g. gaussian) shape, or be optimised for detections of halos with NFW density profile [3].

$$M_{\rm ap}(\boldsymbol{\theta}_0) = \int d^2\theta\, Q(\vartheta)\gamma_t(\boldsymbol{\theta};\boldsymbol{\theta}_0), \quad Q_{\rm NFW}(\theta;\theta_c) \propto \frac{1}{1+e^{6-150\theta}+e^{-47+50\theta}} \frac{\tanh(\theta/\theta_c)}{\theta/\theta_c}$$

We use simulated lensing maps from a large-scale cosmological N-body dark matter simulation as our dataset [4]. These consist of a grid of 2048x2048 "galaxies". For a NFW filter radius of 3 arcsec we must calculate M_{ap} for each point on the grid as the sum over the surrounding 6400 galaxies, requiring a total of 30 billion calculations.

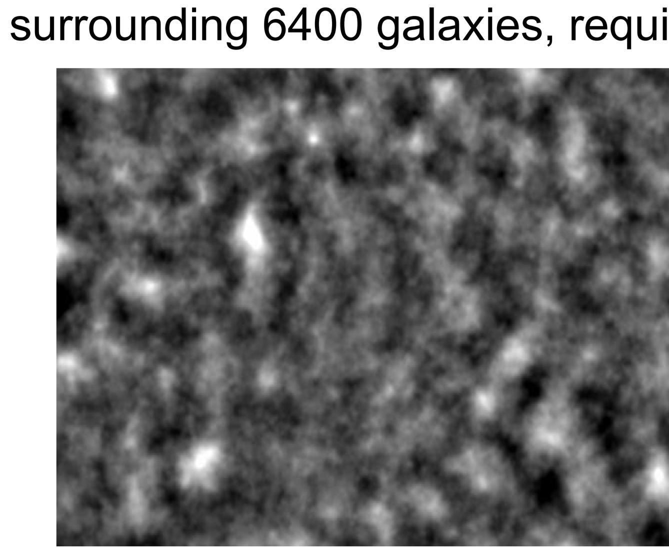

Matter density map from fiducial cosmological simulation, smoothed by a gaussian of radius 3 arcseconds.

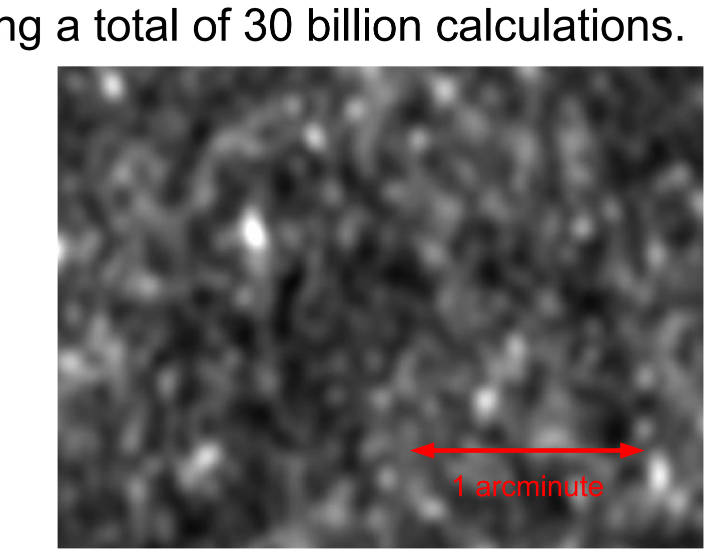

Reconstructed aperture mass map using a NFW filter of radius 3 arcseconds.

GPU Implementation

The vectors of galaxy parameters are loaded into global memory on the GPU. The kernel calculates the contributions of all the surrounding galaxies for each point in parallel, summing the contributions and returning the aperture mass. The number of threads required is equal to the number of "galaxies" in the grid.

	512x512. galaxies	1024x1024 galaxies	2048x2048 galaxies
Compiled C on Intel Xeon 2.53 GHz CPU	1132 s	6810 s	29279 s
CUDA on Nvidia Tesla GPU	4 s	15 s	58 s
Improvement	283 x faster	454 x faster	505 x faster

The speed-up is dramatic, compared to using nested loops in C. With even larger datasets coming soon, this will become an extremely valuable tool for cosmological matter-density calculations.

[1] Landy & Szalay, ApJ 412 **64** (1993)
[2] Huchra *et al,* ApJS (2001)
[3] Schirmer *et al*, A&A 426 (2007)
[4] Kratochvil *et al*, PRD 84 (2010)