COVER PAGE

Beyond Planck – delivering state-of-the-art observations of the microwave sky from 30 to 70 GHz for the next decade

List of participants

Participant No *	Participant organisation name	Country
1 (Coordinator)	Universitetet i Oslo	Norway
2	Universitá Degli Studi Di Milano	Italy
3	Istituto Nazionale Di Astrofisica	Italy
4	Helsingin Yliopisto	Finland
5	Planetek Hellas EPE	Greece
6	Expert Analytics AS	Norway

1. Excellence

1.1 Objectives

The detection of primordial gravity waves created during the Big Bang ranks among the greatest potential intellectual achievements in modern science. During the last few decades, the instrumental progress necessary to achieve this has been nothing short of breathtaking, and we are today able to measure the microwave sky with better than one-in-a-million precision. However, from the latest ultra-sensitive experiments such as BICEP2 and Planck, it is clear that instrumental sensitivity alone will not be sufficient to make a robust detection of gravitational waves. Contamination in the form of astrophysical radiation from the Milky Way, for instance thermal dust and synchrotron radiation, obscures the cosmological signal by orders of magnitude. Even more critically, though, are second-order interactions between this radiation and the instrument characterization itself that lead to a highly non-linear and complicated problem.

In this project, we propose a ground-breaking solution to this problem that allows for joint estimation of cosmological parameters, astrophysical components, and instrument specifications. The engine of this method is called Gibbs sampling, which we have already applied extremely successfully to basic CMB component separation. The new and critical step is to apply this method to raw time-ordered observations observed directly by the instrument, as opposed to pre-processed frequency maps, thereby closing the loop between instrumental characterization, astrophysical component separation and cosmological interpretation. While representing a ~100-fold increase in input data volume, this step is unavoidable in order to break through the current foreground-induced systematics floor.

We will apply this method to observations taken by the Planck Low-Frequency Instrument (LFI) between 2009 and 2013, and deliver new state-of-the-art frequency and component maps to the cosmological community. We will also combine these new data products with similar observations from the Wilkinson Microwave Anisotropy Probe (WMAP) observations and the ground-based C-BASS experiment, and demonstrate consistency and robustness across state-of-the-art experiments.

Thus, building on this base of observations, we will:

1. deliver new legacy Planck LFI 30, 44 and 70 GHz frequency maps.

2. deliver the world's cleanest and most sensitive full-sky estimates of polarized synchrotron emission at CMB frequencies. This new model will form a bed-rock for future CMB B-mode experiments searching for inflationary gravitational waves in the coming decade, as well as for scientists studying the structure and dynamics of the Milky Way.

3. deliver a new likelihood code suitable for large-scale CMB polarization analysis, and use this to derive a new and robust estimate of the optical depth of reionization, one of the most critical parameters in contemporary cosmology.

4. make the software necessary for time-domain analysis available to the community under an Open Science license, allowing other projects and experiments to build on and extend our work.

1.2 Relation to the work programme

The current project is submitted to the Horizon 2020 COMPET-4 program, which forms part of a LEIT (Leadership in Enabling and Industrial Technologies) programme. Specifically, the stated purpose and topic for the current call is to "*support the data exploitation of European missions and instruments, in conjunction, when relevant, with international missions.*" Furthermore, the detailed description of the call states the following (numbers are inserted to facilitate easy reference below):

> *(1) Projects selected under this call may rely on the data available through all the available ESA Space Science Archives when possible or other means (e.g. instrumentation teams). (2) Combination and correlation of this data with international scientific mission data, as well as with relevant data produced by ground-based infrastructures all over the world, is encouraged to further increase the scientific return and to enable new research activities using existing data sets. (3) These activities shall add scientific value through analysis of the data, leading to scientific publications and higher level data products. (4) When possible, enhanced data products should be suitable for feeding back into the ESA archives. (5) Resulting analyses should help preparing future European and international missions. (6) International cooperation is encouraged in particular with countries active in space exploration and science, or where their participation is deemed essential for carrying out the activities of this topic.*

As will be made clear in the following discussion, the proposed project is perfectly aligned with every one of these goals:

(1) The core of the proposed project is the Planck LFI observations, one of ESA's cornerstone missions, and the fundamental data are available through the Planck Legacy Archive (PLA), https://pla.esac.esa.int

(2) We will combine the LFI observations with the international (NASA) WMAP observations as well as the ground-based C-BASS observations, to derive new state-of-the-art estimates of the low-frequency polarized sky. In addition, the full science potential of our new maps will be reached by combining them with external data sets, including both CMB (e.g., Planck HFI, ACT, SPT etc.) and other types of cosmological probes (e.g., Type Ia supernova observations, baryonic accoustic oscillation measurements).

(3) From these observations, we will derive new estimates of the polarized synchrotron sky across the three world-leading full-sky low-frequency CMB experiments. These products and results will be essential both for understanding the properties of the Milky Way (structure, dynamics, magnetic fields, etc.), which is one of the main issues for detecting inflationary gravitational waves with next-generation experiments, and preparing for next-generation polarization B-mode experiments.

(4) Our results will be fed back to the PLA, essentially replacing the currently available products with more robust and sensitive versions, useful for a wider range of scientific applications.

(5) The single most important scientific driver for our project is to inform future searches for inflationary gravitational waves, such as CORE (European-led satellite mission), LiteBIRD (Japanese-led satellite mission), S4 (US-led ground-based experiment) and many others. Indeed, this overall goal is the main reason why it has been possible to establish the necessary team of internationally leading (and very busy) scientists in the first place: We

recognize that this work *has* to happen for the field to go forward, and we are willing to put the necessary amount of time and effort into leading it to a success.

(6) The scope and ambition of this project naturally implies that the team must be international in nature. The necessary know-how, skills and expertise is naturally distributed across several countries. Indeed, the team has been explicitly designed to optimally exploit the strengths and skills of each member, but without adding unnecessary “weight”. Specifically, Oslo will play a leading role in the component separation, Gibbs sampling and integrational aspects of the work; Trieste will lead the LFI time-ordered data processing; Helsinki will lead the map making aspects; Milano will take charge of systematic effects and cosmological interpretation; Athens will take charge of data delivery and archiving. In every single case, these responsibilities are perfectly aligned with the proven expertise of each team; there is literally no other team in the world who are better equipped to do this task than the current, simply because the team has been established strictly according to scientific demands. And as such, international collaboration is an absolute prerequisite.

1.3 Concept and methodology

(a) Concept

i. The hunt for inflationary gravitational waves and the BICEP2 experience

On March 17th 2014, the public relations office at Harvard University issued a press release [1,2] proclaiming that “Researchers from the BICEP2 collaboration announced the first direct evidence for cosmic inflation. Their data represent the first images of gravitational waves, or ripples in space-time.” Few, if any, press releases in modern cosmology have ever generated greater interest. It easily eclipsed the first announcement of our own Planck results [3], and no WMAP [4] press release came even close. Indeed, it was no less visible than the highly celebrated discovery of the Higgs boson in 2012 [5] that was awarded the Nobel Prize of Physics already in 2013. In the 12 months after the BICEP2 release, more than 500 papers appeared on arXiv, the physics preprint server, discussing the BICEP2 results and their implications. In fact, shortly after the release two independent papers appeared back-to-back on the arXiv listing having titles ending with “... in light of BICEP2”.

What caused all this interest and excitement? The answer is simple: if the BICEP2 claims were correct, this would be our first direct view of physics taking place during the creating of the universe, just some 10^{-34} seconds after the beginning of time [6]. It would be our best chance to actually measure physics at the Planck energy scale, and thereby probe quantum gravity. Additionally, it would be the first direct detection of gravitational waves, as predicted by Albert Einstein in 1916 [7], similar in nature to those later observed by LIGO from colliding black holes [8]. It is no exaggeration to say that this would be one of the greatest cosmological discoveries of all time, with far-reaching theoretical implications for both particle physics and cosmology.

The signal reported by the BICEP2 team was massive. Prof. Clem Pryke, co-leader of the BICEP2 team, summarized their findings as follows [1]: “This has been like looking for a needle in a haystack, but instead we found a crowbar.” The black points in the top panel of Figure 1 shows the angular power spectrum the BICEP2 team derived when combining their original measurements with more observations from the Keck telescope (BICEP2's successor; [9]), compared to the expected signal that is caused by normal weak gravitational lensing (red solid line). The excess power corresponds to a best-fit value of the tensor-to-scalar ratio (ie., the relative amplitude of inflationary gravitational waves to the signal expected from standard density perturbations) of r =

0.20 ± 0.06 [2]; before BICEP2, most people expected this ratio to lie somewhere in the range of r = 0.001 to 0.05, if not even lower [6].

Unfortunately, the story was indeed too good to be true. Shortly after the initial release, complementary observations from Planck became available [3]. Since Planck is also observing at both lower and higher frequencies than BICEP2, it is far more sensitive to contamination from astrophysical foregrounds in the form of synchrotron and thermal dust emission. And when comparing Planck's thermal dust map to the measurements from BICEP2 [9], it became evident that the claimed signal was not due to inflationary gravitational waves, but rather contamination from the Milky Way. The bottom panel in Figure 1 shows an overlay of the BICEP2 map on Planck's 353 GHz map, while the blue points in the top panel shows the angular power spectrum after subtracting the best-fit dust contribution. Clearly, the signal vanished, and the hunt for inflationary gravitational waves continues at full strength.

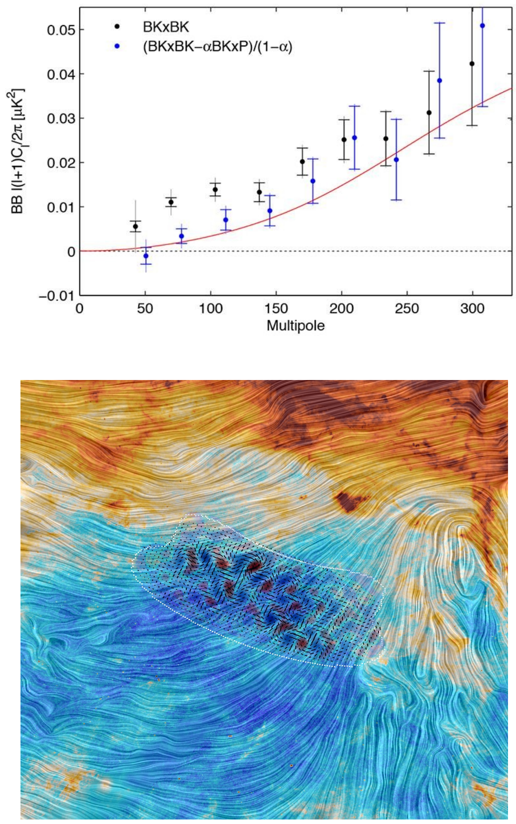

Figure 1: (Top) BICEP2/Keck power spectrum as measured with (blue points) and without (black points) cross-correlation with Planck [9], which is equivalent to marginalize over thermal dust foreground contamination or not [2]. (Bottom) Original BICEP2 field overlaid on the Planck 353 GHz map. Courtesy of Jon Gudmundsson. The Planck map clearly shows strong thermal dust emission in the BICEP2 field.

ii. The next generation – possibilities and challenges

As the BICEP2 results (literally) fell to dust, other competing experiments are currently adjusting their strategies. In particular, after Planck and BICEP2 it has become strikingly obvious that the main challenges ahead will no longer be raw instrumental sensitivity as such, but rather accurately understanding the emission mechanisms in our own Milky Way.

Our main tool for achieving this is frequency coverage. Whereas CMB radiation follows a nearly perfect blackbody with a temperature of 2.7255 K [10], all known astrophysical emission mechanisms have non-thermal spectra [11]. For instance, synchrotron emission has a spectral energy density (SED) that scales as ν^{-3}, and thermal dust emission has a spectrum that may be accurately approximated as a modified black-body spectrum with a slope of β = 1.5 and a temperature of 18-21 K. No other polarized emission mechanisms have been detected in the microwave frequency range as of today, although both free-free, CO and spinning dust are all expected to be polarized at a low level.

To reliably distinguish between the various emission mechanisms, the trick is to observe the microwave sky at many different frequencies between, say, 10 and 1000 GHz. This is of course expensive, and different experiments therefore adopt different optimization strategies, all attempting to maximize their CMB sensitivity in the most cost effective manner. As of February 2017, the nominally most sensitive experiment for large-scale B-mode searches with data on disk is SPIDER [12], a balloon-borne bolometer based experiment launched from McMurdo, Antartica, in January 2016 (see top panel in Figure 2).

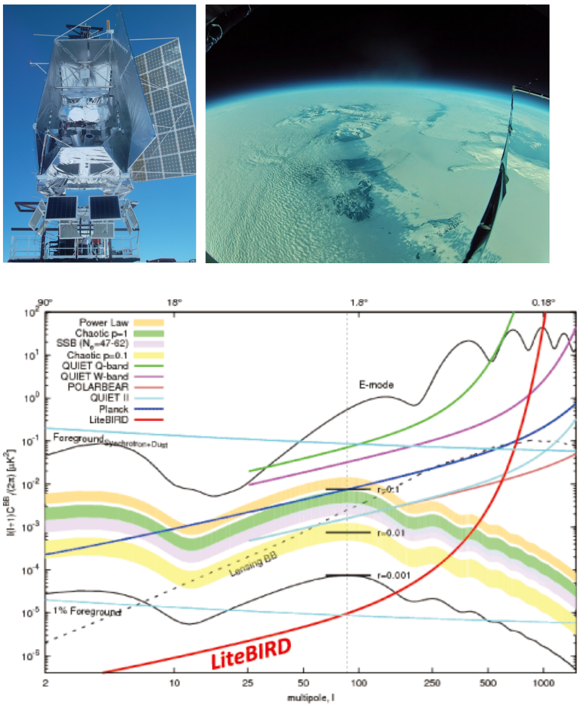

Figure 2: (Top left) The SPIDER payload as launched on January 1st 2016. (Top right) SPIDER view from 36,000 meters above Antartica. (Bottom) Power spectrum forecast for the LiteBIRD satellite concept [13]. The red curve indicates the white noise limits of the experiment, while the colored bands indicate various theoretical expectations. Colored curves indicate the noise limits of a few previous-generation experiments.

Many other experiments are currently being planned, funded and deployed, but due to space limitations we will not even attempt to review them here. However, the most advanced next-generation CMB polarization probe is the Japanese-led LiteBIRD [13] satellite mission. This experiment is currently undertaking its Phase A study. It has a nominal launch target in 2021, and aims to operate for at least one year. According to the latest project description, LiteBIRD will observe at 15 frequencies between 40 and 400 GHz.

