Analyzing the CMB and LSS in the Big Data Era

Alessandro Renzi (INFN Padova, Italy)

Hot topics in Modern Cosmology
Spontaneous Workshop XII 14 - 19 May 2018 Cargèse

Istituto Nazionale di Fisica Nucleare

Outline

- Present and (near) future of Cosmological dataset
- How and (why) analyze CMB and LSS data
- Evolution of computing and its compatibility with current estimators
- Needlets analysis: a possible solution
- Prospect and conclusions

Outline

- **Present and (near) future of Cosmological dataset**
- How and (why) analyze CMB and LSS data
- Evolution of computing and its compatibility with current estimators
- Needlets analysis: a possible solution
- Prospect and conclusions

Cosmological Dataset: Planck CMB

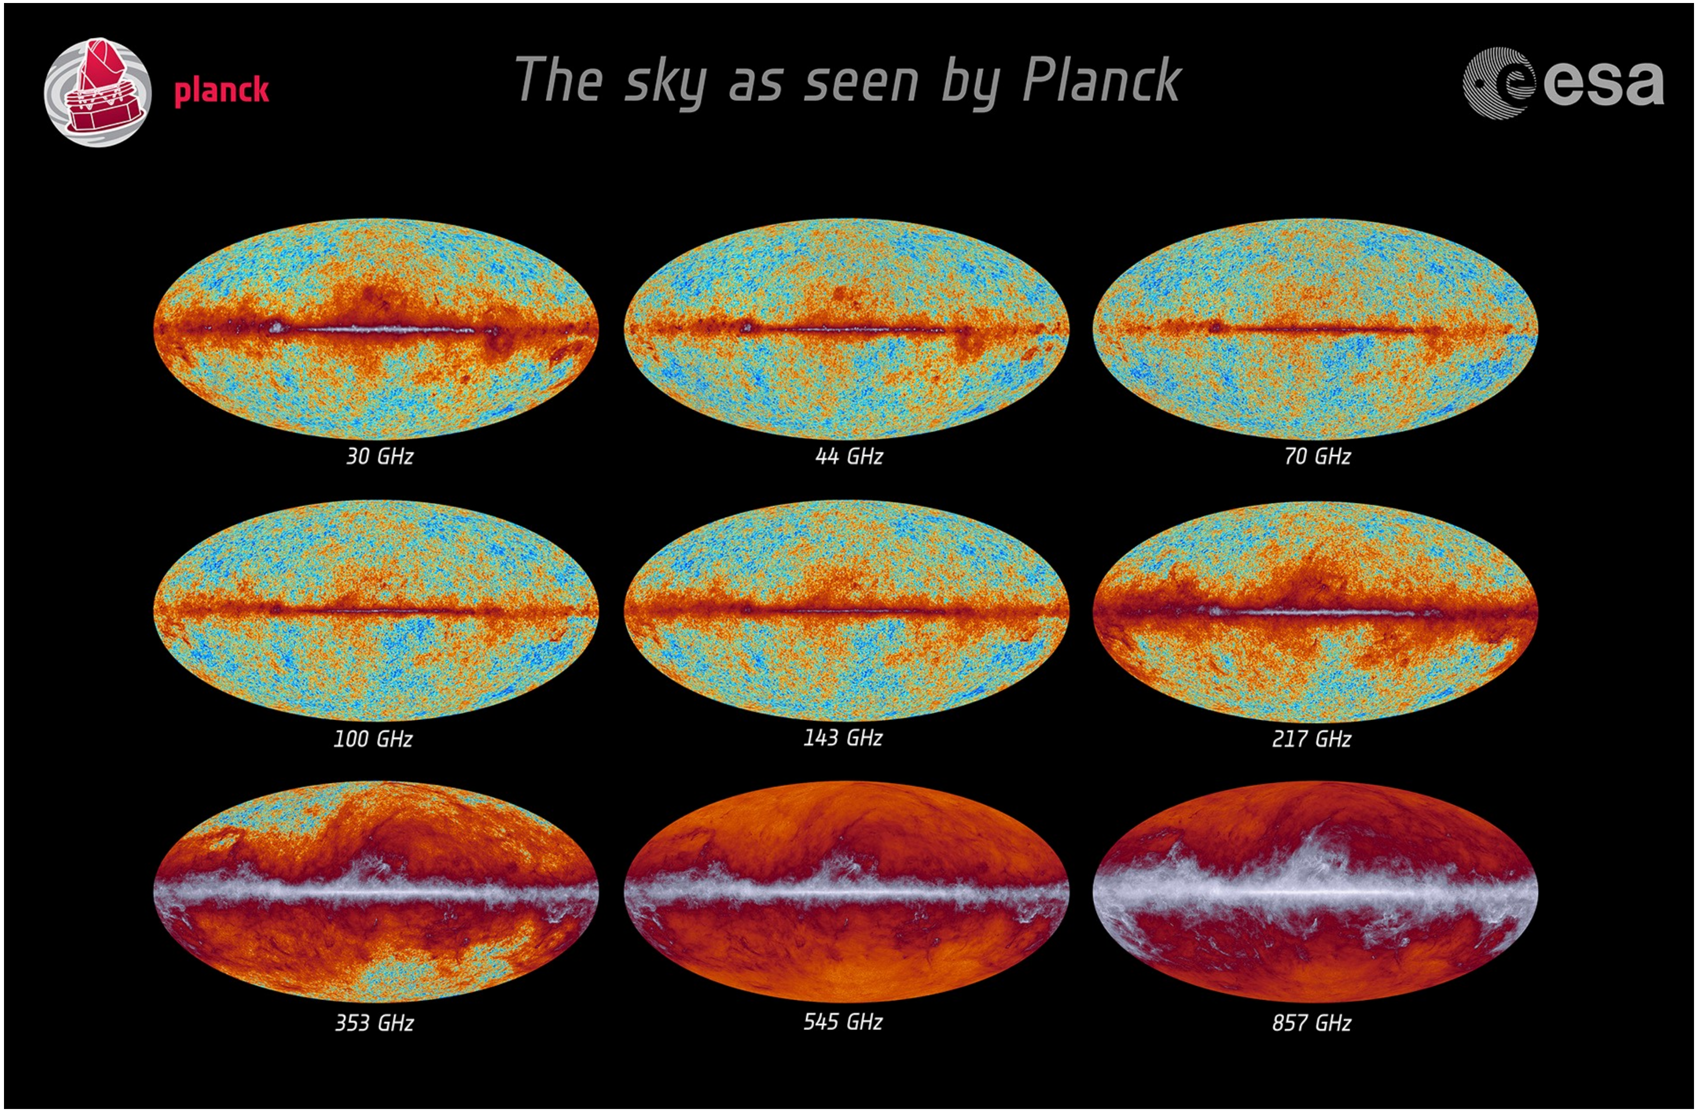

For all the experiments considered in the following the biggest computational effort comes from the simulations (physics [temperature and polarization] + foregrounds + instruments systematics)

Planck ~ O(1M) simulations per release!

Cosmological Dataset: CMB Stage 4

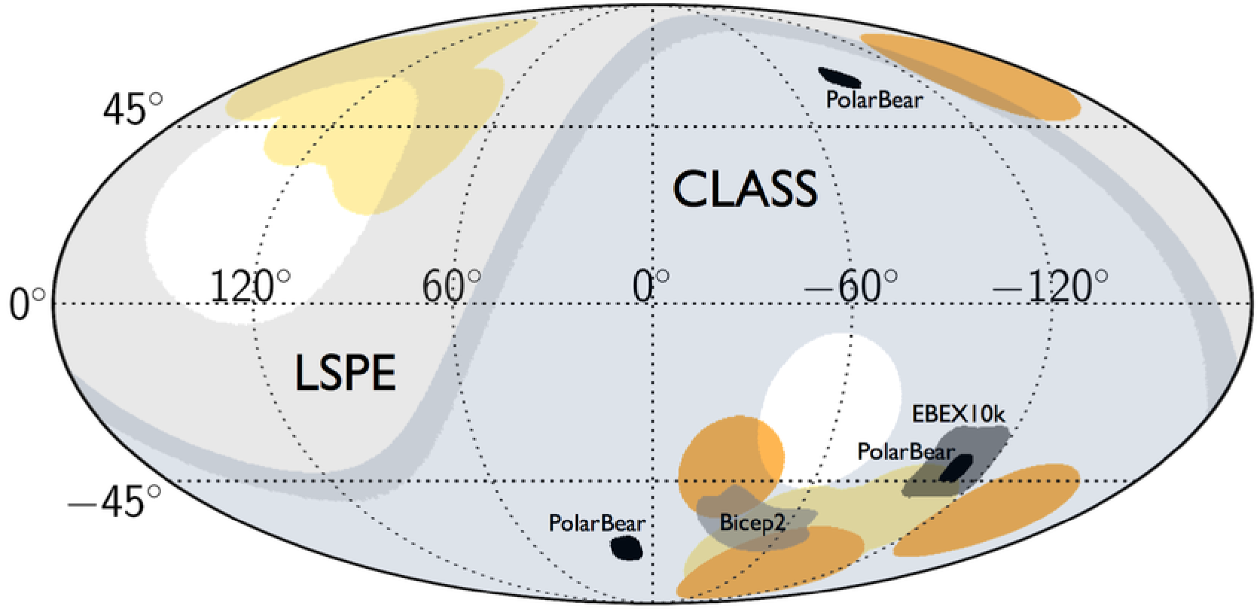

Note that even if the fraction of the sky covered is smaller than Planck, the resolution is higher → #pixels is > than Planck

