Big Bang, Big Data

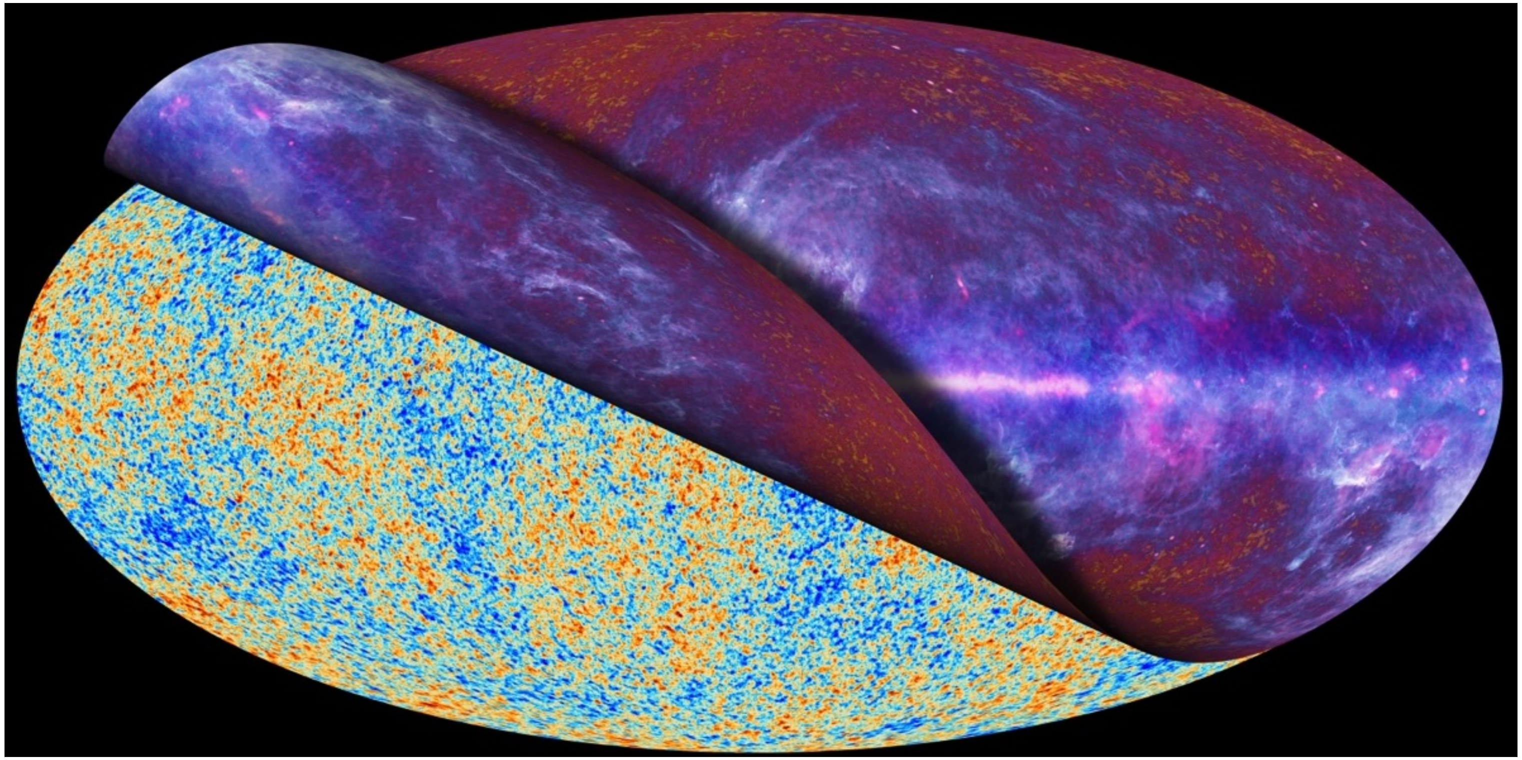

J.-F. Cardoso, CNRS (Télécom-ParisTech, IAP, APC) with the the Planck collaboration

École d'Été de Peyresq. Traitement du Signal et des Images. GRETSI et GdR ISIS. Juin 2014.

“The cosmos, back in the day”. Big Science, big news.

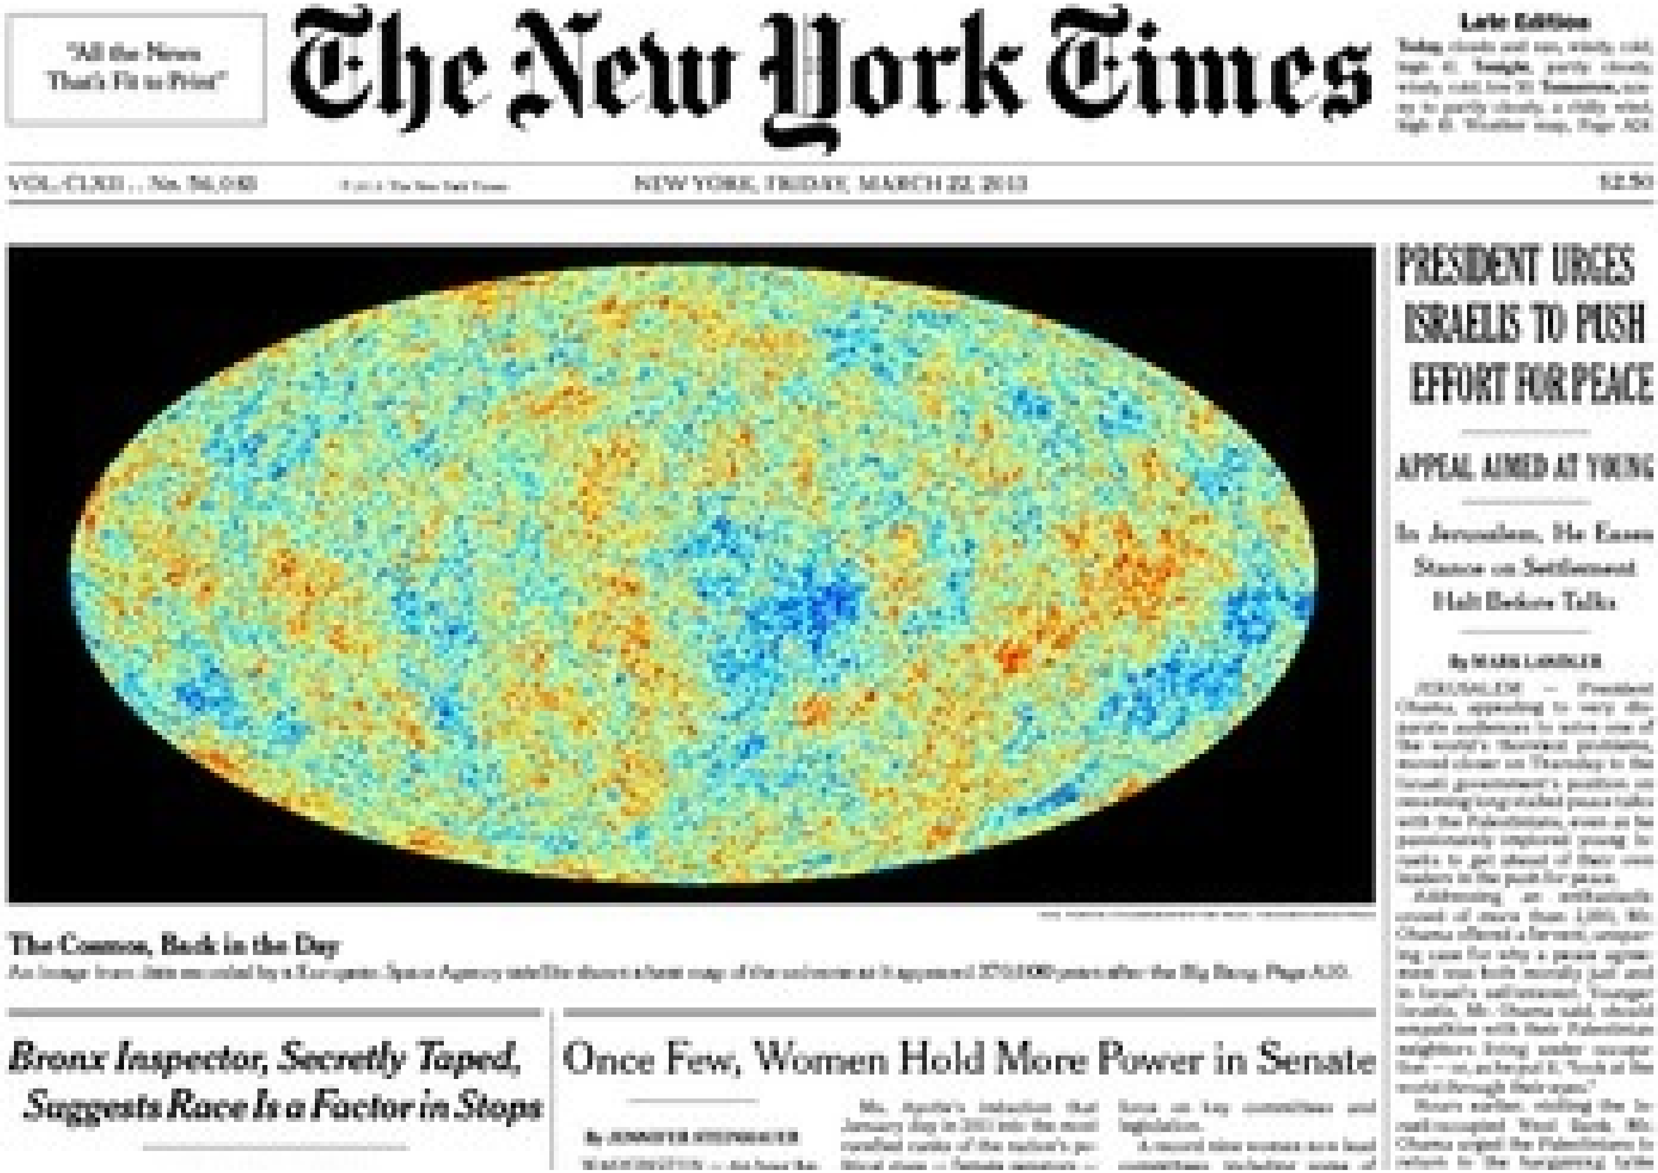
"All the News That's Fit to Print"

The New York Times

Late Edition

NEW YORK, FRIDAY, MARCH 22, 2013

$2.50

The Cosmos, Back in the Day

PRESIDENT URGES ISRAELIS TO PUSH EFFORT FOR PEACE

APPEAL AIMED AT YOUNG

In Jerusalem, He Eases Stance on Settlement Halt Before Talks

Bronx Inspector, Secretly Taped, Suggests Race Is a Factor in Stops

Once Few, Women Hold More Power in Senate

Notre plan

1. **La théorie (le Big Bang)**: un peu de cosmologie, l'histoire de l'Univers, à grands traits.

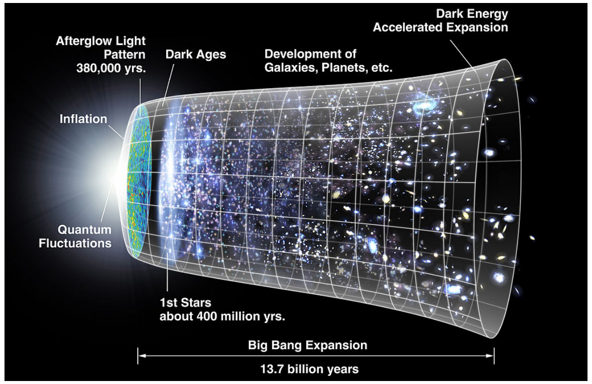

2. **Les observations**: le satellite Planck; ce qu'il a vu, et comment.

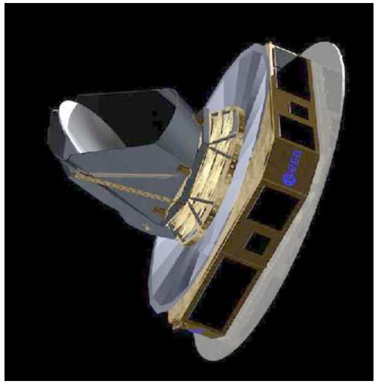

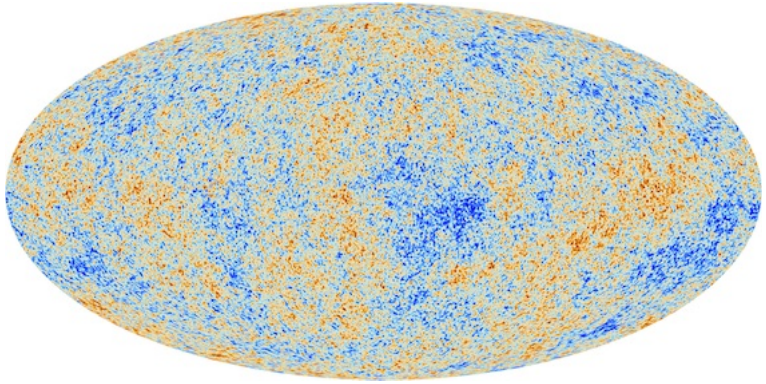

3. **Leur recontre**: extraire la Science des observations.

L'âge de l'Univers est . . .

Préliminaire: des formes et des couleurs

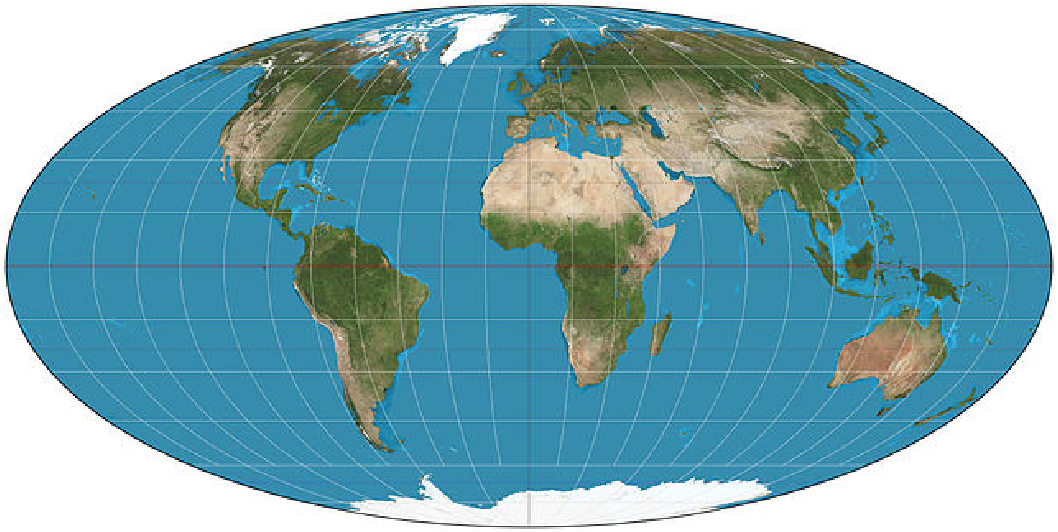 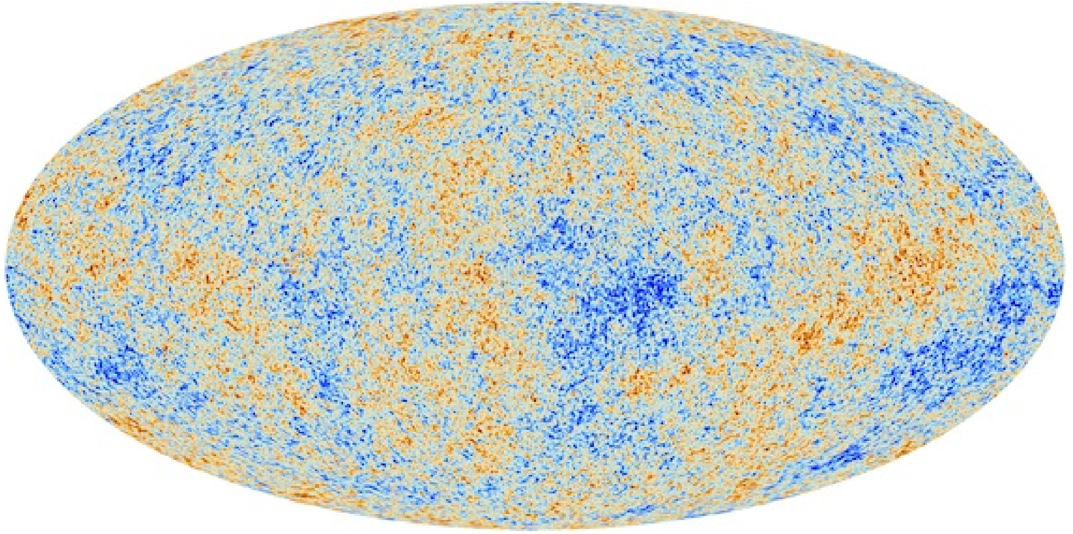

La Terre mise à plat (projection de Mollweide) : un sphéroïde bleu, vert, jaune, blanc,. . . à l'oeil

Le rayonnement fossile, mesuré sur toute la voûte céleste, mis à plat.

Une carte de température en couleurs arbitraires.

Big Bang

L'histoire de l'Univers en une image

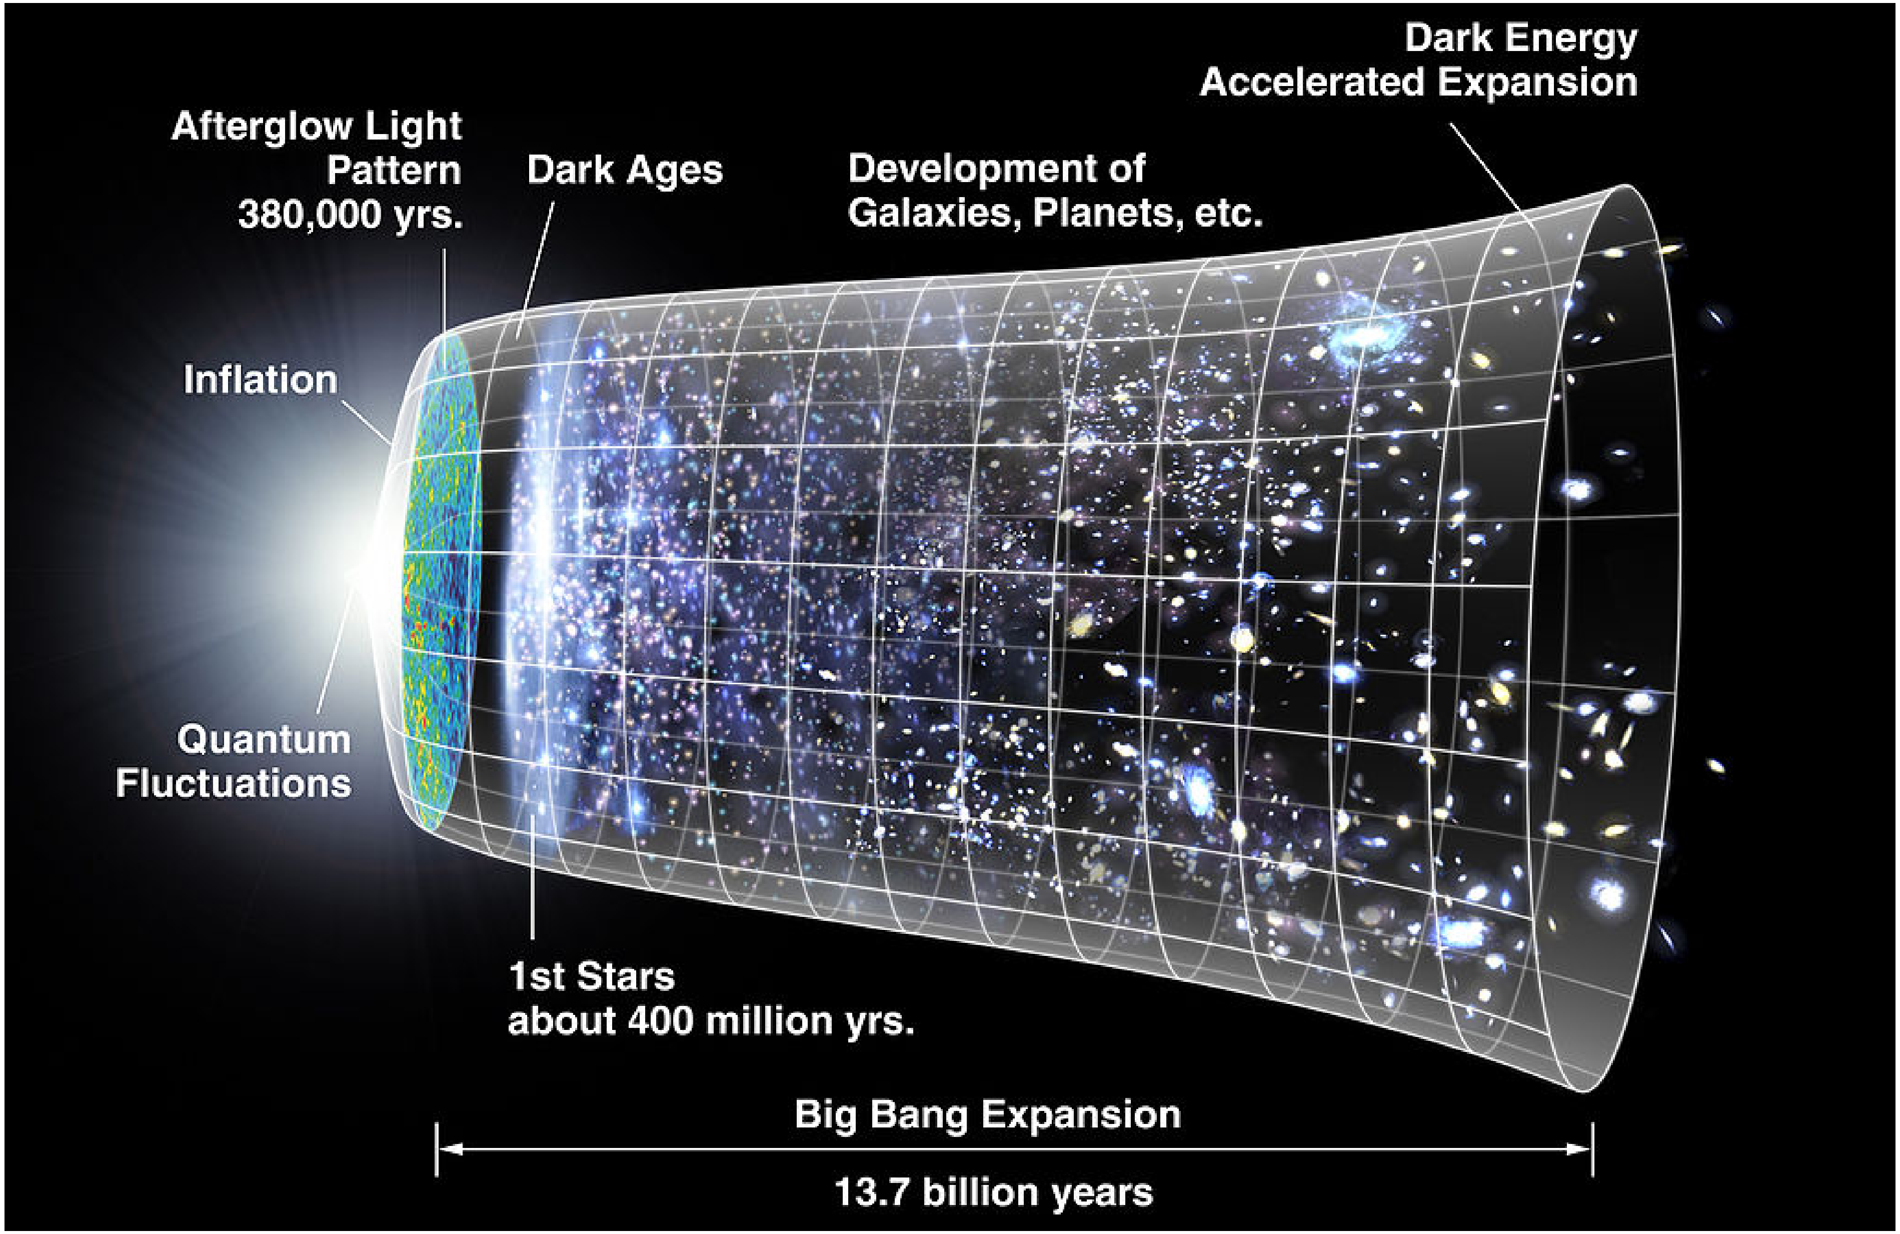

Quand l'Univers devient **transparent**, la lumière se **fossilise**.

