Julian Merten

Supercomputing Techniques in Lensing

Institut für Theoretische Astrophysik
Zentrum für Astronomie
Universität Heidelberg
INAF - Osservatorio Astronomico di Bologna

April 28th, 2010

Lensing in HD

Outline

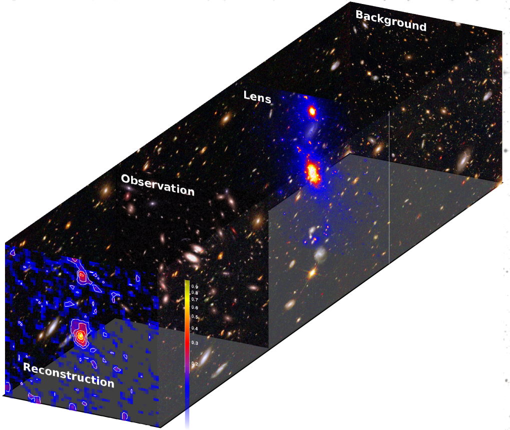

"I am not paid by NVIDIA"

Supercomputing today

Rapidly increased transistor density

Supercomputing today

Special Hardware

- e.g. GRAPE, FPGAs
- Extremely fast
- Extremely expensive, difficult to operate

Massive Parallelisation

- e.g. Infiniband compute clusters
- Fast and relatively cheap
- Extremely flexible

Rapidly increased transistor density

Supercomputing today

GPU Computing

- On-chip parallelisation
- Relatively flexible
- Extremely cheap

Special Hardware

- e.g. GRAPE, FPGAs
- Extremely fast
- Extremely expensive, difficult to operate

Massive Parallelisation

- e.g. Infiniband compute clusters
- Fast and relatively cheap
- Extremely flexible

Rapidly increased transistor density

Single node CPU/GPU Parallelisation

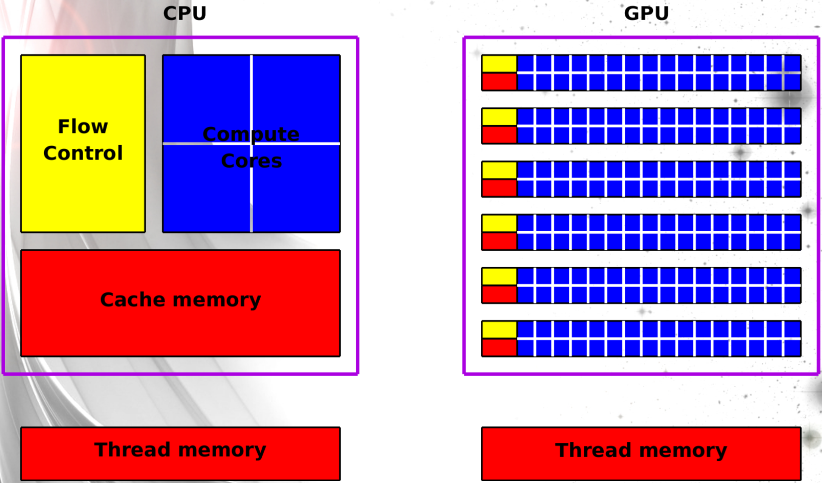

- One single GPU allows for massive parallelisation at $\sim 1/1000$ of the cost, if problem is suited for $\Rightarrow$ **Data-parallel**.

Clusters of puzzles

Density profile

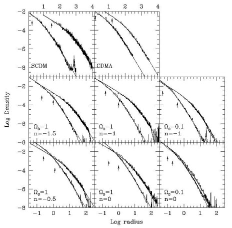

Cool cores

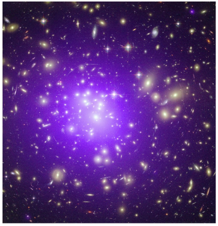

Strong lensing

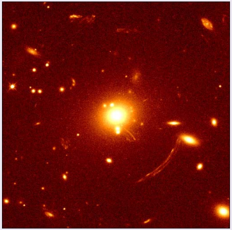

Extreme dynamics

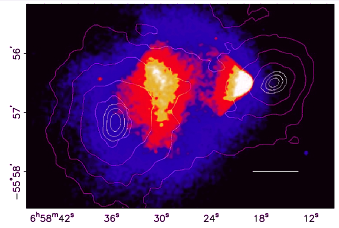

Are those puzzles at all?