In Europe, the biggest effort is organized around a concept called CORE, which in many respects may be thought of as Planck's successor dedicated to polarization observations. Although it has yet to be funded, the enthusiasm in the field is large, and it is safe to say that new iterations of this mission will continue to be developed and proposed in the foreseeable future, simply because there is literally no bigger target in contemporary CMB cosmology than the detection of inflationary gravitational waves.

The main goal of LiteBIRD (CORE) is to constrain the tensor-to-scalar ratio down to $r < 0.001$ (0.0001) without the use of external priors or dependencies. These forecasts are summarized for LiteBIRD in terms of angular power spectra in the bottom panel of Figure 2. The red curve shows the effective CMB noise level as a function of angular scales, after accounting for foreground subtraction and marginalization. The lower black solid curve shows a theoretical B-mode spectrum corresponding to $r = 0.001$, and the dotted line shows the signal expected from weak gravitational lensing. Colored bands indicate a few selected inflationary models, all of which will either be detected or ruled out by LiteBIRD. For comparison, the upper black solid curve shows the (E-mode) polarization signal expected from standard density perturbations. Intuitively, the tensor-to-scalar ratio is essentially the ratio between the lower and upper black solid curves at a pre-defined reference scale.

Within the next five years, all these upcoming B-mode experiments are certain to constrain CMB B-modes to unprecedented levels. However, as discussed in the next two sections, the technical challenges ahead are formidable – and addressing these challenges is exactly what the current project is all about.

iii. Astrophysical component separation, CMB Gibbs sampling and Commander

As described above, after Planck and BICEP2 the topic of astrophysical component separation has taken center stage in current CMB analysis. Of course, this challenge has been anticipated within the community for a long time, and a wide range of algorithmic ideas for how to efficiently deal with it has been developed. The biggest and most coordinated effort has taken place within the Planck collaboration [11,14]. Starting from a field of 10-15 different algorithmic candidates, four fundamentally different methods were adopted for the published Planck analysis. Of these, three (NILC, SEVEM and SMICA; [14]) focused primarily on CMB extraction, while only one consid-

ered the problem from a global point of view, simultaneously addressing both CMB and foreground reconstruction. This latter approach was developed, implemented and applied by Prof. Hans Kristian Eriksen, Dr. Jeffrey Jewell (JPL/NASA) and Dr. Ingunn Kathrine Wehus (Oslo) in a computer code called *Commander* [15,16]; the development of this code was supported by Prof. Eriksen's ERC Starting Grant running from 2011 to 2015.
Commander forms a cornerstone in current Planck analysis, and is involved all the way from low-level gain estimation and map making via intrinsic component separation to power spectrum and cosmological parameter estimation.

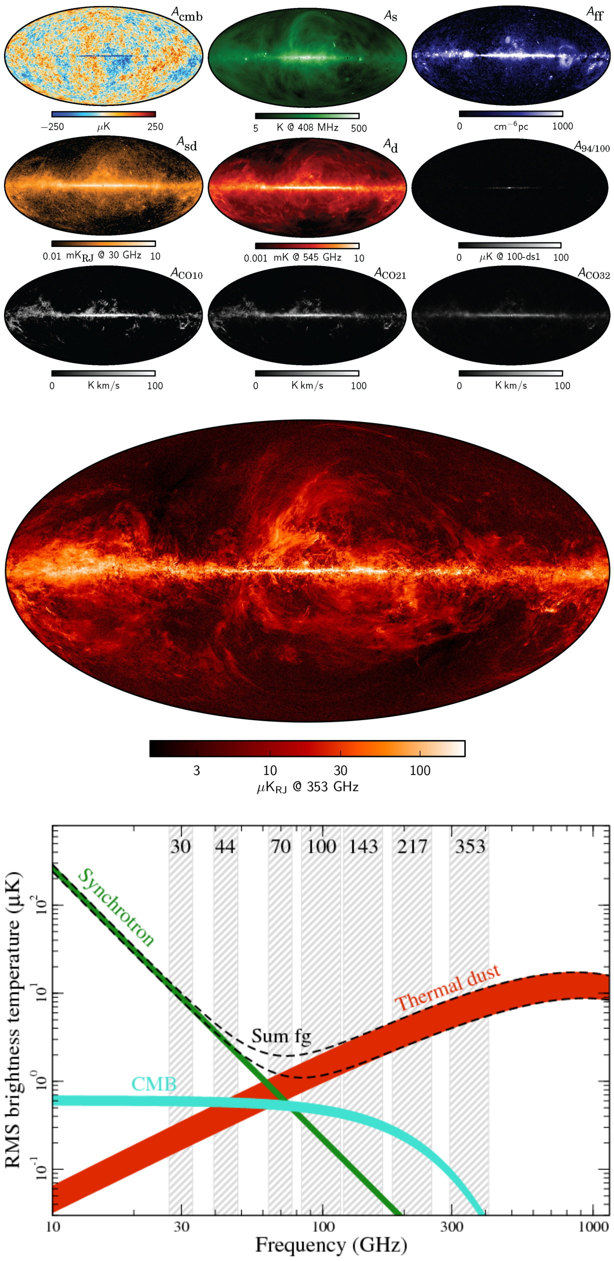

Figure 3: Selected plots from the Planck 2015 astrophysical model [11,14]. This model was in its entirety derived with Commander. (Top) Planck 2015 temperature sky model consisting of CMB, synchrotron, free-free, thermal and spinning dust, and CO and HCN line emission, all separated using frequency information. (Middle) Polarized thermal dust emission. (Bottom) Summary of polarized emission mechanisms as a function of frequency.

The ubiquity of Commander is a direct consequence of its global nature. As sensitivity improves, more and more effects are visible above the noise level, and must be addressed. Usually, implementing "ad-hoc hacks" to address these typically works for a while, but once the complexity of the problem grows, this approach invariably breaks down. And at that stage, only a truly global and exact approach works, and this is what Commander was built for.

In the Gibbs sampling approach we approach the problem of CMB analysis from a Bayesian point of view [16]. That is, we start by writing down an explicit model for all relevant parameters – cosmological, astrophysical and instrumental – and we derive an expression for their joint posterior distribution. As usual in Bayesian analysis, the goal is then to map out this posterior somehow, either in terms of maximum-posterior estimates or confidence intervals. Of course, the CMB posterior typically involves millions of correlated and non-Gaussian parameters, and no explicit sampling or maximization scheme exists for it. The only realistic hope is through Monte Carlo sampling, and in particular *Gibbs sampling*. Gibbs sampling is a well-known statistical technique that allows the user to draw samples from a complicated joint posterior by iteratively cycling through its corresponding *conditional* distributions. Typically, these conditionals are all vastly simpler than the joint posterior, as is also the case for the joint CMB—foreground posterior [16].

With this simple algorithm, members from our team derived a comprehensive signal model for the Planck 2013 and 2015 releases. A few highlights from the latter are shown in Figure 3. The top panel shows our temperature sky model, including CMB, synchrotron, free-free, spinning and thermal dust, CO and HCN emission sky maps. The middle panel shows the iconic Planck polarized

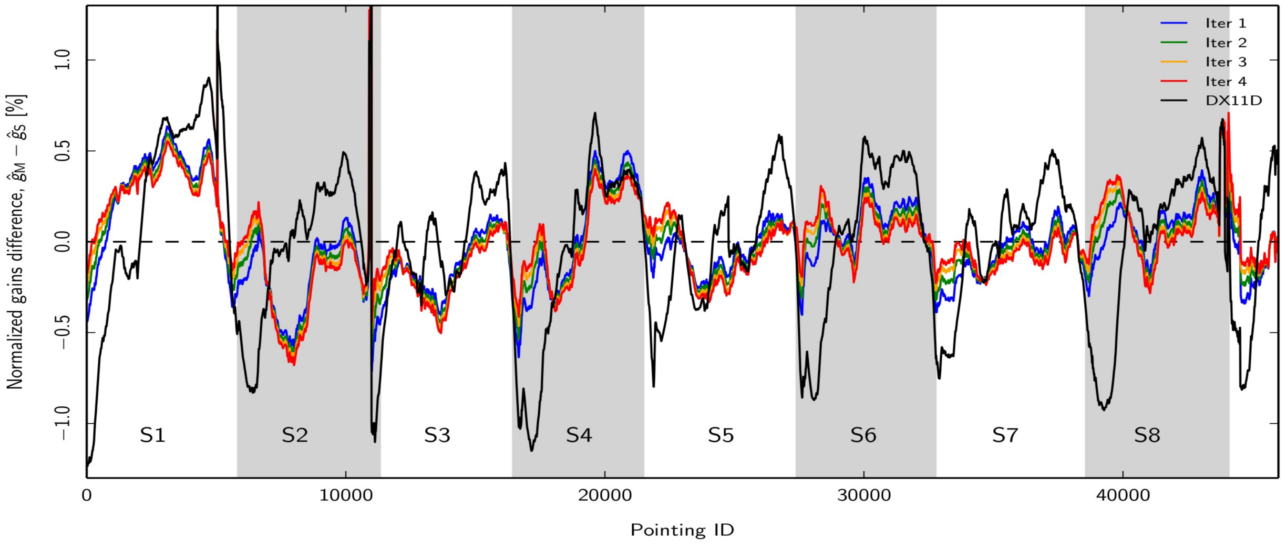

Figure 4: Fractional gain variation as a function of time as estimated for Planck LFI using different polarized foreground assumptions (Planck 2017, in preparation). Uncertainties in the foreground model directly affect the instrumental characterization, leading to a highly non-linear estimation problem. Solving the joint cosmology, astrophysical foreground and instrument estimation problem is the main goal of the current project.

thermal dust emission map. The bottom panel shows the most complete summary of polarized emission mechanisms to date across CMB frequencies.

iv. Lessons learned from Planck: hitting the foreground-induced systematics floor

While the importance of both astrophysical foreground contamination and instrumental systematics has been recognized in the CMB community for a long time, their intimate interplay has not been fully appreciated until recently. For a long time, most of us believed that it would be possible to first handle instrumental systematics, and derive well-behaved frequency maps from "cleaned" time-ordered data. From these maps, one would perform component separation to estimate a CMB map (as described above), before finally making cosmological inferences. Without exceptions, all CMB experiments to date have operated in this linear fashion [3,4].

However, as our understanding of the Planck observations has improved, and we have dug deeper and deeper into the noise floor, it has become painfully apparent that this simplistic approach is no longer appropriate. The problem is simple to formulate: the white noise level of Planck corresponds to about 100 nK per unit area on large angular scales. The only astrophysical calibration source that supports such high accuracy is the Doppler-induced CMB dipole, with an intrinsic amplitude of 3 mK. This number is more than four orders of magnitude larger than the target sensitivity. As a result, it is necessary to achieve a relative calibration stability of better than 10^{-4} in order to eliminate systematic effects that otherwise turn into stripes in the maps [17].

In practice, this is amazingly difficult to achieve for one fundamental reason: in order to reach a calibration precision better than 10^{-4}, it is also necessary to know the *astrophysical sky* to µK levels or better. Otherwise, those astrophysical sources bias the calibrator itself at precisely the same level; after all, the uncalibrated instrument does not observe the CMB dipole alone, but rather the sum of all signals originating from the sky. In other words, it is necessary to know the microwave sky to µK precision before one can characterize the instrument – but it is also necessary to characterize the instrument to 10^{-4} precision before one can estimate the sky! Thus, the problem is no longer linear, but circular.

This has been the crux of the problem for both Planck LFI and HFI ever since our last data release in 2015, and a great amount of effort has been spent in resolving it. For both instruments, we have reached the same basic solution: in order to break through the foreground-induced systematics floor, it is necessary to iterate between instrument characterization, map making and component separation, tying the old linear analysis problem into a closed loop. Unfortunately, because of internal time and resource constraints, leaving no time to rewrite codes from scratch, this process have had to be done “by hand”, feeding data files from one computer to another, and having separate people run their individual codes; one person does gain estimation, another does map making, and a third does component separation.

We have already implemented this manual process for the latest (yet unpublished) Planck 2017 processing, and the colored curves in Figure 4 shows the time-variable gain solutions obtained in four consecutive iterations. For comparison, the black curve shows the solution obtained with no iterations. With no iterations, the relative errors are up to 1%, which is nowhere near sufficient for accurate polarization reconstruction. With the current manual iterations, the relative uncertainties between consecutive iterations are much less than 0.1%. The single biggest problem with this approach is the long turn-over time, resulting in more than one week per iteration due to the manual interaction.

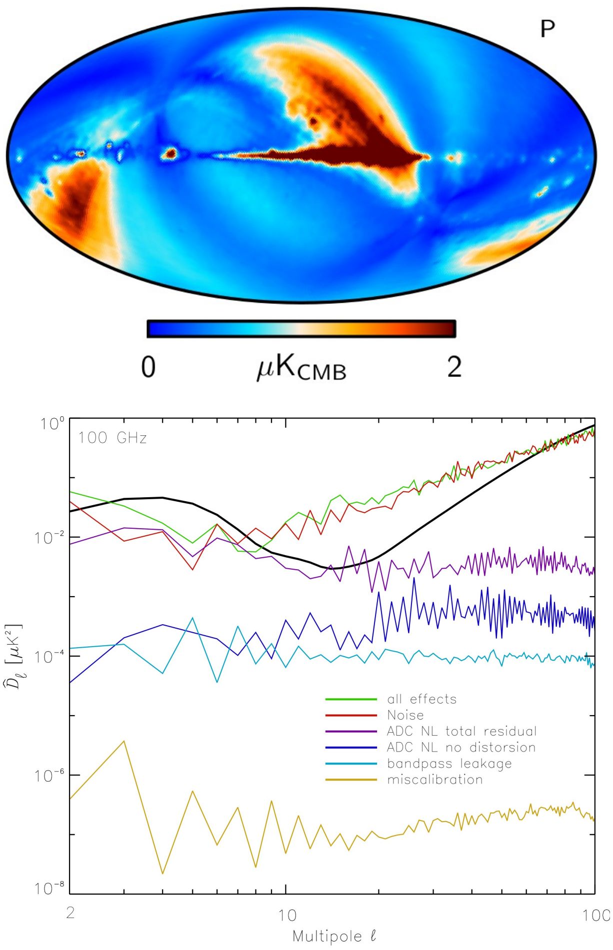

Figure 5: Systematic effects in Planck. (Top) Residual foreground-induced gain template for LFI 70 GHz adopted for (yet unpublished) 2017 analysis. (Bottom) Summary of simulated polarization effects for the pre-2017 HFI analysis [17]. The solid black curve indicates the power spectrum of scalar EE modes. The B-mode spectrum from gravitational waves is (at least) one to three orders of magnitude below this, depending on the value of the tensor-to-scalar ratio, r.