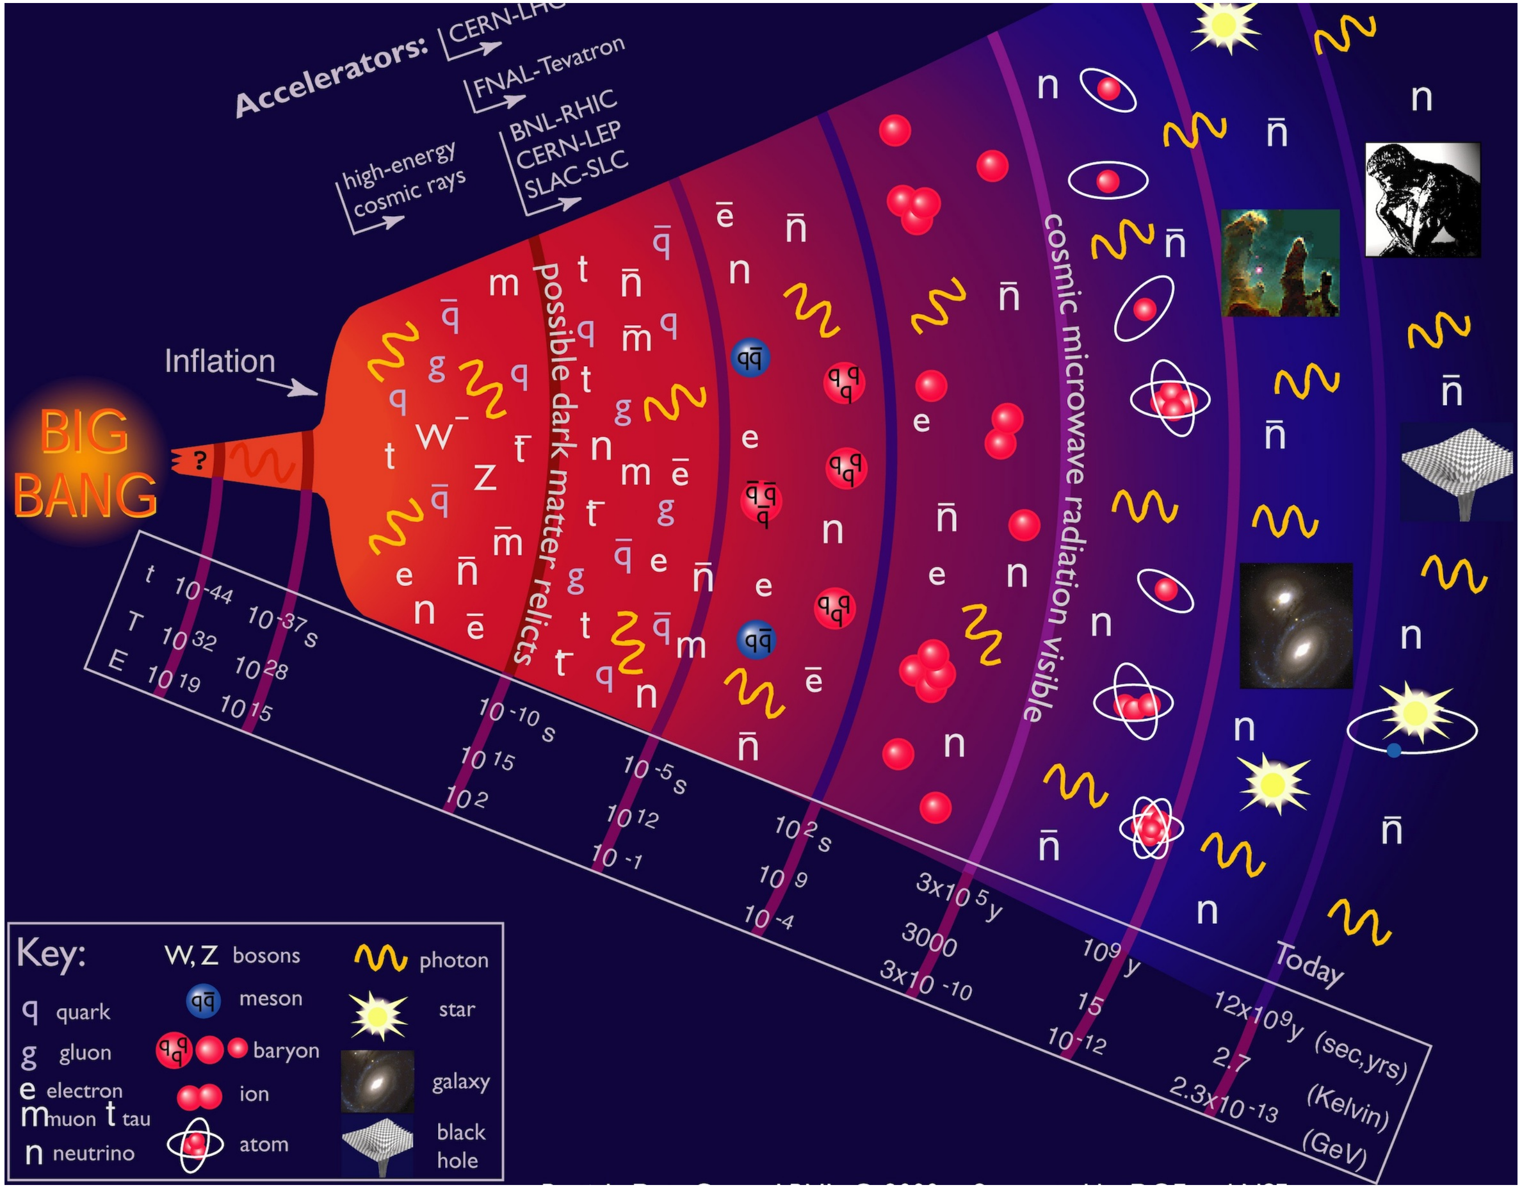

La plus vieille image du monde

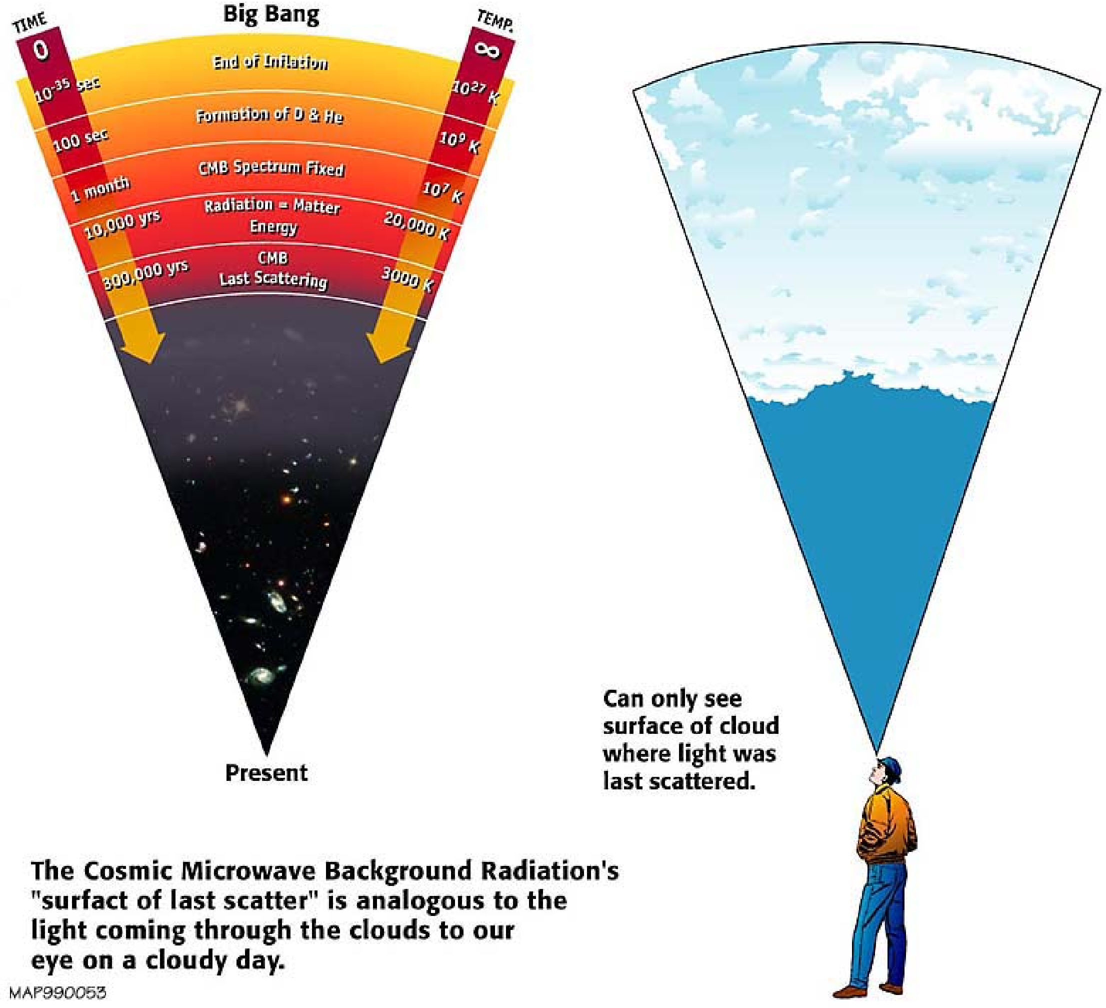

Peut-on réellement percevoir une lumière si lointaine?

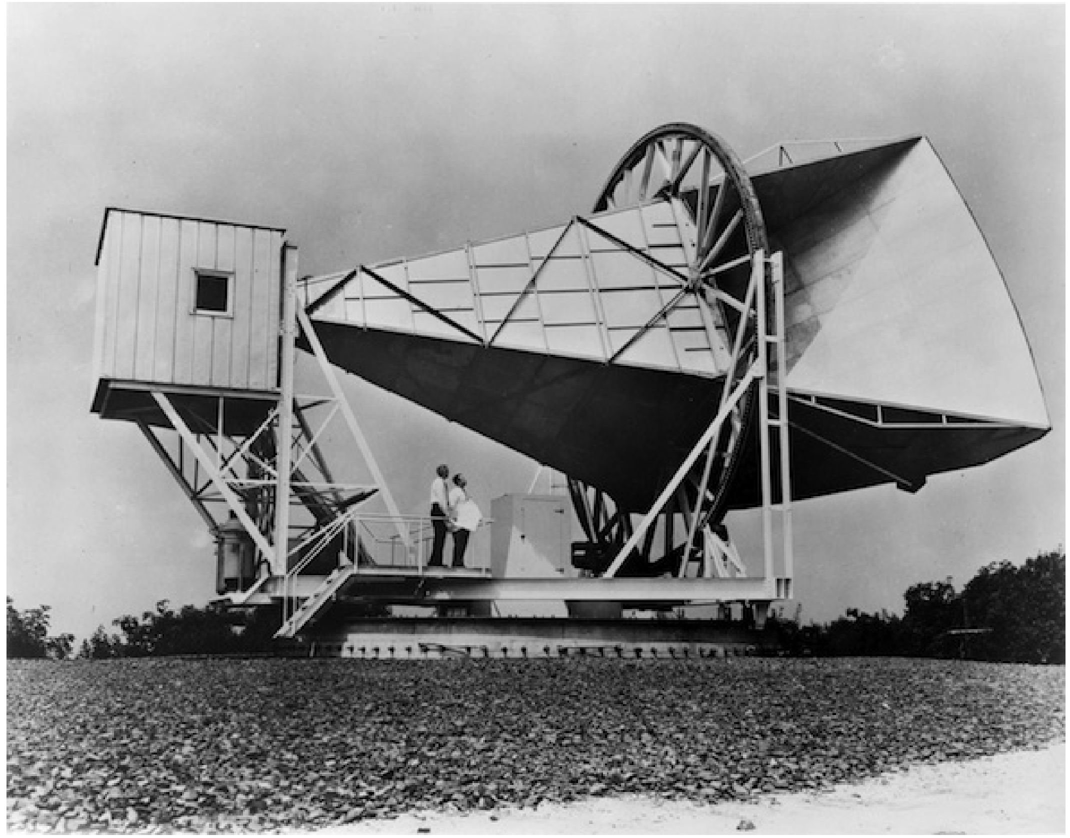

Deux gars l'ont vue, sans faire exprès, en 1965. Nobel pour Penzias et Wilson!

Et ils l'ont trouvée **uniforme** et **froide**: à peu près 3 degrés Kelvin. **C'est-à-dire?**

.

Lumière, matière et température

La théorie (Planck)

La répartition d'énergie en function de la fréquence ν d'un rayonnement de température T est théorisée par la loi de Planck

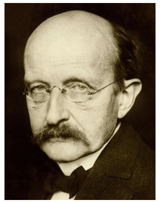

$$I(\nu) = \frac{2h\nu^3}{c^2}\frac{1}{e^{h\nu/kT}-1}$$

avec les honorables constantes:

- c: vitesse de la lumière
- h: constante de Planck
- k: constante de Boltzmann

Les mesures (COBE 1992)

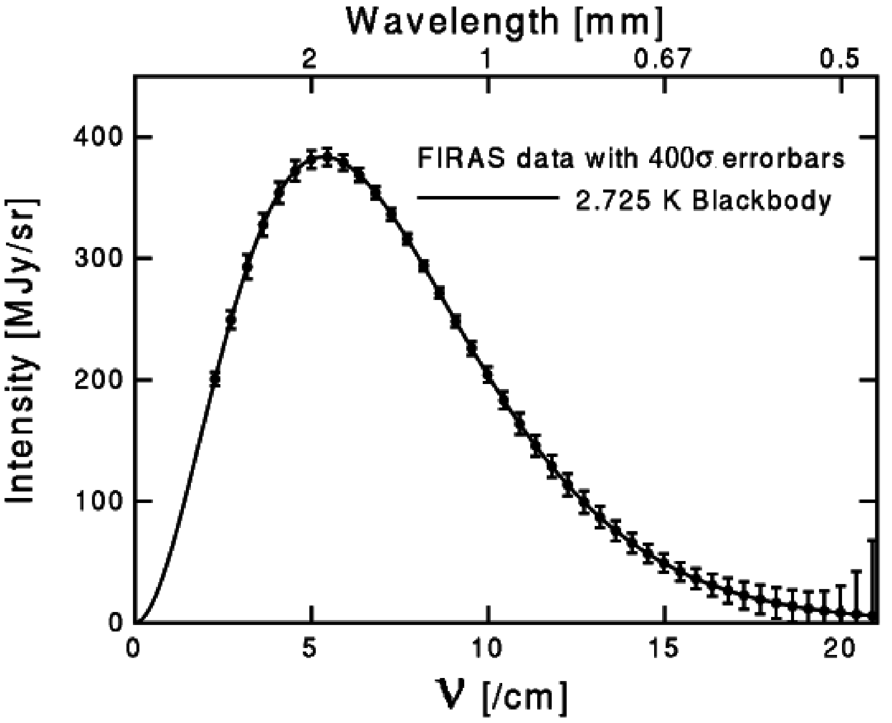

Éblouissant accord entre la mesure et la forme théorique!

L'Univers est rempli de vieux photons froids à 2.725 degrés Kelvin.

Et donc, il s'est dilaté 1000 fois depuis la recombinaison: $z_{\mathrm{rec}} \approx 1000$.

Le rayonnement fossile est isotrope, mais pas trop

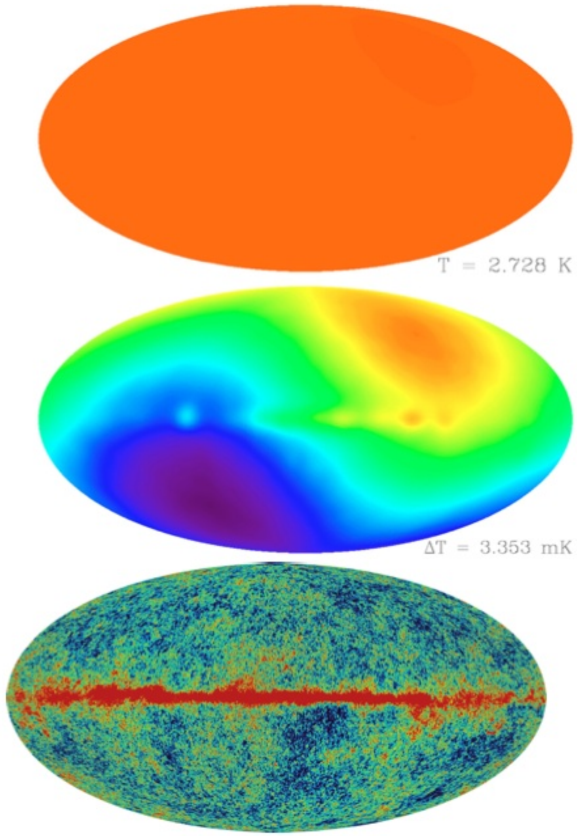

CMB :

isotropic

black-body spectrum

Penzias& Wilson 1965

$T=2.725K$

Dipole :

kinematic of obs. vs CMB

galactic + solar system

COBE 1992

$\Delta T/T \sim 10^{-3}$

CMB Anisotropies (T,P) :

dentity anisotropies in

primordial universe

↔cosmo. parameters

WMAP 2003

$\Delta T/T \sim 10^{-5}$ (I)

Des anisotropies de 0,0001 degrés (Kelvin) et de 1 degré (centigrade).

L'Univers sauvé par les anisotropies

A peine âgé de 380,000 ans, l'Univers est dans une situation délicate:

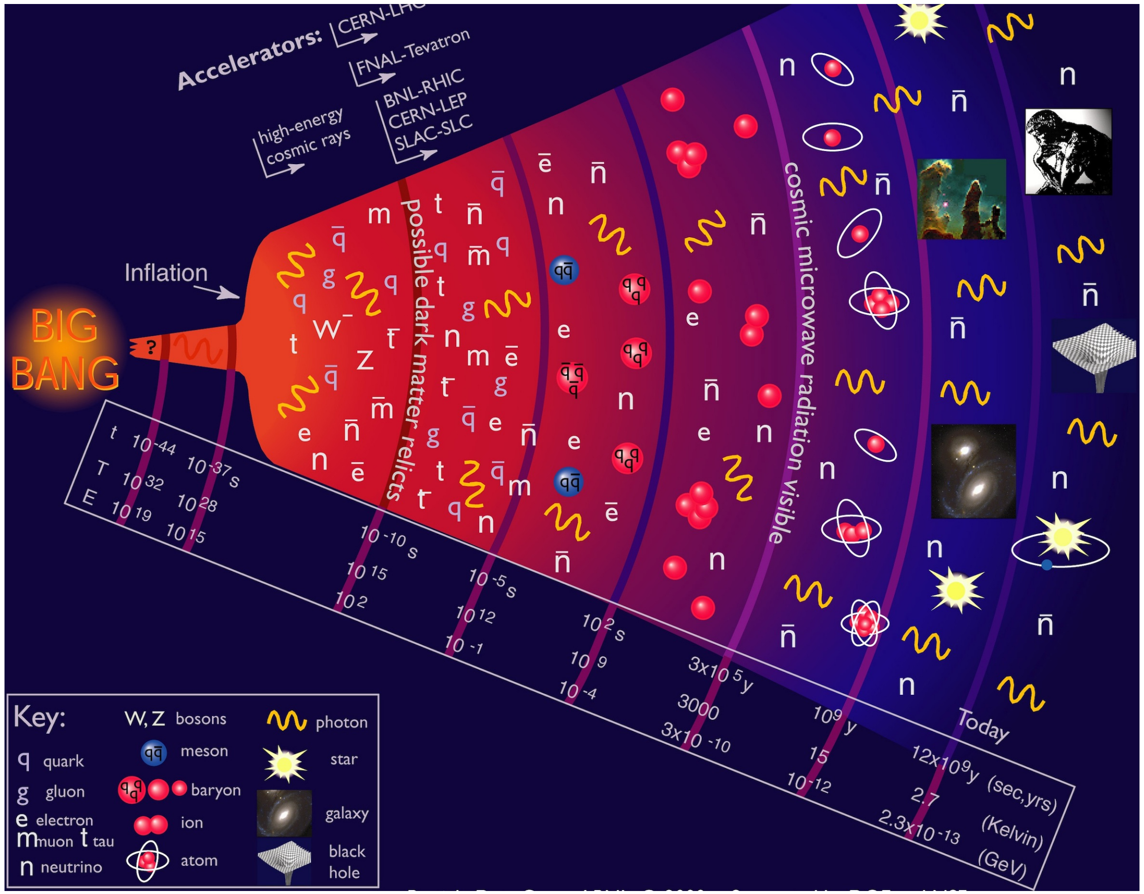

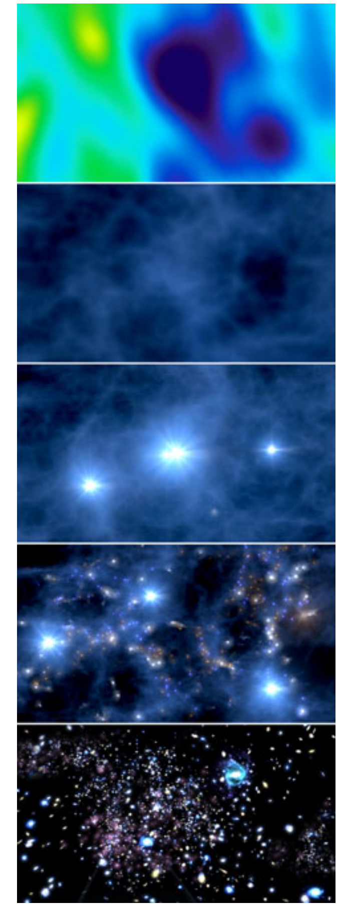

.

Beaucoup reste donc à faire: les étoiles, les galaxies, les planètes, la vie. . .
Tout va-t-il donc partir à vau-l'eau? Il faut initier les grandes structures!