Simulations

- State-of-the-art N-body hydro-sims
- As much physics as possible
    - Cooling
    - Star formation
    - AGN/SN feedback
    - Chemical enrichment
    - ...
- Detailed sims of individual objects
- Cluster populations from cosmological volumes

$\Leftrightarrow$

Observations

- State-of-the-art data
- Joint reconstruction method: lensing, X-ray, dynamics, SZ (JM09+, Bradač05+, Puchwein06+)
- reliable error bars
- large cluster sample

Both sides have to be analysed with the same tools.

Are those puzzles at all?

Simulations

- State-of-the-art N-body hydro-sims
- As much physics as possible
  - Cooling
  - Star formation
  - AGN/SN feedback
  - Chemical enrichment
  - ...
- Detailed sims of individual objects
- Cluster populations from cosmological volumes

$\Leftrightarrow$

Observations

- State-of-the-art data
  - HST/ACS/WFPC3
  - SUBARU/LBT
  - KECK
  - CHANDRA / XMM / SUZAKU
- Joint reconstruction method: lensing, X-ray, dynamics, SZ (JM09+, Bradač05+, Puchwein06+)
- reliable error bars
- large cluster sample

Both sides have to be analysed with the same tools.

Are those puzzles at all?

Simulations

- State-of-the-art N-body hydro-sims
- As much physics as possible
  - Cooling
  - Star formation
  - AGN/SN feedback
  - Chemical enrichment
  - ...
- Detailed sims of individual objects
- Cluster populations from cosmological volumes

⇔

Observations

- State-of-the-art data
  - HST/ACS/WFPC3
  - SUBARU/LBT
  - KECK
  - CHANDRA / XMM / SUZAKU
- Joint reconstruction method: lensing, X-ray, dynamics, SZ (JM09+, Bradač05+, Puchwein06+)
- reliable error bars
- large cluster sample

Both sides have to be analysed with the same tools.

Are those puzzles at all?

Simulations

- State-of-the-art N-body hydro-sims
- As much physics as possible
  - Cooling
  - Star formation
  - AGN/SN feedback
  - Chemical enrichment
  - ...
- Detailed sims of individual objects
- Cluster populations from cosmological volumes

$\Leftrightarrow$

Observations

- State-of-the-art data
  - HST/ACS/WFPC3
  - SUBARU/LBT
  - KECK
  - CHANDRA / XMM / SUZAKU
- Joint reconstruction method: lensing, X-ray, dynamics, SZ (JM09+, Bradač05+, Puchwein06+)
- reliable error bars
- large cluster sample

Both sides have to be analysed with the same tools.

Are those puzzles at all?

Simulations

- State-of-the-art N-body hydro-sims
- As much physics as possible
  - Cooling
  - Star formation
  - AGN/SN feedback
  - Chemical enrichment
  - ...
- Detailed sims of individual objects
- Cluster populations from cosmological volumes

$\Leftrightarrow$

Observations

- State-of-the-art data
  - HST/ACS/WFPC3
  - SUBARU/LBT
  - KECK
  - CHANDRA / XMM / SUZAKU
- Joint reconstruction method: lensing, X-ray, dynamics, SZ (JM09+, Bradač05+, Puchwein06+)
- reliable error bars
- large cluster sample

Both sides have to be analysed with the same tools.

Are those puzzles at all?

Simulations

- State-of-the-art N-body hydro-sims
- As much physics as possible
  - Cooling
  - Star formation
  - AGN/SN feedback
  - Chemical enrichment
  - ...
- Detailed sims of individual objects
- Cluster populations from cosmological volumes

$\Leftrightarrow$

Observations

- State-of-the-art data
  - HST/ACS/WFPC3
  - SUBARU/LBT
  - KECK
  - CHANDRA / XMM / SUZAKU
- Joint reconstruction method: lensing, X-ray, dynamics, SZ (JM09+, Bradač05+, Puchwein06+)
- reliable error bars
- large cluster sample

Both sides have to be analysed with the same tools.