In the following, we propose to solve this problem by developing a statistically coherent framework that allows end-to-end analysis of CMB data completely without intermediate human interaction. This will allow us to run hundreds or thousands of such full analysis cycles, as opposed to just the handful we are currently limited to within Planck, and thereby finally reach full robust convergence and precision.

The top panel in Figure 5 shows the difference between two consecutive iterations for the Planck LFI 70 GHz frequency channel (Planck 2017, in prepration). This residual is comparable in magnitude to the cosmological target signal we aim to measure of about 1μK peak-to-peak, in this case cosmic reionization. Without iterations, the residual systematic would swamp the cosmological signal. Likewise, the bottom panel shows the expected bias in the more sensitive HFI 100 GHz frequency channel after making all the latest corrections [17]. Again, without these types of corrections, the systematic effects drown the signal.

While this discussion has used Planck as a worked example, it is important to realize that the basic problem is fundamental in nature: the only calibrators that are able to support sub-μK and nK precision are astrophysical in nature, but to use these it is in fact necessary to know them to comparable precision. The problem is circular in nature – and it needs to be handled accordingly in order to resolve it.

v. Project concept, deliverables and readiness level

We are now ready to describe the essential concept of the proposed project, both in terms of methods, goals and deliverables.

We propose to build a new CMB Gibbs sampler, following in the footsteps of the successful Commander code, but this time starting from raw time-ordered data, as opposed to co-added frequency maps. Algorithmically, this entails merging instrument characterization, map making and component separation into one coherent framework, such that the results from one step can feed seamlessly into the next. This is thus fully equivalent to what we have already done by hand in Planck, but by automating the process, we will be able to run hundreds and thousands of cycles, as opposed to just a handful.

We will then use this new end-to-end code to reprocess the Planck LFI data, going directly from time-ordered data to astrophysical component maps and cosmological products. We will also combine the Planck LFI data with the best complementary data products available anywhere for the same frequency range, namely the WMAP sky maps and C-BASS. WMAP is already available both in terms of raw time-stream data and pixelized sky maps, whereas C-BASS is expected to be released within the next year.

Based on this combined data set, we will deliver a new model of polarized synchrotron emisssion in the microwave frequencies. This model will be essential for planning, forecasting and analyzing future B-mode polarization data. While many experiments are planning to provide new information regarding thermal dust polarization in the coming years, the combination of Planck LFI, WMAP and C-BASS is very likely to remain unchallenged on full-sky synchrotron modelling in the coming decade, and it is therefore of the greatest importance to optimally extract all relevant information from these data sets at the present time. This is the primary scientific goal of the current proposal.

In the process, we will also produce a great amount of secondary science products, including, but not limited to, new constraints on the optical depth of reionization, one of the most important cosmological parameters; the structure and dynamics of the galactic field of the Milky Way; new beam deconvolved Planck LFI frequency maps; new estimates of low-frequency point sources etc. All these results will be published in leading international journals.

In addition to these basic scientific results, which will be released both in the form of peer-reviewed publications and machine readable data products (sky maps, power spectra, likelihoods etc.), we also aim to transform how CMB science is done, such that all experiments can benefit from our experience, both from this project and from Planck. An important step towards this goal is to adopt a strict focus on Open Science. In particular, all source code developed in this project will be made available in open-access repositories (GitHub), and all data products, parameter files etc. will be made available on public servers.

By providing a well-defined list of deliverables to the cosmological community, this project positions itself near the center of a spectrum spanning from “lab to market”. We will take raw data produced by one of ESA's leading cosmology missions, and deliver easily usable products to the community at large. In terms of Technical Readiness Level, we consider all individual components to be at TRL7 (“S*ystem prototype demonstration in operational environment*”), considering that each component has already been employed extensively in Planck analysis for the last decade. The only new step is to integrate these into a single code that will not require human interaction between each step, and we therefore consider the entire project to be at TRL6 (“*technology demonstrated in* *relevant environment*"). We do not foresee any major complications in order to reach TRL9 within the first year of the project.

(b) Methodology

i. Time-domain Gibbs sampling – end-to-end analysis within a single coherent framework

As discussed above, the essential problem with the current state-of-the-art approach to end-to-end CMB analysis is the lack of feedback in the process. It is essentially a linear process starting with low-level data (data selection, gain calibration, noise estimation etc.) going into frequency maps, which then are converted into component maps, from which cosmology finally is being derived. However, with the high sensitivity of the Planck observations, this is no longer suitable, because uncertainties in the astrophysical sky bias the low-level instrumental characterization at a level comparable with the desired cosmological results. It is therefore necessary to close the loop, and feed the final products back into the instrument calibration process, and iterate some significant number of times until convergence is reached.

This process may be mathematically formulated in terms of a modern statistical technique called Gibbs sampling [15,16,29,30]. In this method, a complicated multivariate distribution is broken down into a simpler set of conditional distributions through iterative sampling. In our case, we will be interested in the joint distribution describing all relevant cosmological (CMB map, power spectrum and parameters etc.), astrophysical (synchrotron, thermal dust emission etc.) and instrumental parameters (gain, noise, beam characterizations etc.).

As in any Bayesian analysis, the first step is to write down an explicit parametric model of the data. For the latest Commander-based analysis performed in Planck, we adopted the following model [11,14],

$$d_{p,\nu} = g_\nu \sum_{i=1}^{N_{\mathrm{comp}}} F^i_{p,\nu}(\theta) a^i_p + n_{\nu,p}$$

where the sum runs over all relevant sky signals, both cosmological and astrophysical; a^i_p denotes the amplitude of component i in pixel p; F describes the frequency dependence of component i, the so-called mixing matrix, which depends on some set of spectral parameters θ, for instance physical temperatures or spectral indices; g_ν is a calibration uncertainty per frequency; and $n_{p,\nu}$ denotes instrumental noise. (In addition, we introduced parameters to quantify uncertainties in the detector frequency response of each detector, but for notational simplicity we include these among θ for now.) In addition, we are interested in the angular power spectrum, C_l, which is defined through the signal covariance matrix of the CMB component, most typically adopting a spherical harmonic basis for the CMB map,

$$S(C_\ell) = \left\langle a^{\mathrm{cmb}}_{\ell m} a^{\mathrm{cmb},t}_{\ell m} \right\rangle$$

Note that the connection to cosmological parameters, such as the tensor-to-scalar ratio is made via this power spectrum. The full set of free parameters is then $\{a^i_p, \theta, g_\nu, C_l\}$, and our primary goal is to quantify the corresponding joint posterior distribution, $P(a^i_p, \theta, g_\nu, C_l| d_{p,\nu})$. Of course, this distribution involves millions of parameters (remember that each a and θ is a full parameter sky map), many of which are both non-Gaussian and strongly correlated. Obviously, it may not be written in a closed analytic form. In order to map it out, we therefore resort to Gibbs sampling, and set up the following sampling scheme,

$$
\begin{array}{rcl}
a_p^i & \leftarrow & P(a_p^i | d_{p,\nu}, \theta, g_\nu, C_\ell) \\
\theta & \leftarrow & P(\theta | d_{p,\nu}, a_p^i, g_\nu, C_\ell) \\
g_\nu & \leftarrow & P(g_\nu | d_{p,\nu}, a_p^i, \theta, C_\ell) \\
C_\ell & \leftarrow & P(C_\ell | d_{p,\nu}, a_p^i, \theta, g_\nu)
\end{array}
$$

The left-arrow symbol indicates over-writing the variable on the left hand side with a sample drawn from the distribution on the right-hand side. The full Gibbs sampling chain simply consists of looping through these steps until convergence, and marginal distributions for each parameter may be computed from the resulting ensemble of samples.

Each of the conditionals in the above cycle has a relatively simple analytic form that allows efficient sampling [16]. For instance, the first and third distributions are multivariate Gaussians, while the second may be mapped out in terms of a simple χ^2. The fourth is an inverse Wishart distribution. All of these are implemented with efficient algorithms in Commander.

In our existing implementation, however, the input data are considered to be pixelized sky maps, $d_{p,\nu}$, derived through some pre-processing step (instrument characterization and map making) that is in principle unknown to our Gibbs sampler. As discussed extensively above, while this was good enough for all experiments until Planck, it is not sufficient for experiments that require sub-µK and nK sensitivity. For these, it is necessary to perform a joint analysis of cosmological, astrophysical and instrumental parameters. And since the instrument can only really be described properly at the time-level, we are forced to adopt a much more basic data model, namely

$$d_{t,\nu} = g_{\nu,t} P_{t,p}^{\nu} \sum_{i=1}^{N_{\mathrm{comp}}} F_{p,\nu}^i(\theta) a_p^i + n_{\nu,t}$$

This expression incorporates two important changes. First, the data vector is now dependent on time rather than pixel. This time variable corresponds to the index of the time-ordered data recorded by the instrument, and runs over billions of samples, as opposed to tens of millions of pixels. In order words, this seemingly minor change represents an increase in input data volume of two orders of magnitude or more. Second, a new projection operator, P, has been introduced that describes the pointing of the instrument. For any given time sample, this operator encodes which pixel the instrument is currently observing.

The important difference with respect to the standard pixel-based approach concerns the instrumental characterization. When using pre-pixelized sky maps, it is only possible to change overall instrumental parameters, such as the absolute calibration or the mean noise level. It is not possible to model time-dependent changes in these – but it is precisely such variations in time that creates the problems discussed above. However, with this new time-dependent approach we have a new and unique handle to tackle these issues. For instance, we can now attach a time variable to the calibration, $g_\nu(t)$, and estimate this ourselves internally in the Gibbs chain. Explicitly, step 3 in the above prescription takes over the task of gain estimation, which now may be done with the full sky signal as a calibrator, since that is modeled in step 1, and not just the CMB dipole.

In general, the complexity of the model and the number of free parameters needed will of course depend strongly on the experimental setup. However, the quintessential point is that by considering time-domain observations as input data, as opposed to pixelized sky maps, one gains access to the correct “knobs” required to tune the full system; it is precisely in this domain that the instrument lives its life. And this is therefore in this space possible to model instrumental systematic effects and uncertainties to the necessary level, whether it is propagation of beam asymmetries, corrections for ADC non-linearity, or deconvolution of detector transfer functions.

So far, our discussion has mostly revolved around the importance of removing *biases* from systematic effects. However, once that is done, it is also critical to propagate the corresponding *uncertainties* to final cosmological results. For most experiments, including BICEP2, Planck and WMAP, this is done exclusively through forward Monte Carlo simulations [21], in which one establishes a simple model of the systematics, process these into frequency maps, derive corresponding cosmological parameters, and show (and pray) that the uncertainties due to systematic effects are small compared to instrumental noise. However, this is unlikely to be adequate in the future, simply because the instrumental noise is an increasingly small part of the full error budget. Instead, most of the errors will come from uncertainties in the instrumental characterization and component separation. With the new end-to-end sampling scheme presented in this proposal, this problem will be effectively resolved. Both instrumental and astrophysical uncertainties are seamlessly propagated to the final cosmological results through the Gibbs scheme.

Implementing the analysis pipeline described above from scratch would clearly be far beyond the scope of a 2-year project with a limited budget. The only reason this is possible, is that all individual components already exists among the partners of this project, and the task is primarily one of organization and pipeline, rather than re-implementing and developing algorithms from first principles. For the same reason, we know that the project is computationally very tractable. The final analysis will have a total computational cost bounded above by that of producing a set of end-to-end Planck LFI simulations, a task which we already do at a regular basis.

ii. Step-by-step overview of the analysis pipeline

The previous section focuses on the overall global analysis process from a high-level point of view, within the framework of an end-to-end Gibbs sampler. However, a Gibbs sampler consists of nothing but a series of conditional distribution samplers ("estimators"). In this section we therefore provide a brief breakdown of each of the most important sampling steps, emphasizing how these all tie together. Each step will mirror one Work Package, and is as such described in greater detail in Section 3. Here, though, our main goal is to provide intuition rather than details.

Overall, the work may be described in terms of the following main components:

- **Gibbs sampling infrastructure –** As described above, a Gibbs sampler is nothing but a chain of consecutive conditional sampling steps, and is as such often very simple to implement; the actual work lies in writing the modules that implements each conditional distribution, not in the actual Gibbs sampler. However, in this particular project, which deal with TB data objects, it will be of utmost importance to achieve optimal IO, memory bus bandwidth and parallelization performance. The first step in the development will therefore be to write the raw infrastructure that allows fast low-level data operations, shifting data from one module to another with minimal use of IO. This work will be led by Prof. Hans Kristian Eriksen (Oslo), who has played an internationally leading role in applications of Gibbs sampling for cosmological purposes.

- **Instrument characterization –** The first module that must be implemented is instrument characterization, corresponding to Step 3 in the above Gibbs chain. Within this larger module, there will be many separate sub-modules, but the two first to be implemented are time-variable gain estimation with the CMB dipole as primary calibrator, as well as a noise estimation module, deriving so-called 1/f parameters for each segment of time. This work will be led by Dr. Samuele Galeotta (INAF), who carries a main responsibility for gain estimation and low-level time-ordered data processing in Planck LFI.

- **Map making –** Once the low-level infrastructure is up and running, the most critical step is to implement the sampling step corresponding to frequency map making. We will implement two versions of this, one that employs so-called "destriping" [24], in which case the correlated part of the instrumental noise model is approximated in terms of a small number of simple basis functions, and one that employs so-called "deconvolution map making". The former has been widely used in Planck, and will be useful in particular during code development and debugging due to its faster speed, while the latter is more experimental, and is more critical for small-scale analysis. This work will be led by Dr. Elina Keihänen (Helsinki), the main author of the MADAM code used for the official Planck LFI map making.

- **Component separation –** Having frequency maps ready at hand, the next step is to derive estimates of the astrophysical sky in terms of physical components. This work will be done with the existing Commander implementation, which itself is a fully-functional Gibbs sampler. The computational cost of this step is small compared to the instrumental characterization and map making. Typical output products are shown in Figure 3, which summarizes Commander products from the last Planck data release. This work will be led by Dr. Ingunn Kathrine Wehus (Oslo), who also has been leading this work within the official Planck processing.

- **Power spectrum and cosmological parameter estimation –** With astrophysical component maps in hand, the final step is cosmological interpretation. At its most basic, this process may be thought of as deriving estimates of the angular CMB power spectrum and cosmological parameter estimation from cleaned CMB maps, and provides the connection between data and physics. It thus serves as the ultimate end stage of the entire process. However, within the scope of the current project, it also serves another extremely important purpose, namely as an internal consistency and quality control check. It is only at this very late stage that we get a global overview on whether each sub-step makes sense or not. This work will be led by Drs. Loris Colombo (Milano) and Eirik Gjerløw (XAL), who have been responsible for this type of analyses in Planck LFI for more than 10 years between them.