Planck

Pour mieux y voir: la mission Planck

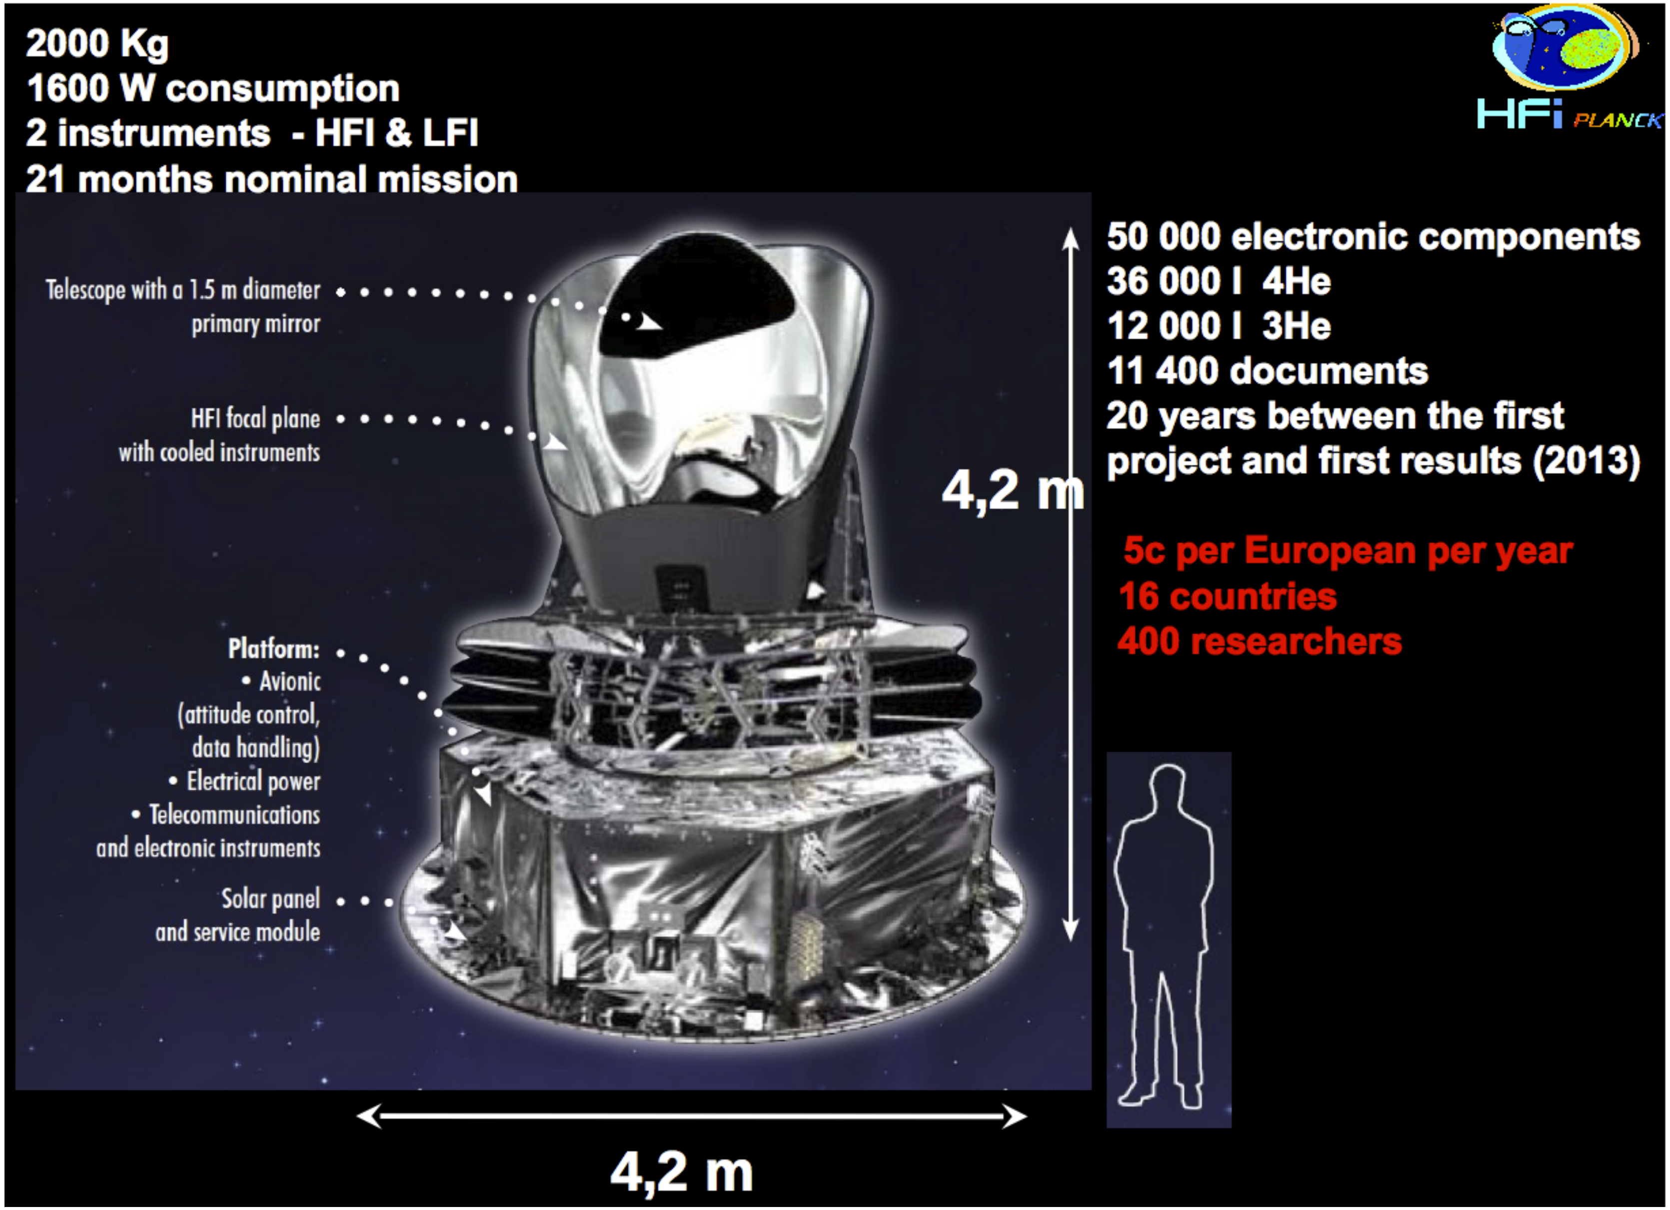

En route pour le deuxième point de Lagrange!!

Les cosmologistes adorent le 2ème point de Lagrange Terre-Soleil

Scanning the sky

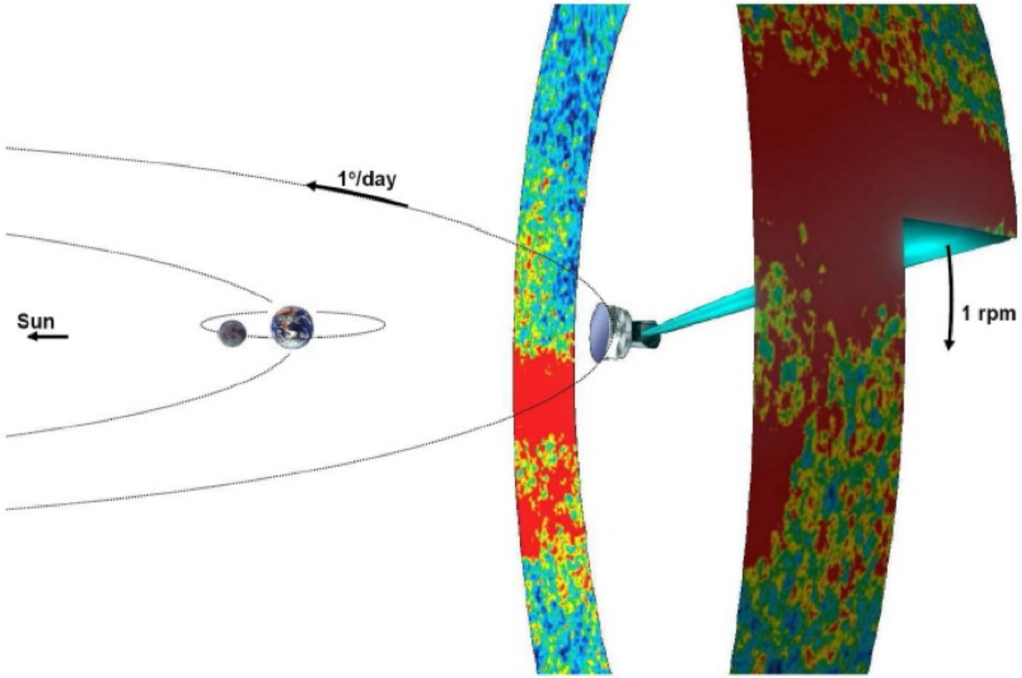

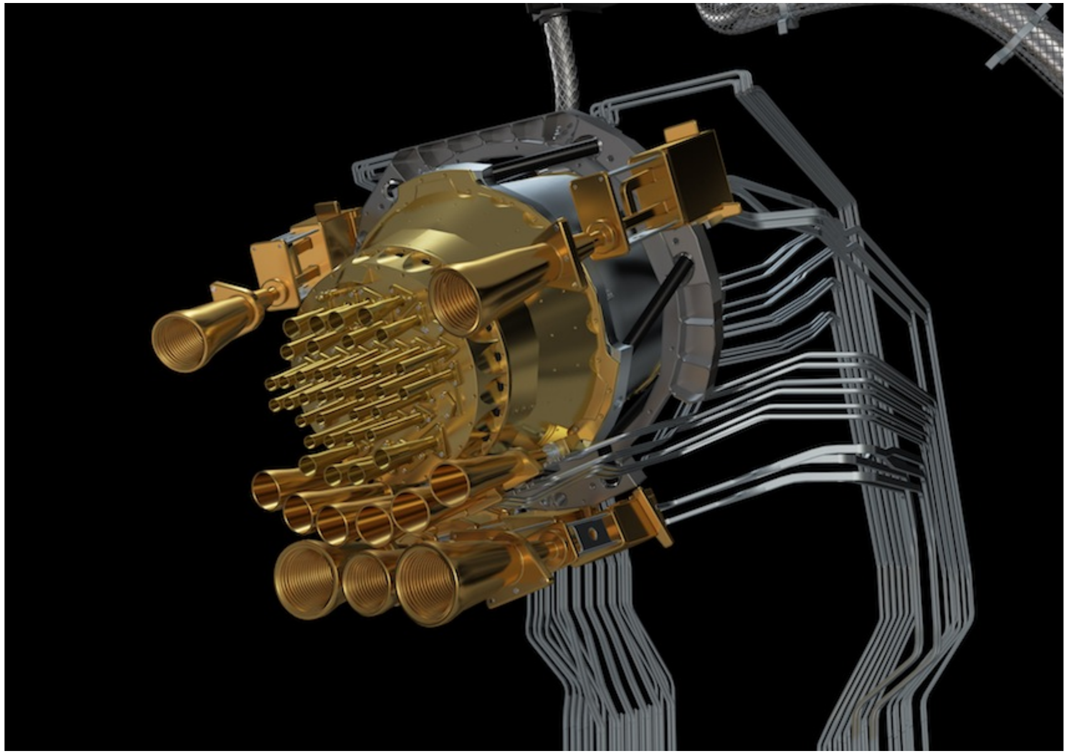

Le téléscope focalise la lumière vers 52 cornets (HFI) qui la filtrent et la guident vers les bolomètres refroidis à 0.1 degrés Kelvin.

Notez les tuyaux: Planck est aussi un exploit frigorifique.

Le plan focal de Planck HFI

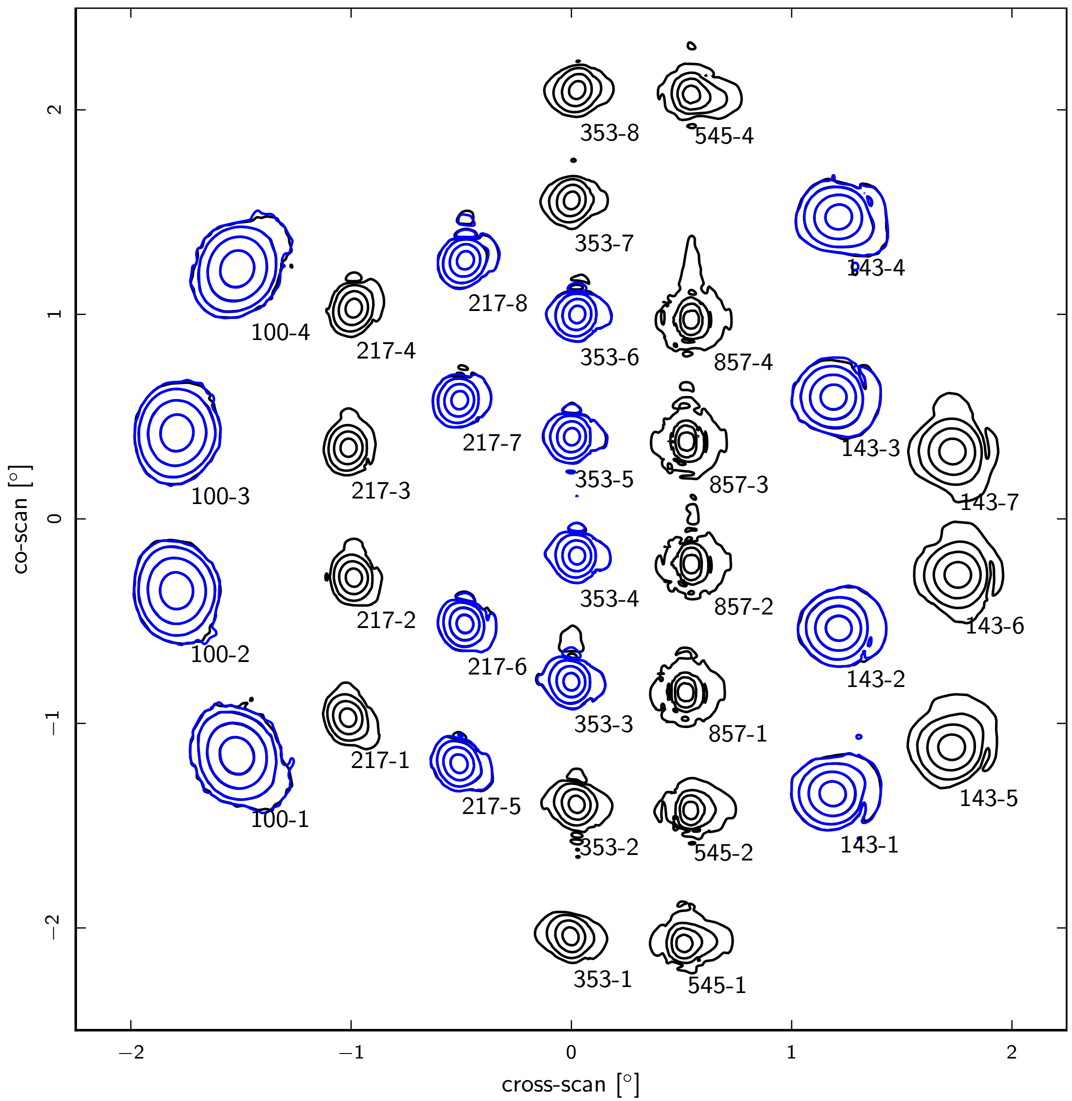

5 arc-minute resolution in best channels → $50\,10^6$ pixels over the sky.

Les canaux spectraux moyens de Planck HFI

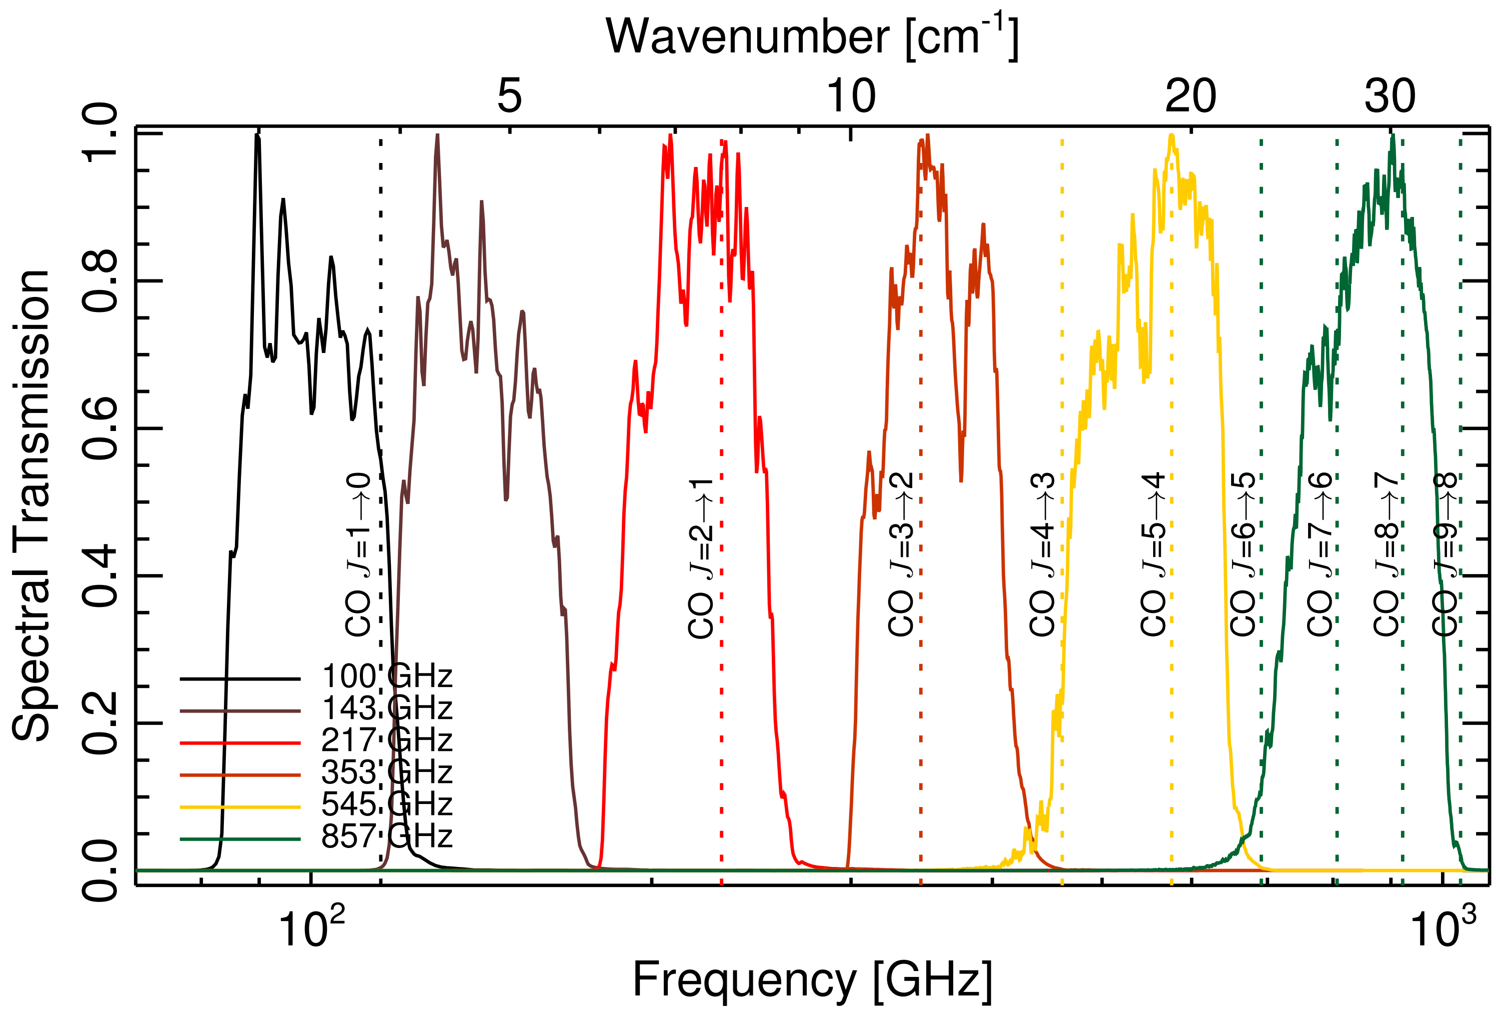

Signals from a 143 GHz, a 545 GHz, a dark bolometer

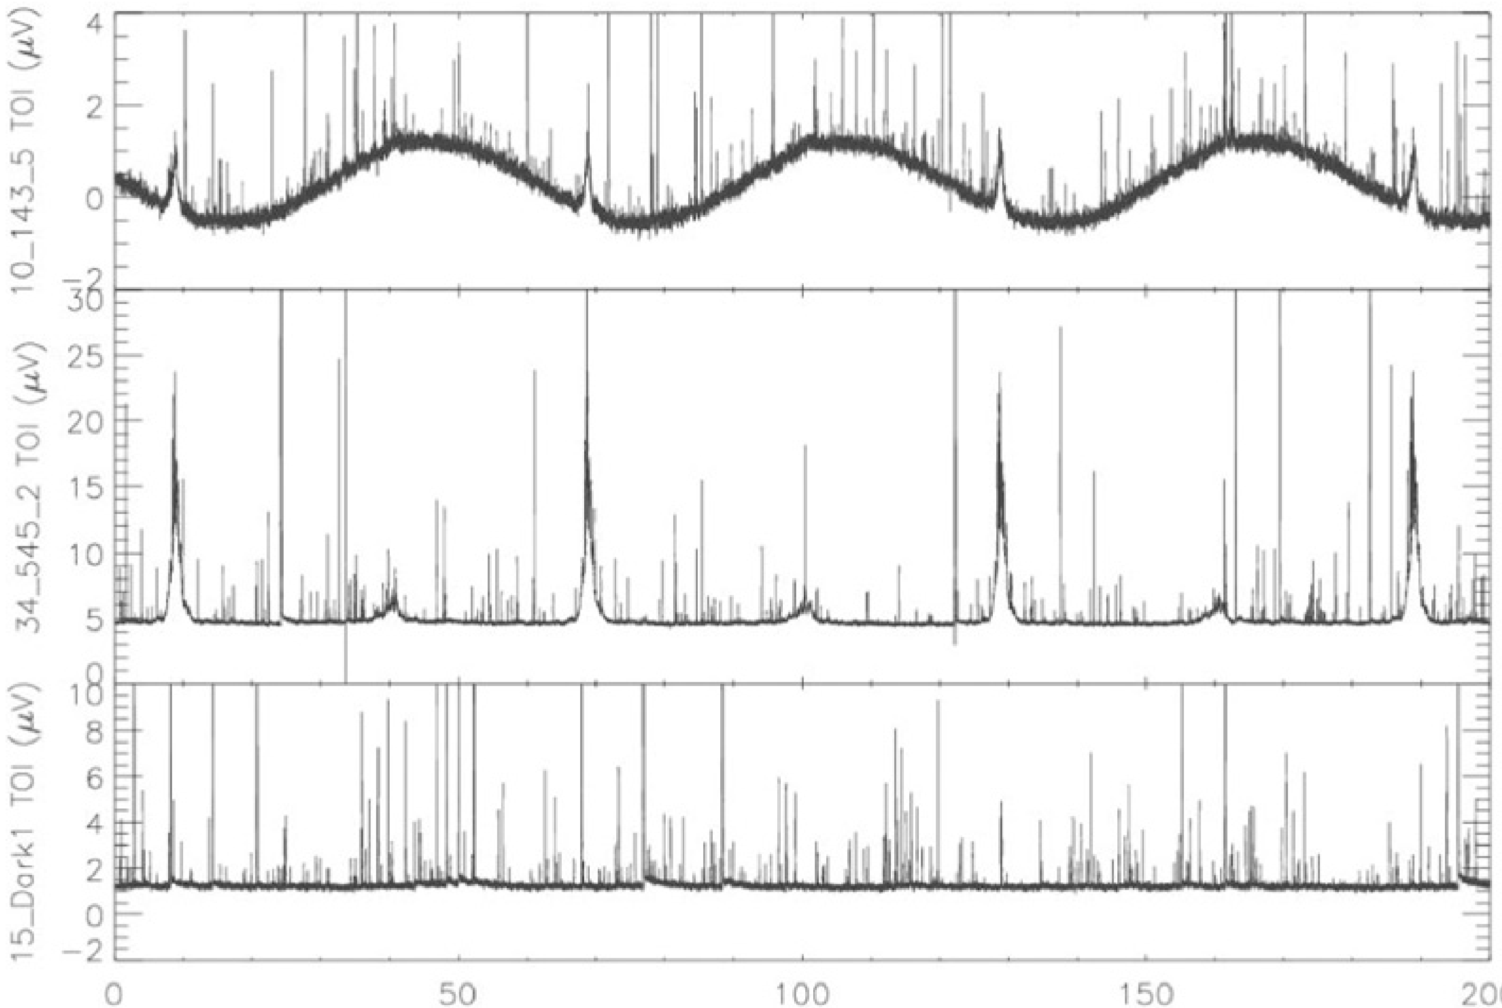

A serious case of glitchopathy.