Cosmological Dataset: CMB Time Ordered Data

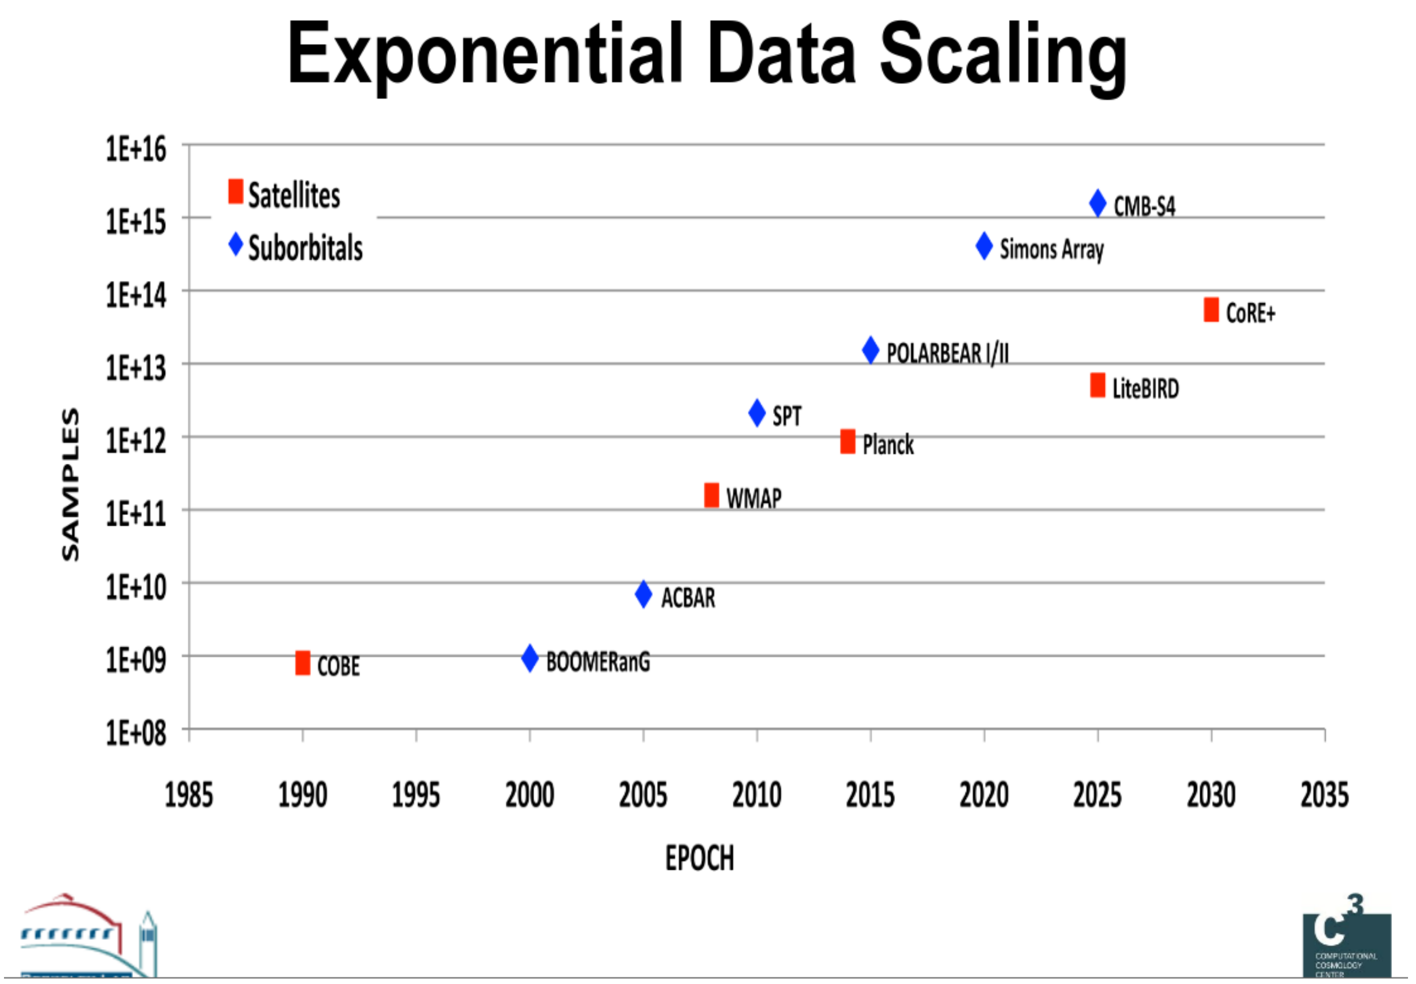

Figure taken from Julian Borrill presentation (Planck 2014, Ferrara)

Cosmological Dataset: Euclid (et al.)

LSS requires N-body simulations to properly account for non-linearities!

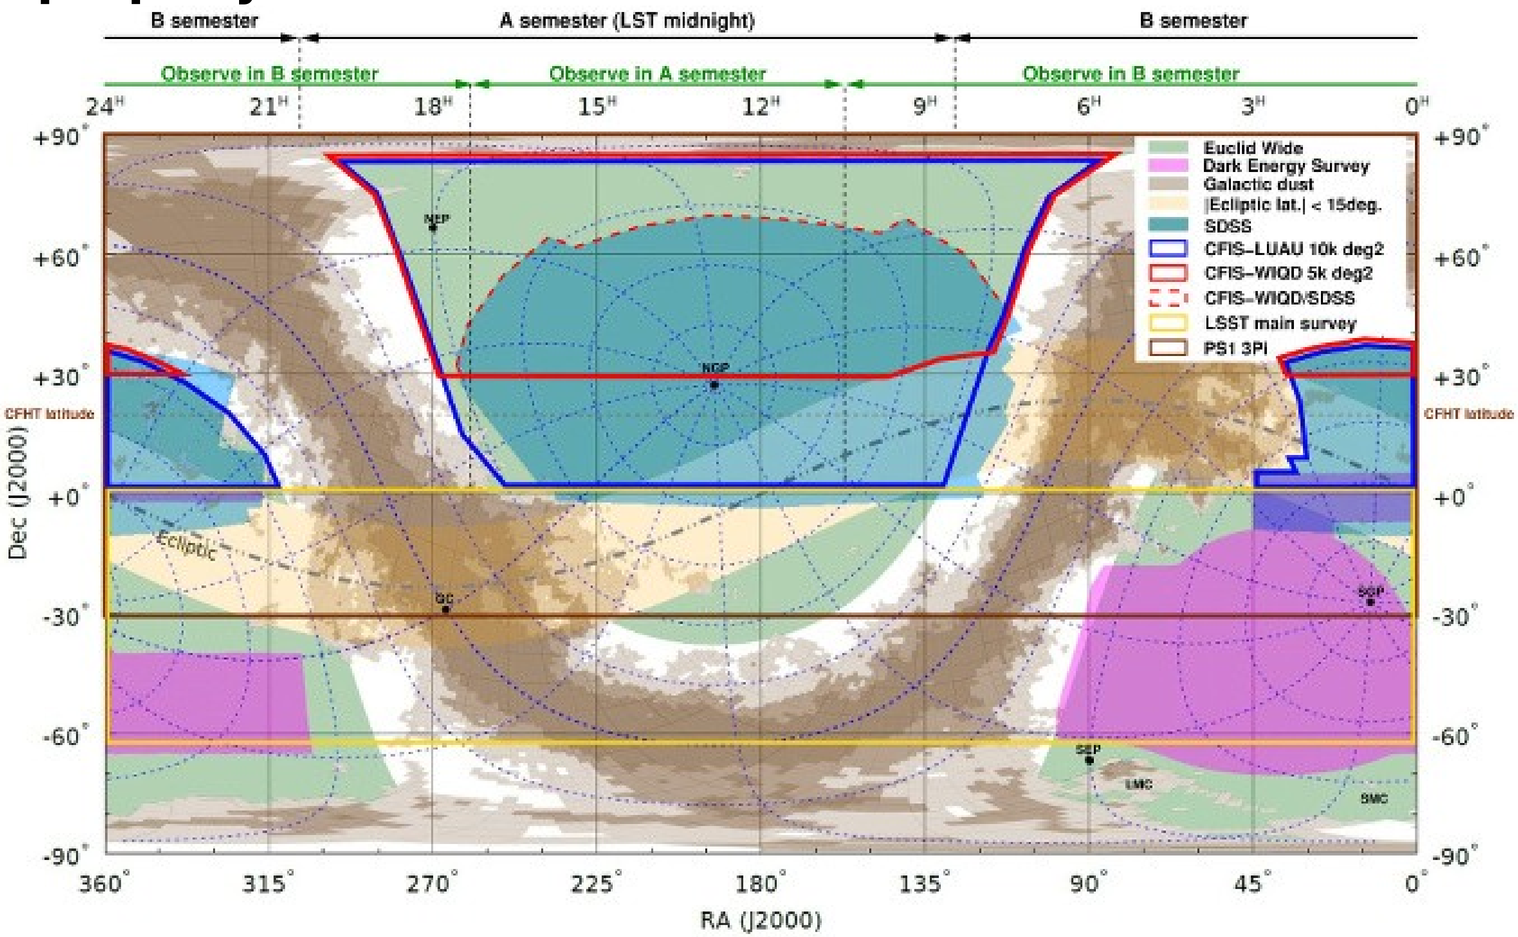

External datasets are much larger than Euclid alone!

And Euclid is larger than Planck!

O(1G) photometric redshift

O(10M) spectroscopic redshift

Euclid (and suborbital) fsky is large!

Cosmological Dataset: computational forecast

	[1603.09303] PLANCK	[1603.09303] CMB-S4 Next Generation CMB Experiment	[1701.08158] euclid
HPC Computation (cores x hours / year)	~ 10M – 100M	**~ 1G – 10G**	**~ 100M – 1G**
Storage for Products	~ 1 PB	**~ 10P – 100P**	**> 100 PB**

With respect to Planck, next generation cosmology experiments will require:

- **order of magnitude** in CPU-hour needs
- **orders of magnitude** in Disk Space needs

Take-home messages

Outline

- Present and (near) future of Cosmological dataset
- **How and (why) analyze CMB and LSS data**
- Evolution of computing and its compatibility with current estimators
- Needlets analysis: a possible solution
- Prospect and conclusions