- **Systematic error assessment –** While the fundamental goal of the current project is to resolve the most important remaining systematic errors in Planck LFI at the level of time-ordered data, there will always be some residuals left in the data at the end of the analysis. To quantify these, it is generally necessary to establish low-level simulations of each associated effect, and process these simulations through the same software as used for the main analysis. For these simulations to be representative of the real world, though, an intimate knowledge of the instrument is required. This work will therefore be led by Prof. Marco Bersanelli (Milano), the Planck LFI Instrument Scientist.

- **Product distribution –** This project aims to provide a community service by delivering the best frequency sky maps and astrophysical component maps published to date. This will be done through the existing Planck Legacy Archive (PLA) infrastructure, which already is committed to serve the public into the foreseeable future. Integration into this infrastructure will be handled by Planetek Hellas, who already has been developing the PLA for two years. In addition, all codes will be tested and documented by Planetek software scientists before distribution under OpenSource licenses.

1.4 Ambition

Already in their current form, the Planck LFI observations represents the state-of-the-art in terms of full-sky microwave observations below the foreground minimum at 70 GHz, and thereby feed broadly into both cosmology and general astrophysics. However, as discussed above, the currently available data products suffers from several known short-comings in the form of astrophysically-induced instrumental systematic errors, and these limit the use of the current data products. Furthermore, while these issues will be described in detail in next official (and final) Planck release, they will – because of time and funding limitations – not be resolved. Our best hope to actually produce final and systematically clean Planck LFI sky maps on all angular scales is the current proposal, which will provide the necessary funding to actually resolve these problems.

It should therefore be obvious that the products from the current project will naturally define a new state-of-the-art once brought to completion. However, as described above, we do not only aim at providing the world's best microwave sky maps between 30 and 70 GHzin this project, but more generally to define a new baseline for how to deal with cosmological data at a very profound level: we aim to transform the very approach with which all future CMB experiments reduce their data, emphasizing the need for global end-to-end analysis with iterative feedback. This issue is not specific for Planck LFI or HFI, but rather concerns any CMB experiment with sufficiently high signal-to-noise ratio: once the sensitivity becomes large enough, the limiting factor will be interplay between instrumental systematics and the astrophysical sky, and then it becomes essential to consider the entire problem jointly. Our analysis will be the world's first analysis of this type, and it will set the standard for all subsequent CMB analyses. Furthermore, the codes we develop will be made publicly available through OpenSource licensing, allowing other users to benefit from our work.

We conclude by noting that while our primary scientific target in this current proposal is the LFI observations, we will in the future be most interested in applying this same framework to other data sets, and in particular the Planck HFI observations. However, a full analysis of that data set, which is considerably more complex than Planck LFI in terms of instrumental uncertainties and overall noise levels, is clearly beyond the scope of the current budget. We leave this analysis for future applications.

[1] https://www.cfa.harvard.edu/news/2014-05
[2] BICEP2 Collaboration 2014, Phys. Rev. Lett. 112, 241101
[3] Planck Collaboration 2016, A&A, 594, A1
[4] C. Bennett et al. 2013, ApJS, 208, 20
[5] CMS Collaboration, 2012, Phys. Lett. B, 716, 30
[6] D. H. Lyth & A. Riotto 1998, Phys. Rept. 314, 1
[7] A. Einstein 1916, Annalen der Physik, 49, 769
[8] LIGO Collaboration 2016, Phys. Rev. Lett., 116, 061102
[9] BICEP2/Keck/Planck Collaborations, 2015, Phys. Rev. Lett. 114, 101301
[10] D. Fixsen, 2009, ApJ, 707, 916
[11] Planck Collaboration 2016, A&A, 594, A9
[12] J. Filipini et al. 2011, SPIE, [1106.2158]
[13] http://litebird.jp/eng
[14] Planck Collaboration 2016, A&A, 594, A8
[15] H. K. Eriksen et al. 2004, ApJS, 155, 227
[16] H. K. Eriksen et al. 2008, ApJ, 676, 10
[17] Planck Collaboration 2016, A&A, 596, A107
[18] Planck Collaboration 2016, A&A, 594, A16
[19] H. K. Eriksen et al. 2004, ApJ, 605, 14
[20] M. Tegmark 1997, ApJL, 480, 87
[21] Planck Collaboration 2016, A&A, 594, A12
[22] O. G. King et al. 2014, MNRAS, 438, 2426
[23] R. Genova-Santos et al. 2015, MNRAS, in press, [1501.04491]
[24] M. A. J. Ashdown et al. 2009, A&A, 493, 753
[25] E. Hivon et al. 2002, ApJ, 567, 2
[26] A. Lewis and S. Bridle 2002, Phys. Rev. D. 66, 103511
[27] B. Racine et al. 2016, ApJ, 820, 31
[28] G. Prezeau et al. 2010, ApJS, 190, 267
[29] J. B. Jewell et al. 2004, ApJ, 609, 1
[30] B. D. Wandelt et al. 2004, Phys. Rev. D, 70, 083511

vii. References

2. Impact

2.1 Expected impacts

According to the H2020 2016-2017 work plan, the expected impact requirements for COMPET-4 2017 read as follows:

> *A higher number of scientific publications based on Europe's space data, high-level data products made available through appropriate archives, and tools developed for the advanced processing of data. Proposals are also expected to add value to existing activities on European and international levels, and to enhance and broaden research partnerships.*

As described above, we propose to derive new legacy data products from Planck LFI data, following closely in the footsteps of the extremely successful Planck mission. The proposal team consists of leading Planck LFI scientists with a strong scientific desire to see this program succeed, and have a major personal commitment to the project. The overall goal of this work is precisely to derive world-leading science results, and publish this in top quality journals. For reference, the Planck collaboration as a whole has produced more than 150 papers based on the full Planck data set. For our project, we foresee at least 9 papers resulting directly from our processing, and many more exploiting our products. The core set of papers will describe at least the following main topics:

1. Overview paper, providing an overview of the entire project (Lead: HKE)
2. End-to-end Gibbs sampling; technical algorithm paper (Lead: HKE)
3. Low-level time-ordered data processing and frequency maps (Leads: AZ, SG, EK)
4. Low-frequency astrophysical foregrounds in temperature and polarization (Lead: IKW)
5. CMB maps and their characterization in temperature and polarization (Leads: IKW, LC)
6. Cosmological parameters (Lead: LC)
7. Statistical isotropy on large angular scales as measured with polarization measurements (Lead: HKE)
8. End-to-end simulations and systematic uncertainties (Lead: MB)
9. Reproducible research (Lead: EG)

In addition to the actual scientific publications, the most long-lasting impact of the project will be the actual re-derived Planck LFI data products, including at least the following:

1. Cleaned time-ordered data with associated flagging, gain and noise estimation
2. Frequency maps clean of instrumental systematics in both temperature and polarization, with and without explicit beam symmetrization
3. New astrophysical sky maps for synchrotron, free-free and spinning dust, as well as a new compact source catalog for sources between 30 and 70 GHz

4. New CMB maps in both temperature and polarization; these maps will be the first full-sky polarization maps with full support to the very largest angular scales, down to a multipole of $l = 2$.

From these products, we will also derive a new estimate of the optical depth of reionization that is robust with respect to both astrophysical foregrounds and instrumental systematics, and we will study the statistical properties of the universe on the very largest angular scales with CMB polarization data. For instance, there have been several claims in the literature regarding potential violations of statistical isotropy on large scales with temperature data, and our maps will be the first that allows a detailed study of these issues with CMB polarization data; until now, such studies has been prohibited by the very same residual instrumental systematics that we aim to resolve.

In addition to producing new state-of-the-art science products, we also aim to transform the way CMB science is done across the field, by providing the first real-world example of true end-to-end data processing. As discussed at length above, we believe this is the only way to break through the current error floor defined by the interaction between uncertainties in the instrumental characterization and the astrophysical sky. The current project will thus serve as a pathfinder for the next-generation CMB polarization experiments, including ESA's CORE concept, JAXA's LiteBIRD mission, and a vast range of ground- and sub-orbital based experiments.

Not only will our project serve as a practical example, but we will also release our codes under Open Source licenses, allowing other experiments to benefit directly from our work. And this brings us to the final major ground-breaking aspect of this project: we will adopt a strong focus on *reproducible research.* In parallel with the dedicated cosmologists and astrophysicists working on the data processing, we will have another dedicated team of computing scientists, and therefore “non-experts” with respect to cosmology, who will both replicate the analyses internally, and be directly in charge of communication with external users. Adding this extra “non-expert layer” between the astrophysicists and the users will maximize the external usability of our products, because all implicit and internal assumptions must by construction be made explicit and documented before final delivery. A total of 20'% of the requested budget is dedicated to these aspects of the project. We believe that the success of this approach may play a transformational role in the field, and become a standard requirement for future projects of this type.

We consider any risks regarding the extent to which each of the primary expected impacts will be achieved to be small. The members of our team have long experience working with the Planck LFI data, and we understand very well the challenges that are involved. The sole reason this work has not already been undertaken is exclusively one of limited time and resources; the true underlying explanation for the systematic errors we see in the data did not become clear until about 6-12 months ago, after which it was too late to change the analysis pipeline within the official Planck collaboration, which is scheduled to conclude its work in the first half of 2017. The current proposal aims to rectify this problem by providing the necessary time and funds to complete the work.

2.2 Measures to maximise impact

a) Dissemination and exploitation of results

In this section we outline a draft *Plan for the Exploitation and Dissemination of Results* (PEDR), which will serve as a useful monitoring and guiding strategic document throughout the project period. However, before presenting the actual PEDR, we discuss the overall philosophy for dissemination and quality control that we adopted for our project, which revolves around the concept of reproducible research.

For many years scientists have been using electronic tools to collect, manipulate and publish their results. With the increase in high quality and easy-to-use collaboration and computational tools comes the advantage of creating highly customizable research workflows that unfortunately, for most of the times, remain undocumented. The output of the research work will be abstracted, introduced, analyzed, defended and cited in a very detailed way in the final research paper, but the scientific and computational workflow will often be omitted, or in the best case barely mentioned in a footnote.

There are two basic reasons why a reproducible scientific workflow is of benefit to both the producers and the consumers of the scientific workflow. The first is to provide an easy way to validate the correctness of the scientific results. Descriptions and verbal explanations of the processes used to produce results, are rarely sufficient to convince the sceptical readers of such work. Having a well documented and reproducible path that anyone that is curious or has doubts about the results in the papers can retrace on their own, will increase the validity and trustworthiness of the work that the authors present.

The second reason is to make the presented research more approachable and extendable for other scientists, allowing them to easily reuse existing methods and results and take the given research a little further. Having a way to reuse or extend the given workflow, the research can be easily applied to analyze a similar or complementary area of scientific research.

We will address these issues by defining a dedicated work package for reproducible research and user communication. The goal of the this work package will be to align the scientific research performed in the scientific Work Packages with the five major categories of the computational reproducibility spectrum as defined in the ICERM workshop "Reproducibility in Computational and Experimental Mathematics", which are

1. **Reviewable Research**. The descriptions of the research methods can be independently assessed and the results judged credible. (This includes both traditional peer review and community review, and does not necessarily imply reproducibility.)
2. **Replicable Research**. Tools are made available that would allow one to duplicate the results of the research, for example by running the author's' code to produce the plots shown in the publication or recreating the output files from the research performed. (Here tools might be limited in scope, e.g., only essential data or executables, and might only be made available only upon request.)
3. **Confirmable Research**. The main conclusions of the research can be attained independently without the use of software provided by the author. (But using the complete description of algorithms and methodology provided in the publication and any supplementary materials.)

4. **Auditable Research**. Sufficient records (including data and software) have been archived so that the research can be defended later if necessary or differences between independent confirmations resolved. The archive might be private, as with traditional laboratory notebooks.

5. **Open or Reproducible Research**. Auditable research made openly available. This comprised well-documented and fully open code and data that are publicly available that would allow one to (a) fully audit the computational procedure, (b) replicate and also independently reproduce the results of the research, and (c) extend the results or apply the method to new problems.

To realize this potential, the actual work of this work package will include at least the following items:

- Work closely with the scientists of the scientific work packages to identify the list of existing workflows and tools they are using to perform their work.
- Collate the results from the collected data and define the best tool from each area to use in a proposed reproducible pipeline.
- Attempt to implement an online service that provides instructions and hooks to external services that fulfill the five major categories of the Reproducibility motto and capture the workflow of each scientific package of this proposal
- Assume the role of a third party sceptical observer and recreate the scientific work produced by the scientific work packages

Of course, there are some risks associated with this type of work, such as

- inability to make parts of the computational efforts reproducible. Due to the fact that some of the computational efforts require distributed or computational grids in order to achieve acceptable execution times, finding alternative execution environments might be prohibitively expensive.
- the resulting scientific output, might be extremely large in terms of data storage (potentially in the order of Terabytes). This will make sharing and transferring these files a difficult procedure.

Even if these risks are experienced in the progress of the project and although we might not be able to meet all the original goals of the project, still a majority of the workflow should be able to still be carried out at the cost of small gaps in the final automated pipeline.

Finally, WP9 will also devote some of its effort to a more experimental approach, aiming to replace some of the expensive parts of the calculations with faster and cheaper GPU calculations, with the goal of allowing external users to reproduce our results with cheaper computational equipment. We emphasize, however, that this component is considered experimental, and does not lie in the critical path of the project as a whole.

Plan for the Exploitation and Dissemination of Results

What kind of needs does the project respond to?

The aim of this project is to provide to the cosmological community state-of-the-art measurements and models of the microwave sky between 30 and 70 GHz. These types of measurements form a cornerstone in modern cosmology, and have during the last decades allowed us to probe the origin and evolution of the universe to unprecedented accuracy, from a fraction of a second after the Big Bang until today.

The next frontier in observational cosmology is the search for inflationary gravitational waves. According to current theories, these may be found in the large angular scale perturbations of the CMB polarization field. However, the amplitude of the signal is minute, possibly several orders of magnitude smaller than obscuring radiation from synchrotron and thermal dust radiation in the Milky Way, and interactions between this radiation and the instrumental characterization renders the overall analysis problem non-linear and highly convolved.

In this project, we aim to provide the world's best model of polarized synchrotron emission on large angular scales in the frequency range required for gravitational wave measurements by reanalyzing the state-of-the-art Planck LFI observations in combination with external data sets, such as Planck HFI, WMAP and C-BASS. In the process, we will deliver legacy Planck LFI sky maps and associated products to the cosmological community, and we will demonstrate a new approach to CMB data analysis that we believe will form a reference standard for future experiments.

What kind of problem the proposed solution will solve and why this solution will be better than existing ones?