Les mêmes, après deglitching

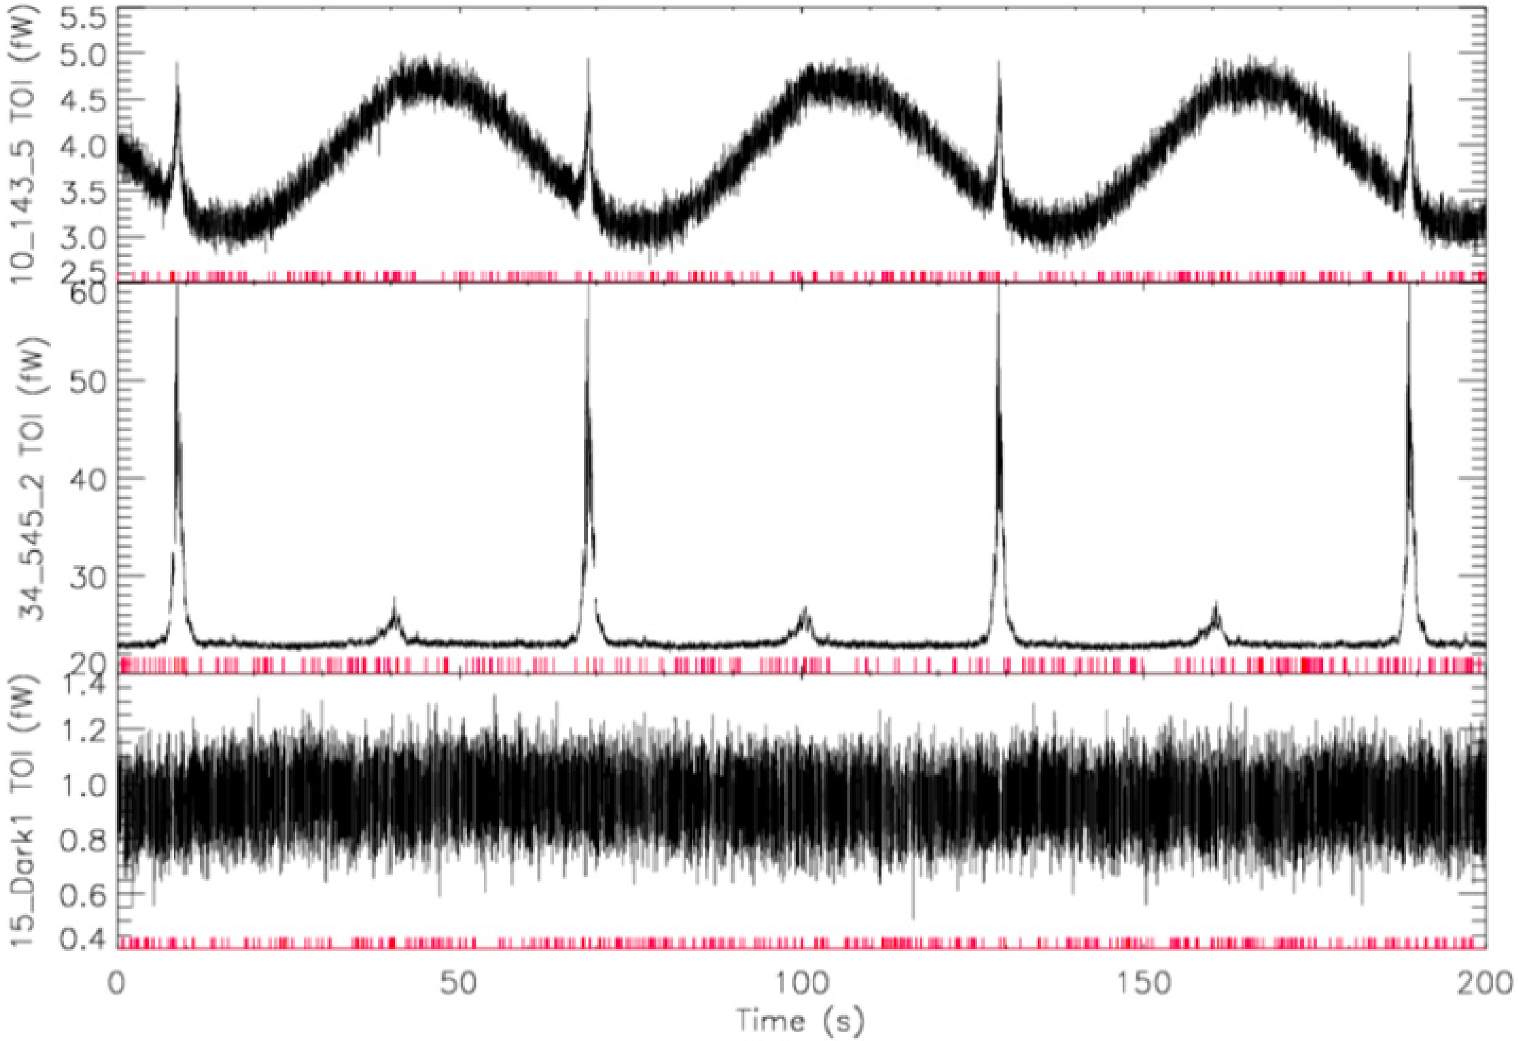

After some glitchotherapy (thanks to recent advances in glitchology).

Map Making needed

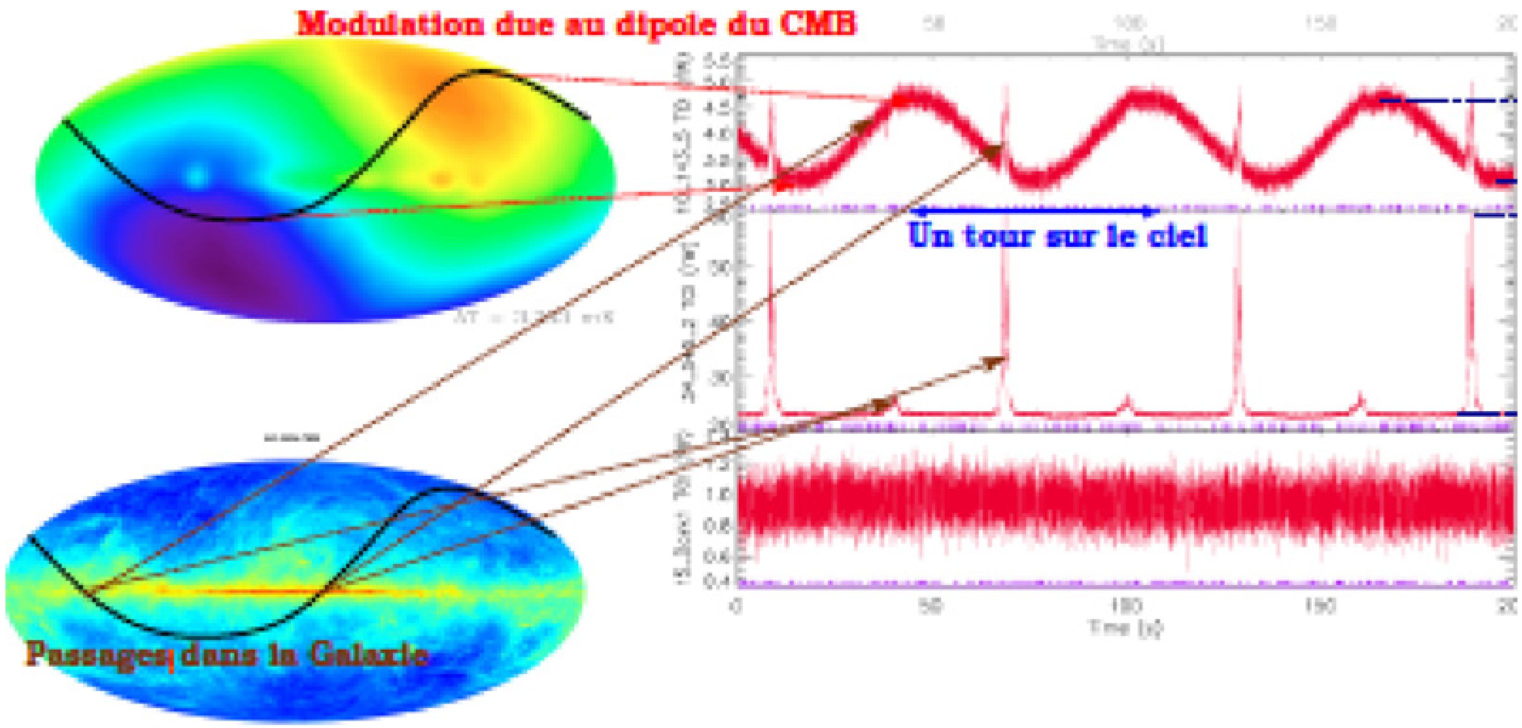

The cosmological signal is still under the noise.

How to go from noisy time lines to less noisy spherical maps: another adventure in big data.

The sky as seen by Planck

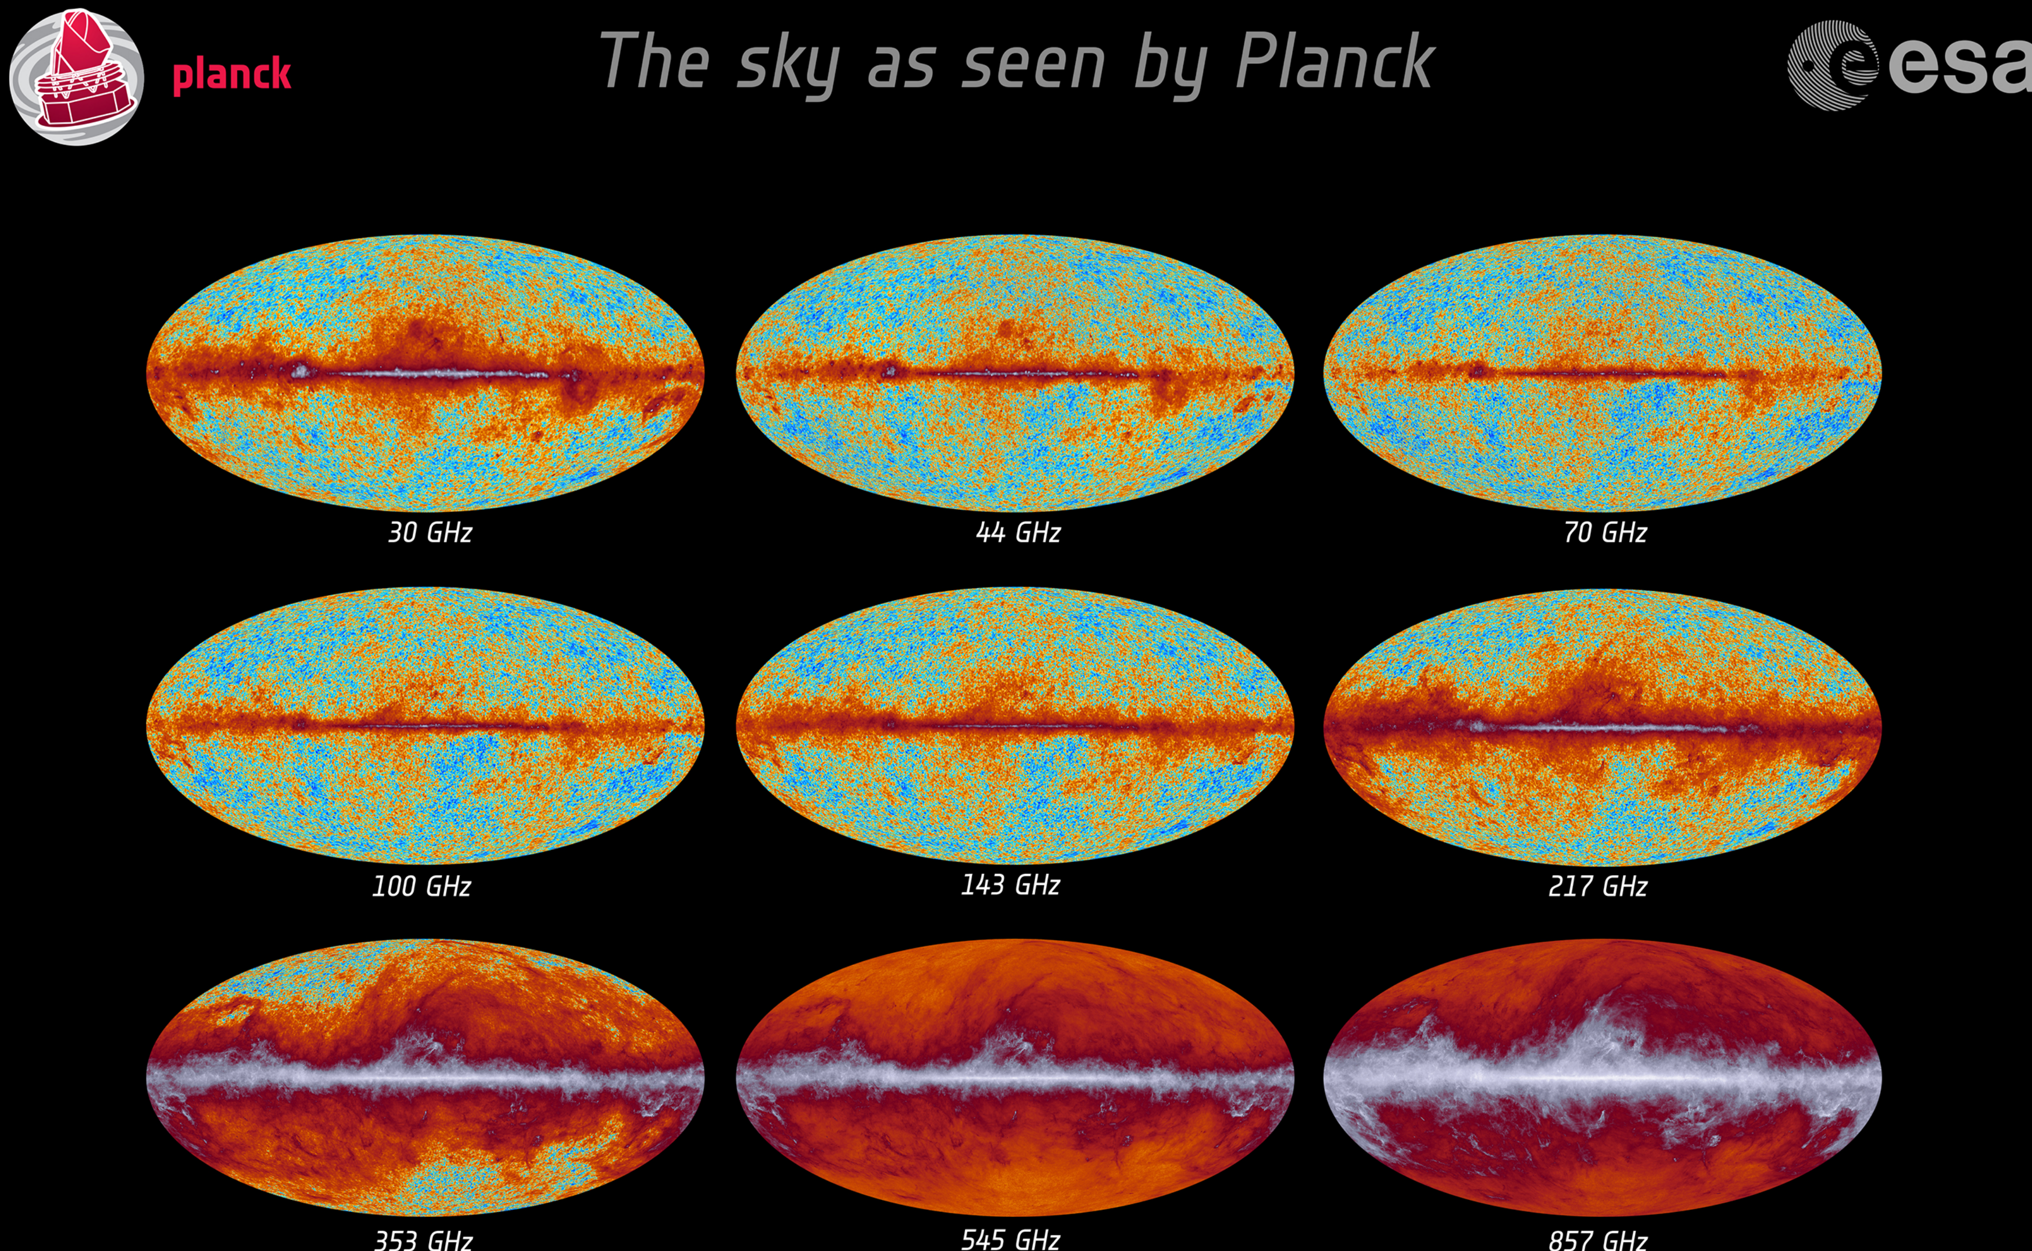

Combinaison des 9 canaux Planck pour extraire le rayonnement fossile

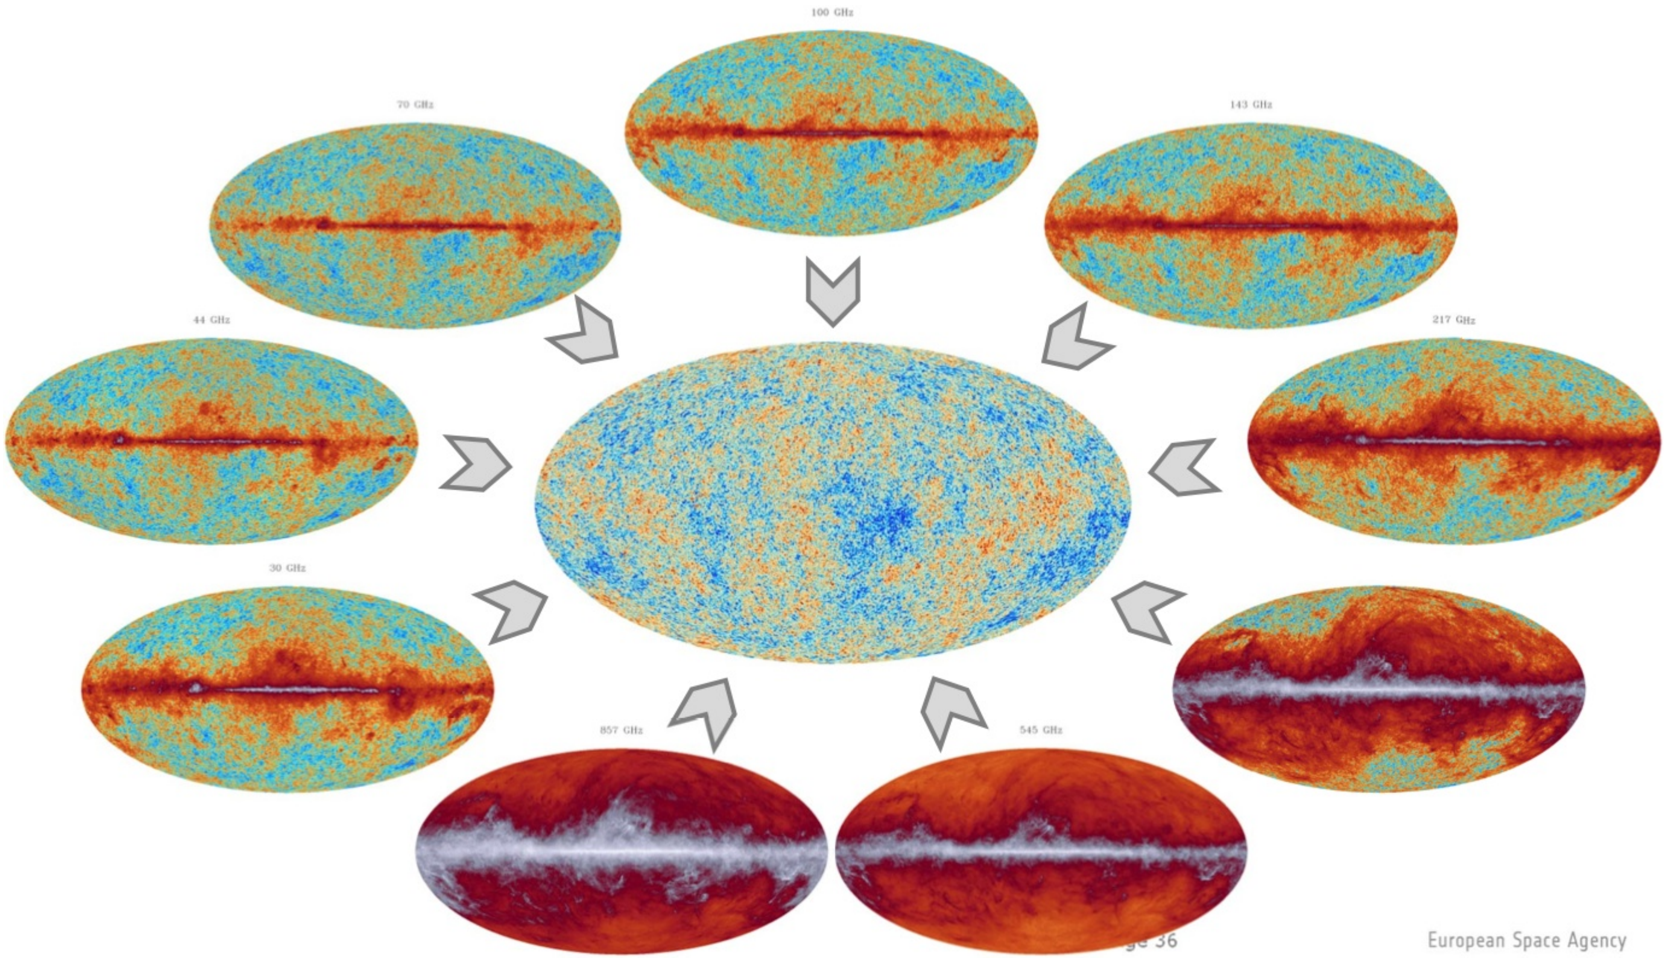

Credit: ESA, FRB

La plus vieille image du monde, par Planck

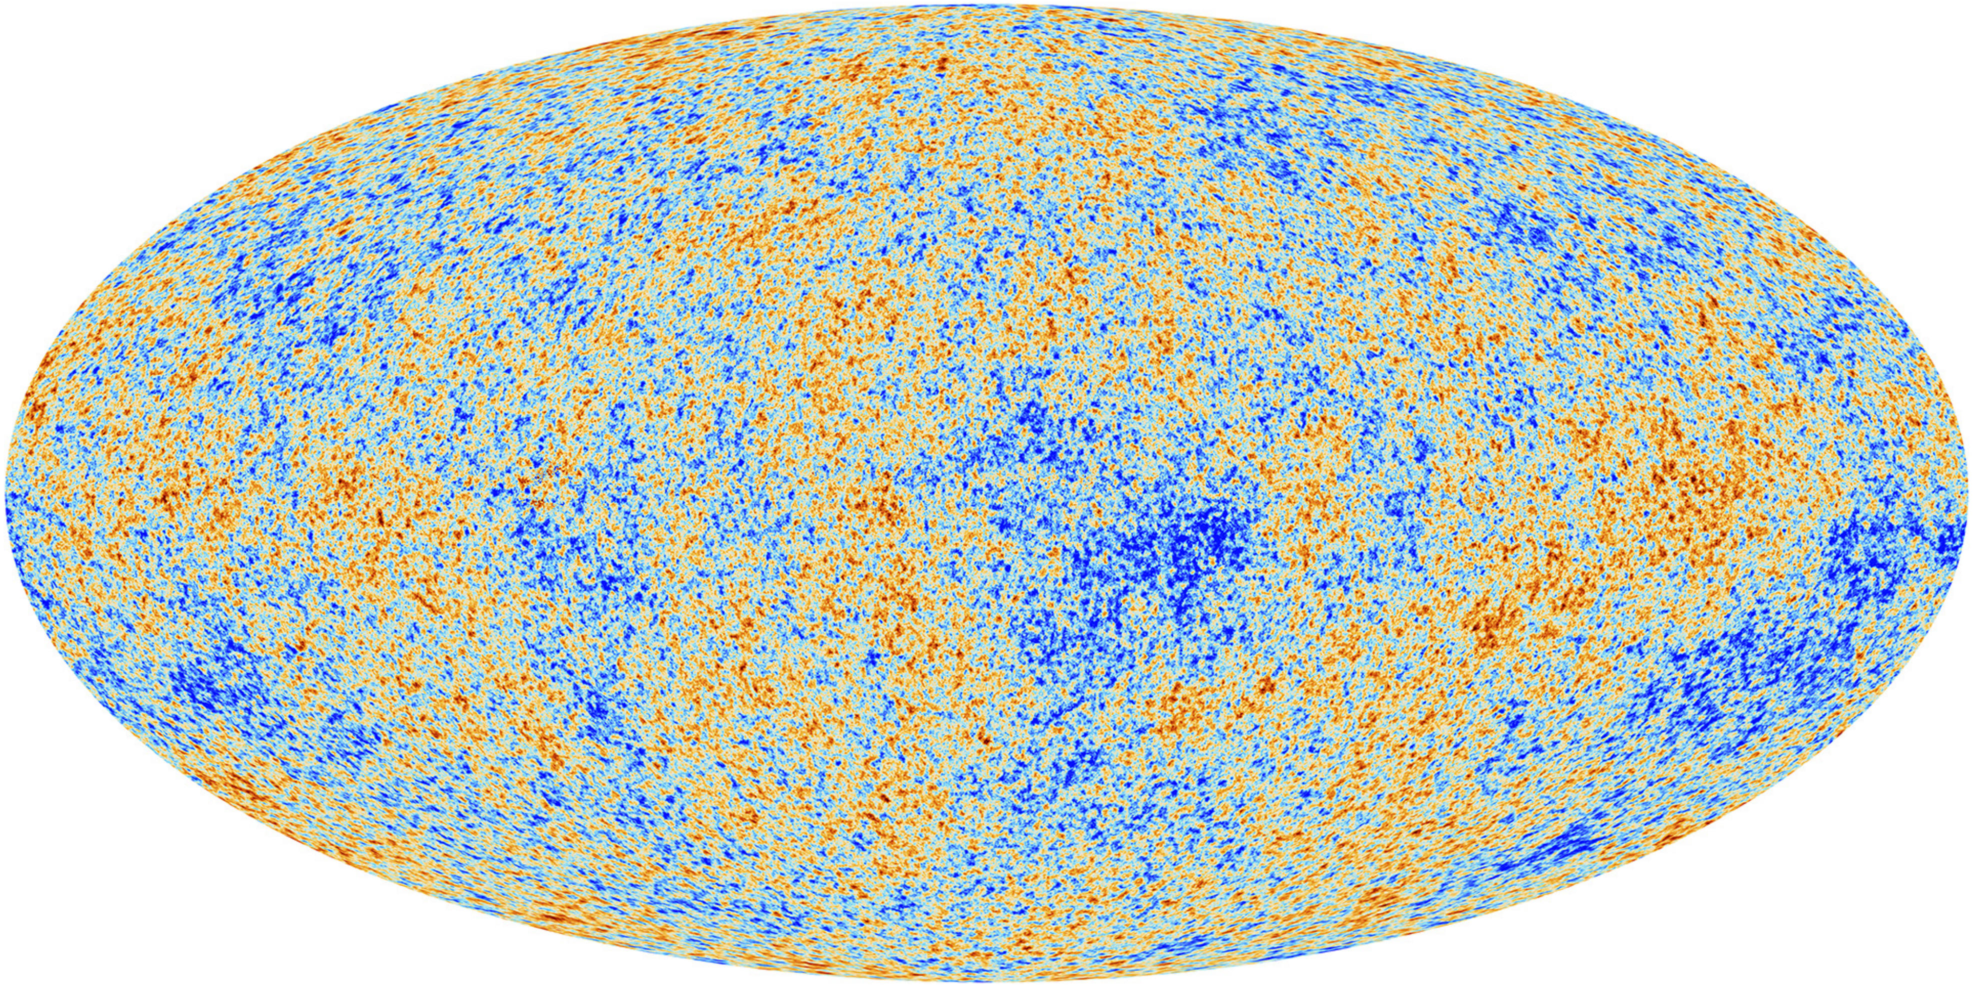

Echelle de couleur: $\pm$ 300 millionièmes de degré.