Analysis: CMB data

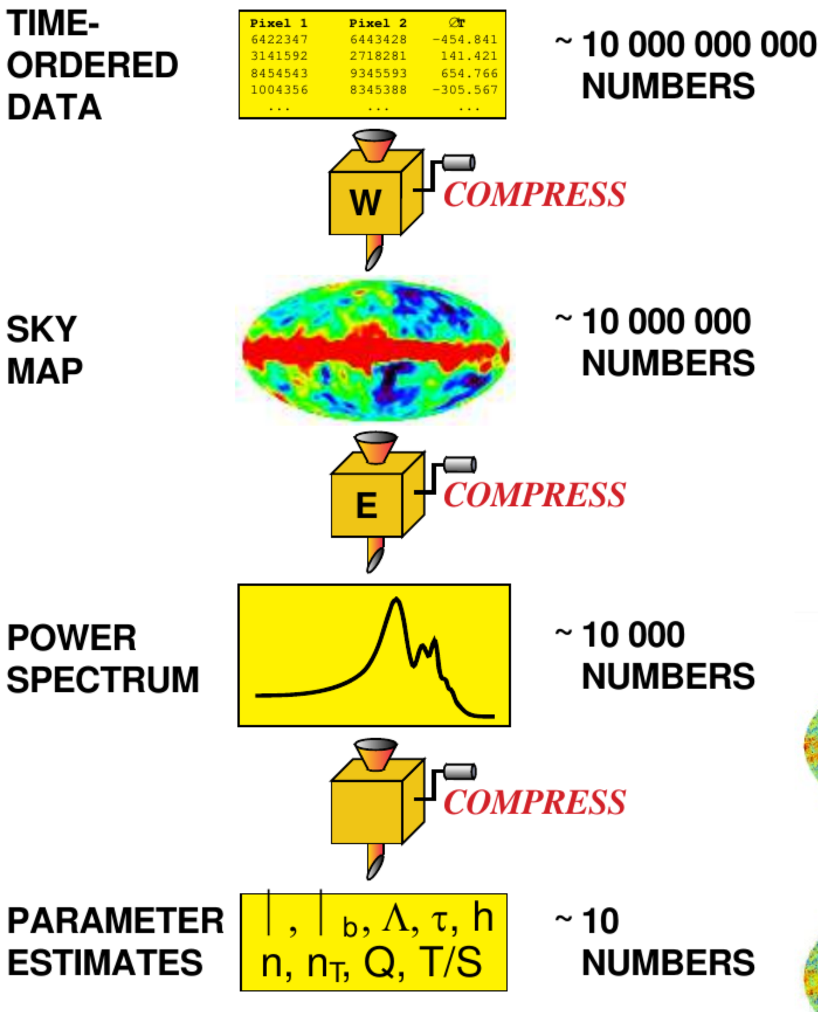

Figure taken from Tegmark

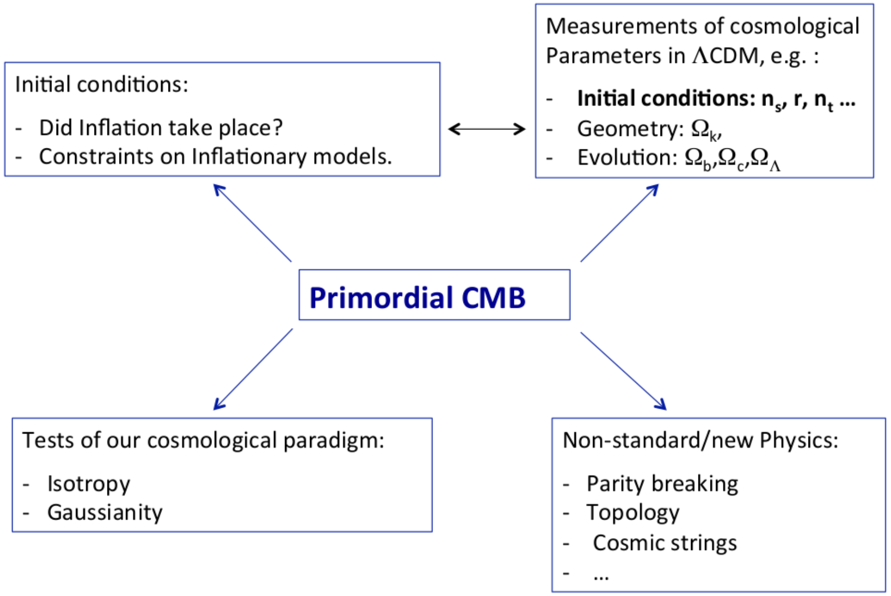

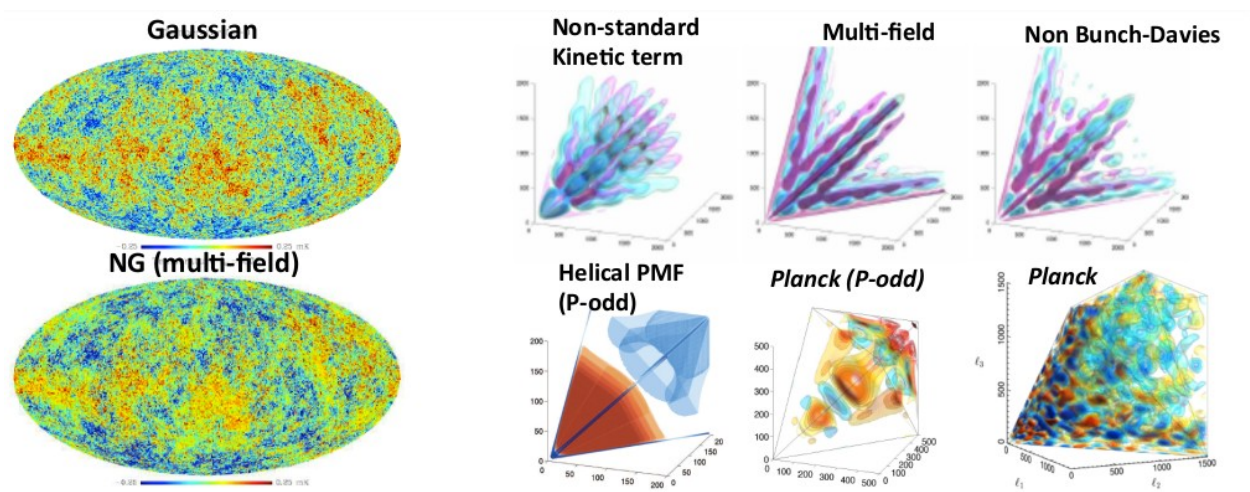

Analysis: LSS data

Figure taken from Tegmark

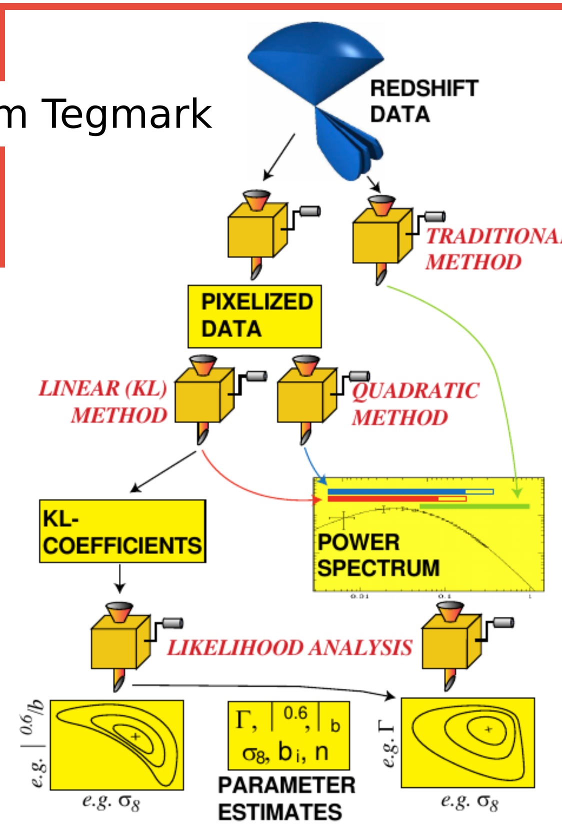

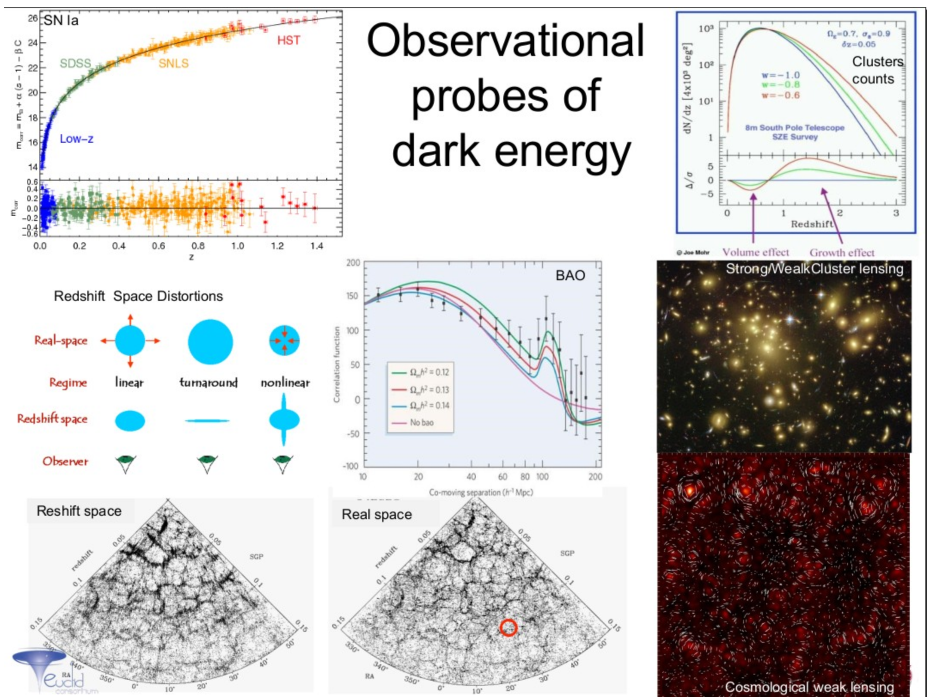

Figure taken from J. Mellier (Euclid PI) presentation

Astroparticle:

- Neutrino masses
- Nature of Dark Matter.

Measurements of cosmological Parameters

Large Scale Structure

Study formation, growth and clustering of cosmic structures

Understanding acceleration:

- Dark Energy
- Modified Gravity

An example of LSS N-body simulation computational requirements

- **Simulations** are the main tool to investigate evolution of the universe at small scale where **non linear effects become important**
- C. Carbone, M. Petkova, K. Dolag [1605.02024]
- 4 simulations with different Σm_ν [0, 0.17, 0.3, 0.53 eV]
- 62 sampling in time
- 3D grid dimensions: $L_{box} = 2h^{-1}$Gpc and a mesh of 4096^3 cells
- For every time sampling save a snapshot of:
  - CDM particles
  - Neutrinos
  - 3D grid of the gravitational potential
  - 3D grid of the derived gravitational potential
- Resources for <u>**one**</u> simulation:
  - **~ 1M CPU-hours**
  - **~ 90 TB**