Are those puzzles at all?

Simulations

- State-of-the-art N-body hydro-sims
- As much physics as possible
  - Cooling
  - Star formation
  - AGN/SN feedback
  - Chemical enrichment
  - ...
- Detailed sims of individual objects
- Cluster populations from cosmological volumes

$\Leftrightarrow$

Observations

- State-of-the-art data
  - HST/ACS/WFPC3
  - SUBARU/LBT
  - KECK
  - CHANDRA / XMM / SUZAKU
- Joint reconstruction method: lensing, X-ray, dynamics, SZ (JM09+, Bradač05+, Puchwein06+)
- reliable error bars
- large cluster sample

Both sides have to be analysed with the same tools.

CLASH: An HST/MCT Programme

- PI: M. Postman
  (Co-PIs including Matthias, Arjen, T. Broadhurst, O. Lahav, A. Riess, P. Rosati,...)
- Target: 25 well-chosen X-ray clusters
- Goal: Density profile of clusters, high-z Universe

CLASH Facts

- 524 orbits
- ACS + WFC3 obs.
- 14 wave bands
- wide follow-ups with SUBARU

Our pipeline I: **SkyLens** *(M. Meneghetti, P.Melchior, JM)*

Name	Description
D	aperture diameter
g	detector gain
A_{pix}	pixel area
$F(\lambda)$	used filter
$M(\lambda)$	mirror filter curve
$O(\lambda)$	optics filter curve
$C(\lambda)$	CCD filter curve
FoV	total field-of-view
RON	detector readout-noise
f	flat-field accuracy
a	residual flat-field error
PSF	PSF model
t_{exp}	exposure time
$A(\lambda)$	atmospheric extinction
m_{a}	airmass
SED_{sky}	sky-background emission
SED_{gal}	background population
α	deflection angle map

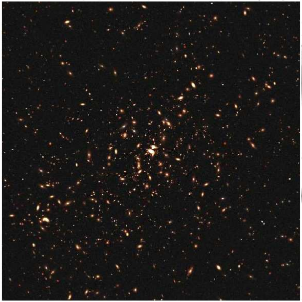

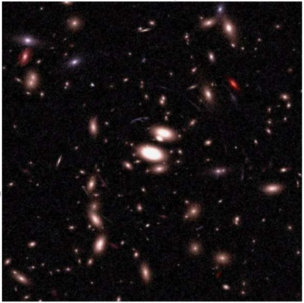

Parallelisation strategy:

Ray-tracing

- From snapshots to lensing:

Deflection angle

$$\hat{\boldsymbol{\alpha}}(\boldsymbol{\xi}) = \frac{4G}{c^2}\int d^2\xi' \, \Sigma(\boldsymbol{\xi}')\frac{\boldsymbol{\xi}-\boldsymbol{\xi}'}{|\boldsymbol{\xi}-\boldsymbol{\xi}'|^2}.$$

- Calculated on a pixelised grid
- Usual approach: Barnes-Hut tree codes
- Better approach: direct N-Body summation

Parallelisation strategy:

Ray-tracing

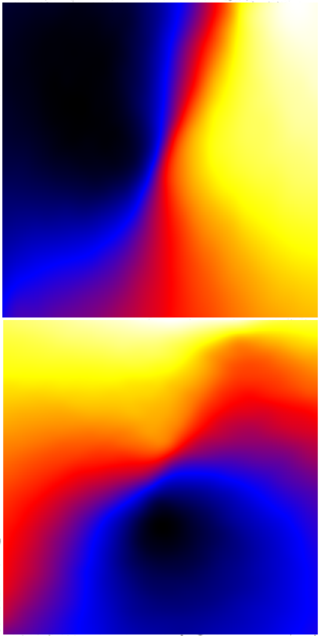

Our pipeline III: **SaWLens** *(JM09/10)*

Features

- Fully nonparametric
- Stat. grid-based approach
- AMR implementation
- Computationally rather demanding