COBE 1992, WMAP 2001, Planck 2013

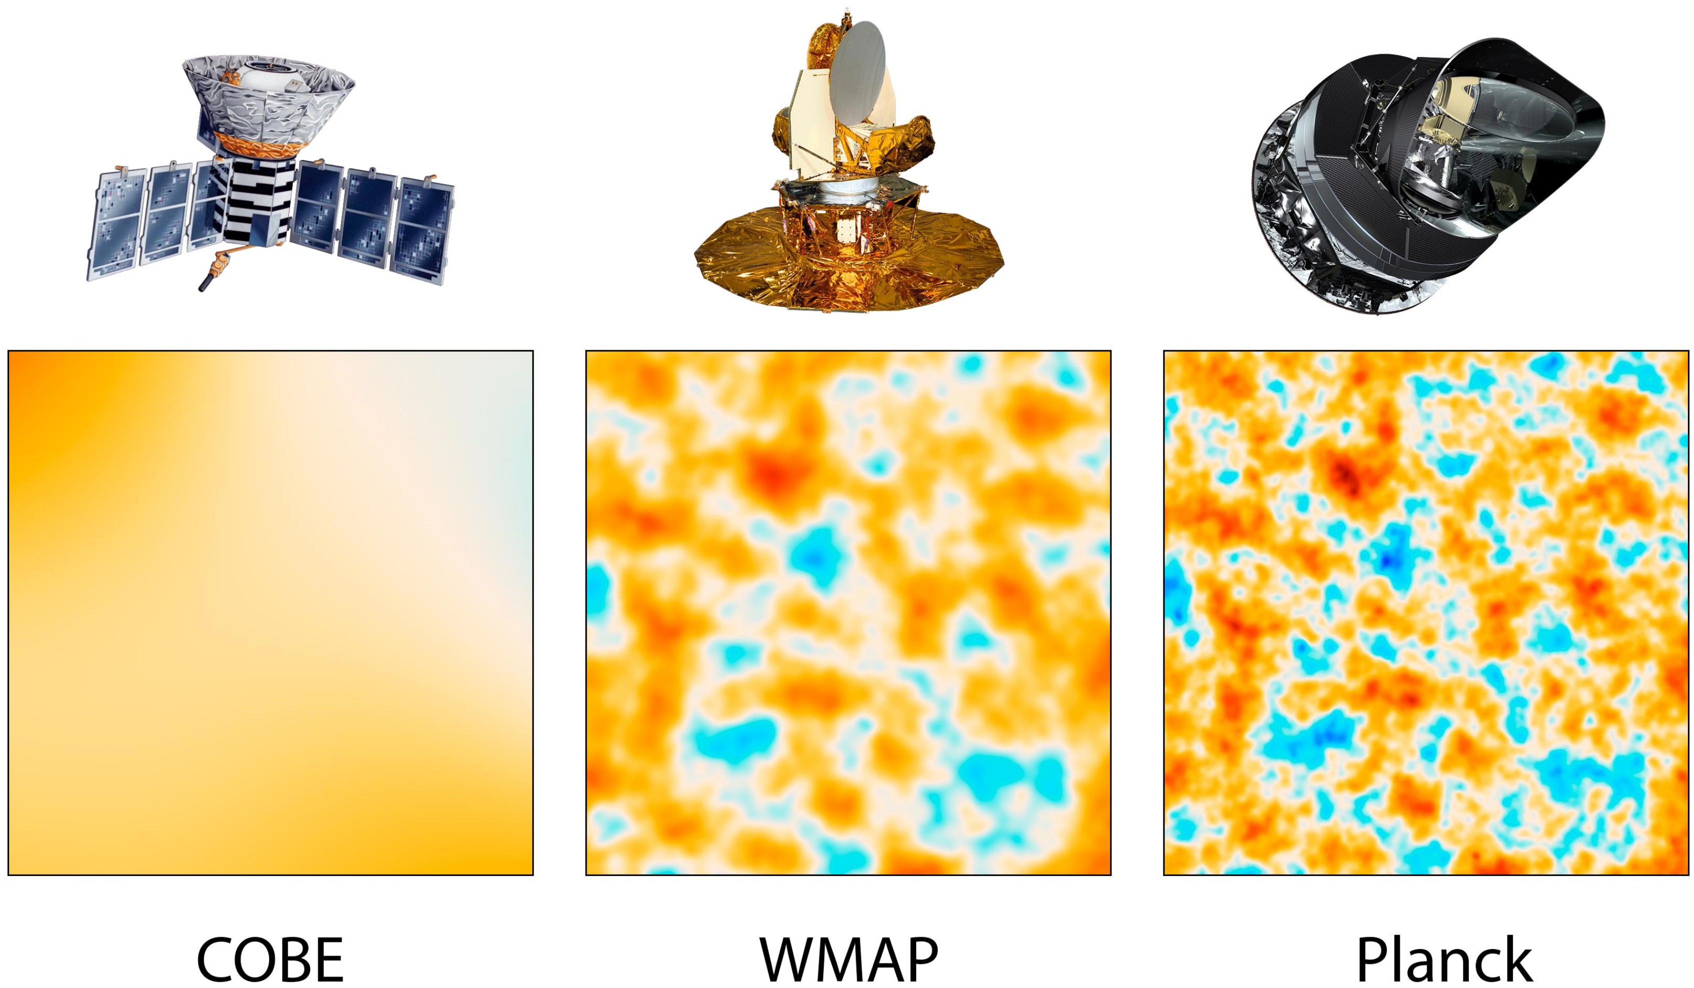

Fine, but how does one do Science with such a boring image ?

Cosmological inference on the sphere

Météo marine et analyse spectrale

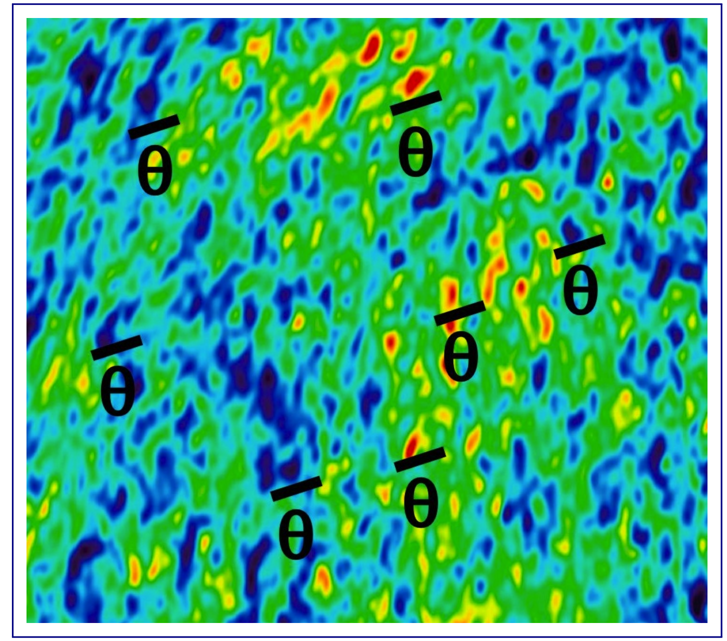

A gauche: Vent de nord-ouest, mer calme.

A droite:
Anisotropies de la température du rayonnement fossile:
taille angulaire 1 degré, amplitude: 0,0001 degré Kelvin.

Pourquoi nous voulons voir l'univers vibrer.

En route vers le quantitatif.

Toute l'information (ou presque) est contenue dans le **spectre angulaire**.

Il mesure la distribution de l'amplitude des fluctuations selon leur échelle.

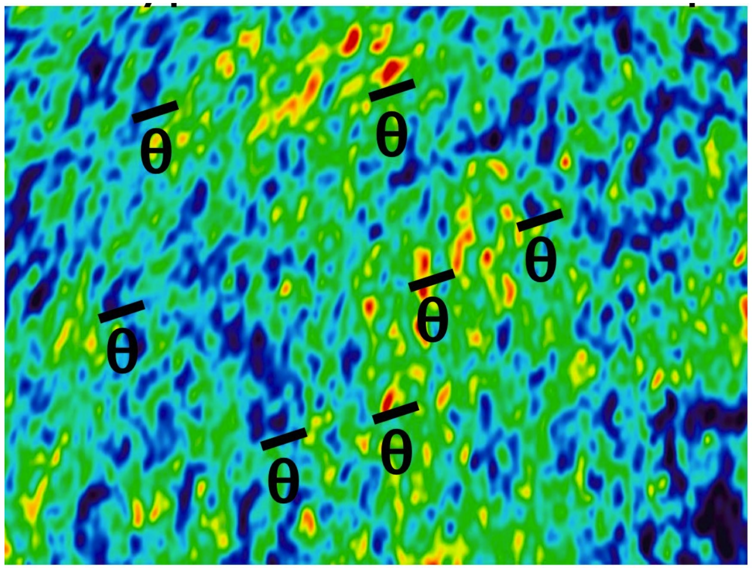

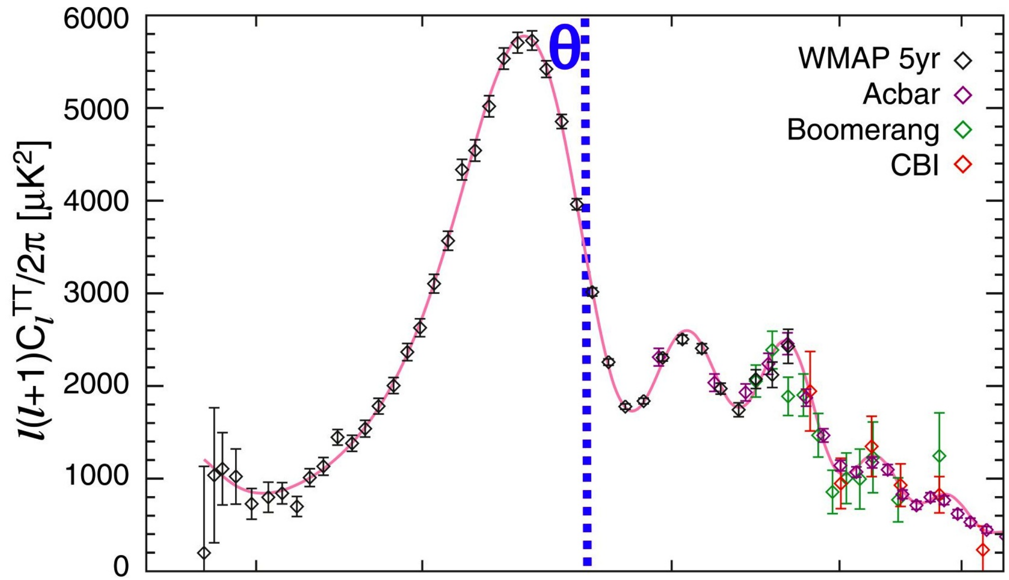

La physique du plasma primordial se lit aux échelles plus fines que 1 degré.

So, we need to go spherical, don't we ?

- No universally satisfying pixelization.
  In particular, no easy translation by almost any amount.

- Rotations and translations are the same thing.

- The rotation group for the sphere $SO(3)$ is not commutative.

- There is no convolution on the sphere (although. . .)

- There is no dilation without distortion (wavelets?).

- We need a 'good' basis for the space of functions defined on the sphere.
  Here 'good' means: with the same symmetries as the sphere itself.

- So, what is the Fourier basis?

- And what is a "spherical frequency" in the first place?

Multipoles and angukar frequencies

A spherical field $X(\theta, \phi)$ can decomposed into 'harmonic' components:

$$X(\theta, \phi) = \sum_{\ell \geq 0} X^{(\ell)}(\theta, \phi) \qquad [\theta, \phi] = [\text{(co)latitude, longitude}]$$

called monopole, dipole, quadrupole, octopole, . . . , multipole. Visually:

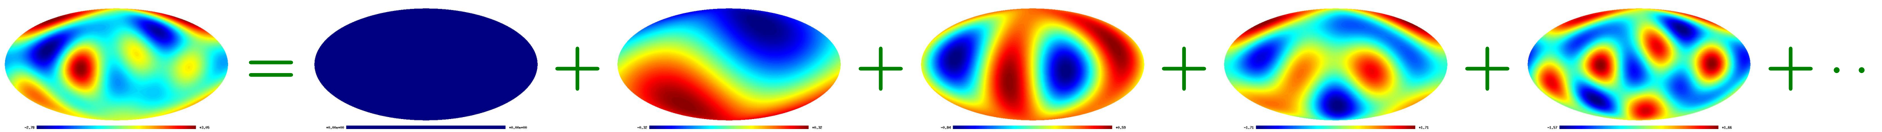

Each multipole component can be characterized as the orthogonal projection onto the ℓ-th eigenspace of the spherical Laplacian $\Delta = \partial^2/\partial x^2 + \partial^2/\partial y^2$:

$$\Delta X^{(\ell)}(\theta, \phi) = -\ell(\ell + 1)\, X^{(\ell)}(\theta, \phi)$$

Compare to the circular case: $[\partial^2/\partial\theta^2]\, e^{im\theta} = -m^2\; e^{im\theta}$.

The ℓ-th eigen-subspace of the spherical Laplacian,

- contains all functions at <u>angular frequency</u> (or multipole) ℓ,
- is globally invariant under rotation,
- has dimension $2\ell + 1$,
- has a nice orthonormal basis

Spherical harmonics

- An ortho-basis for spherical fields: the spherical harmonics $Y_{\ell,m}(\theta,\phi)$:

$$X(\theta,\phi) = \sum_{\ell\geq 0} \sum_{m=\ell}^{m=-\ell} a_{\ell,m}\; Y_{\ell,m}(\theta,\phi) \quad \longleftrightarrow \quad a_{\ell,m} \;= \int_\theta \int_\phi Y_{\ell,m}(\theta,\phi)\; X(\theta,\phi)$$

- Better than simply orthogonal: it factors as $Y_{\ell,m}(\theta,\phi) \;=\; e^{im\phi}\; P_{\ell,m}(\cos\theta)$.

- So we have semi-fast spherical harmonic transforms. . .

. . . for good sampling schemes.

- Think of the double index (ℓ, m) as a spherical version of $\left(\sqrt{k_x^2 + k_y^2},\; k_y\right)$.

- The decomposition into multipoles:

$$X^{(\ell)}(\theta,\phi) = \sum_{m=\ell}^{m=-\ell} a_{\ell,m}\; Y_{\ell,m}(\theta,\phi)$$

Angular spectrum and likelihood

- The spherical harmonic coefficients $a_{\ell,m}$ of a stationary random field are uncorrelated with variance C_ℓ called "the angular power spectrum" :

$$\mathsf{E}\Big(a_{\ell,m}\ a_{\ell' m'}\Big) = C_\ell\ \delta_{\ell\ell'}\ \delta_{mm'}$$ distribution of power vs angular frequency ℓ.

- Thus, for a Gaussian field, the empirical spectrum

$$\widehat{C}_\ell = \frac{1}{2\ell+1} \sum_{m=\ell}^{m=-\ell} a_{\ell,m}^2 \qquad \mathsf{Var}\Big(\widehat{C}_\ell / C_\ell\Big) = \frac{2}{2\ell+1}$$ (Cosmic variance)

is a sufficient statistic, compressing all information since the likelihood reads:

$$-2\log P(X|\{C_\ell\}) \ = \ \sum_{\ell\geq 0} (2\ell+1) \left(\frac{\widehat{C}_\ell}{C_\ell} + \log C_\ell\right) + \mathsf{cst}$$ **Super compression!**

- It is also a spectral mismatch:

$$-2\log P(X|\{C_\ell\}) \ = \ \sum_{\ell\geq 0} (2\ell+1)\ k(\widehat{C}_\ell / C_\ell) + \mathsf{cst'} \qquad k(u) \stackrel{\mathsf{def}}{=} u - \log u - 1$$

Angular spectrum and likelihood

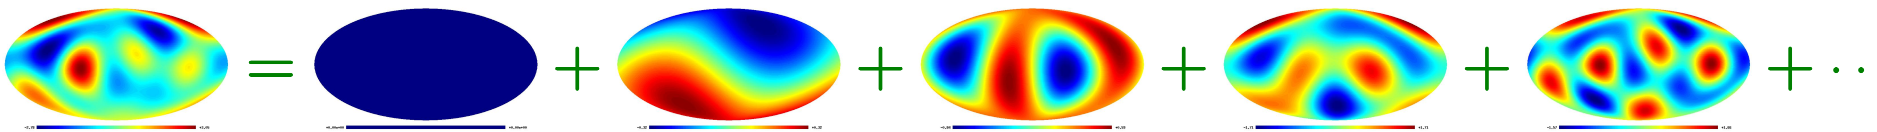

$$X(\theta,\phi) = X^{(0)}(\theta,\phi) + X^{(1)}(\theta,\phi) + X^{(2)}(\theta,\phi) + X^{(3)}(\theta,\phi) + \cdots$$

The empirical angular spectrum $\quad \widehat{C}_\ell \stackrel{\text{def}}{=} \|X^{(\ell)}\|^2/(2\ell+1)$

The "true" angular spectrum $\quad C_\ell \stackrel{\text{def}}{=} E\,\widehat{C}_\ell$

We have only one sky, one observable Universe, one set of $\widehat{C}_\ell$.

The "true" angular spectrum C_ℓ exists only in the theory.

The likelihood of our Universe

- Cosmologists build physical models predicting a Gaussian stationary CMB sky with an angular spectrum depending on fundamental cosmologic parameters:

$$C_\ell = C_\ell(\alpha) \qquad\qquad \alpha = (\Omega_\Lambda, \Omega_m, H_0, \ldots)$$

- Instrumentalists painfully measures the angular spectrum $\widehat{C}_\ell$ of the CMB sky.

- Statisticians know that $\widehat{C}_\ell$ then is a sufficient statistic: the likelihood reads:

$$-2 \log p(\mathrm{CMB}|\alpha) = \sum_{\ell \geq 0} (2\ell + 1) \left(\frac{\widehat{C}_\ell}{C_\ell(\alpha)} + \log C_\ell(\alpha) \right) + \mathrm{cst.}$$

- In real life, things (the likelihood, the estimate $\widehat{C}_\ell$) are much more complicated, but we still match a model spectrum to an empirical spectrum.

Theoretical angular spectrum of the CMB

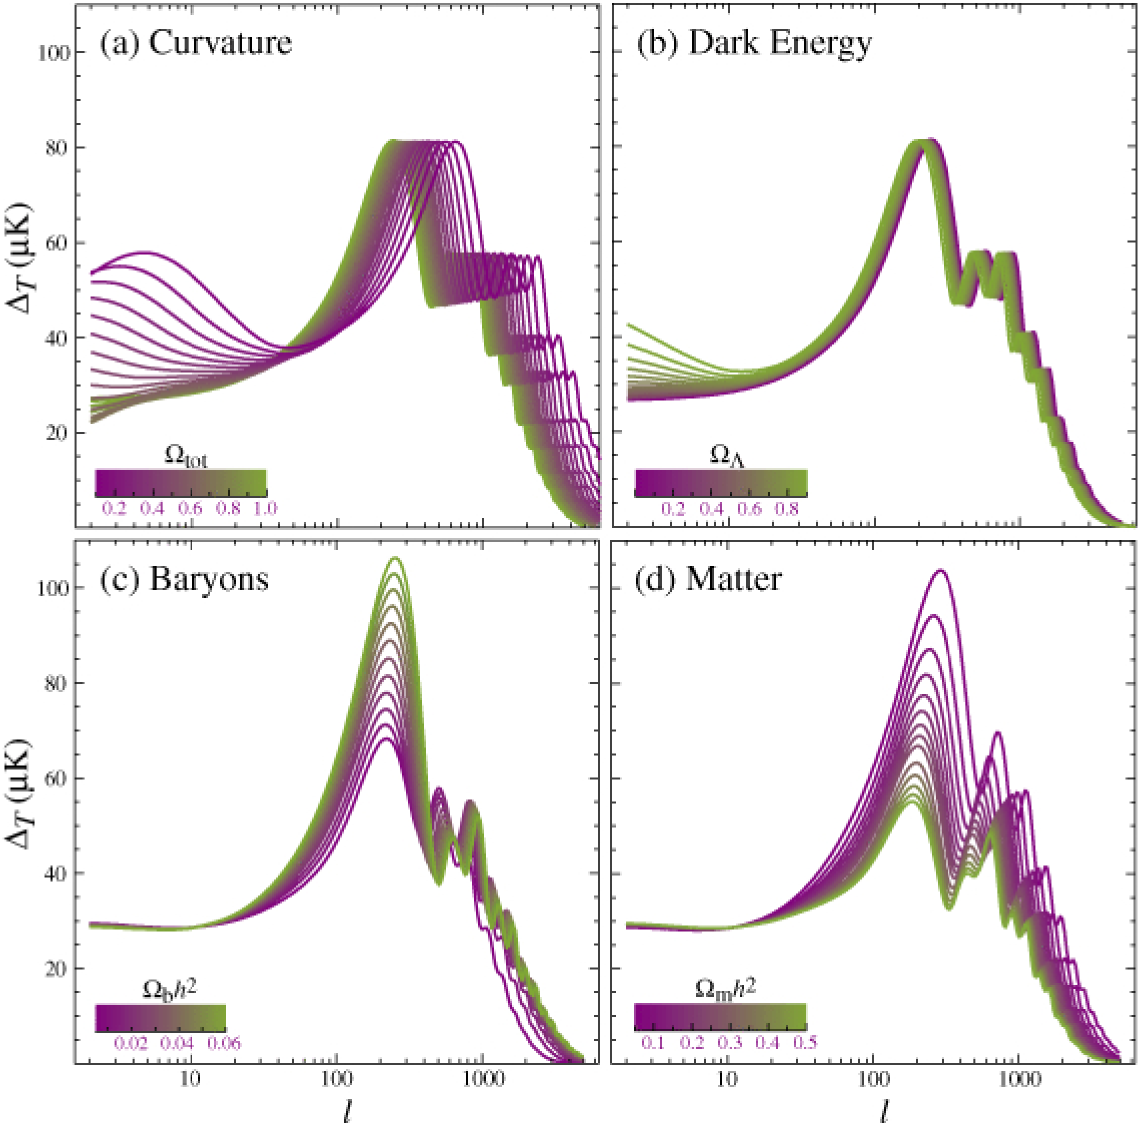

A cosmological model has to predict the angular spectrum of the CMB as a function of "cosmological parameters".