Analysis: putting all together

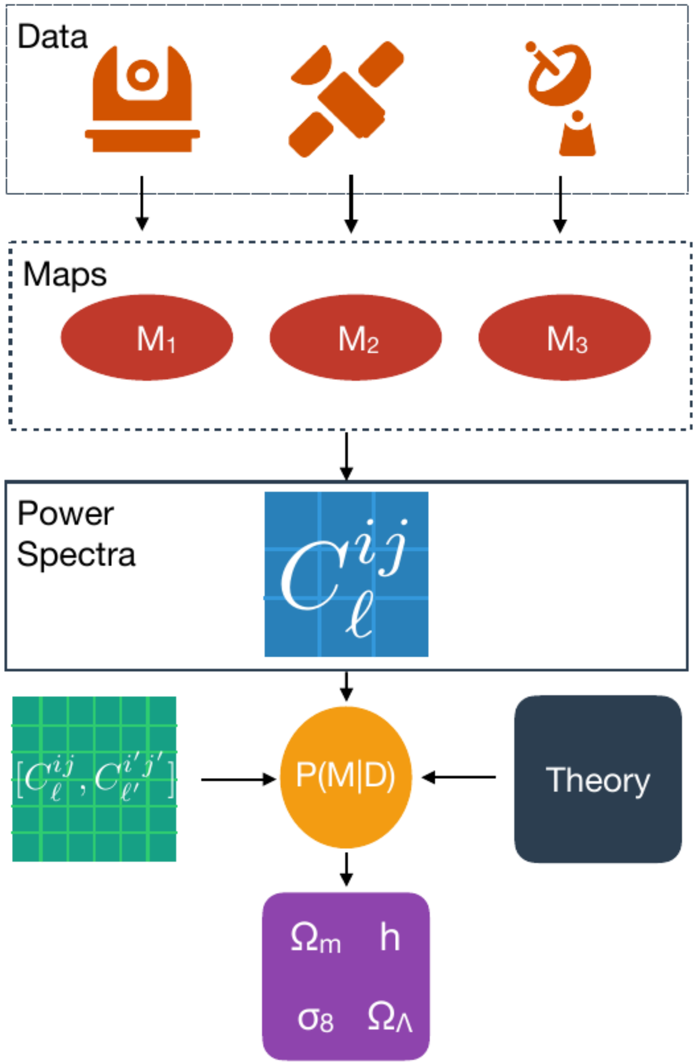

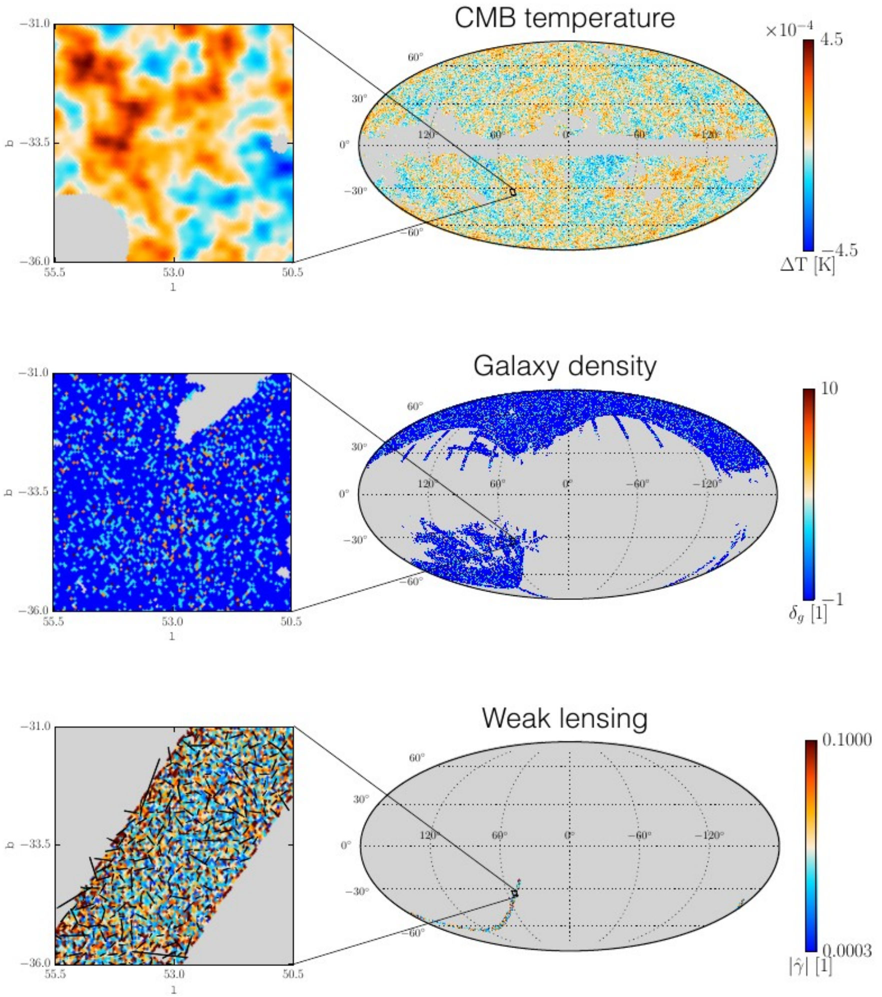

Both figures are taken from: arXiv:1607.01014 Nicola etal

Outline

- Present and (near) future of Cosmological dataset
- How and (why) analyze CMB and LSS data
- **Evolution of computing and its compatibility with current estimators**
- Needlets analysis: a possible solution
- Prospect and conclusions

Status of computing: present

Rank	Site	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
1	National Supercomputing Center in Wuxi China	**Sunway TaihuLight** - Sunway MPP, Sunway SW26010 260C 1.45GHz, Sunway NRCPC	10,649,600	93,014.6	125,435.9	15,371
2	National Super Computer Center in Guangzhou China	**Tianhe-2 (MilkyWay-2)** - TH-IVB-FEP Cluster, Intel Xeon E5-2692 12C 2.200GHz, TH Express-2, Intel Xeon Phi 31S1P NUDT	3,120,000	33,862.7	54,902.4	17,808
3	DOE/SC/Oak Ridge National Laboratory United States	**Titan** - Cray XK7 , Opteron 6274 16C 2.200GHz, Cray Gemini interconnect, NVIDIA K20x Cray Inc.	560,640	17,590.0	27,112.5	8,209
4	DOE/NNSA/LLNL United States	**Sequoia** - BlueGene/Q, Power BQC 16C 1.60 GHz, Custom IBM	1,572,864	17,173.2	20,132.7	7,890
5	DOE/SC/LBNL/NERSC United States	**Cori** - Cray XC40, Intel Xeon Phi 7250 68C 1.4GHz, Aries interconnect Cray Inc.	622,336	14,014.7	27,880.7	3,939
6	Joint Center for Advanced High Performance Computing Japan	**Oakforest-PACS** - PRIMERGY CX1640 M1, Intel Xeon Phi 7250 68C 1.4GHz, Intel Omni-Path Fujitsu	556,104	13,554.6	24,913.5	2,719

Cores ~ O(1M)

Tflops ~ O(few * 10k)

Status of computing: the problem

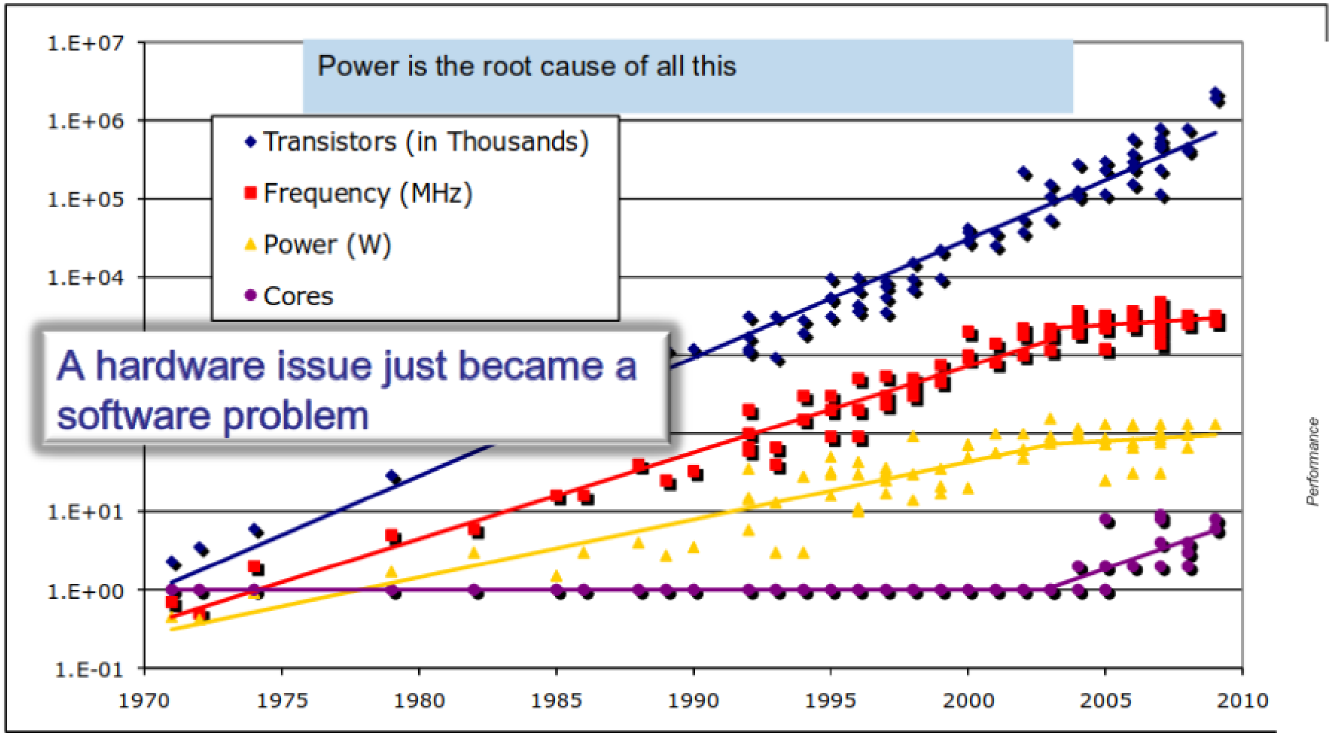

Data from Kunle Olukotun, Lance Hammond, Herb Sutter,
Burton Smith, Chris Batten, and Krste Asanoviç
Slide from Kathy Yelick

Status of computing: some (rough) definitions

- Some definitions will be useful for the following slides, with some rough approximation:

  - **Core** - a single hardware computing unit
  - **Node** - a physical object that contains one or a collection of cores along with the memory that the cores can *directly* access.
  - **Node communication** – the cores in a node **can** access the memory of another code thought a fast(?) physical inter-nodes connection
  - **I/O communication** – the cores in a node can access the physical disk memory thought *very slow* I/O operations (potentially parallel)
  - **Thread** - a single logical computational job, for simplicity we will assume that every core could process only a single thread
  - **Task** - a task is composed by a single thread or more parallel threads, for simplicity we will assume that a node can process only a single task, tasks can communicate between them

  - **HPC** - High Performance Computing, main characteristics: top computational speed, limited amount of nodes and cores, top memory per node, top node communication, very expensive. Very difficult to program efficient parallel algorithm for those machines.
  - **HTC** – High Throughput Computing, main characteristics: general commodity hardware computational speed, enormous amount of nodes, low memory per node, none or slow node communication, very cheap. Easy algorithmic development.

→ The key point: **optimal performance are obtained by exploiting data locality**

Status of computing: the solution

Power Cost of Frequency

- Power ∝ Voltage^2 x Frequency (V^2F)
- Frequency ∝ Voltage
- Power ∝ Frequency^3

	Cores	V	Freq	Perf	Power	PE (Bops/watt)
Superscalar	1	1	1	1	1	1
"New" Superscalar	1X	1.5X	1.5X	1.5X	3.3X	0.45X
Multicore	2X	0.75X	0.75X	1.5X	0.8X	1.88X

(Bigger # is better)

50% more performance with 20% less power
Preferable to use multiple slower devices, than one superfast device

Go many cores! **Exascale computing**

And memory remain nearly constant in dimension and speed!

- Number of cores per chip doubles every 2 year, while clock speed decreases (not increases!) → Future generation (we!) will have billions of threads!

Status of computing: missed predictions for Exascale computing

Exascale systems are likely feasible by 2017± 2

Potential System Architectures

Systems	2015	2018-2020
System peak	100-200 Pflop/s	1 Eflop/s
System memory	5 PB	10 PB
Node performance	200-400 Gflop/s	1-10 Tflop/s
Node memory bandwidth	100 GB/s	200-400 GB/s
Node concurrency	O(100)	O(1000)
Interconnect bandwidth	25 GB/s	50 GB/s
System size (nodes)	O(100,000)	100,000-1,000,000
Total concurrency	O(50,000,000)	O(1,000,000,000)
Storage	150 PB	300 PB
IO	10 TB/s	20 TB/s
MTTI	days	O(1 day)
Power	~10 MW	~20 MW

Note!

The problem of current estimators: harmonic (and/or) Fourier analysis and its relation with data locality

- Data analysis is based on signal processing, and the cornerstone of that is Fourier decomposition (or Spherical Harmonics decomposition for spherical symmetric fields like in cosmology)

- Fourier decomposition impose a complete space localization in real space and a complete de-localization in momentum space → to know the value of one real sample, we need to now *all* the Fourier coefficients

$$\text{(Fourier Analysis)} \qquad a_k = \frac{1}{N} \sum_{n=-N/2+1}^{N/2} f(\theta_n) \exp^{-ik\theta_n}, \qquad (1)$$

$$\text{(Fourier Synthesis)} \qquad f(\theta_n) = \sum_{k=-N/2+1}^{N/2} a_k \exp^{ik\theta_n}, \qquad (2)$$

- The implications of that is: **if we cannot store all the coefficients in memory at once, we have to rely on some sort of communications, namely I/O or Node communications**.

- I/O communication is the slowest possible communication, and the most inefficient solution for data analysis

- Fast Node communication is fast (by definition :D), but not as fast as the use of memory in node

- Both solutions require a dedicated (somehow complex) optimization strategy and usually the only viable solution to that it is the use of HPC systems (with all the related problematic of: cost of resources and complex optimization)

Friendly advice: An overlooked goldmine

- Most of data analysis algorithms in physics (Fourier-related) are required to be run in HPC systems

- Since HPC systems are able to run efficiently even HTC-like algorithms, the interest in HTC farms are reduced in the data analysis physics community.