Possible input

- Shear/Flexion
- Multiple images + Critical curve estimators
- (Cluster dynamics)
- (ICM-tracers)

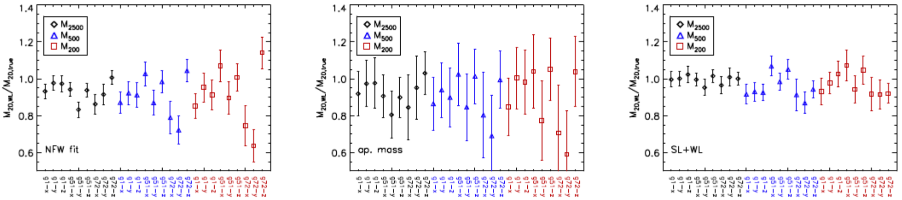

Parallelisation strategy:

Independence of grid cells, matrix summation schemes

Programming GPUs: `C for CUDA`

Basic workflow

- Host-code: "Any language"
- Device-code: `C for CUDA` or `CUDA driver API`.
- Host-code calls the device if necessary
- `MPI` + `CUDA` possible
- Several `GPUs` in a code also possible
- Wide range of tools in the `CUDA SDK`.

`C for CUDA`

- High-level language
- `C++` syntax with `C` functionality
- Several levels of memory addressing, including on-chip memory
- Easy thread-indexing
- Device-code objects created with `nvcc` compiler
- Some libraries: FFTW, BLAS, basic math functions

GPU workflow

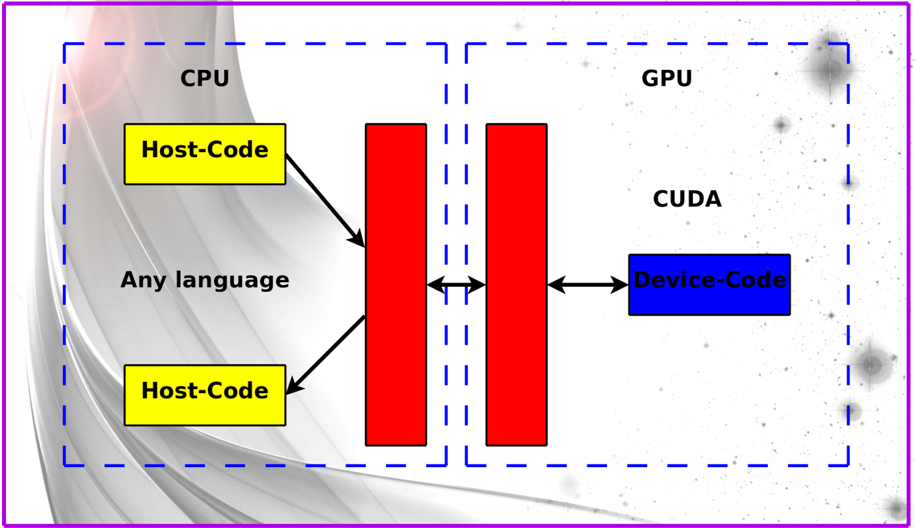

GPU-Systems `Jabba the Hutt` (BO) & `Kolob` (MA)

NVIDIA Tesla C1060

- 240 streaming cores
- 4 GB DDR3 GPU memory
- 933 GFLOPS peak performance
- Upcoming `Fermi` cards

First results

- The toy problem:
  - Simulate a typical SaWLens problem
  - Calculate a typical coefficient matrix

$$\mathcal{B}_{lk} = a_i b_j \mathcal{C}_{ij} \mathcal{D}_{ik} \mathcal{E}_{jl},$$

    while using Albert's sum convention.
  - Dimensions:

$$l, k \in [0, ..., 2499], \quad i, j \in [0, .., 15]$$

- Competitors:
  1. Jabba's CPU: Intel XEON quadcore @ 2.5 GHz, one core used
  2. Jabba's GPU: NVIDIA Tesla C1060 @ 1.2 GHz, 240 cores used
- The runtime:
  1. CPU: `82.3 s`
  2. GPU: `1.03 s`

SaWLens runtime will be reduced to $\mathcal{O}$(min).