Left: examples of the dependance of the spectrum on some parameters of the $\Lambda - \text{CDM}$ model.

Important note : we plot $\ell^2 C_\ell$.
Large scales dominate the power.

W-MAP 5-year angular spectrum.

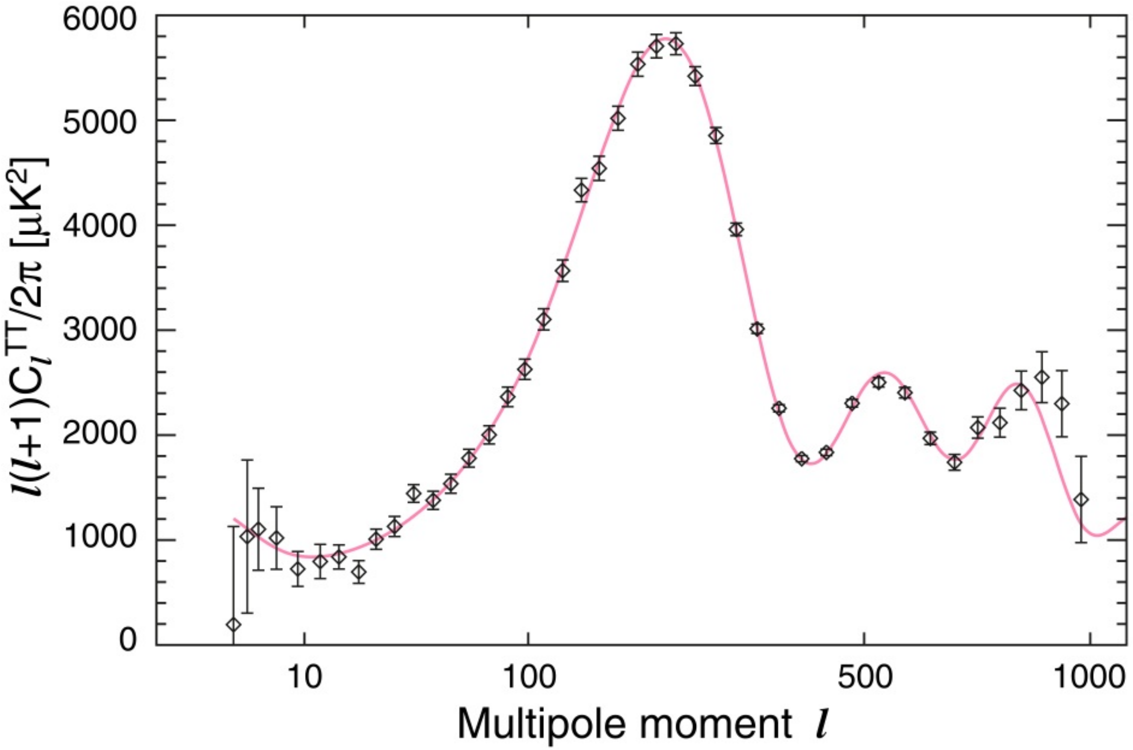

Plot of rescaled spectrum:

$$D(\ell) = C(\ell) \times \frac{\ell(\ell+1)}{2\pi}$$

Black ink: measures.

Red ink: best-fit theory spectrum.

Three acoustic peaks: congrats, W-MAP!

Note the puzzingly low quadrupole (harmonic frequency $\ell = 2$)

Cosmic variance dominates for $\ell \leq 540$;
instrumental noise dominates at higher multipoles.

Component separation

Extracting the CMB from the 9 Planck frequency channels

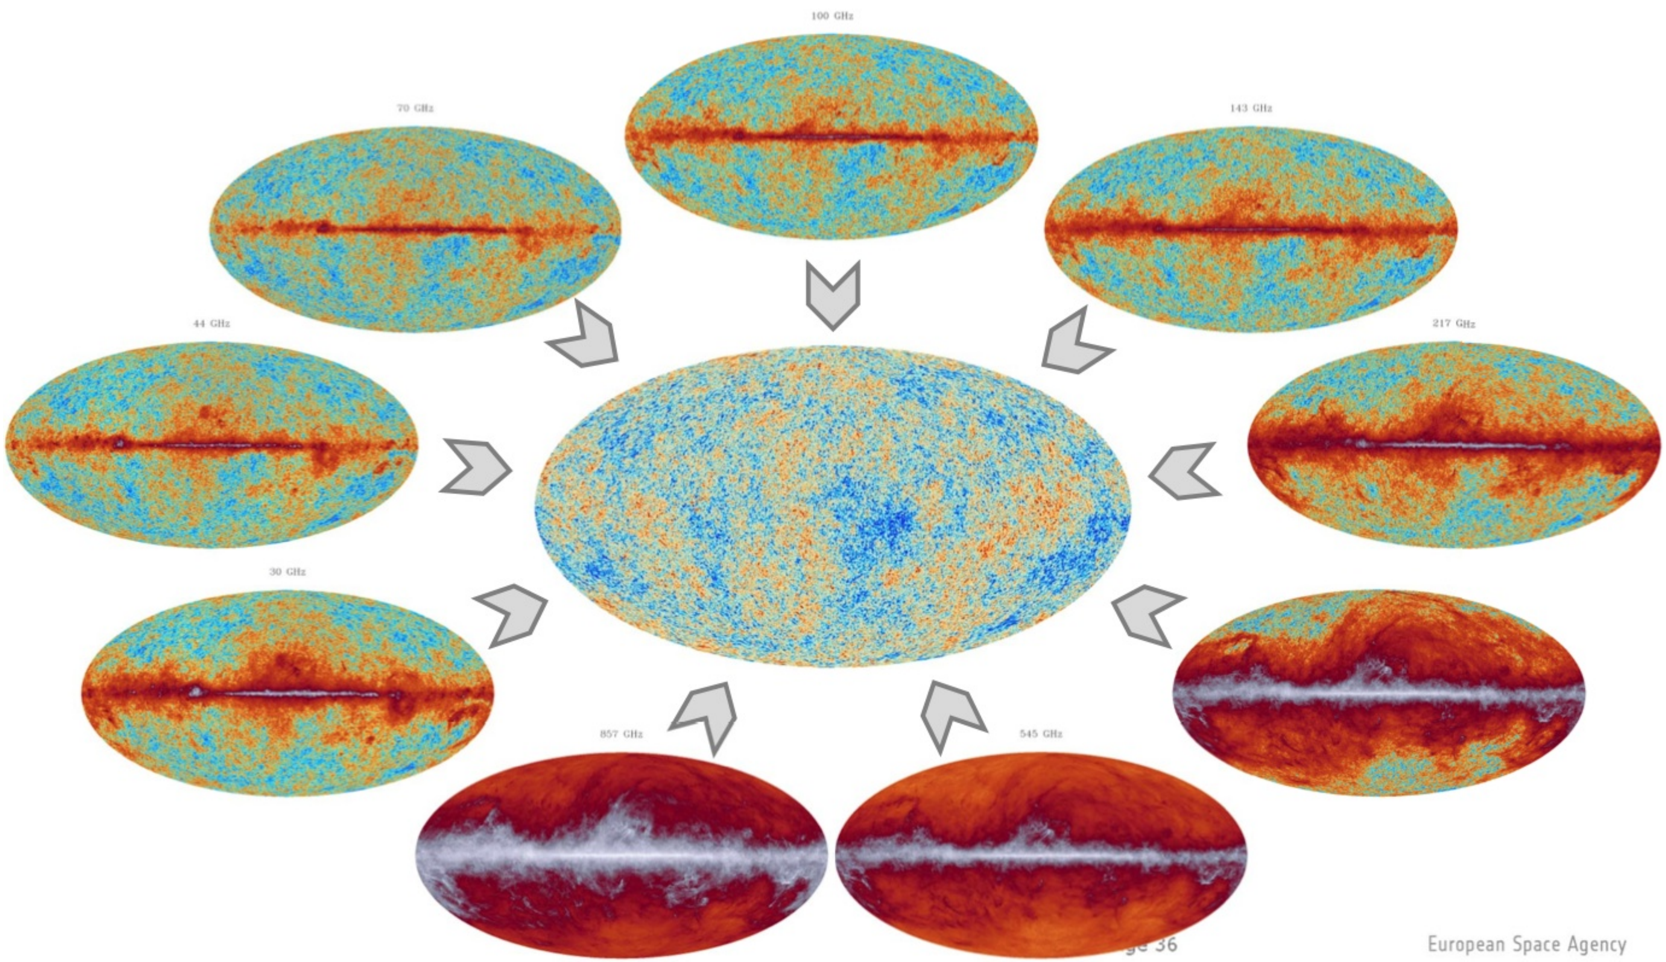

Color scale: hundreds of micro-Kelvins.

Credits: ESA, FRB.

Four CMB anisotropy maps delivered to the Planck Legacy Archive

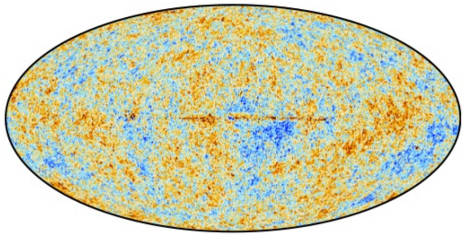
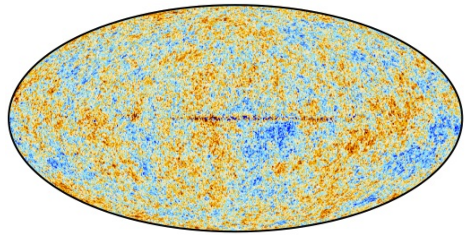
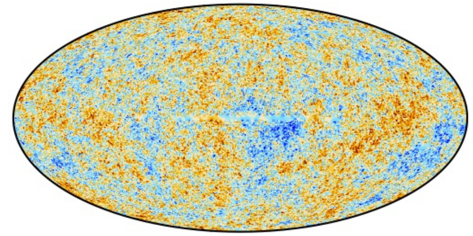
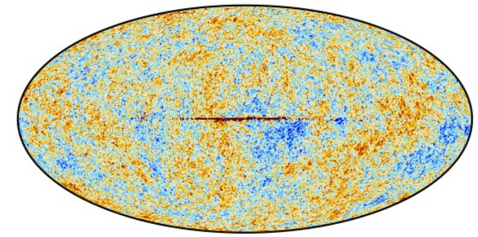

NILC	SEVEM	SMICA	C-R
$\ell_{\mathrm{SNR}=1} = 1790$	$\ell_{\mathrm{SNR}=1} = 1790$	$\ell_{\mathrm{SNR}=1} = 1790$	$\ell_{\mathrm{SNR}=1} = 1550$
Wavelet space	Wawelet-like	Harmonic space	Pixel space
non-parametric	non-parametric	semi-parametric	parametric

- Various assumptions about the foregrounds.

- Various filtering schemes (space-dependent, multipole-dependent, or both).

- The SMICA (Spectral Mismatch ICA) method selected for
the 'Main product' for the Planck CMB map.

Some requirements for producing a CMB map

- The method should be accurate and high SNR (obviously).

- The method should be linear in the data:

  1. It is critical not to introduce non Gaussianity

  2. Propagation of simulated individual inputs should be straightforward

- The result should be easily described (e.g. map=beam*sky+noise) with a well defined transfer function.

- The method should be fast enough for thousands of Monte-Carlo runs.

- The method should be able to support wide dynamical ranges, over the sky, over angular frequencies, across channel frequencies.

Wide dynamics over the sky

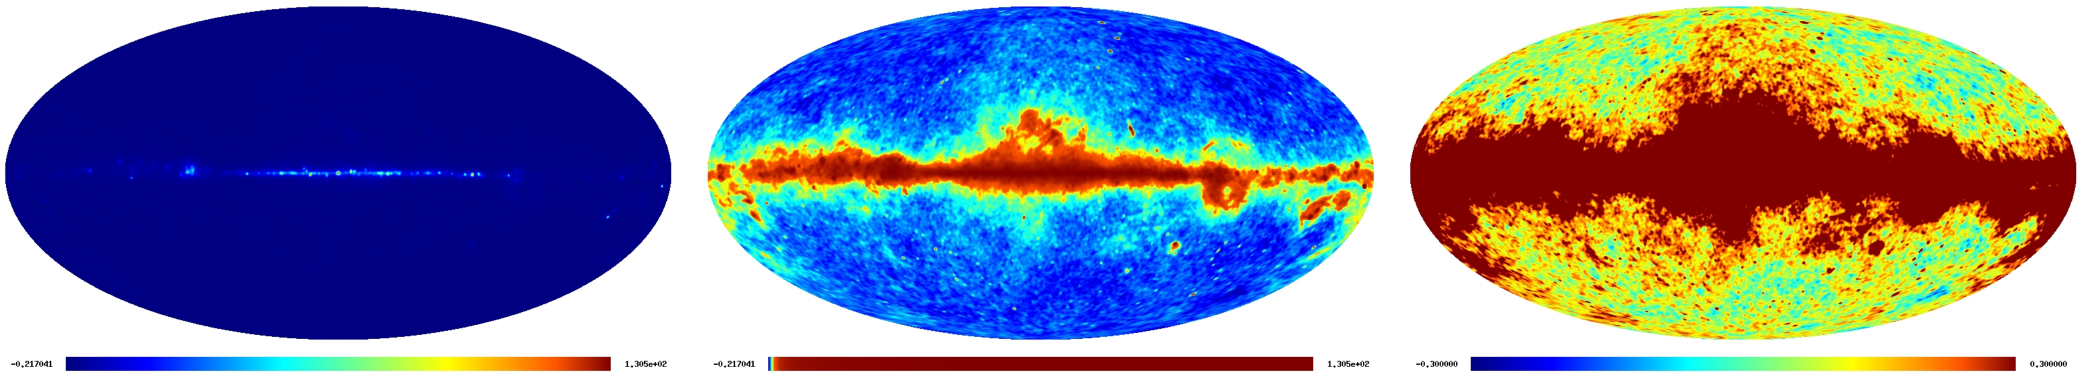

Left: The W-MAP K band. Natural color scale [-200, 130000] μK.
Middle: Same map with an equalized color scale.
Right: Same map with a color scale adapted to CMB: [-300, 300] μK.

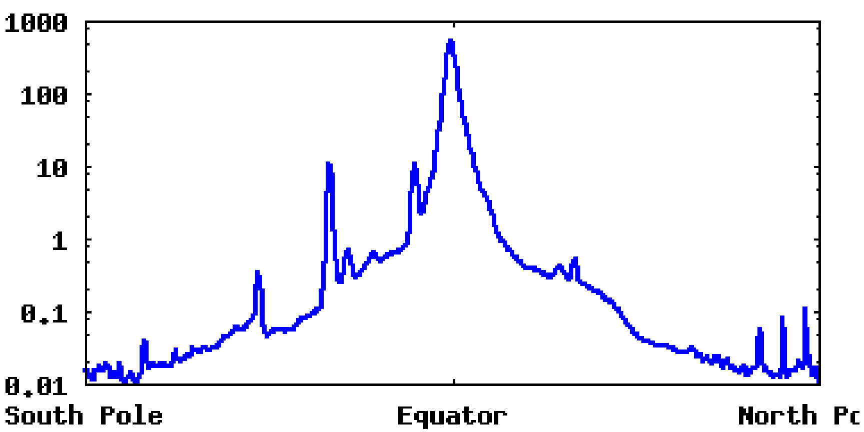

Average power as a function of latitude on a log scale for the same map.

Do we really want to estimate covariance matrices over the whole sky?

Wide spectral dynamics, SNR variations

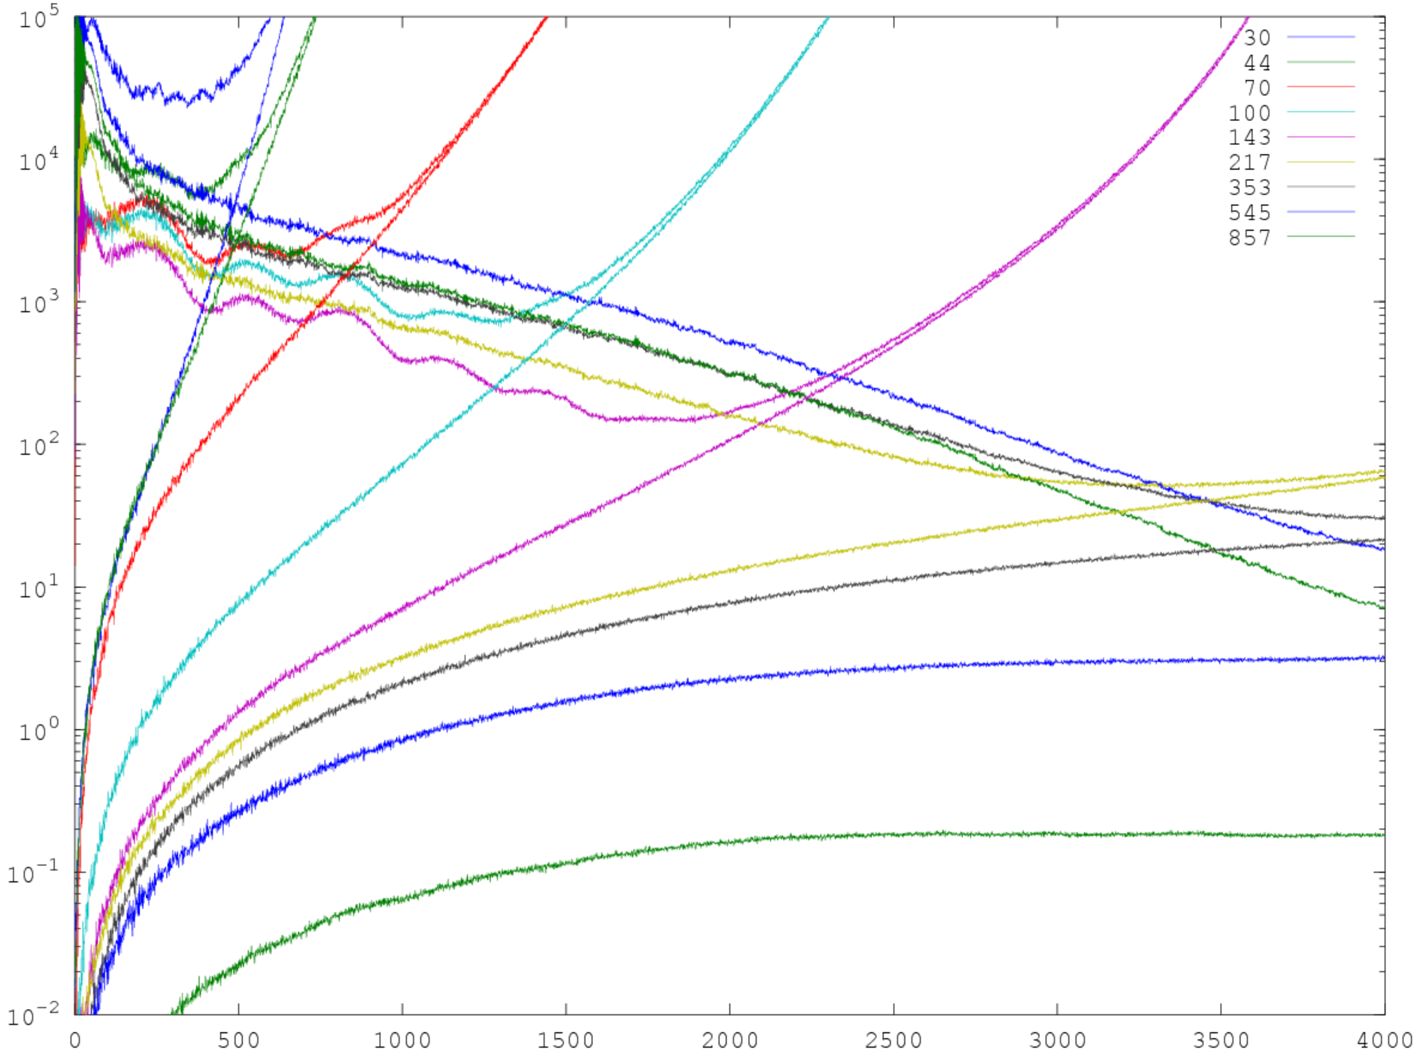

$\widehat{C}(\ell) \cdot \ell(\ell+1)/2\pi$ in $[\mu K_{\mathrm{RJ}}]^2$ for $f_{\mathrm{sky}} = 0.99$.

And what about the noise ?

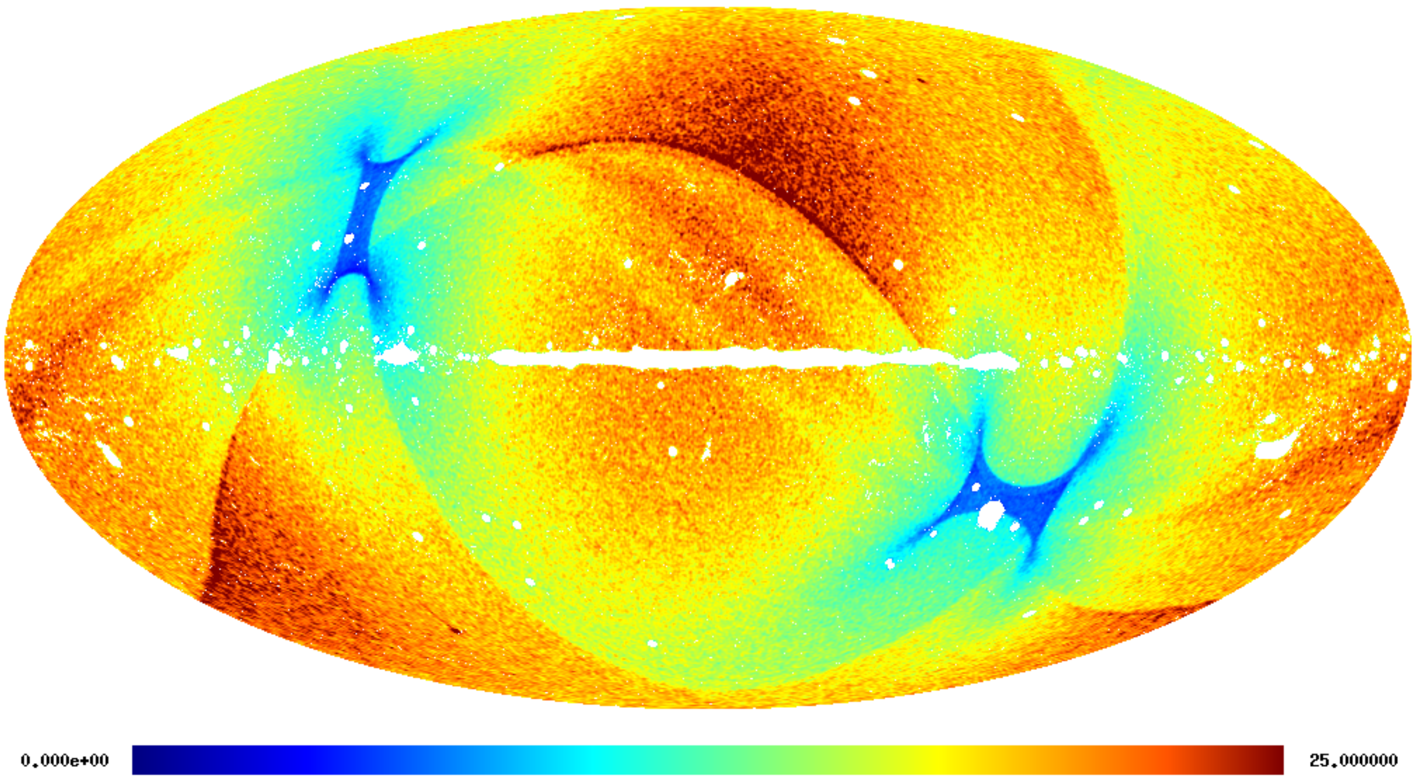

Noise RMS in μK in the SMICA map.