- We are overlooking the biggest HTC farm ever build for physics: **the LHC Grid**

- Just to fix some numbers (top500 list):

  - CINECA Marconi supercomputer (configuration A2 – Intel Xeon Phi)
    - #14 in the top500 supercomputer
    - 3600 nodes,
    - 68 x 1.4GHz core per node with about 250 threads
    - 16GB or 94GB ram per node, for a grand total of about 250k cores
      - To efficiently use the 250 core, highly tuned parallel codes based on hybrid threads/tasks parallelization must be developed

  - The Worldwide LHC Computing Grid:
    - 2M tasks run every day
    - 750k computer cores
    - ~ Gbyte/second of calculations

Solution to the exascal computing paradigm

- (**Optimization**) Dedicate your scientific career to became a Computer Engineer with the only objective of optimizing code(s) for a particular HPC cluster
- (**Algorithm development 1**) Develop approximated solutions targeted to reduce the computational effort
- (**Algorithm development 2**) find new way to analyze your data (e.g. AI revolution or wavelets analysis)

Outline

- Present and (near) future of Cosmological dataset
- How and (why) analyze CMB and LSS data
- Evolution of computing and its compatibility with current estimators
- **Needlets analysis: a possible solution**
- Prospect and conclusions

Algorithm development with Needlets

- Needlets are introduced in CMB analysis here: arXiv:0707.0844 (however they can be used in LSS as well, see arXiv:1607.05223)

- Needlets are a particular wavelet base with the following properties:

  - Quasi-exponential localization in pixel space
  - Bounded support on multipoles
  - Tight frames: exact reconstruction formula
  - Minimal correlation in harmonic and real domain
  - Needlets coefficients are uncorrelated at higher needlets frequencies
  - (But most importantly, for the purpose of this talk) doubly localized in pixel and momentum space

$$\text{(Needlet Analysis)} \quad \beta_j(\xi_{jt}) = \sum_{n=-N/2+1}^{N/2} f(\theta_n)\Psi_j(\theta_n,\ \xi_{jt}), \tag{3}$$

$$\text{(Needlet Synthesis)} \quad f(\theta_n) = \sum_{j=0}^{j_{\max}} \sum_{t=-N/2+1}^{N/2} \beta_j(\xi_{jt})\Psi_j(\theta_n,\ \xi_{jt}), \tag{4}$$

$$\text{(Needlet Base)} \quad \psi_j(\theta_n,\ \xi_{jt}) = \sqrt{\lambda_{jt}} \sum_{k=-N/2+1}^{N/2} b\left(\frac{k}{B^j}\right) \exp^{-ik(\theta_n-\ \xi_{jt})}, \tag{5}$$

Needlets in momentum space

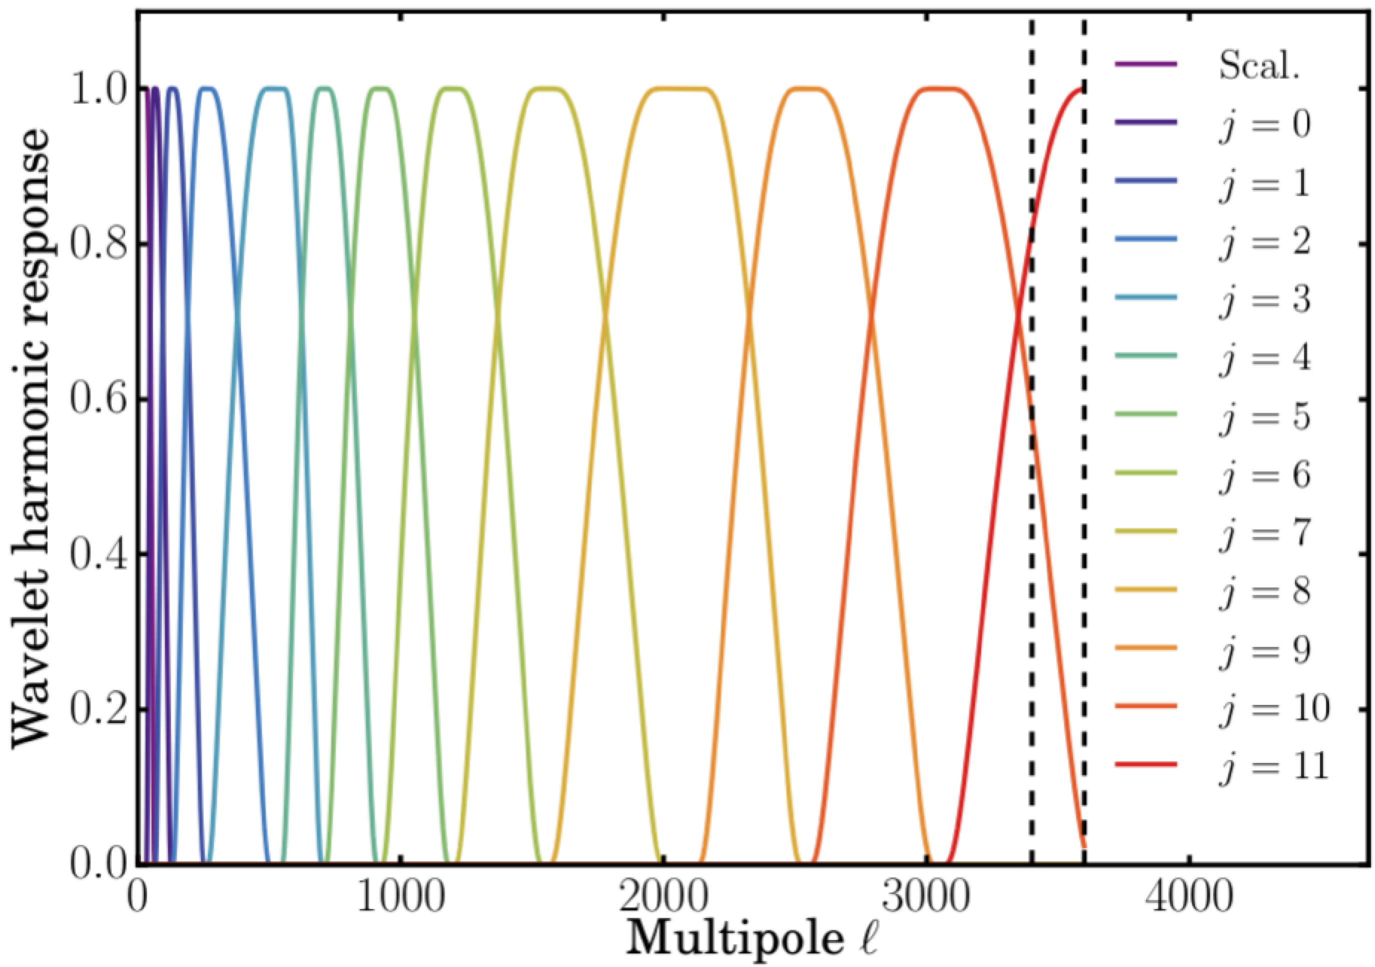

Needlets in real space

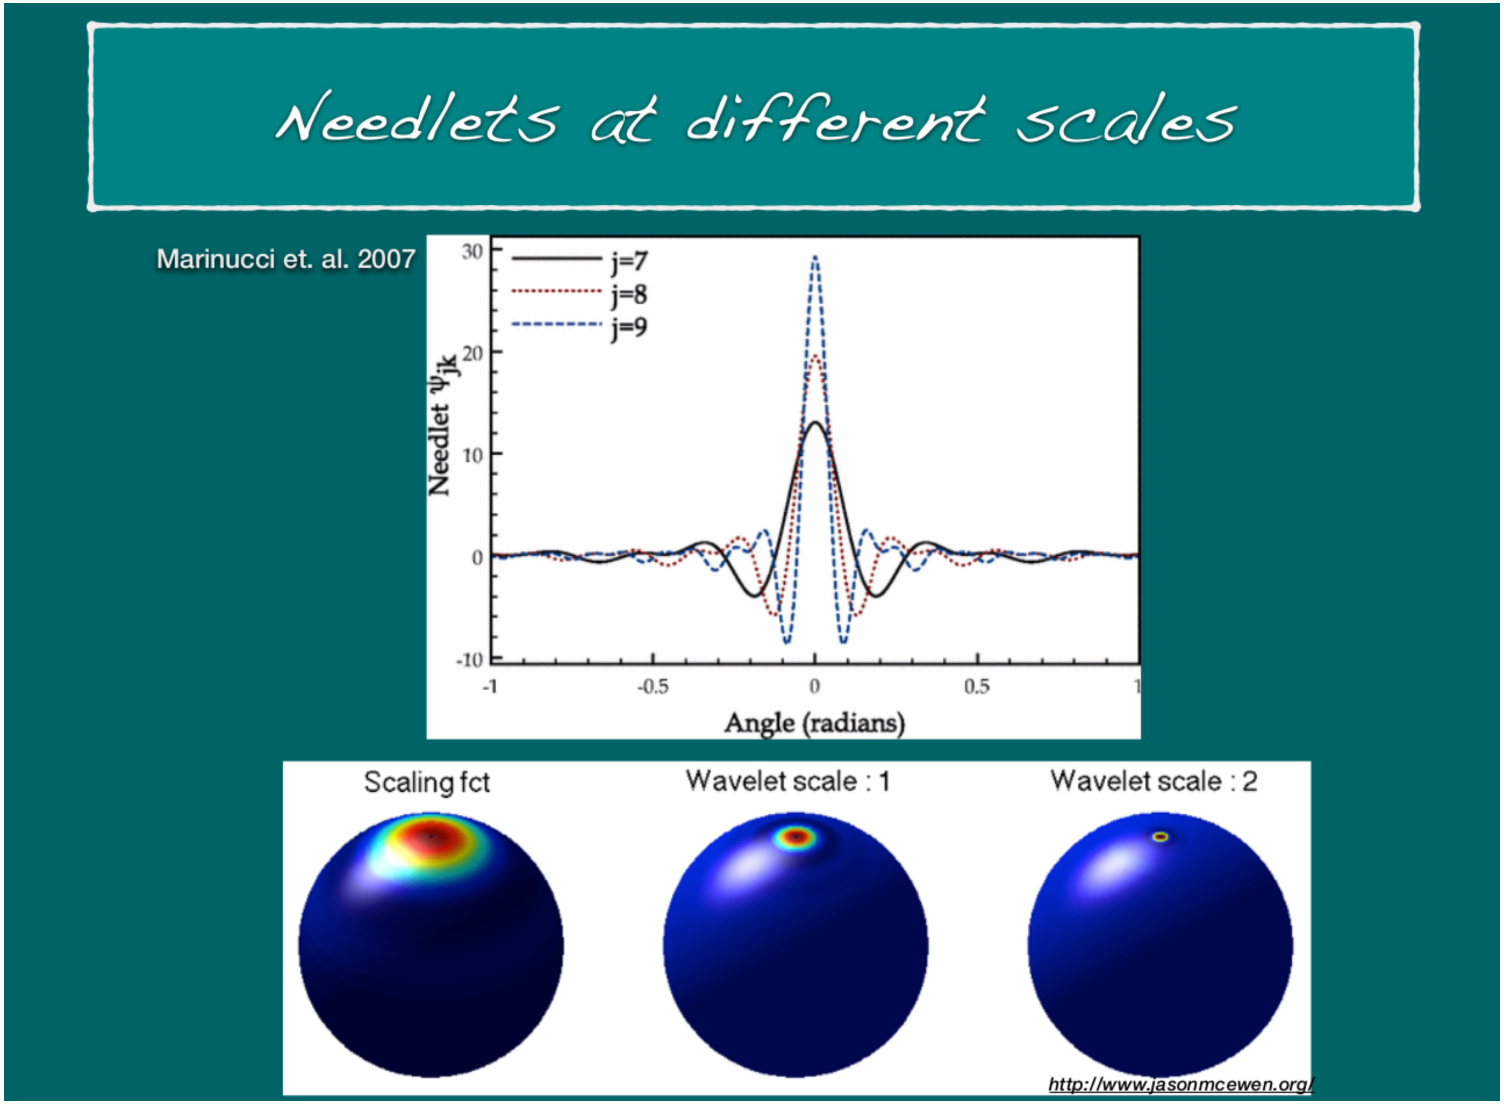

3D Needlets

See arXiv:1408.1095

Radial 3D needlets – Setup

Our method envisages a data collection environment in which an observer located at the center of the ball is surrounded by concentric spheres with the same pixelization at different radial distances, for any given resolution.