The single most important outstanding challenge regarding the existing data products are instrumental systematics induced by the presence of astrophysical foregrounds. Such systematics can only be mitigated by considering the entire analysis problem globally, fitting both astrophysical and instrumental parameters jointly. This has never been attempted to this date, due to lack of expertise, funding and time. Exploiting the unique expertise represented within the team proposing this project, and leveraging the resources made available through the current proposal, we will be the first group to succeed in this task.

What new knowledge and results the project will generate?

The project will generate new knowledge on many fronts:

- In terms of hard data products, we will produce new state-of-the-art maps and models of the microwave sky between 30 and 70 GHz, a critical frequency range for modern cosmology, including both cosmological and astrophysical components.
- From these, we will derive new constraints on cosmological parameters and characteristics from large-scale polarization measurements that are robust with respect to instrumental and astrophysical systematics, including new measurements of the optical depth of reionization as well as limits on statistical anisotropy on large angular scales.
- The novel end-to-end methods we develop will be essential for next-generation CMB B-mode experiments searching for inflationary gravitational waves. Our analysis will serve as a real-world test-bed for experiments such as CORE, LiteBIRD, S4 and many others.

Who will use these results?

The primary users of these results are cosmologists and astrophysicists at the highest international level. These products will serve as standard references for the entire microwave astronomy community until the next CMB satellite mission flies, which may be a decade or longer.

These products will also be used in public outreach and science communication to the public. Maps of the cosmic microwave background and the Milky Way appear everywhere from high-school physics textbooks via popular science, all the way to popular culture; for instance, the CMB temperature field appears regularly in the “Big Bang Theory” sitcom in the form of a beach ball.

What benefits will be delivered and how much benefit?

Because of the minute amplitude of the signals in question here, the main concern of most CMB cosmologists before Planck revolved around random instrumental noise. The emphasis was constantly on producing more and more sensitive maps. However, after the last generation of ultra-sensitive experiments such as Planck and BICEP2, the field has undergone a critical transition: the main challenge is now longer raw sensitivity, but rather astrophysical contamination and instrumental systematics. What we offer in this project is a viable algorithmic path through this problem, and as such open up the road to continued cosmological progress for the coming decade. Ultimately, this may lead to one of the biggest discoveries in modern cosmology, namely the detection of gravitational waves created during the Big Bang.

How will end users be informed about the generated results?

End-users will be informed through a six-level dissemination structure:

1. All main results will be published in international journals with a “Gold Open Access” policy, ensuring free access to everybody.

2. All data products will be published through open long-term data repositories, most notably the Planck Legacy Archive (PLA), funded and maintained by ESA.

3. All software products will be accessible through an open GitHub repository, allowing everybody to download, reuse and extend our work.

4. We will create a dedicated web page for the project as a whole, describing its background, purpose and status, including all necessary references to results, products and software.

5. Near the end of the project period, we will host an international conference dedicated to the discussion and dissemination of these results.

6. Throughout the project period, we will be pro-active in terms of seeking exposure in general media, both in terms of popular science and radio/TV shows.

All results and products will be available indefinitely through their respective dissemination channel. Once published, none of the adopted channels require continued funding from the current project. Furthermore, all data products will be released using industry-standard formats. For instance, all sky maps will be provided in HEALPix FITS format; all time-ordered data products will be released in HDF format; and all tabulated data will be provided in standard ASCII formats. No internal or proprietary formats will be employed, ensuring easy use for external users.

b) Communication activities

As described in the PEDR, the main dissemination strategy will build on six pillars:

1. Peer-reviewed publications in international journals

As is customary in astrophysics, all main results will be published in international journals. The preferred journal for our purposes is *Astronomy and Astrophysics* (A&A), a peer-reviewed top-tier journal adminstrated by the European Southern Observatory (ESO). All publications from the Planck Collaboration have been published in A&A, and we will continue this tradition in the current project, maintaining the same rigorous quality levels and style guidelines as for the official Planck publications. Likewise, all Planck publications were published under a “Gold Access” license, ensuring open access to all external readers, and we will maintain the same policy in the current project.

In addition to A&A, all publications will be posted on the arXiv preprint server, allowing early access to results and ensuring maximal exposure throughout the community. This is where most cosmologists read new results, and we will take advantage of this well-established structure.

2. Data deliveries through the Planck Legacy Archive

The Planck Legacy Archive (PLA; https://pla.esac.esa.int/pla) was created as part of the first Planck data release, providing direct access to all main data products. It is funded by ESA, and will serve the cosmological community into the foreseeable future. We will take advantage of this existing infrastructure, and feed our final results back into the PLA near the end of the project period. Two of our project partners, Planetek and XAL, have been working on the PLA for the last two years, and are intimately familiar with its organization and structure.

3. Direct software access through open GitHub accounts

By now, GitHub has become an effective industry standard in terms of software version control for large-scale projects, combining development and distribution into one convenient infrastructure. We will adopt this approach in our project, maintaining a semi-open software repository for all codes used. Specifically, all internal users will have full push and pull (write and read) rights, whereas external users will only have pull (read) rights. In addition to the continuously accessible software repository, we will provide a self-contained software package with necessary scripts (Makefiles, autoconf files etc.) to allow for easy compilation on most systems.

4. Dedicated project web page

We will create a dedicated web page for the project (hosted by the University of Oslo to ensure long-term support), providing easy access and overview of the project to external users. This web page will include both a general introduction to cosmology for the general public, as well as detailed information regarding data and software access required for scientific purposes. It will also provide direct links to all associated repositories.

5. International conference hosted at the end of the project

To ensure effective distribution of our results, we will host a major international conference dedicated to the dissemination and use of the derived results. This will follow in the footsteps of previous Planck release conferences, and the invited speakers will comprise a mix of internal researchers and external experts in the field. The exact time, location and duration will be decided at a later time, but one possible scenario is a one-week conference in Norway in late 2019 or early 2020.

6. Public outreach through general media

Cosmology and astrophysics are generally very popular topics in the general media, and we will exploit this interest to present our results to the public. As a simple example of such efforts, HKE has during the last year acted as an effective “house astronomer” for one of the biggest radio shows in Norway (“Norgesglasset”, P1, NRK), appearing often several times a month to discuss recent astronomical developments. The current project will be of great interest to such audiences. Other examples of such activities include popular science articles, general news articles and giving high-school lectures.

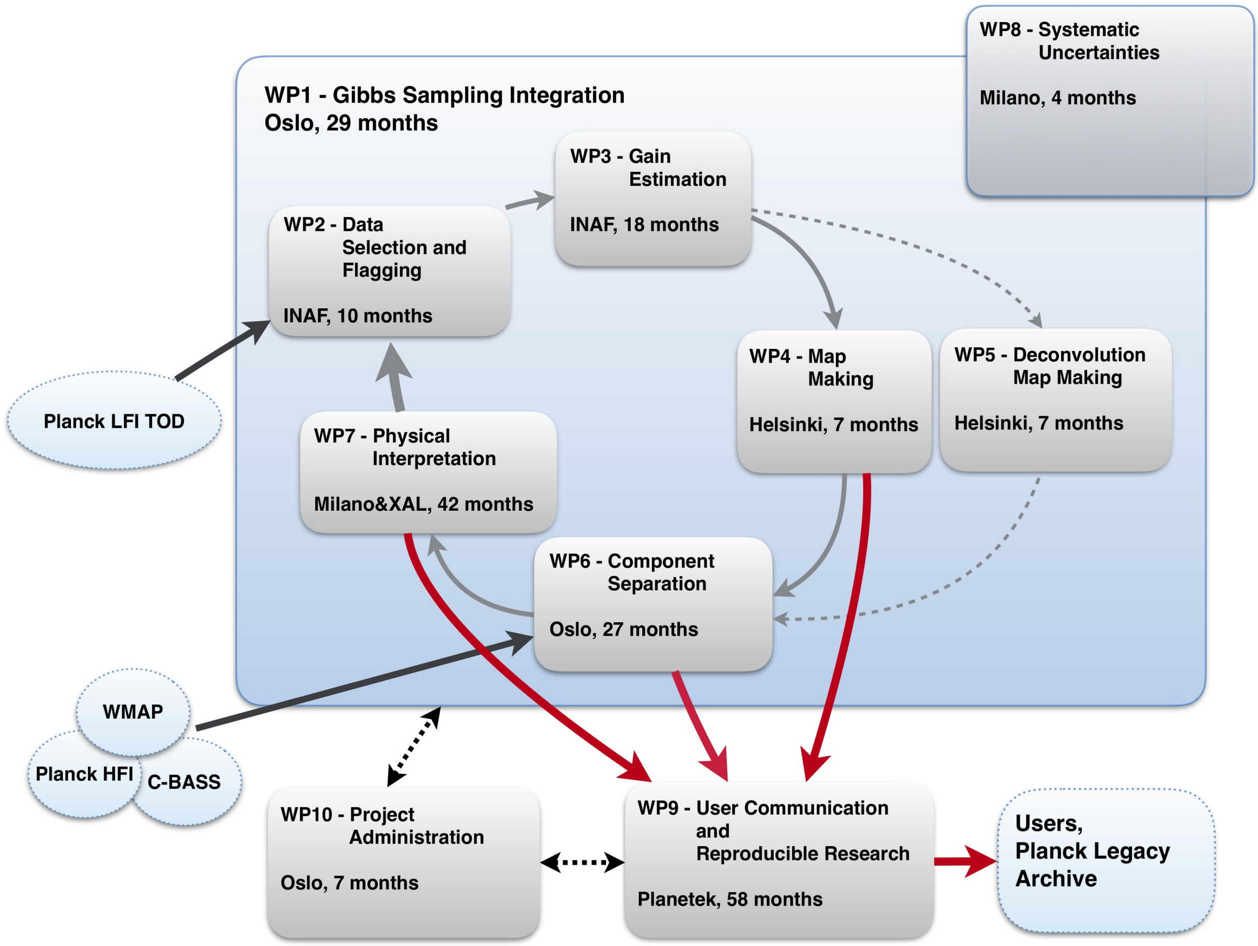

Figure 6: Schematic overview of work packages and their inter-dependencies. In this diagram "Months" correspond to person-months, not WP duration.

3. Implementation

3.1 Work plan — Work packages, deliverables

We now turn our attention to the implementation and organization of the project, and start with the overall work plan as defined in terms of work packages (WPs).

As described in Section 1b, the core of the current project revolves around developing a single end-to-end analysis framework that converts raw bits from the Planck LFI instrument into final science products in the form of sky maps, cosmological parameters and astrophysical models. This overall machine may be implemented in terms of a Gibbs sampler, and is as such an intrinsically iterative approach.

The work flow corresponding to this approach is schematically illustrated in Figure 6 in terms of WPs and their inter-relationships. Going through this diagram in chronological order, we start the process by inserting observations from Planck and external experiments (ellipses on the left-hand side in Figure 6) into the overall Gibbs sampler (WP1). The raw Planck LFI time-ordered data are first processed through a Data Selection module (WP2) that identifies and removes (ie., flags) any bad data segments from further processing. Note that this process can be made significantly more accurate and efficient by using prior information regarding what the sky actually does look like.

The second step in the processing is Gain Estimation (WP3), which translates measured voltages to astrophysically relevant sky temperatures. The main calibrators for this purpose are the Doppler-induced CMB dipole due to the Earth's motion around the Sun, and the corresponding dipole due to the Sun's motion with respect to the background. However, in order to reach the necessary sub-µK precision level required for this project, it is essential to also account for radiation from the Milky Way.

Having clean and calibrated data at hand, the third step is to produce clean frequency maps through Map Making (WP4/5). This corresponds essentially to performing a weighted average for each sky pixel over all samples falling within that pixel. We will implement two different approaches for this purpose, namely standard destriping, as implemented in the state-of-the-art MADAM code (WP4), but also so-called beam-deconvolved map making (WP5). Since the instantanous instrumental response function, often called “point-spread function” or “beam”, of a CMB detector is generally azimuthally asymmetric, the measured output from the detector depends on the specific orientation of the detector with respect to the sky. This effect can, however, be accounted for by taking into account the known beam response profile and the detector orientation through so-called beam deconvolution. This approach results in maps with simpler signal properties than standard maps, but more complicated noise properties. As a result, the optimal method depends on the particular application the maps will be used for. We will implement both, but note that the standard destriper approach will be our main algorithm of choice.

The fourth step is Component Separation (WP6), and for this we will adopt the well-established Bayesian Commander framework discussed above. This WP will be responsible for deriving astrophysical component maps (CMB, synchrotron, free-free, spinning and thermal dust emission etc.) from the frequency maps, as well as combining the Planck LFI maps with external data, whether they come from WMAP, C-BASS, Planck HFI or any other source.

Fifth and finally, the analysis chain is completed by Physical Interpretation (WP7). This WP derives cosmological and astrophysical science from the component maps generated in WP6, and as such serves as the highest-level analysis step in the pipeline. Products from this WP includes, but are not limited to, angular CMB power spectra and cosmological parameters.

So far, the analysis pipeline outlined above follows a very conventional procedure, with a linear progression from raw data to final science products. However, the fundamentally new step in our procedure is to close the loop, indicated by a bold arrow from WP7 back to WP1, indicating that this will indeed by a circular and iterative process, as opposed to the standard one-shot linear process. The main task of WP1 (Gibbs sampling) is to provide the computational infrastructure that binds this entire process together.

In addition to the main Gibbs-related WPs, we also define two external special-purpose WPs. The first of these is called Systematic Errors (WP8), which carries the responsibility of understanding and quantifying residual systematic errors in the final data products after full processing. This work will partly take place inside the main infrastructure and partly outside. Generally speaking, this work will revolve around understanding the behaviour of the instrument in light of the most up-to-date models, and determine how and why they differ. Thus, the typical mode of operation will be to establish an imperfect model of a given instrumental effect; project this into time-ordered data; process those data through the pipeline; and quantify the residuals.

The ninth WP is called User communication and reproducible research (WP9), and is responsible for all aspects of user communication. As already discussed extensively in Section 2.2, we adopt an ambitious philosophy based on reproducible research concepts, ensuring that all deliveries are well understood, documented and accessible by external users. This work package is responsible also for delivery of all products to external repositories, including to the Planck Legacy Archive.

Finally, the 10^{th} and last WP covers all administrative and non-scientific aspects of the project, including budgeting, audits, meeting organization etc.

In Table 3.1a we provide a detailed description of each WP, while Table 3.1b gives an overview of all WPs. Table 3.1c lists all deliverables, and, finally, in Figure 7 we provide a Gantt chart summarizing the timeline of the project.