Simple CMB cleaning by "template removal"

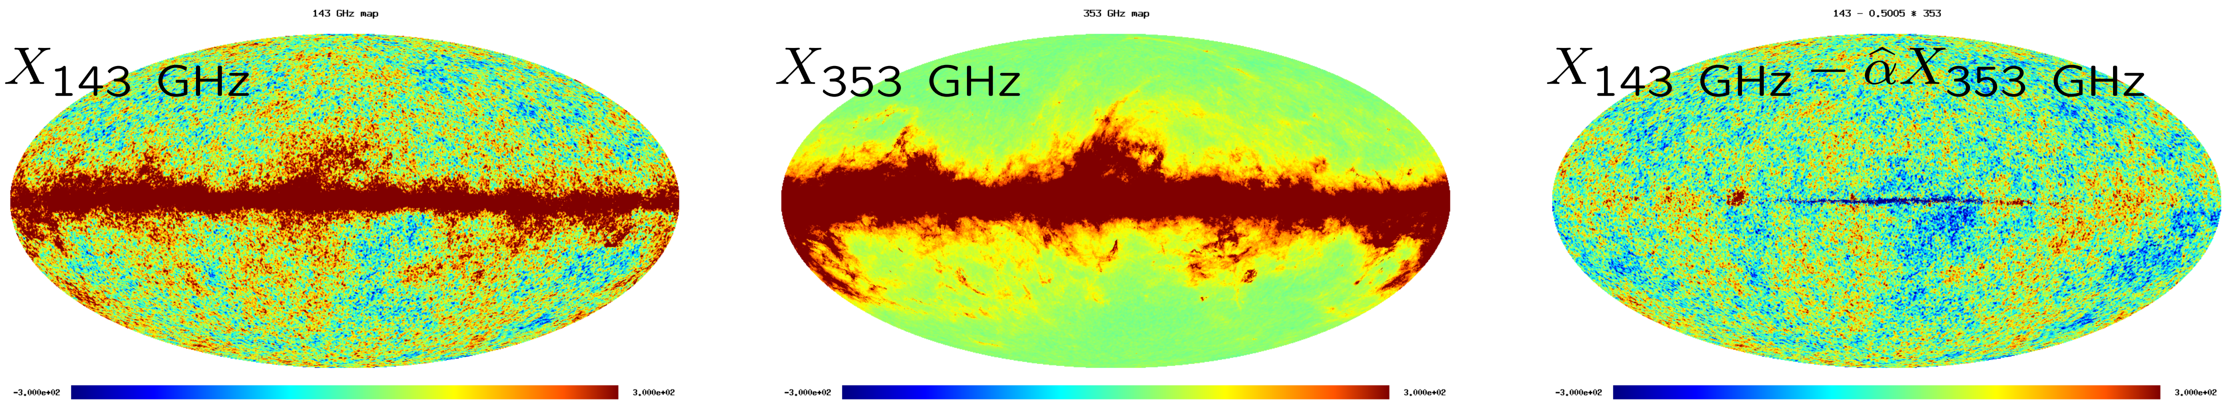

Assume that the 353 GHz channel sees only dust emission
and that the 143 GHz channel sees CMB plus a rescaled dust pattern:

$$X_{143} = \mathsf{CMB} + \alpha\ X_{353}$$

Find α by correlation: $\hat{\alpha} = \langle X_{143} X_{353} \rangle / \langle X_{353} X_{353} \rangle$ and a clean (?) CMB map as

$$\widehat{\mathsf{CMB}} = X_{143} - \frac{\langle X_{143} X_{353} \rangle}{\langle X_{353} X_{353} \rangle}\ X_{353}$$ where $\langle \cdot \rangle$ denotes a pixel average

The result (top right) does not look bad, but it is !

Note: By construction $\langle \widehat{\mathsf{CMB}}\ X_{353} \rangle = 0$.

Combining all 9 Planck channels, non parametrically: the ILC

Stack the 9 Planck channels into a data 9×1 vector $\mathbf{d} = [d_{30}, d_{44}, \ldots, d_{545}, d_{857}]^\dagger$ and estimate the CMB signal $s(p)$ in pixel p by weighting the inputs:

$$\widehat{s}(p) = \mathbf{w}^\dagger \mathbf{d}(p) \qquad p = 1, \ldots, N_{\mathsf{pix}}$$

At frequency ν, the CMB signal $s(p)$ has amplitude a_ν and contaminated by $f_\nu(p)$

$$d_\nu(p) = a_\nu \, s(p) \; + \; f_\nu(p) \qquad \text{or} \qquad \mathbf{d}(p) = \mathbf{a} \, s(p) + \mathbf{f}(p)$$

The best ($\mathsf{min}_w \langle (s - \mathbf{w}^\dagger \mathbf{d})^2 \rangle_p$) unbiased ($\mathbf{w}^\dagger \mathbf{a} = 1$) estimator is:

$$\mathbf{w} = \frac{\widehat{\mathbf{C}}^{-1} \, \mathbf{a}}{\mathbf{a}^\dagger \widehat{\mathbf{C}}^{-1} \, \mathbf{a}} \qquad \text{with} \qquad \widehat{\mathbf{C}} = \langle \mathbf{d}\mathbf{d}^\dagger \rangle_p, \quad \text{the sample covariance matrix}$$

That is known as ILC (Internal Linear Combination) in CMB circles, as MVBF (Minimum Variance Beam Former) in array processing, otherwise elsewhere.

Looks good: linear, unbiased, minimum MSE, very blind, very few assumptions: knowing $\mathbf{a}$ (calibration) and the CMB uncorrelated from the rest (very true).

Is the ILC good enough for Planck data ?

A simulation result:

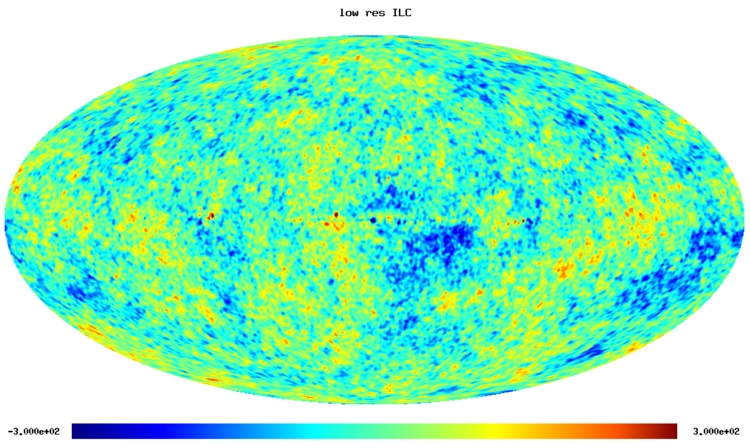

$\longleftarrow$ ILC map on a $\pm 300\mu K$ color scale

Error on a $\pm 50\mu K$ color scale $\longrightarrow$

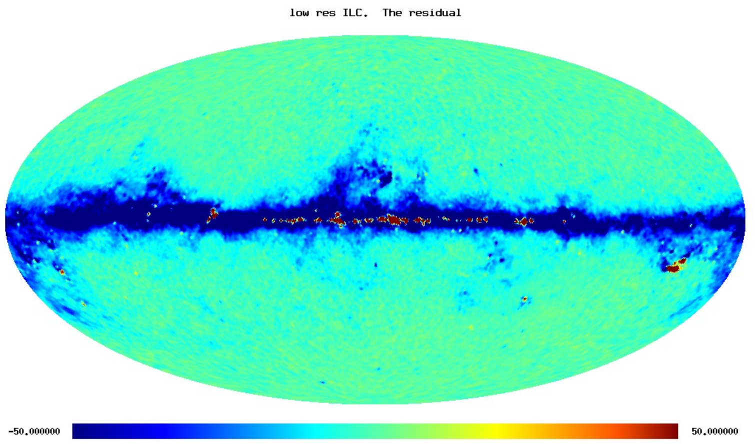

ILC looked promising, but something went wrong.

Actually two things, at least, need fixing:

- harmonic dependence and
- chance correlations.

SMICA: Linear filtering in harmonic space

Since resolution, noise and foregrounds vary (wildly) in power over channels and angular frequency, the combining weights should depend on ℓ.

The SMICA CMB map is synthesized from spherical harmonic coefficients $\widehat{s}_{\ell,m}$, obtained as linear combinations:

$$\widehat{s}_{\ell,m} = \mathbf{w}_\ell^\dagger \, \mathbf{d}_{\ell,m} \quad \text{with, again,}$$

$$\mathbf{w}_\ell = \frac{\mathbf{C}_\ell^{-1}\mathbf{a}}{\mathbf{a}^\dagger \mathbf{C}_\ell^{-1}\mathbf{a}} \qquad \mathbf{C}_\ell = \mathrm{Cov}(\mathbf{d}_{\ell,m})$$

- At high ℓ, the (spectral) covariance matrices $\mathbf{C}_\ell$ are well estimated by their sample counterparts
- At lower ℓ, we need to get smarter.

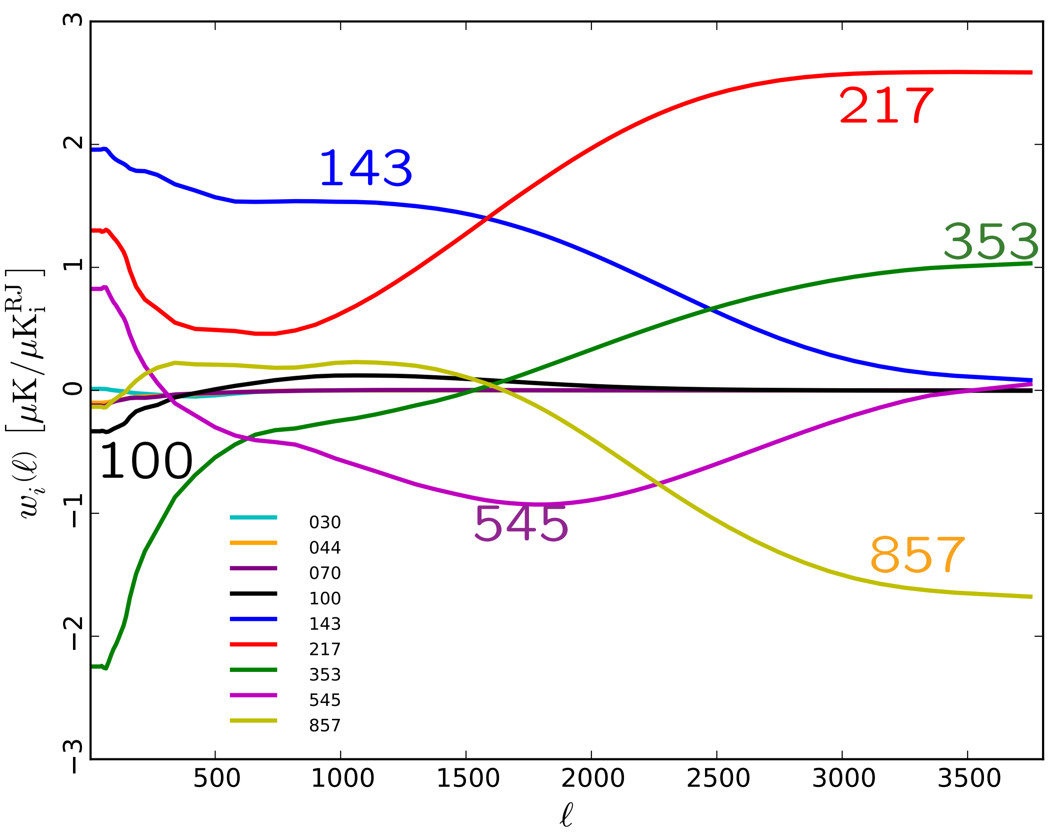

Note: spectral localization is a must. Spatial localization do not seem critical (See NILC perf.).

Foregrounds and how to get rid of them (at low ℓ) ?

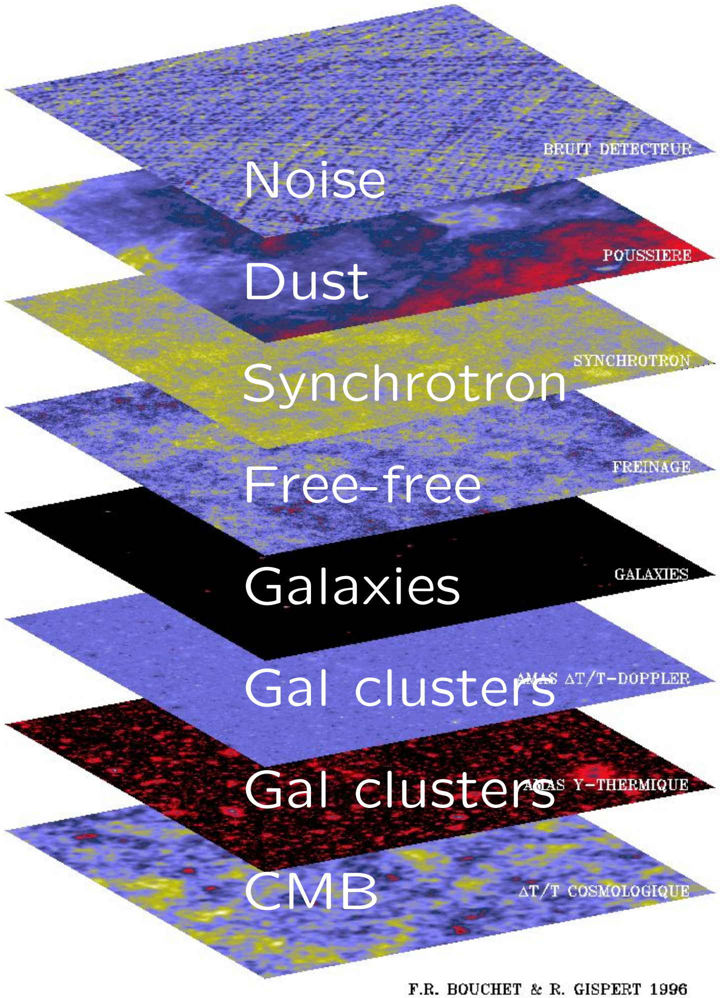

Various **foreground** emissions (both galactic and extra-galactic) pile up in front of the CMB.

But they do so additively !

Even better, most scale rigidly with frequency: each frequency channel sees a different mixture of each astrophysical emission:

$$\mathbf{d} = \mathbf{A}\mathbf{s} + \mathbf{n} \quad \text{(data = mixture x sources + noise)}$$

Such a linear mixture can be inverted . . . if the mixing matrix $\mathbf{A}$ is known. How to find it or do without it ?

1 Trust astrophysics and use parametric models, or

2 Trust your data and the power of statistics.

Foregrounds, physical components and the mixing matrix

At low ℓ (i.e. large angular scales), there are less Fourier modes available for estimating spectral statistics $\widehat{\mathbf{C}}_\ell$: the variability (chance correlations) must be decreased by fitting a model $\mathbf{C}_\ell(\theta)$ to $\widehat{\mathbf{C}}_\ell$.

- Mixing matrix. The 9 Planck channels as noisy linear mixtures of components:

$$\mathbf{d}_{\ell,m} = \mathbf{A}(\theta)\ \mathbf{s}_{\ell,m} \ + \ \mathbf{n}_{\ell,m}$$

- Some models for the mixing matrix $\mathbf{A} = \mathbf{A}(\theta)$:

Type	Mixing matrix	parameters θ
physical, fixed	$\mathbf{A} = [\mathbf{a}_{\text{cmb}}\ \mathbf{a}_{\text{dust}}\ \mathbf{a}_{\text{CO}}\ \mathbf{a}_{\text{LF}}]$	$\theta = [\ ]$
physical, parametric	$\mathbf{A} = [\mathbf{a}_{\text{cmb}}\ \mathbf{a}_{\text{dust}}(T)\ \mathbf{a}_{\text{CO}}\ \mathbf{a}_{\text{LF}}(\beta)\ ]$	$\theta = (T, \beta)$
non-parametric ($\sim$ ILC)	$\mathbf{A} = [\mathbf{a}_{\text{cmb}}\ \mathbf{B}]$ (a square matrix)	$\theta = \mathbf{B}$
semi-parametric, SMICA	$\mathbf{A} = \mathbf{A}$ (any tall matrix)	$\theta = \mathbf{A}$

- Which model, which fitting criterion? See next the SMICA case.

SMICA semi-parametric model

- SMICA models the 9 Planck channels as noisy linear mixtures of CMB and 6 "foregrounds":

$$\begin{bmatrix} d_1 \\ d_2 \\ \vdots \\ \vdots \\ d_9 \end{bmatrix} = \begin{bmatrix} a_1 & F_{11} & \dots & F_{16} \\ a_2 & F_{21} & \dots & F_{26} \\ \vdots & \vdots & \dots & \vdots \\ \vdots & \vdots & \dots & \vdots \\ a_9 & F_{91} & \dots & F_{96} \end{bmatrix} \times \begin{bmatrix} s \\ f_1 \\ \vdots \\ f_6 \end{bmatrix} + \begin{bmatrix} n_1 \\ n_2 \\ \vdots \\ \vdots \\ n_9 \end{bmatrix} \quad \text{or} \quad \mathbf{d}_{\ell,m} = [\,\mathbf{a} \mid \mathbf{F}\,] \begin{bmatrix} s_{\ell,m} \\ \mathbf{f}_{\ell,m} \end{bmatrix} + \mathbf{n}_{\ell,m}$$

- SMICA only assumes <u>decorrelation</u> between foregrounds and CMB.

The foregrounds must have 6 (say) dimensions but are otherwise completely unconstrained: they may have any spectrum, any color, any correlation...

So the data model is very blind: all non-zero parameters are <u>free</u> !

$$\mathrm{Cov}(\mathbf{d}_{\ell,m}) = [\,\mathbf{a} \mid \mathbf{F}\,] \begin{bmatrix} C_\ell^{\mathrm{cmb}} & 0 \\ 0 & \mathbf{P}_\ell \end{bmatrix} [\,\mathbf{a} \mid \mathbf{F}\,]^\dagger + \begin{bmatrix} \sigma_{1\ell}^2 & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & \sigma_{9\ell}^2 \end{bmatrix} = \mathbf{C}_\ell(\mathbf{a}, C_\ell^{\mathrm{cmb}}, \mathbf{F}, \mathbf{P}_\ell, \sigma_{i\ell}^2).$$

- Blind identifiability: can it be done? Maths say: yes!

- Fit by $\min_\theta \sum_\ell (2\ell+1) \left[\mathrm{trace}\ \widehat{\mathbf{C}}_\ell \mathbf{C}_\ell(\theta)^{-1} + \log\det \mathbf{C}_\ell(\theta)\right]$ = Gaussian stationary likelihood.

- Only $\mathrm{Span}(\mathbf{A})$, the foreground subspace, is needed to suppress the foregrounds.

It is collectively determined by all the multipoles involved in the fit.

Why the stationary Gaussian likelihood is OK (and sparsity useless)

Consider the noise-free, square case : $\mathbf{d} = [\,\mathbf{a} \mid \mathbf{F}\,] \begin{bmatrix} s \\ \mathbf{f} \end{bmatrix}$ with known $\mathbf{a}$.

For any matrix $\mathbf{G}$, the 'preprocessor' $\mathbf{T} = [\,\mathbf{a} \mid \mathbf{G}\,]^{-1}$ ensures

$$\mathbf{TA} = [\,\mathbf{a} \mid \mathbf{G}\,]^{-1}\,[\,\mathbf{a} \mid \mathbf{F}\,] = \begin{bmatrix} 1 & \alpha^\dagger \\ 0 & \mathbf{Y} \end{bmatrix} \quad \text{so that} \quad \widetilde{\mathbf{d}} = \mathbf{Td} = \begin{bmatrix} s + \alpha^\dagger \mathbf{f} \\ \mathbf{Y}\,\mathbf{f} \end{bmatrix}$$

Hence, the first pre-processed channel $\widetilde{\mathbf{d}}_1$ contains the signal of interest s contaminated by a linear combination of the other observed channels $\widetilde{\mathbf{d}}_2, \widetilde{\mathbf{d}}_3, \cdots$.

• Statistical foreground models **not** needed: they are deterministically observed.

• Hence, a statistical model is needed only for the CMB. Since it is Gaussian and stationary, the likelihood has a simple expression which is readily maximized in the cleaning coefficients. That justifies the use of a Gaussian stationary likelihood for fitting the SMICA model.

More explicitly

Assuming without loss of generality that $\mathbf{Y} = \mathbf{Id}$ and denoting $d = \tilde{\mathbf{d}}_1$, the model is, in each pixel p

$$d(p) \;=\; s(p) + \alpha^\dagger \mathbf{f}(p)$$

with $\mathbf{f}(p)$ observed. For the best cleaning, we need the optimal estimate of α.

But the likelihood is trivial in harmonic space, with the CMB as 'noise' (!)

$$-2 \log P(d|\alpha) = \sum_\ell (2\ell + 1) \; \frac{(d_{\ell,m} - \alpha^\dagger \mathbf{f}_{\ell,m})^2}{C_\ell} + \mathrm{cst}$$

The (trivial) solution corresponds to combining the inputs with weight vector

$$\mathbf{w} = \frac{\widehat{\mathbf{C}}^{-1}\,\mathbf{a}}{\mathbf{a}^\dagger \widehat{\mathbf{C}}^{-1}\,\mathbf{a}} \qquad \text{i.e. an ILC with} \qquad \widehat{\mathbf{C}} = \sum_\ell \sum_m \mathbf{d}_{\ell,m} \mathbf{d}^\dagger_{\ell,m} / C_\ell$$

This is **also** the SMICA solution in the same context:
chance correlation is optimally mitigated in the spectral domain.

Results

The big pipeline picture, from time series to cosmology, ideally

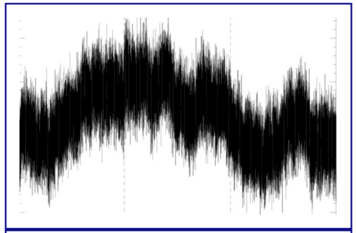

$\leftarrow$ SCANNING THE SKY.
Make sure you capture those μKs in your time lines.
Deglitch, flag, deconvolve, calibrate...

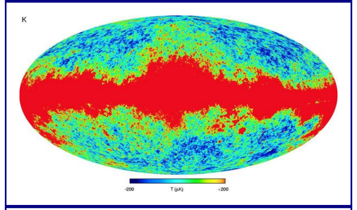

$\leftarrow$ MAP MAKING: from time lines to spherical maps.
Here, the microwave sky at 23 GHz seen by W-MAP.

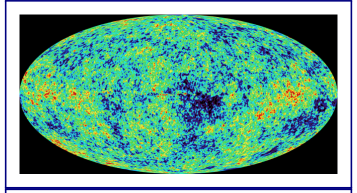

$\leftarrow$ COMPONENT SEPARATION: from several frequency channels maps to a component map. Here, pure (?) CMB from WMAP.

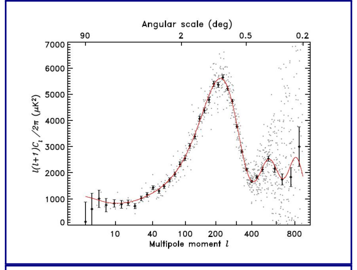

$\leftarrow$ SPECTRAL ESTIMATION: a bumpy 'angular spectrum'.

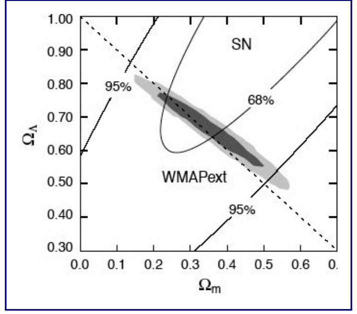

$\leftarrow$ LIKELIHOOD ANALYSIS. Here, likelihood of matter Ω_m and vacuum Ω
energy densities in front of CMB data (and supernovae).

$\rightarrow$ "Thus", the Universe is flat and 13.7 ± 0.2 billions years old (says WMAP)...

Spectre angulaire de la carte Planck sur 89% du ciel.