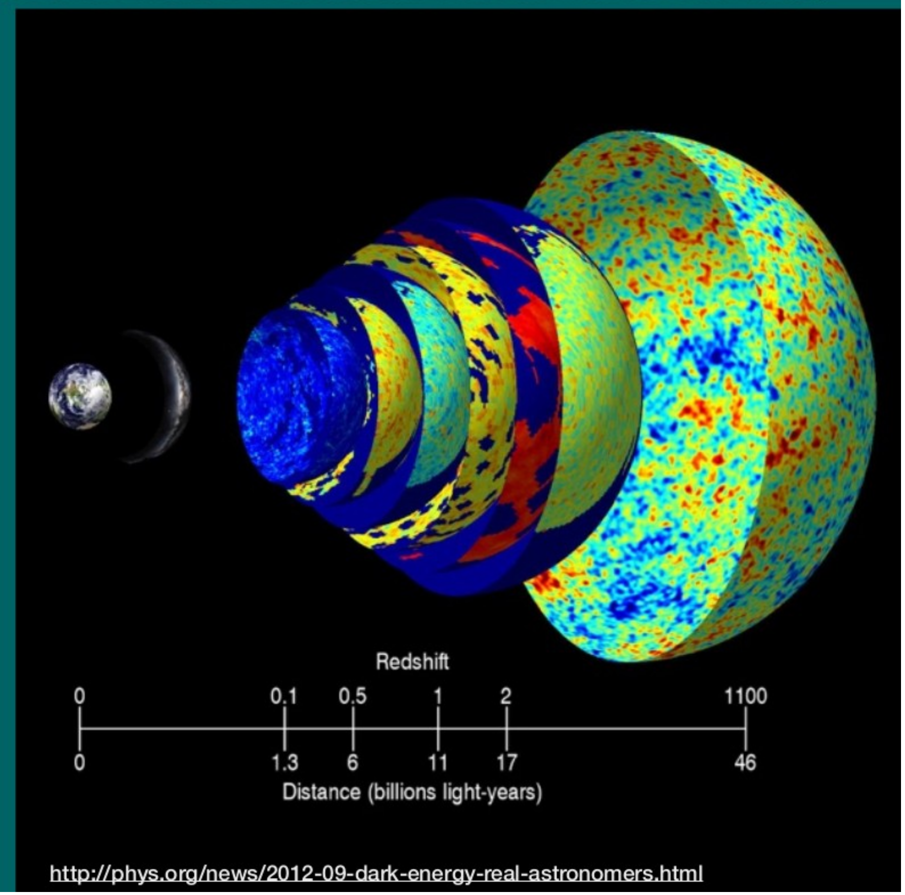

Needlets can be used to directly "sample" the physics at a particular scale

$$\beta_{jk} = \int_{S^2} T(\hat{\gamma})\psi_{jk}(\hat{\gamma})d\Omega$$

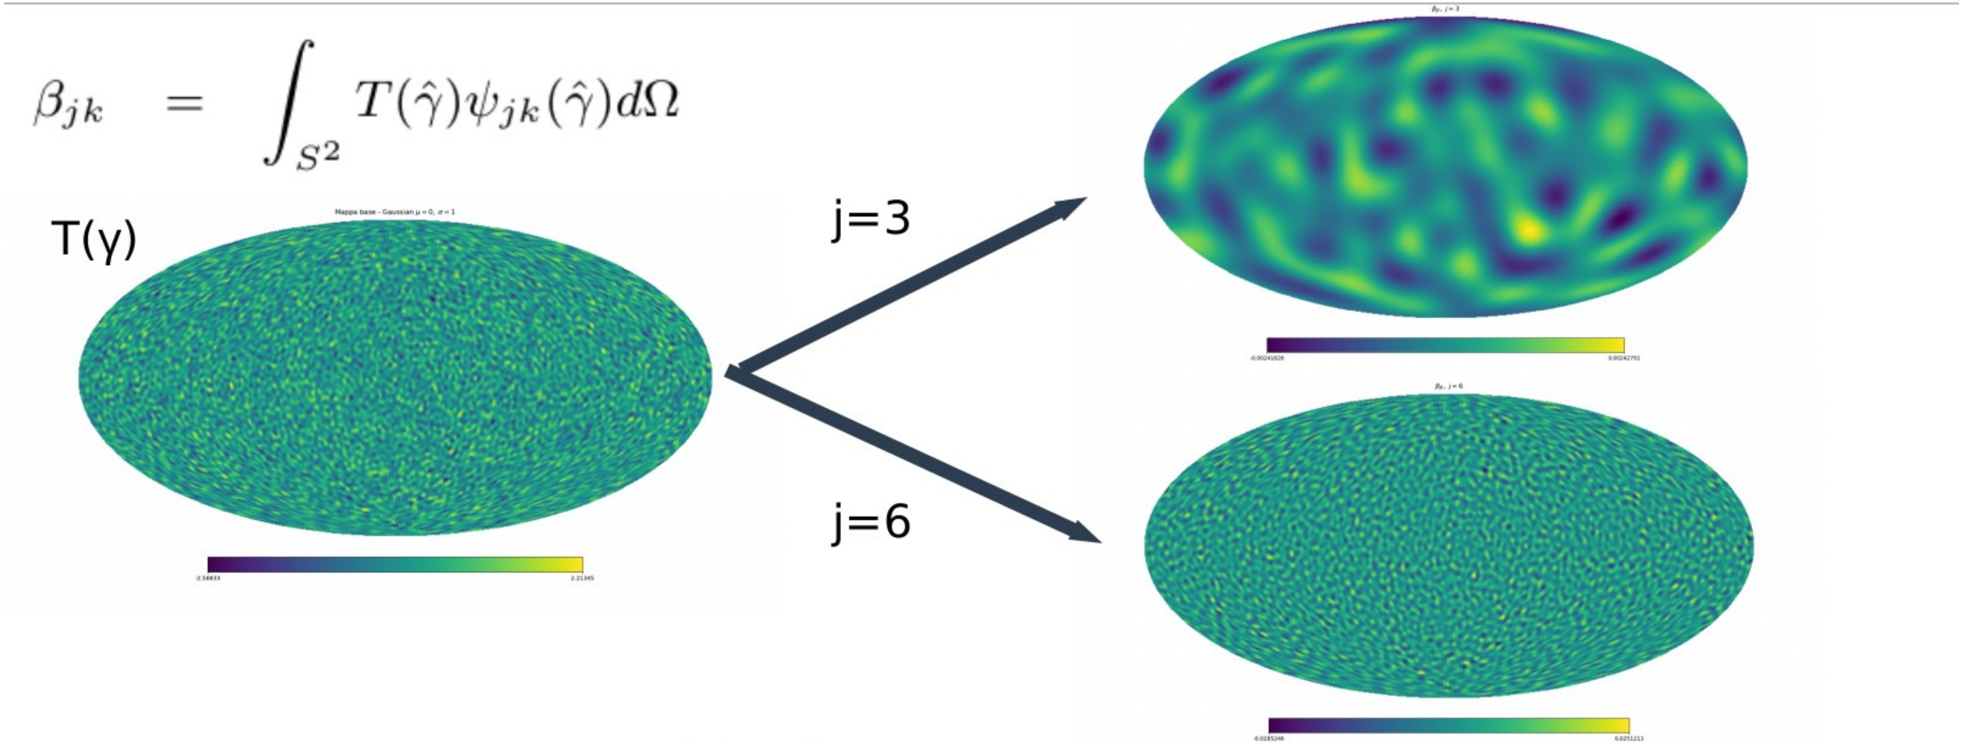

$$\psi_{jk}(\hat{\gamma}) = \sqrt{\lambda_{jk}}\sum_{\ell} b(\frac{\ell}{B^j}) \sum_{m=-\ell}^{\ell} \overline{Y}_{\ell m}(\hat{\gamma})Y_{\ell m}(\xi_{jk});$$

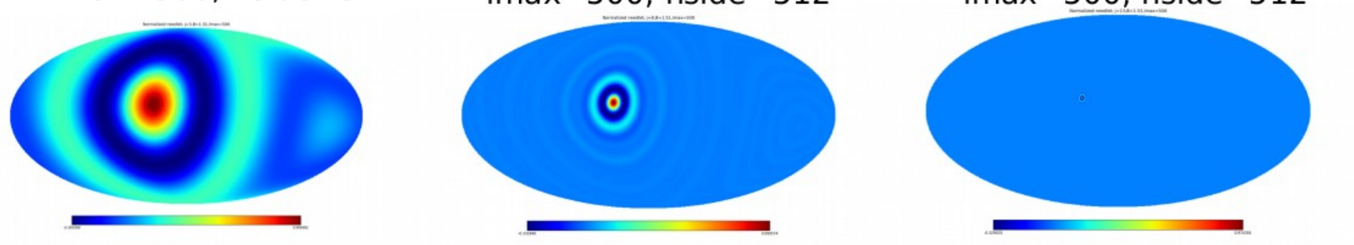

Space locality is data locality – Needlets as a viable algorithmic solution!