Table 3.1a: Work package descriptions

Work package number	**1**		**Lead Beneficiary**				**Oslo**
Work package title	Gibbs sampling integration						
Participant number	1	3					
Short name of participant	Oslo	INAF					
Person months per participant:	27	2					
Start month	1			**End month**	24		

Objectives This WP serves as the framework that provides contact between the various sub-modules, feeding data and partial results from one operation to the other. For now, this task is performed by human interaction, limiting the number of iterations to a handful. The main objective of this WP is to automate and streamline this process, such that hundreds of iterations may be run completely without human intervention. After cross-module integration, this work package will also be responsible for delivering the final joint products.

Description of work The work will consist of two phases, namely infrastructure construction and module population. In the first phase, we will implement a placeholder pipeline that is able to input raw time-ordered data, pass them on to simplified sub-modules, and then output final maps and deliveries to disk. This process will be coordinated by HKE (with responsibility for the Gibbs sampling aspects of the work) and SG (with responsibility for the TOD access aspects), while most of the work will be done by PD1; we foresee this work to require continuous attention for the first 12 months of the project. In the second phase, each of the sub-modules will be replaced with fully functional modules; these include gain estimation and data flagging (SG), map making (EK, PD3), component separation (IKW, PD2) and cosmological interpretation (LC).

Deliverables The deliverables from the WP are:

- A prototype pipeline, useful for allowing other partners to start their work; due in Month 3
- A fully functional pipeline, useful for full analysis; due in Month 12.
- Final scientific data products, due in month 21.

Work package number	**2**		**Lead beneficiary**			**INAF**	
Work package title	Data Selection and Flagging						
Participant number	3						
Short name of participant	INAF						
Person months per participant:	10						
Start month	1			**End month**	6		

Objectives This WP serves as starting point to select which data, at the timelines level, will be used in our analysis. It will, based on pre-defined criteria, flag the data that should be excluded like manoeuvre period, gain changes in the data acquisition electronics that cause saturation, abrupt changes in voltage outputs caused by gain fluctuation (to determined using cross correlation with House Keeping) etc . The main objective of this WP is to give to the subsequent WPs a more instrument effect clean data input.

Description of work The work will consist of two phases. In the first phase we will define the requirements to be used to exclude (flag) the data (roughly estimation is to at least be able to use 80% of the data) and verify if the manoeuvre period (about 8%) can be used for scientific exploitation. This process will be coordinated by AZ and MM (requirements collection) and should be concluded in 1 months. The second phase will consist in the application of the requirements in a semi-automatic pipeline coordinated by SG and DT with the goal to made available a first guest after Month 3. Feedback from Science analysis will tune the requirement and changes in the process with the scope to release the final flagged timelines at Month 6.

Deliverables The deliverables from the WP are:

- Prototype version of data flagging module, due in Month 3
- Final version of data flagging module, due in Month 6

Work package number	**3**		**Lead beneficiary**		**INAF**		
Work package title	Gain Estimation Module						
Participant number	3	1					
Short name of participant	INAF	Oslo					
Person months per participant:	17	1					
Start month	4			**End month**	15		

Objectives This WP serves as the framework that provides conversion of time-ordered streams of voltages into time-ordered streams of thermodynamic temperatures with all the know instrument systematics removed. This WP will be then integrated in an iterative pipeline as described in the Gibbs sampling integration WP.

Description of work The work can be split in two phases. The first one is the collection of all the elements required for the thermodynamic calibration like beam model and Cosmological dipole signal (MM), foreground model to be used for initialization (IKW), and identification of the main instrument systematics to be removed at timelines level (DT). The second step will be the creation of a stand-alone pipeline that from raw time order data using the previous described input will be able to calibrate in thermodynamic unit the Planck LFI timeline. This pipeline will be implemented in dedicated and self-contained modules to allow a quick integration in an iterative procedure. This work will be coordinated by AZ and SG, and realized by SG and GM.

Deliverables The deliverables from the WP are:

- A prototype gain module estimation, consequently calibrated timeline, useful for starting the iterative foreground removal, is due in Month 9
- Fully functional pipeline and final calibrated timelines will be released at Month 12 .

Work package number	**4**	**Lead beneficiary**					**Helsinki**
Work package title	Map making module						
Participant number	4						
Short name of participant	Helsinki						
Person months per participant:	7						
Start month	4		**End month**	18			

Objectives

Integrate the Madam mapmaking code into the pipeline.

Description of work

Map-making step if performed with the Madam mapmaker code (Keihänen et al, A&A 510:A27 (2010)). The code is already fully functional and is the main map-making tool of LFI DPC. The work package consists of integrating the code with the pipeline in such a way that it can be run repeatedly as part of the automatized pipeline, testing and validation.

The mapmaking module is run at two phases. As input the code takes the calibrated timelines. When run as part of the Gibbs sampler loop, the main output consists of timelines that that cleaned of correlated noise, and which are then used as input for the next iteration. At the end of the loop, map-making module will construct sky maps of the final timelines.

The work is done by an unnamed Postdoc NN under the supervision of EK.

Deliverables The deliverables from the WP are:

- Prototype MADAM module integrated in Gibbs sampler, due in Month 6
- Tuned MADAM module integrated in Gibbs sampler, due in Month 12

Work package number	**5**	**Lead beneficiary**			**Helsinki**		
Work package title	Beam deconvolution map making						
Participant number	4						
Short name of participant	Helsinki						
Person months per participant:	7						
Start month	10		**End month**	24			

Objectives

Production of beam-deconvolved maps.

Integration of the necessary tools with the pipeline.

Description of work

Use the ArtDeco deconvolver code (Keihänen and Reinecke, A&A 548, (2012)), to produce beam-deconvolved maps, where the effective beam is symmetric, and leakage from temperature to polarization due to asymmetric beams has been eliminated. The exact role of the deconvolved maps in the pipeline is yet to be determined. At minimum the are used for cross-check and validation purposes. The mentioned code package involves tools related to forward beam convolution, and associated analysis. These are integrated with the pipeline as need arises.

The work is done by a yet unnamed Postdoc under the supervision of EK.

Deliverables The deliverables from the WP are:

- Beam deconvolution map maker module in Gibbs sampler, due in Month 15
- Associated low-resolution noise covariance matrix module, due in Month 18

Work package number	**6**		**Lead beneficiary**		**Oslo**		
Work package title	Component separation						
Participant number	1						
Short name of participant	Oslo						
Person months per participant:	27						
Start month	1		**End month**	24			

Objectives Produce astrophysical component maps from frequency maps

Description of work The work in this WP may be split into two main components, namely code development and data analysis. The first stage will consist of integrating the existing Commander component separation code into the new and larger Gibbs sampler. This will mostly involve interfacing existing routines with the new infrastructure. The second part will be to actually run the code, improving existing foreground models. This work package will also be responsible for integration and analysis of external data, including WMAP, C-BASS and Planck HFI. This work will be carried out by a postdoc under the supervision of IKW.

Deliverables The deliverables from the WP are:

- First iteration of astrophysical sky maps, needed for initialization, due in Month 3. Based on existing Commander sky model
- Modularized Commander code, suitable for insertion into main Gibbs sampler; due in Month 3
- First end-to-end astrophysical sky maps (CMB, synchrotron, free-free and spinning dust in temperature, and CMB and synchrotron in polarization) from new Gibbs sampler, due in Month 12
- Final release candidate maps, due in Month 21

Work package number	**7**		**Lead beneficiary**			**Milano**	
Work package title	Physical interpretation						
Participant number	2	6					
Short name of participant	Milano	XAL					
Person months per participant:	24	18					
Start month	1			**End month**	24		

Objectives The current science analysis of Planck maps is limited by an imperfect understanding of instrumental systematics and a simplified component separation approach. This WP aims at two main objectives: 1) development of tools for the likelihood analysis of the improved maps, which are able to fully propagate the uncertainties from residual instrumental systematics and component separation all the way to the cosmological parameters, 2) perform the final cosmological analysis of the maps, either alone or in conjunction of other cosmological datasets.

Description of work This WP consists of two main phases, tools development and science exploitation. The first phase involves the development of a likelihood module to be integrated in the main Gibbs pipeline. This module will take maps (either single frequency or component separated) as inputs and produce power spectra and posterior distributions for the cosmological parameters of interest. This activity will start immediately, with the goal of producing a final module for the full pipeline integration in Month 12. Once the full pipeline is completed, the focus of the WP will shift to the scientific characterization of the maps, and the production of the final science results. Given the nature of LFI data, the focus will be on the analysis of large angle features, while the analysis of the small angle features will be used mainly as a consistency check. Preliminary results are expected after 6 months of analysis, with final results completed by Month 21, leaving 3 months for the preparation of the related scientific publications. The main responsibility for this WP will be with LC, with support from EG.

Deliverables Deliverables from the WP are:

- Cosmological interpretation module for integration in the main pipeline (Month 12);
- Scientific characterization of maps, including power spectra and cosmological parameter constraints (Month 21);
- Scientific papers for publication in peer-review journals (Month 24).

Work package number	**8**	**Lead beneficiary**				**Milano**
Work package title	Systematic Uncertainties					
Participant number	3					
Short name of participant	Milano					
Person months per participant:	4					
Start month	1		**End month**	24		

Objectives

This WP covers the analysis of systematic effects affecting the LFI maps produced in the new pipeline developed in this project.

Description of work

Our new pipeline will improve the quality and self-consistency of the Planck/LFI polarization maps, with minimal contamination from instrumental and astrophysical effects. However, deviations from the ideal map reconstruction will always be present at some level, due to a combination of residual foreground radiation and systematic effects. In the proposed optimisation process, it will be crucial to understand the origin of remaining imperfections, identify their source, and control the convergence of the iteration process. Building on our 25-years-experience of design, development, testing, flight operation and data analysis of the Planck/LFI instrument, our group at Milano University will support the interpretation of residual instrumental effects in the LFI map produced by the pipeline.

The work will be split in the following tasks:

- Methodical inspection of different generations of maps produced by the pipeline, for different versions and at increasing iteration levels
- Develop and implement an optimal scheme for null-test analysis to highlight residual anomalies in the data
- Interpretation of the source of systematic effects in terms single instrumental features (such as stray light from far sidelobes, bandpass effects, ADC non-linearity, spikes in time domain, thermal effects, etc.)
- Run simulations, where necessary, to confirm the above interpretation
- Compare Planck/LFI maps with external data, in particular with WMAP maps and Planck-HFI maps

Deliverables Deliverables from the WP are all due in Month 21, and include:

- Quantitative systematic error estimates on all main parameters
- Residual error simulations (sky maps) processed through the pipeline

Work package number	**9**	**Lead beneficiary**					**Planetek**
Work package title	User communication and reproducible research						
Participant number	5						
Short name of participant	Planetek						
Person months per participant:	58						
Start month	1		**End month**	24			

Objectives The main objectives of this WP are as follows:

- Disseminate the results of the rest of the work packages
- Investigate how the scientific work performed in the previous work packages can be reproducible.
- Implement a system that will be able to capture and recreate the scientific operations performed in this project.

Description of work

This work package can be sub-divided into three main classes. The first class concerns code organization and distribution. According to the OpenSource philosophy of this project, all source codes will be made publicly available through a GitRepository. The main work in this class is to set up and maintain this repository throughout the project period, and develop a user-friendly webpage that explains the purpose and functionality of the project. The second and most work-intensive class concerns *reproducible research*. In order to ensure that external users will be able to access, reuse and reproduce the codes developed in the project, Planetek software scientists (who are not themselves cosmologists) will run the codes externally, as if they were external users, and they will be in charge of developing suitable documentation. By having non-cosmologists performing this work, the end-products will be far more user-friendly. For full discussion of the topic of reproducible research, see Section 2.2a. Third and finally, this WP is responsible for integrating with and delivering final products through the Planck Legacy Archive. With the aim to facilitate reproducibility of parallel algorithms, this WP will also investigate techniques to replace expensive computing grid calculcations with low cost local or remote GPU based environment by converting suitable code into low level representation for GPU execution.

Deliverables Deliverables from this WP include:

- A skeleton GitHub repository for internal use (Month 1)
- First public project web page (Month 6)
- Report on the methodologies, tools and input files required for the completion of each scientific WP (Month 6)
- Report on the selected tools and libraries to be used in the implementation of the Reproducibility Framework (Month 12)
- Reproducibility Framework tool and documentation (Month 24)
- Product delivery to Planck Legacy Archive (Month 24)

Work package number	**10**	**Lead beneficiary**				**Oslo**
Work package title	Administration					
Participant number	1	2	3	4	5	6
Short name of participant	Oslo	Milano	INAF	Helsinki	Planetek	XAL
Person months per participant:	2	1	1	1	1	1
Start month	1			**End month**	24	

Objectives This WP covers all non-scientific administrative aspects of the project.

Description of work The work in this WP includes meetings with the H2020 organization; providing periodic reports; internal budgeting; initializing audits; organizing work meetings and release conference etc. This work will be led by HKE in close collaboration with the EU team at the Faculty of Mathematics and Natural Sciences, University of Oslo,

Deliverables The main deliverables from the WP are:

- PEDR, due in Month 6
- Periodic reports
- Final report, due in Month 24

Table 3.1b: List of work packages

Work package No	Work Package Title	Lead Participant No	Lead Participant Short Name	Person-Months	Start Month	End month
1	**Gibbs sampling integration**	1	Oslo	29	1	24
2	**Data Selection and Flagging**	3	INAF	10	1	6
3	**Gain Estimation**	3	INAF	18	4	19
4	**Map making**	4	Helsinki	7	4	18
5	**Deconvolution map making**	4	Helsinki	7	10	24
6	**Component separation**	1	Oslo	27	1	24
7	**Physical Interpretation**	2	Milano	42	1	24
8	**Systematic Uncertainties**	2	Milano	4	1	24
9	**User communication and public communication**	5	Planetek	58	1	24
10	**Administration**	1	Oslo	7	1	24
				209		

Table 3.1c: List of Deliverables

Deliverable (number)	Deliverable name	Work package number	Short name of lead participant	Type	Dissemination level	Delivery date (in months)
9.1	GitHub repository	9	Planetek	OTHER	PU	1
1.1	Proto Gibbs sampler	1	Oslo	DEM	PU	3
2.1	Prototype flagging module	2	INAF	DEM	PU	3
6.1	Initial astrophysical sky model	6	Oslo	DEM	PU	3
6.2	Commander module	6	Oslo	OTHER	PU	3
2.2	Final flagging module	2	INAF	OTHER	PU	6
4.1	Prototype MADAM module	4	Helsinki	DEM	PU	6
10.1	Data management plan	10	Oslo	R	PU	6
9.2	Project web page	9	Planetek	DEC	PU	6
9.3	WP methodology report	9	Planetek	R	PU	6
3.1	Protoype gain estimation module	3	INAF	DEM	PU	9
3.2	Final gain estimation module	3	INAF	OTHER	PU	12
4.2	Tuned MADAM module	4	Helsinki	OTHER	PU	12
1.2	Operational Gibbs sampler	1	Oslo	OTHER	PU	12
6.3	First astrophysical sky model products	1	Oslo	DEM	PU	12
7.1	Cosmological interpretation module	2	Milano	OTHER	PU	12
9.3	Reproducible research methodology report	9	Planetek	R	PU	12
5.1	Beam-deconvolved map maker module	4	Helsinki	OTHER	PU	15
5.2	Low-resolution noise covariance matrix module	4	Helsinki	OTHER	PU	18

Deliverable (number)	Deliverable name	Work package number	Short name of lead participant	Type	Dissemination level	Delivery date (in months)
1.3	Final Gibbs/posterior products	1	Oslo	OTHER	PU	21
6.4	Final astrophysical sky model products	1	Oslo	OTHER	PU	21
7.2	Scientific characterization	2	Milano	OTHER	PU	21
8.1	Systematic error assessment	2	Milano	OTHER	PU	21
8.1	Systematic error simulations	2	Milano	OTHER	PU	21
9.4	Reproducibilty framework tool	9	Planetek	OTHER	PU	24
9.5	Product delivery	9	Planetek	OTHER	PU	24
7.3	Paper release	2	Milano	R	PU	24
10.2	Final report	10	Oslo	R	PU	24

Type:

R: Document, report (excluding the periodic and final reports)
DEM: Demonstrator, pilot, prototype, plan designs
DEC: Websites, patents filing, press & media actions, videos, etc.
OTHER: Software, technical diagram, etc.