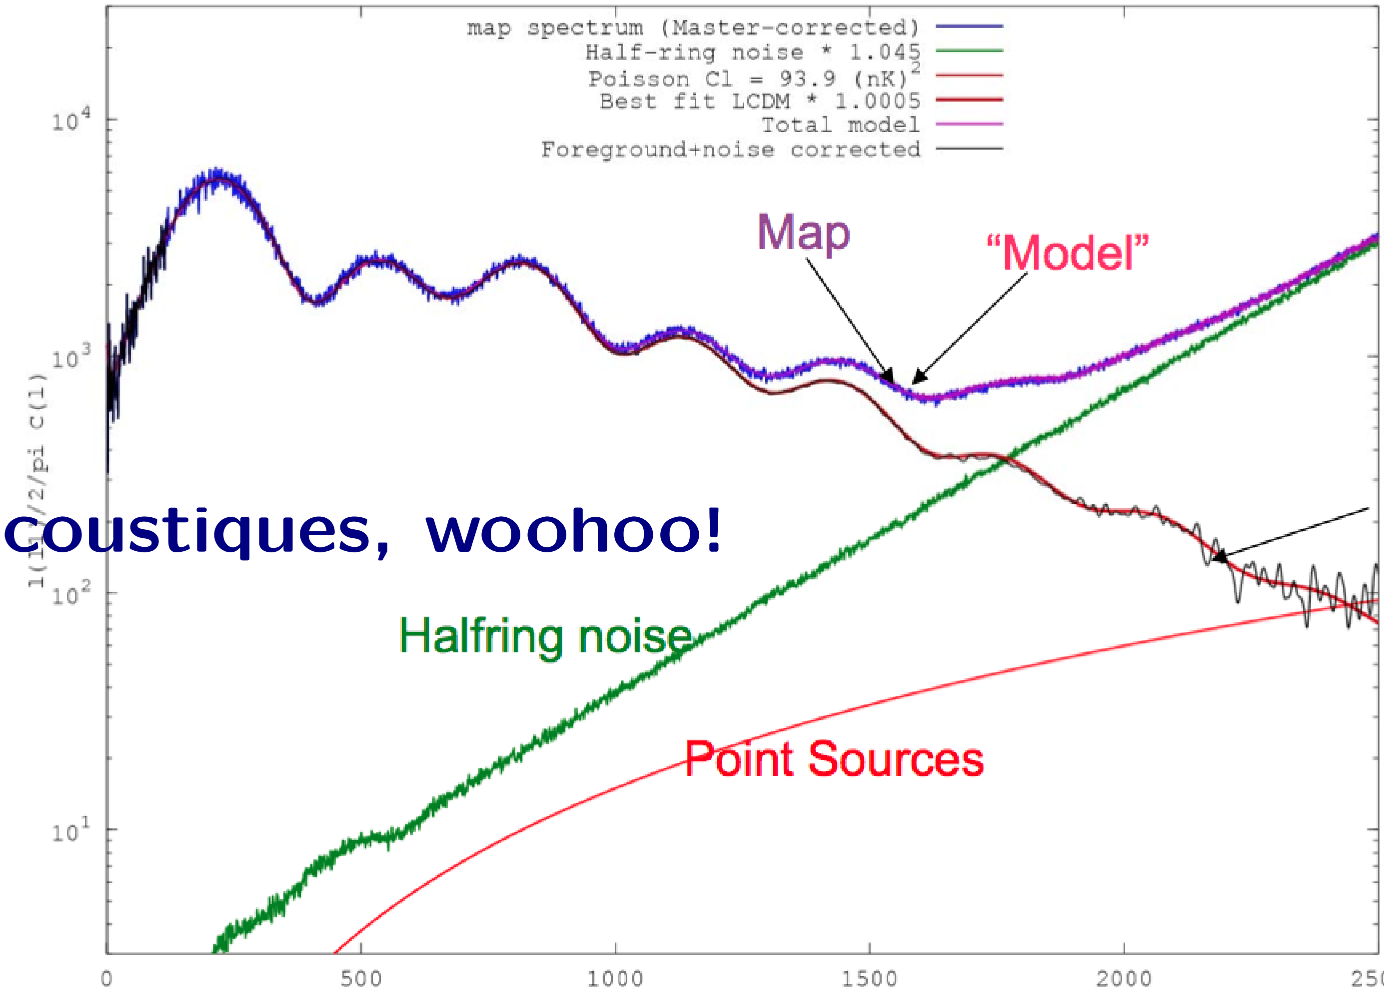

Sept pics acoustiques, woohoo!

Les erreurs sont dominées par la variance cosmique jusqu'à $\ell = 1500$ (disons).

Le spectre angulaire de Planck: contact entre théorie et observations

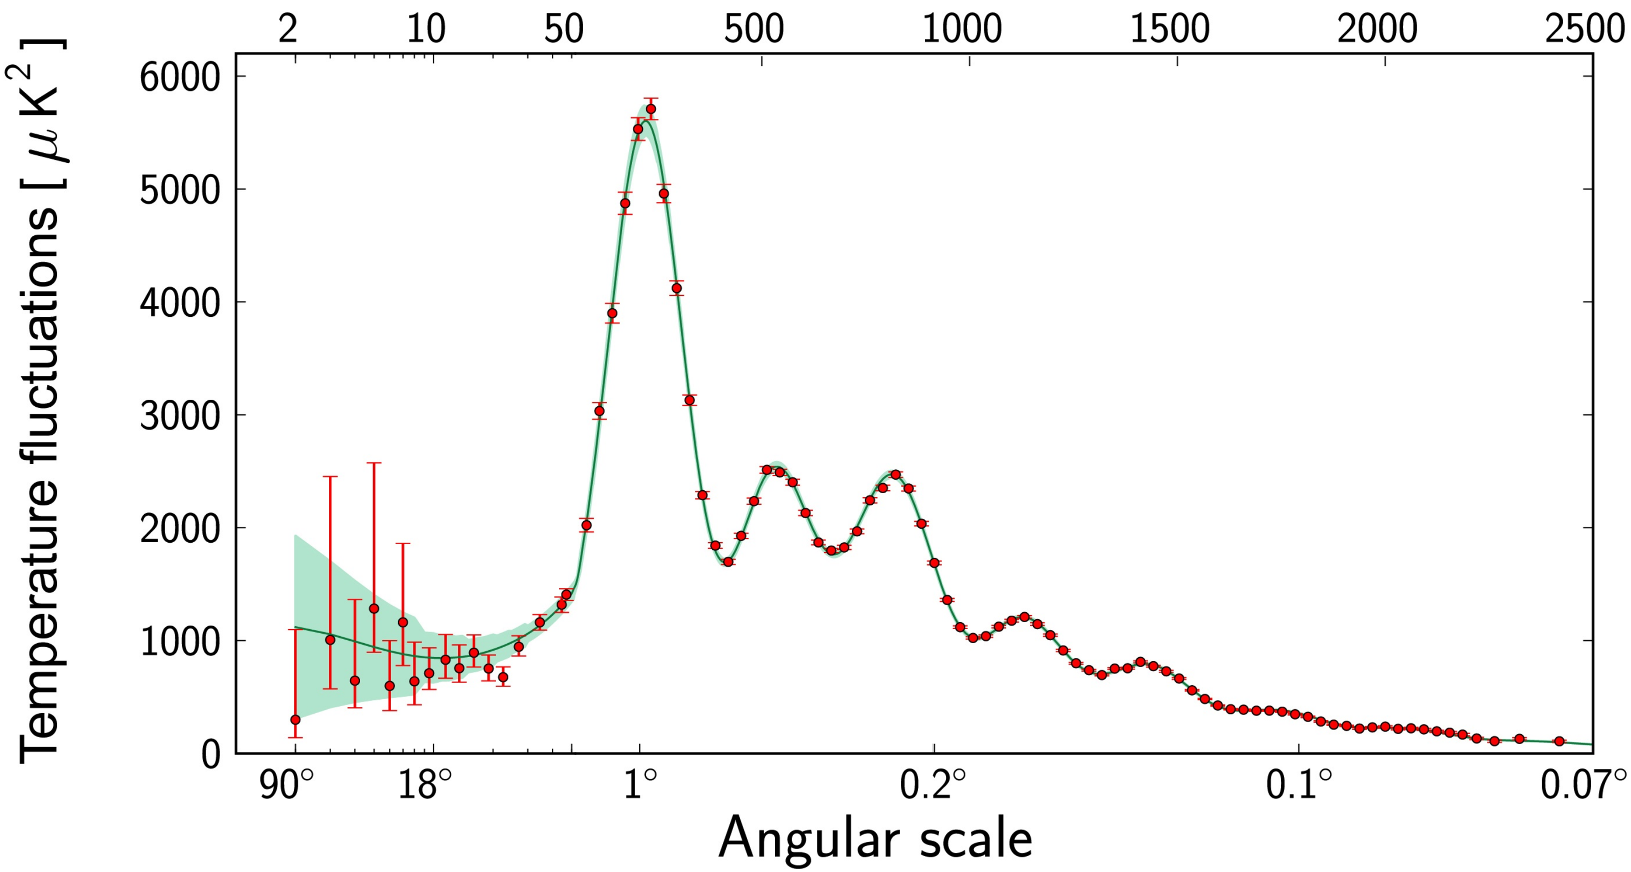

Un superbe ajustement des mesures par les prédictions du plus simple des modèles cosmologiques: le modèle $\Lambda - \mathrm{CDM}$ à 6 paramètres.

Le modèle standard du Big Bang

Un scénario à 6 paramètres: Le modèle Λ-CDM.

1. Amplitude A des fluctuations primordiales
2. Leur indice spectral n_s
3. Densité de matière noire Ω_d
4. Densité d'énergie noire Ω_Λ
5. Taux d'expansion H_o de l'Univers
6. Profondeur optique de ré-ionisation τ

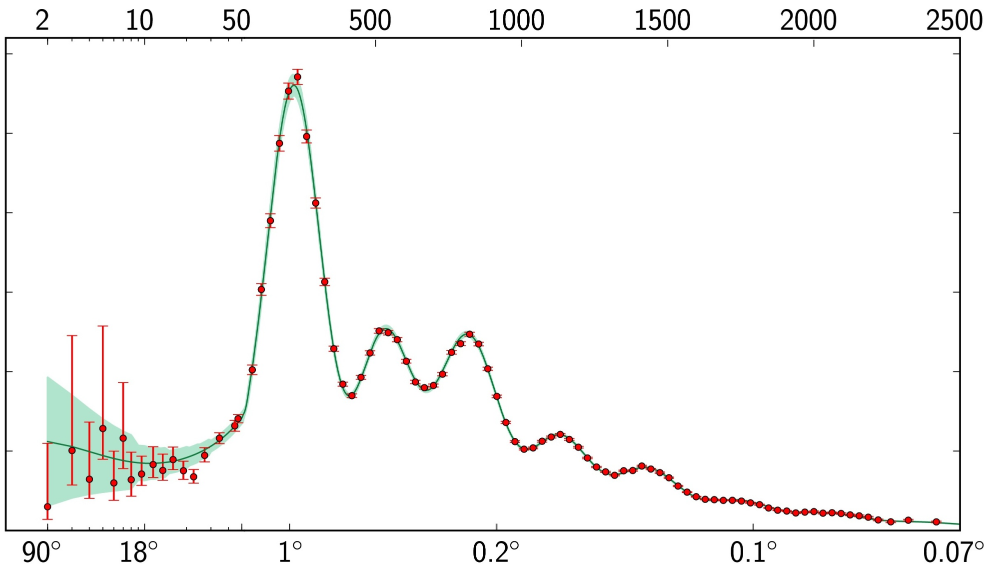

.

Les 2 premiers paramètres A et n_s décrivent le spectre primordial $\rightarrow$ inflation.

Les trois suivants H_o, Ω_d, Ω_Λ contrôlent sa "mise en forme".

Quelques résultats

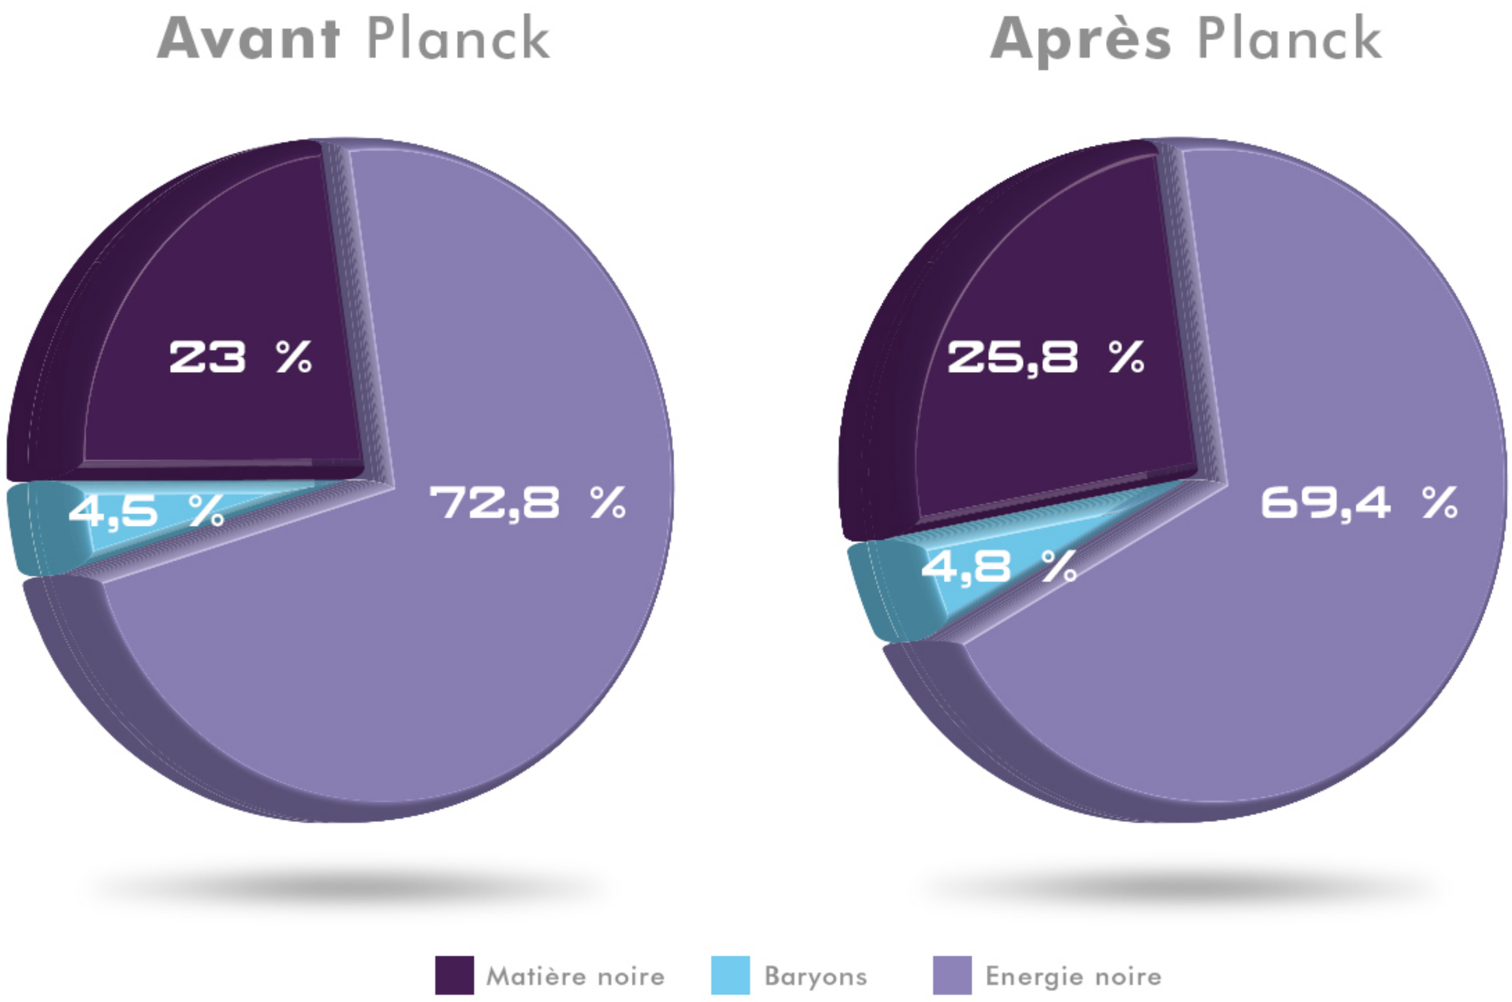

Un taux d'expansion H_o de 67,15 km/s/Mpc et un âge de 13,8 milliards d'années.

Planck results

- La plus grosse surprise: pas de grosse surprise.
- Excellente prédiction des observations avec un modèle simple...
  mais un peu bizarre: Λ-CDM = inconnu + inconnu !
- Sans parler de quelques anomalies marginales...
- Beaucoup d'autres façons d'exploter les données Planck,
  tant en Cosmologie qu'en Astrophysique.
- Magnifique succès scientifique, mais ce n'est pas fini...

Plus de science: lentillage gravitationnel

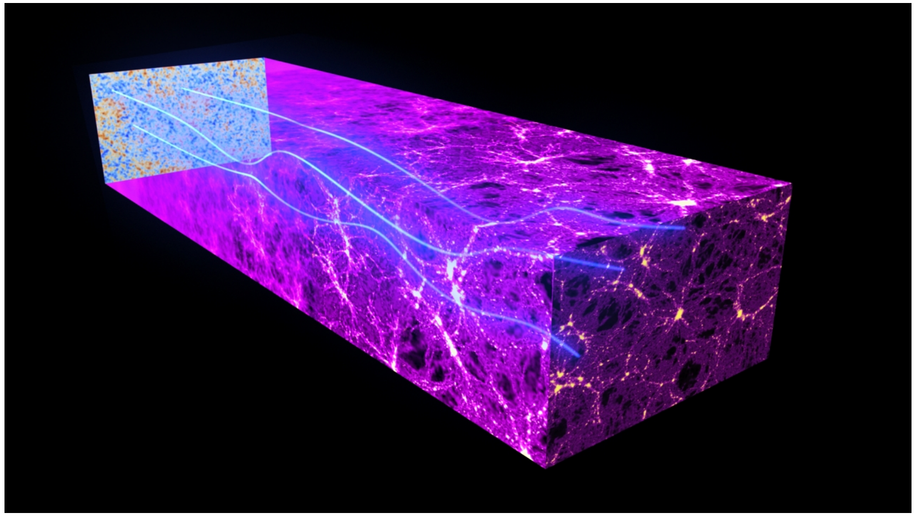

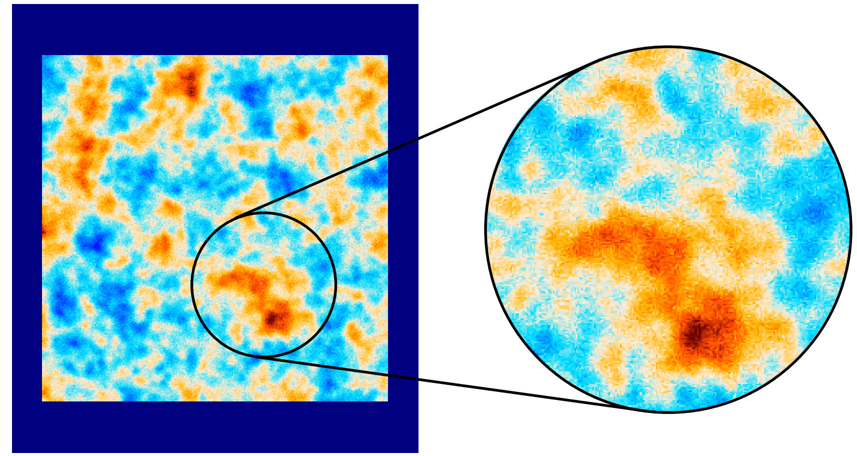

A movie

Plus de science: lentillage gravitationnel

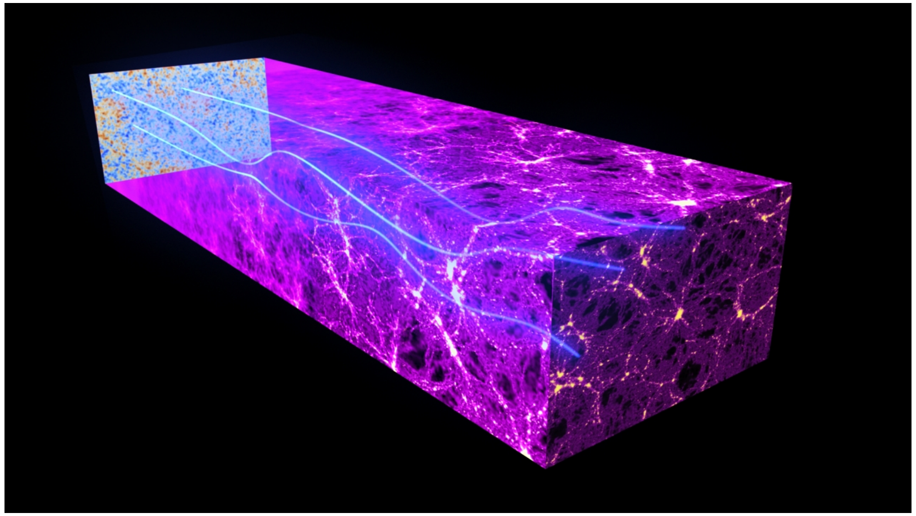

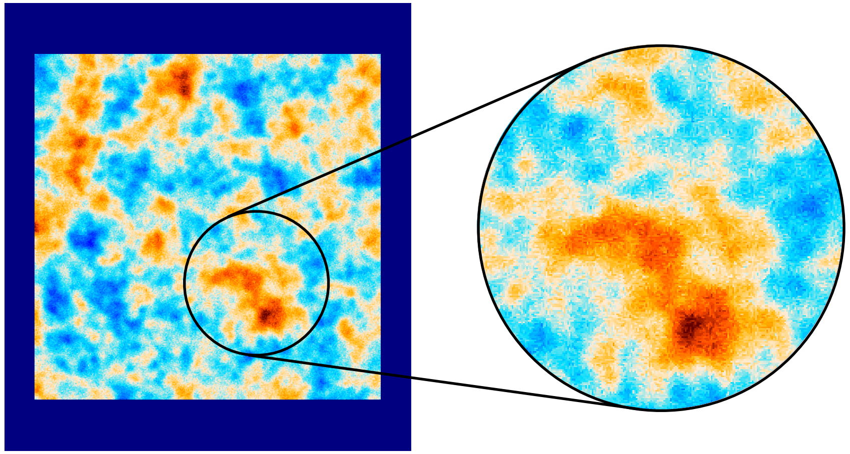

A movie

Big Data ? Who has the biggest?

For HFI: 52 bolometers (2 broken) plus 2 ”darks” (plus 16 ancillary channels: thermometers,. . .)

Mission lasts as long as Helium: our fridge achieved $>$ 2.5 years = 5 sky (about 1000 days).

1000 days of 24 hours of 60 minutes of 60 secs of 180 samples from 52 bolometers $\sim 8 \cdot 10^{11}$ bytes.

- Compression with various kinds of redundancy:

From timelines to rings: a factor of ≈ 40
From rings to spherical maps: about 5000 rings per sky, maps of $12 \cdot 2048^2 \sim 50 \cdot 10^6$ pixels
From detector maps to frequency channel maps (several bolometers for a frequency band)
From 9 frequency channel maps to one CMB map
From one map to a angular spectrum: all multipoles have SNR > 1 up to $\ell \approx 1800$
From the angular spectrum to 6 cosmological parameters
$\rightarrow$ Ultimate data size: 6 bytes.

- Bad/sad reasons for dumping data:

Instability after depointing: 1 minute per ring
Bad compression tuning in front of bright sources (was quickly fixed)
Dump data if dark bolos appear inconsistent (for whatever reason)
Nasty comsic rays give nasty glitches.
$\rightarrow$ In the end 30% of the scientific data had to be discarded.

Thanks and get ready for polarization!

The scientific results that we present today are a product of the Planck Collaboration, including individuals from more than 100 scientific institutes in Europe, the USA and Canada

Planck is a project of the European Space Agency, with instruments provided by two scientific Consortia funded by ESA member states (in particular the lead countries: France and Italy) with contributions from NASA (USA), and telescope reflectors provided in a collaboration between ESA and a scientific Consortium led and funded by Denmark.