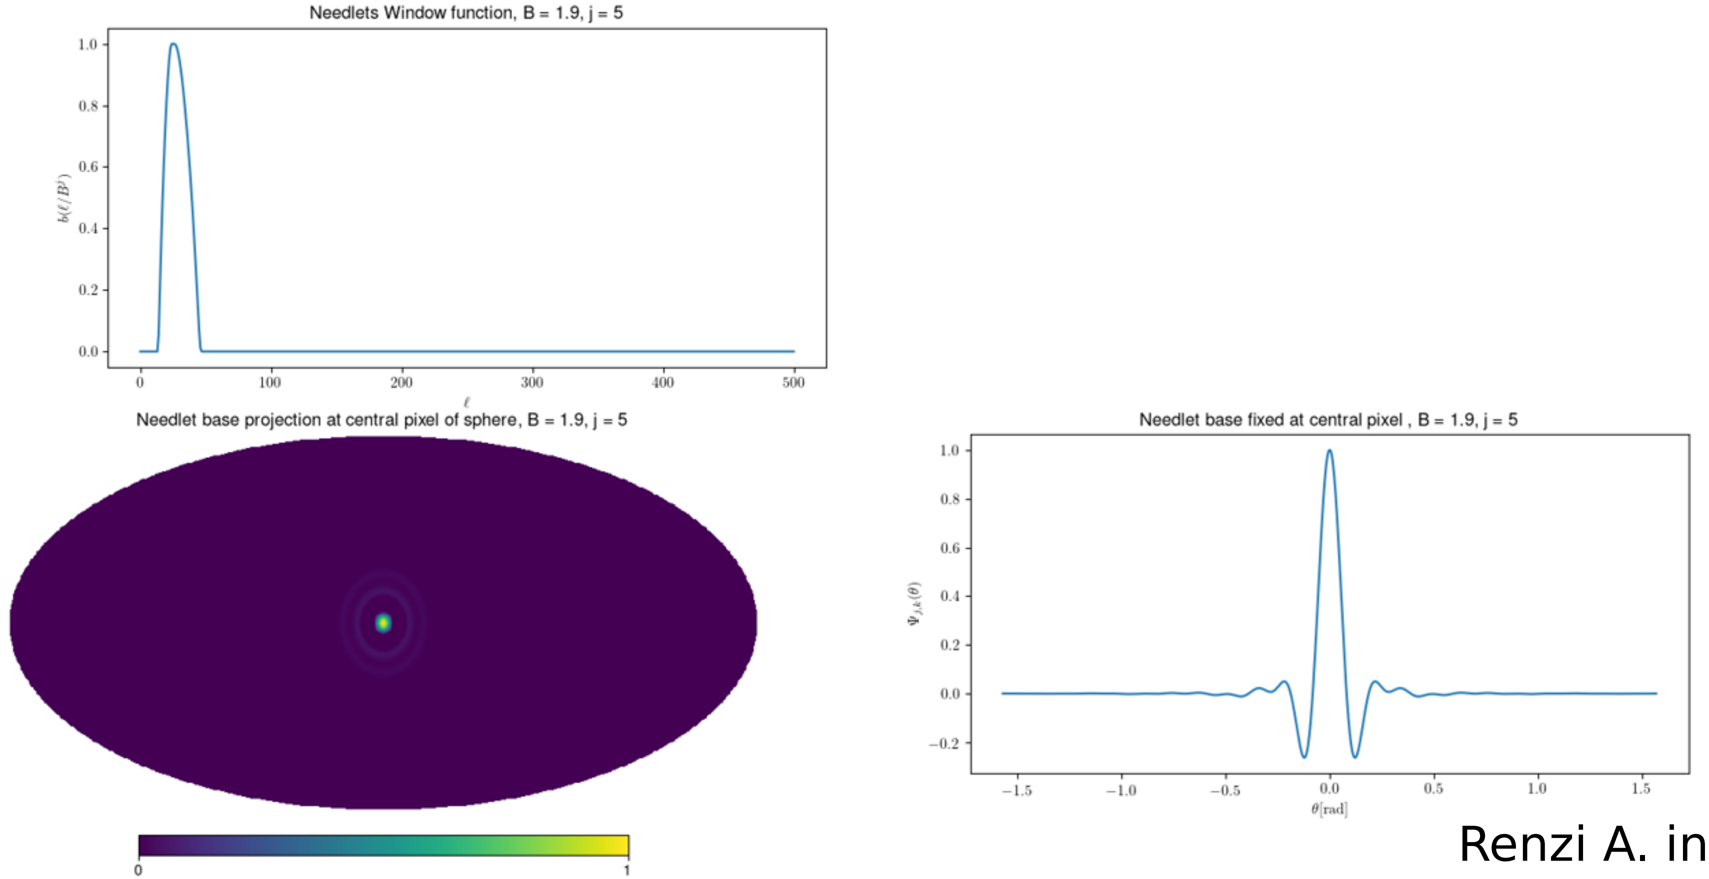

Renzi A. in prep.

Fig. 3. The figure presents the needlets localization and bounded support properties. We choose as example the needlet parameters B = 1.9 and jmax = 10, showing the mean needlet frequency j = 5. In the top left panel, the bounded support and harmonic space localization properties are depicted by the needlet window function plot, while the real space localization is visible from the bottom left panel where we plot the (adimensional) projection of the needlet base function on the sphere (placed at the central pixel), and then shows its 1d equivalent in the bottom right panel. For larger j the window function support enlarge, increasing the real space localization, the contrary happens for lower j where, however, a reduced resolution of the real space map is suitable in needlets analysis.

Exploiting Needlets double-localization property

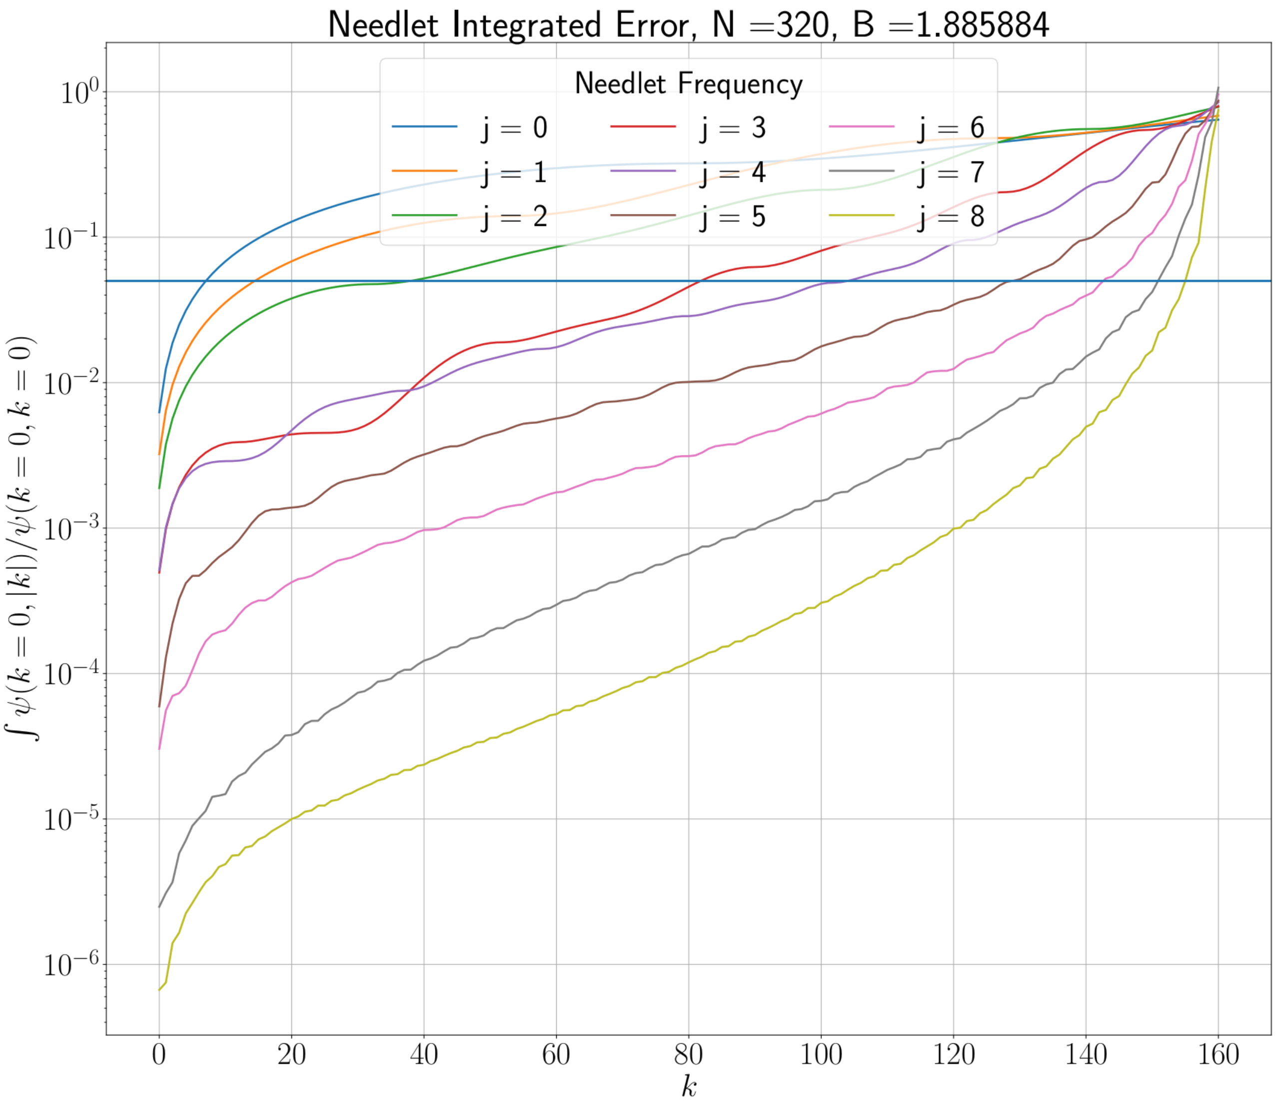

With the knowledge of you computational hardware parameters, you can simply tune your code to be efficient on that HTC system, without complex targeted code optimization!

Outline

- Present and (near) future of Cosmological dataset
- How and (why) analyze CMB and LSS data
- Evolution of computing and its compatibility with current estimators
- Needlets analysis: a possible solution
- **Prospect and conclusions**

You can always use needlets and gain some more (CMB)

Planck 2015 results. XVII. Constraints on primordial non-Gaussianity (arXiv:1502.01592)

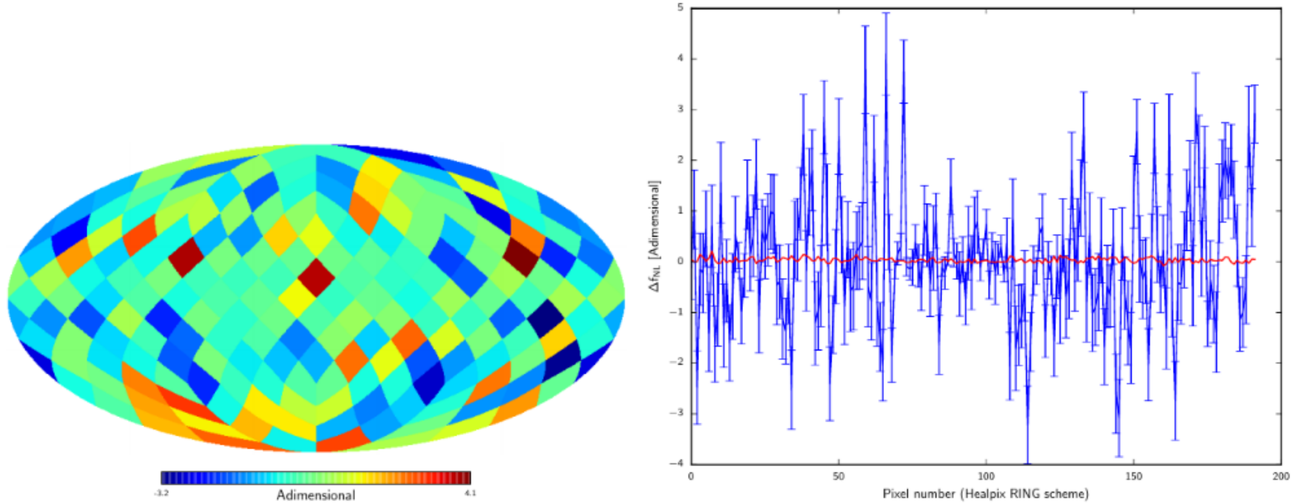

Fig. 19. Temperature only, local f_{NL} directional contributions from SMICA. As explained in the text, summing over all the pixel values would give the full sky f_{NL} needlet estimator result. The left panel displays the directional f_{NL} map. On the right, the blue points represent the f_{NL} contibution for each direction (i.e., for each pixel in the directional map), with Monte Carlo error bars. The red line is the average from simulations, which is consistent with zero.