Dissemination level:

PU = Public, fully open, e.g. web
CO = Confidential, restricted under conditions set out in Model Grant Agreement
CI = Classified, information as referred to in Commission Decision 2001/844/EC.

Delivery date

Measured in months from the project start date (month 1)

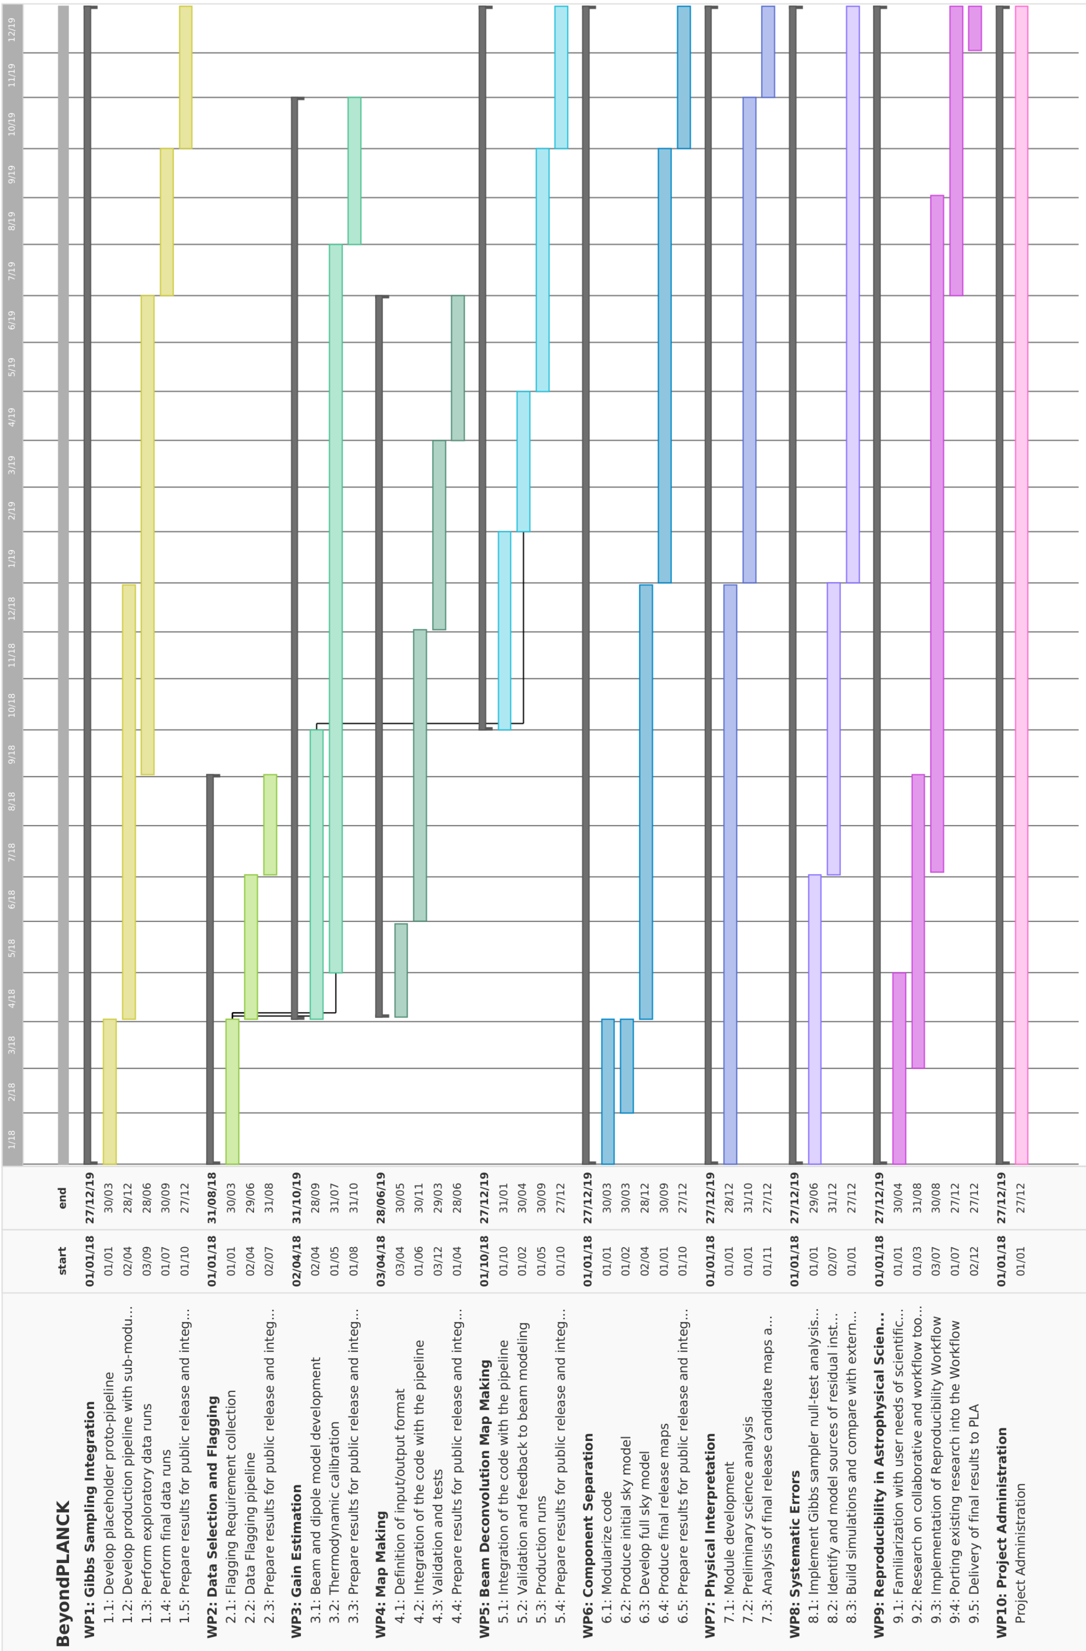

Figure 7: Project time line, indicating the start and end times for each WP.

3.2 Management structure, milestones and procedures

To understand the formal management structure of the project, it is useful to first review the origins of the underlying collaboration. This project derives directly from the efforts of the existing Planck LFI collaboration, which has already existed and operated for more than two decades. As such, the overall management routes and decision-making structures of this project are well established, in that all key players know each other very well through many years of close and productive collaboration.

The existing Planck LFI collaboration comprises more than 150 scientists distributed across 13 countries, spanning a wide range of research interests and fields. It is a big and complex organization, sometimes with long turn-over times, often taking many months from a decision has been made until it is implemented in practice. We do not aim to reproduce this structure in the current project, but rather aim for a much smaller and more tuned approach. Specifically, the philosophy adopted when defining the current project team has been to set up the smallest possible team that still covers all critical topics. Furthermore, recognized leaders within the collaboration (as measured in terms of production and proven delivery, not formal position or status) have been identified and invited to lead the respective working package. This definition process has led to the following leadership team:

- WP1 – Gibbs sampling integration: Hans Kristian Eriksen (Oslo; project coordinator)
- WP2 – Data selection: Samuele Galeotta/Andrea Zacchei (INAF)
- WP3 – Gain estimation: Samuele Galeotta/Andrea Zacchei (INAF)
- WP4 – Map making: Elina Keihänen (Helsinki)
- WP5 – Deconvolution map making: Elina Keihänen (Helsinki)
- WP6 – Component separation: Ingunn Kathrine Wehus (Oslo)
- WP7 – Physical interpretation: Loris Colombo (Milano)
- WP8 – Systematic errors: Marco Bersanelli (Milano)
- WP9 – User communication: Efstratios Gerakakis (Planetek)
- WP10 – Administration Hans Kristian Eriksen (Oslo)

Within this leadership group, and within the collaboration as a whole, we will adopt an overall a flat and “democratic” decision making structure. Thus, while HKE formally serves as the project coordinator and formal leader, the entire team will participate in all major decisions. Of course, the reason this is possible is precisely because the members know each other very well already after years of working together, and have established excellent working relations, both professionally and personally. Operationally, the main inter-project communication channels will be weekly telecons and frequent interaction through email lists. Both will be open to everybody in the collaboration, and all relevant information will be distributed widely.

In short, we will base the management and organization structure on those already operating in Planck, but we will improve on them in terms of transparency and openness. With such a small team as ours, where everybody knows everybody, and each knows what her or his tasks are,, there is no need for multiple hierarchical communication layers. Complex communication structures are more likely to introduce misunderstandings than avoid them.

With this background in mind, we now turn to the formal aspects of the management structure.

3.2.1 Formal management structure

The consortium comprises 6 organizations, and is coordinated by the University of Oslo and Prof. Hans Kristian Eriksen. The consortium partners agrees to the management structure outlined below. This structure as well as the decision-making procedures is established according to the following objectives:

- consideration of the equality and collective responsibility of all participants,
- efficiency and transparency of the overall management,
- ensuring of compliance with all relevant regulations of the European Commission (EC),
- realization of sound monitoring and professional administration to avoid time and cost escalation,
- realization of effective quality management and respect time schedule, including milestones and deliverables,

The consortium is fully aware that management activities are extremely important for the successful implementation of the project as well as for a transparent accountability of the EC contribution. Fully accomplishing the rationale of Horizon 2020 and its focus towards innovation fully integrated with research, the management structure has been designed to boost the process of foreground transfer, especially by enabling and fostering the transfer of complementary expertise among partners. A pragmatic organizational structure has been adopted, with attention paid to both the clear distribution of responsibilities and the organized flow of information. The governance structure of the collaboration has three common organizational levels:

1. The project board consists of one representative from each project partner and WP. It is the project's highest authority and acts as the central forum for discussing the project and making decisions on its status and progress.
2. The coordination level oversees the project's overall progress and coordinates its activities and the interaction between partners, and between the project and the EC. This includes the coordinator, who is responsible for the overall coordination of the project's scientific aspects.
3. The Coordination Committee Team (CCT) consists of all work leaders, and forms the main operational level of the project. It reports to the coordinator.

The tasks of the coordination team include, but are not limited to:

1. to strategically guide the project choices, including the orientation and focusing of different scientific tasks, including their deviations from the work plan, when necessary,
2. to ensure interdisciplinary exchange and continuous flow of information between the WPs; and to coordinate the relations between the WP Leaders and management of the interdependencies between various tasks;
3. to coordinate the conduct of project activities and the implementation of the work plan, including the dissemination activities and the interactions of the activities among the Work-packages;
4. to review, enforce and monitor continuously the implementation of the activity scheduling, included the production of the deliverable and achievement of milestones;
5. to supervise the organisation and implementation of the scientific and management events planned throughout the project lifetime;
6. to establish and maintain a complete record of the material produced by the project (working documents and internal reports, workshop presentations and proceedings, deliverables,

progress and management reports, etc.), to be posted on the project website and regularly updated.

7. monitoring work progress and assessing the fulfillment of tasks and deliverables;
8. establishing procedures to ensure that the partners work is performed according to the work plan;
9. evaluating scientific results achieved by WT leaders;
10. promoting integration among the consortium different expertise;
11. organizing, if necessary, site visits, extra meetings, or workshop within the WP;
12. agreeing upon press releases and joint publications regarding the project;
13. making proposals on dissemination of results and IPR-related matters;
14. assisting the Coordinator in the preparation of scientific reports.

The EU team at the University of Oslo will support the coordinator on issues related to accountancy, financial controlling, legal advice and support on contract management and protection of intellectual property rights led by a financial expert. This team has wide experience in the financial management of national, European and international projects.

The highest decision-making body of the project will be the Project Board. The project board is responsible for the overall management of the WPs, for the correct communication among partners, for the timely production of deliverables, and for the integration of the work of partners and WPs. Members of the project board have a strong background knowledge of scientific, technical and managing dealt within this project.

The project board is chaired by the coordinator. Decisions are taken by a majority of two-thirds (2/3) of the votes, unless otherwise provided in the Consortium Agreement. The board meets at least three times over the life span of the project: at the beginning (Kick-off Meeting) and at the end of each project year. Additional meetings may be organised, subject to the availability of resources and necessity.

Work-package Leaders share responsibility with the coordination team for the timely and effective implementation of the activities planned in each WP of the project. Main activities cover:

- ensuring performance and progress of the activities with regard to the deliverables and project milestones;
- coordination and monitoring on a day-to-day basis of the progress of the Work Packages, with a particular attention to the activities carried out within the other WPs;
- ensuring communication between members of the WP and to the CCT of any plans, deliverables and information concerning the work packages;

Notwithstanding the above, from the administrative and financial perspective, all Partners bear the same obligations towards the Commission. Moreover this co-responsibility vision to the implementation of the project activities is an indication of all participants' strong engagement in the project and their high commitment. In details, main activities cover:

- substantive contribution to the scientific coordination of the project through the active participation of one representative of the partner to the PA;

- efficient implementation of the tasks within their own Work packages;
- specific support to the WP leaders for what concerns the preparation of the EC periodic reports;
- basic support to the PC for what concerns the administrative and financial matters, in particular concerning the period cost statements.