arXiv:1407.0624 The Needlet CMB Trispesctrum;
arXiv:1310.8617 Constraining the WMAP9 bispectrum and trispectrum with needlets;
arXiv:1202.1478 On the linear term correction for needlets/wavelets non-Gaussianity estimators;
arXiv:0910.4362 Foreground influence on primordial non-Gaussianity estimates: needlet analysis of WMAP 5-year data;
arXiv:0907.4443 Non-Gaussianity in WMAP 5-year CMB map seen through Needlets;
arXiv:0906.3232 Directional Variations of the Non-Gaussianity Parameter f_NL;
arXiv:0905.3702 Needlet Bispectrum Asymmetries in the WMAP 5-year Data;
arXiv:0901.3154 An Estimate of the Primordial Non-Gaussianity Parameter f_NL Using the Needlet Bispectrum from WMAP;
arXiv:0802.4020 The needlets bispectrum

You can always use needlets and gain some more (LSS)

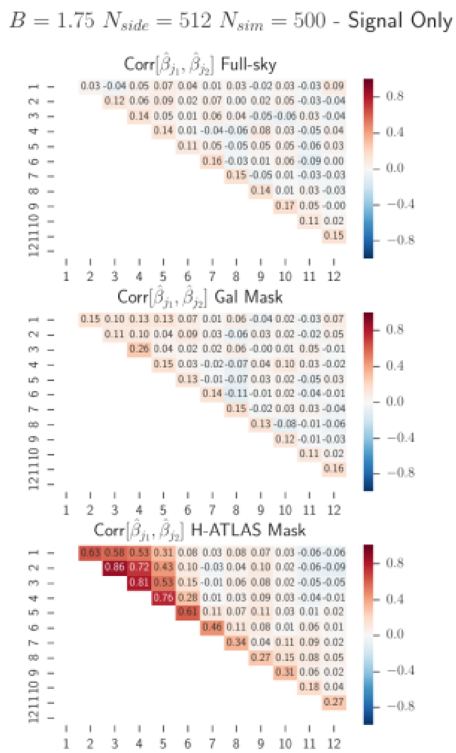

Figure 3. Cross-correlation coefficient matrices, defined as $\mathrm{Corr}_{ij} \equiv \mathrm{Cov}_{ij}/\sqrt{\mathrm{Cov}_{ii}\mathrm{Cov}_{jj}}$, about the needlet space estimator. From top to bottom we show results for the full-sky, galactic mask, and H-ATLAS mask cases respectively.

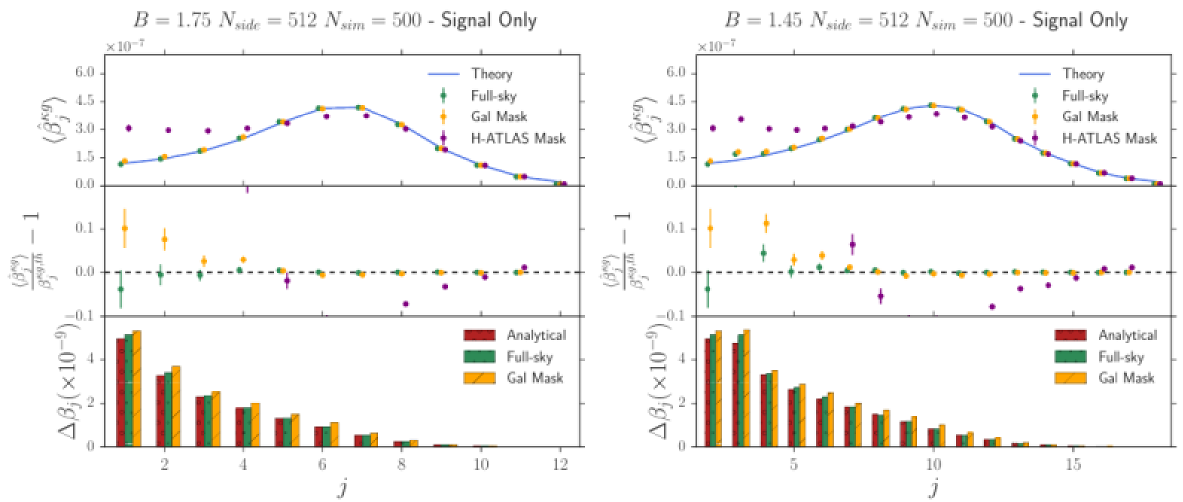

Figure 4. *Upper panel*: Recovered mean needlet cross-power spectrum between correlated CMB convergence and galaxy density maps for different masks and width parameter ($B = 1.75$ and 1.45 on the left and right parts respectively). Green, yellow and purple bandpowers represent full-sky, galactic mask (with $f_{\rm sky} = 0.65$) and H-ATLAS mask (with $f_{\rm sky} = 0.013$) cases respectively. Solid blue line is the generative theoretical input cross-power spectrum. Error bars shown are the diagonal components of the covariance matrices (defined in eq. 4.1), properly scaled by $\sqrt{N_{\rm sim}}$. Note that reconstructed mean needlet power spectra $\langle\hat{\beta}_j^{\kappa g}\rangle$ are corrected for the observed sky fraction using eq. (3.4). *Central panel*: Fractional difference between mean recovered and theoretical needlet cross-spectra for the cases shown in the upper panel. *Lower panel*: Error bars comparison for the cases shown in the upper panel.

arXiv:0910.4362 Foreground influence on primordial non-Gaussianity estimates: needlet analysis of WMAP 5-year data
arXiv:astro-ph/0611797 Dark Energy Constraints from Needlets Analysis of Wmap3 and NVSS Data
arXiv:astro-ph/0606475 Integrated Sachs-Wolfe effect from the cross correlation of WMAP3 year and the NRAO VLA sky survey data: New results and constraints on dark energy

New directions of data analysis using needlets

- **Foreground analysis**
  - exploiting sparsity in the framework of wavelet thresholding
- **Parameter estimation**
  - Uncorrelation at high frequencies is a nice feature for statistical analysis, can this be exploited to improve accuracy in parameter estimation?
- **Flat-sky/small-patch analysis of new generation sub-orbital experiments**
  - Reduce mask issues
  - No need of tangent-plane approximation
  - Mapping between different spherical projections
- **Make predictions with Needlets?**
  - In Needlets analysis, low and high frequency signals are always "distinct" by constructions. Could this be exploited at theoretical level?

Conclusion

- In this talk I have tried to convince you that, due to the evolution of computational systems, and the increase of dataset dimensions to be analyzed, the importance of computation is increasingly critical

- There is the needs of dedicated people working on that, but more importantly there is the needs of dedicated people to improve algorithms without relying (only) on code optimization (your data analysis problems won't be solved simply by employing *real* programmers)

- Needlets are a different way to analyze your data. You can use them in place of Fourier analysis and gain something more from that (scale and patches analysis)

- Needlets, as well as Fourier, could be used to make predictions too. Can we gain something from that?

- There are convincing results that needlets analysis could be optimally applicable in the new Exascale computational systems (or LHC machines!)

The End

Thanks for your attention!

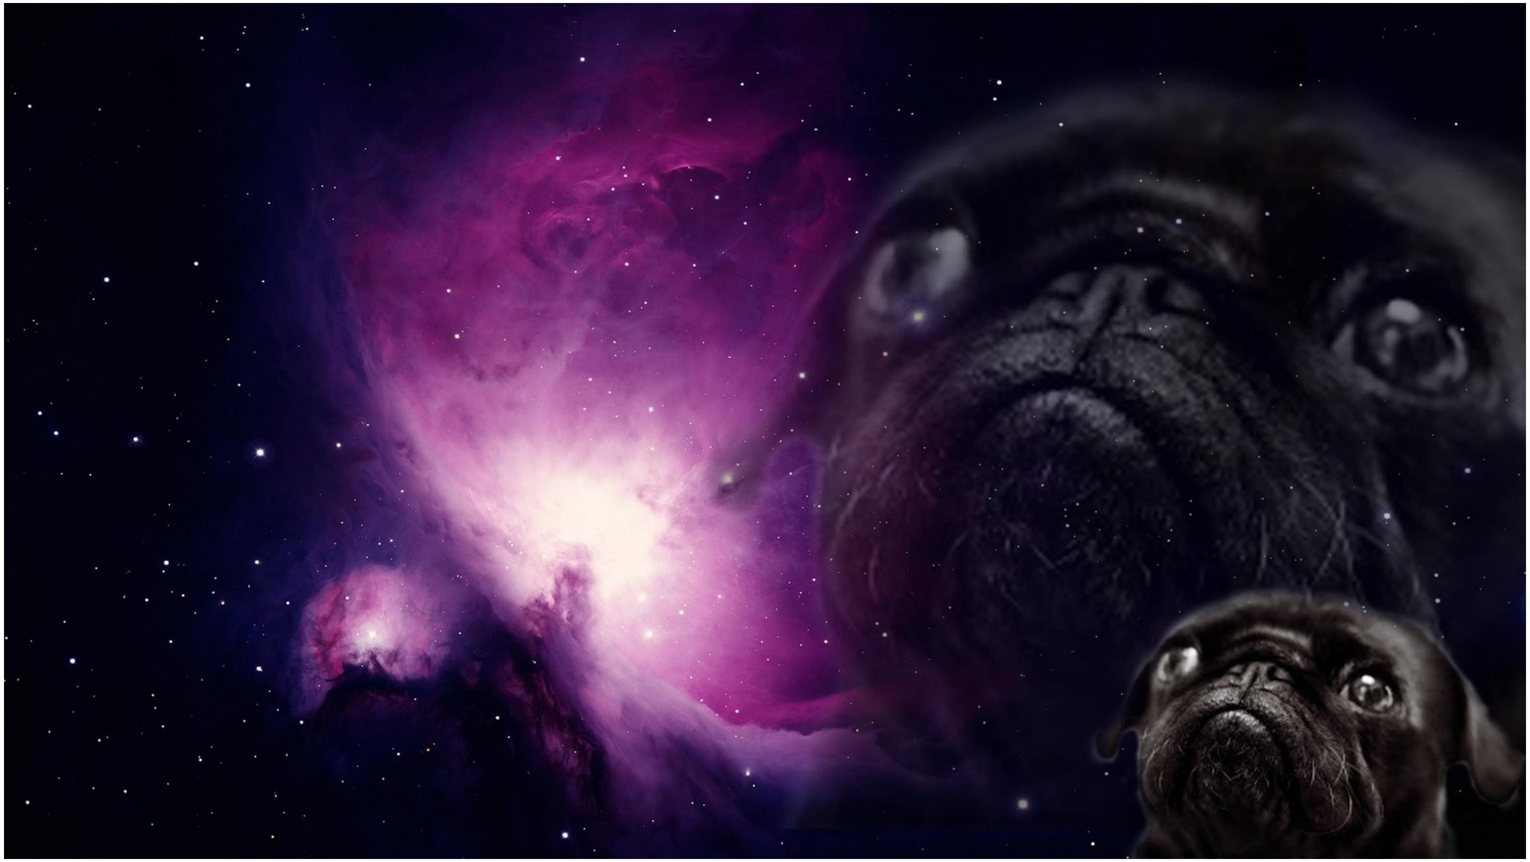

The future of cosmology is bright, we have only to look at that with the right eyes!