3.2.2 Meetings

Regular Project Meetings will be organized twice per year. The purposes of these meetings are to verify the achieved results and milestones, to exchange scientific and technical information about the activities in the project. Specific meetings of the CCT will be organized within Project Meetings. Additionally, regular web- meetings and workshops of the CCT will be used with flexibility during the project to plan, execute and summarize specific tasks, to discuss and promote integration among the different consortium expertise.

The Final Meeting will be held in Brussels (to facilitate participation and to increase the project impacts) and will be divided into two parts: one will consist of scientific presentations and discussions within the Consortium, and the other one will be a seminar dedicated to enterprises, policy makers, EU officers, coastal authorities and end- users for demonstrating the tools and the integrating web-platform, for disseminating the adaptation guidelines useful and the best practices that may be considered for policy implementation purpose.

Internal meetings: Project internal meetings by telecons typically takes place every week. Project internal meetings agenda include the project state, the progress of activities being implemented, actions and problems. Project internal meetings are recorded in meeting minutes.

3.2.3 Reporting

The project will be reporting regularly on its development and results at three different levels:

1. Management Reports. Every 6 months the coordinator will collect, via WP leaders, a concise management report following the meetings of the CCT. The report will include: i) a summary of all activities carried out in the preceding period and a review of progresses; ii) plans for the following 6 months detailed at task level including, if necessary, adjustments proposal; iii) minutes of the CCT meetings, iv) updated list of publications submitted and published and presentations of project results. All project publications will expressively indicate the EU funding (contract number, project title and acronym).
2. Annual scientific reports. Produced by all participants and integrated by the Coordinator into a coherent document according to H2020 guidelines, will provide a synthesis of the results and deliverables produced in the previous year.
3. Final report. A final project report will be delivered in compliance with the general conditions of the Grant Agreement.

3.2.4 Conflict management

Identification of any conflicts which arise in the project is the responsibility of all project participants. Any signs of disagreement between project participants should be notified to the work package leader or coordinator (as appropriate), who should then instigate the conflict resolution procedure:

- The coordinator should separately contact all parties either in person or by telephone, to identify the different viewpoints (it is important not to use email: that medium very often leads to a rapid escalation of disagreements). Based on a clarification of viewpoints, the manager should try to propose a solution. If one is achieved, it should be recorded in a short report; if not, the problem should be escalated.
- If level 1 fails, the matter should be taken up by the CCT (at a special meeting, if need be). At this level, all work should be in writing. If conflicts relate to matters which would normally be assessed as part of the annual reviews by the Commission, the views of the Commission should be sought.

3.2.5 Project re-planning and change management

In an ambitious and dynamic project of this kind, changes to customer requirements are expected and will generate changes to the project plans. Handling changes in project plans will therefore be regarded as a normal part of project management, to be carried out without undue formalities. Project progress will be continuously monitored, and where discrepancies between plans and progress are observed (or predicted), corrective actions will be initiated. In particular, the CCT will carry out risk assessment at their regular meetings. This involves identifying project risks and assessing their probability and the nature of the consequences should the risk be incurred. If the risk level is judged to be high, changes in project planning may be necessary. A set of project risks has already been identified (see Section 3.2.3). It will serve as the basis for risk assessment at the first meeting of the Coordination Committee Team, and will be continuously updated thereafter. Decisions on any necessary replanning of detailed tasks at the work package level will be made by the Work Package Leader, in consultation with all partners involved in the work package. Results should be reported to the Coordinator. Project level changes will be the responsibility of the CC (except in the case of major changes). In addition to any reviews arising from regular risk assessment, the detailed project plan will be reviewed at least once per year, and revised if necessary.

Certain types of re-planning may require the approval of the Commission, according to the terms of the Grant Agreement. It will be the responsibility of the Project Coordinator to contact the Commission regarding the matters. Project re-planning which results in changes deemed to be major must be handled by the General Assembly, using voting procedures. Changes will be deemed to be major if any one partner protests about a proposed change, or automatically if the change involves modifications to the Consortium Agreement or to the management structures and principles, problems with the performance of any partner, partner request to leave the Consortium or re- allocation of budget between work packages and/or partners. Implementation of major changes may necessitate a change in the overall project plan, detailed project plans or the work breakdown structure of the project. As explained above, the management structure of the project essentially follows the work breakdown structure of the project. The management structure can therefore adapt to changes in the work breakdown structure.

3.2.6 Effective innovation management.

Innovation management in this collaboration is considered as an issue of highest priority in the project management activities and in the implementation of the work plan. The Consortium considers innovation as a task for the whole project which needs systematic approaches and strategies to bring innovations into practice. Innovation management is a process which requires an understanding of the market aspects and technical aspects of the project, with the goal to convert the project results in sustainable long-term sources of value. Therefore it requires the contribution from research/academic entities, end-users and industrial suppliers, all well represented in the consortium Innovation management is not a standalone track in the project: in fact it is tightly intertwined will activity streams of all work packages and will make a large section of the exploitation plan.

3.2.7 Critical risks, relating to project implementation,

Risk management procedures will be implemented in WP1 to ensure a smooth progress of the project. Major risks are identified in table 3.2b. Continuous monitoring of risks and of their impact will be established by the CC within the project. Additional risks will be identified by the individual WP Leaders immediately after appearance and mitigation measures will be adopted. Table 3.2b will be updated in the annual report.

Tables for section 3.2

Table 3.2a: List of milestones

Milestone number	Milestone name	Related work package(s)	Due date (in month)	Means of verification
1	GitHub repository	9	1	Up and running
2	Proto-type Gibbs sampler	1	3	Up and running
3	Commander module	6	3	Available in GitHub, and running
3	Data flagging module	2	6	Available in GitHub, and running
4	Web page	9	6	Open for public
5	PEDR	10	6	Posted on web page
6	Gain module	3	12	Available in Github, and running
7	First end-to-end Gibbs sampler	1	12	Available in GitHub, and running
8	Interpretation module	7	12	Available in GitHub, and running
9	First internal data release	1	12	Data products distributed internally
10	Alpha release of reproducibility tool	9	12	Distributed internally
12	Deconvolution map maker	5	15	Available in GitHub, and running
13	Beta release of reproducibility tool	9	18	Distributed internally
14	Final data release products	1	21	Data products distributed internally
15	Public data release	All	24	All products and papers available at public repositories
16	Final release of reproducibility tool	9	24	Publicly available
17	Final report	10	24	Report submitted to EU

Table 3.2b: Critical risks for implementation

Description of risk (indicate level of likelihood: Low/Medium/High)	Work package(s) involved	Proposed risk-mitigation measures
C-BASS observations not released in time (Low)	WP6	The C-BASS observations are not essential to the success of the project, but rather a useful bonus in terms of effective signal-to-noise for the synchrotron spectral index. If not available in time, the analysis will include only LFI and Planck data.
Not able to identify good PhD and postdoctoral candidates for new positions (Low)	WP1, WP4, WP5, WP6, WP8	Five WPs will hire new manpower for their tasks, and this always involves a certain level of risk. However, given the timing of this project, there will be many candidates on the job market from Planck, and we consider the likelihood of securing strong candidates as excellent. We will advertise the open positions immediately after the project is approved, to allow for maximum time for identifying good candidates.
Computer system failures (Low)	All	There is always a risk involved in high-computing projects in that computing facilities may be down for maintenance or due to failures. We will therefore maintain two independent and fully functional copies of the pipeline, both in Oslo and Trieste, ensuring that analysis may take place at any time.
Difficulty in reproducing computational efforts (Medium)	WP9	Inability to make parts of the computational efforts fully reproducible (either automatically, or even manually). Due to the facthat some of the computational efforts require distributed or computational grids in order to achieve acceptable execution times, finding alternative execution environments might be prohibitively expensive thus making the reproducibility of these particular scripts very difficult. The possible inability to recreate computationally expensive operations might be leveraged by our investigation into specific code transformation techniques for exploiting local or remote multi-core CPUs or GPUs without the need for these expensive computational grids.

Institution	Member	WP(s)	Self-financed FTE Months	H202 financed FTE months
Oslo	Hans Kristian Eriksen	WP1, WP10	6	
	Ingunn Kathrine Wehus	WP6	3	
	Postdoc 1	WP1		24
	Postdoc 2	WP6		24
Milano	Marco Bersanelli	WP8, WP10	2	
	Loris Colombo	WP7		24
	Davide Maino	WP8	1	
	Aniello Mennella	WP8	1	
	Maurizio Tomasi	WP8	1	
INAF	Samuele Galeotta	WP1,WP2,WP3		10
	Gianmarco Maggio	WP3	1	
	Michele Maris	WP2,WP3		10
	Daniele Tavagnacco	WP3	2	
	Andrea Zacchei	WP2,WP3, WP10		7
Helsinki	Elina Keihänen	WP4,WP5, WP10		3
	Postdoc 3	WP4,WP5		12
Planetek	Efstratios Gerakakis	WP9, WP10		24
	Computer scientist 1	WP9		24
	Computer scientist 2	WP9		11
XAL	Eirik Gjerløw	WP7, WP10		19
Sum FTE Months = 209			17	192

Table 3.3a: Overview of all consortium members and their funding sources

3.3 Consortium as a whole

We now describe the consortium as a whole in terms of individual members, the responsibilities of each person, and how they inter-relate. First, Table 3.3a provides an overview of each member, their main WP responsibility, and their funding source. Corresponding CVs are provided in Section 4.

First, we note that the project is by far mostly dependent on H2020 funding, accounting for a total of 192 out of 209 FTE months, or 92%. This is precisely why this application is absolutely critical for this essential work to take place. The remaining 8% of the budget is covered by the Universities of Oslo and Milano through other funding sources.

Looking at this list of scientists and their CVs in detail, one can easily recognize the manifestation of the design philosophy discussed above, in that every member plays a leading role in the existing Planck LFI consortium within her or his topic. In the following overview, the main WP responsibility is indicated in paranthesis for each person, while the corresponding description provides a short statement of prior role in the Planck LFI collaboration:

- Hans Kristian Eriksen (WP1 Gibbs sampling, WP10 Administration) is the main developer of the Commander analysis software package, and one of the pioneers in introducing Gibbs sampling in modern cosmology, and applying this to Planck observations.
- Ingunn Kathrine Wehus (WP6 Component separation) leads the astrophysical component separation work in Planck, both for LFI and HFI, and has been responsible for delivering the diffuse component products in both Planck releases.
- Marco Bersanelli (WP8 Systematic errors) has served as Planck LFI Instrument Scientist for more than two decades, and has overseen the design, construction and deployment of the LFI instrument from the beginning to the end of the mission.
- Loris Colombo (WP7 Physical Interpretation) has carried a main responsibility for developing, characterizing and validating the low-l likelihood and parameter estimation for Planck LFI.
- Davide Maino (WP8 Systematic errors) has has been responsible for the implementation of the second stage of the LFI pipeline, covering the entire signal processing from single detectors time-order information to calibrated frequency maps.
- Anniello Mennella (WP8 Systematic errors) has led the systematic error analysis for the last Planck LFI release, establishing high-fidelity end-to-end simulations of all known sources of instrumental errors.
- Maurizio Tomasi (WP8 Systematic errors) has has been responsible for the in-flight calibration of the LFI data, based on the CMB dipole modulation, including assessment of absolute and relative uncertainties.
- Samuele Galeotta (WP1 Gibbs sampling, WP2 Data flagging, WP3 Gain estimation) has been the main responsible for the low-level LFI data processing from raw maps to final frequency maps.
- Michele Maris (WP2 Data flagging, WP3 Gain estimation) has carried a main responsibility for modelling the relationship between the Planck LFI gain calibration and the instrumental response function, including 4π beam and sidelobe modelling.
- Daniele Tavagnacco (WP2 Data flagging, WP3 Gain estimation) has been responsible for producing low-resolution maps and noise covariance matrices for Planck LFI, as well as for removing low-level artifacts (spikes etc.) from the time-ordered Planck LFI data.
- Andrea Zacchei (WP2 Data flagging, WP3 Gain estimation) has for more than a decade served as the Planck LFI Data Processing Center (DPC) manager, and has carried the main responsibility for delivering final LFI products to the public.
- Elina Keihänen (WP4-WP5 Map making) is the main developer of the MADAM destriping code used for Planck LFI map making.
- Efstratios Gerakakis (WP9 User communication) has carried a main responsibility for developing the high-level aspects of the Planck Legacy Archive Added Value Interface (PLAAVI) during the last two years.
- Eirik Gjerløw (WP7 Physical interpretation) has carried a main responsibility for developing the low-level algorithmic aspects of the Planck Legacy Archive Added Value Interface (PLAAVI) during the last two years.

It should be clear from this overview that each member is uniquely suited to complete her or his tasks, and it is our assertion that this team is optimally designed for the tasks at hand in terms of individual and complementary expertises and responsibilities.

Table 3.4a: Summary of staff effort

	WP1	WP2	WP3	WP4	WP5	WP6	WP7	WP8	WP9	WP10	Total Person-Months per Participant
1. Oslo	27		1			27				2	57
2. Milano							24	4		1	29
3. INAF	2	10	17							1	30
4. Helsinki				7	7					1	15
5. Planetek									58	1	59
6. XAL							18			1	19
Total	29	10	18	7	7	27	42	4	58	7	209

3.4 Resources to be committed

A full break-down of the H2020 budget sought for the current project is provided in the budget table in Section 3 of the Administrative proposal form, Part A. The total amount is €1.499.647,50, covering a total of 193 FTE months plus travel and operating costs; 14 additional FTE months are covered by the University of Oslo and the University of Milano. Table 3.4a summarizes the total number of FTE months per institution and work package.

For most institutions, the personnel costs represent more than 85% of the total budget. The only two exceptions from this are the University of Oslo and INAF. For Oslo, additional expenses come from serving as the coordinator, and it will therefore cover publication costs and meeting expenses. For INAF, the additional expenses are due to maintaining the Planck computer cluster. These expenses are summarized in Table 3.4b.

Table 3.4b: ‘Other direct cost’ items (travel, equipment, other goods and services, large research infrastructure)

1. Oslo	Cost (€)	Justification
Travel	53,000	This cost covers travel expenses for four persons (HKE, IKW, 2 PD) for two years at €4k/person/year = €32k, as well as expenses for two internal working meetings at €3k/meeting = €6k, and one release conference at €15k.
Equipment		
Other goods and services	20,000	Dissemination costs (publication charges, web hosting etc.)
Total	73,000	

3. INAF	**Cost (€)**	**Justification**
Travel	40,000	This cost covers travel expenses for five persons for two years at €4k/person/year = €40k.
Equipment		
Other goods and services	30,000	Maintainance of existing dedicated Planck cluster (dedicated electricity contract).
Total	70,000	