Project number: 776282

Project acronym: BEYONDPLANCK

Project title: BeyondPlanck – delivering state-of-the-art observations of the microwave sky from 30 to 70 Ghz for the next decade

Periodic Technical Report

Part B

Period covered by the report: from 01.03.18 to 28.02.19

Periodic report: 1st

1. Explanation of the work carried out by the beneficiaries and Overview of the progress

1.1 Objectives

As stated in Section 1.1 of Annex 1, the main objectives of the BeyondPlanck project are the following:

1. To deliver new legacy Planck LFI 30, 44 and 70 GHz frequency maps
2. To deliver the world's cleanest and most sensitive full-sky estimates of polarized synchrotron emission at CMB frequencies.
3. To 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.
4. Make the software necessary for time-domain analysis available to the community under an Open Source license.

As detailed further in the following, the BeyondPlanck consortium has during the last 12-month period made excellent progress towards reaching all of these objectives. Indeed, we believe it is fair to say that the achieved results exceed our original expectations, and we believe that they will, once publicly released, have an even greater impact in the community than originally anticipated.

This rapid progress has been made possible through a two-track research program. First, in accordance with the original proposal, we have by Month 12 successfully established a fully functional time-line Gibbs sampler for Planck LFI data by integrating existing codes into a single analysis framework, running without intermediate human interaction at the computing facilities at the University of Oslo. This has already reduced the wall clock time per iteration from about 2 weeks to about 2 days, corresponding to a factor of seven in run time. Further optimization is currently on-going, and we consider it likely that another factor of 4-10 may be possible, at which point the entire loop will be run in 5-10 hours. The entire software suite is currently available in the BeyondPlanck GitLab repository (for details, see previous deliverable descriptions), and documentation is currently being developed.

The second main development track that has guided our efforts during the last 12 months is based on the same manual Gibbs sampling approach as was carried out within the Planck LFI data processing center, and that established the starting point for the BeyondPlanck project. In this mode, intermediate data products are transferred between the calibration, mapmaking and component separation codes by hand. A total of four such cycles were completed within the official Planck LFI processing, and the end product of this process eventually defined the official Planck LFI products.

After the Planck LFI DPC processing was officially completed, and the final Planck products were frozen, we continued this manual processing in parallel with the construction of the automatic pipeline. The main goal of this track is to understand the issues that cause the systematic residuals observed in the data at a fundamental level, and to fix them at a basic level. This approach is thus complementary to the automation process, in that it aims to improve each of the sub-components in the pipeline, rather than simply make them run without human interaction.

A major breakthrough has taken place as a direct consequence of this work after the official completion of the Planck project. Specifically, in collaboration with Dr. Reijo Keskitalo at NERSC, Lawrence Berkeley National Laboratory, US, we have been able to significantly improve the overall quality of the Planck data set, by applying the same manual Gibbs sampling approach described above to both LFI and HFI. This is important, because even though HFI is not formally a primary component of the BeyondPlanck project, these data are essential in order to understand the contribution from both CMB and thermal dust emission at LFI frequencies. With the previously available data, systematic errors in the HFI maps could easily end up being a limiting factor for BeyondPlanck. However, with the new joint manual processing of the full Planck data set, we are now confident that HFI systematic errors are well under control for the purposes of BeyondPlanck.

Indeed, this new end-to-end processing of the full Planck data set performed in collaboration between Dr. Keskitalo and the BeyondPlanck team has been so successful, that a new official Planck data release is currently under preparation under the leadership of Dr. Keskitalo. This release will be called «NPIPE», which simply means «NERSC pipeline», recognizing the computer facilities employed for the final processing. The name is chosen intentionally to be politically non-controversial, in order not to cause internal conflicts within the Planck consortium. However, in practice NPIPE is a manual implementation of the BeyondPlanck concept, applied jointly to both the LFI and HFI data sets. Furthermore, several BeyondPlanck codes are used directly as critical sub-components in NPIPE, most importantly the Madam destriper and the Commander component separation code. Furthermore, all NPIPE component separation work has been performed within the BeyondPlanck project.

The NPIPE data products are scheduled to be released in mid-2019, and will represent the first publicly available products from BeyondPlanck. Furthermore, these products will be recognized as official Planck products, including authors from the Planck consortium, although with a clear acknowledgment of their BeyondPlanck heritage and financial support. However, it is important to note that NPIPE is still a manually operated analysis pipeline. As we go into the final 12 months of the BeyondPlanck project, and the automatic pipeline is finally running, our goal is to combine the fundamental understanding of the data set gained from NPIPE with the computational power of the native BeyondPlanck pipeline.

Thus, main scientific highlights resulting from the above efforts are the following:

- We have established new Planck LFI sky maps that show both greater internal consistency and better agreement with WMAP, as well as new HFI maps with significantly lower systematics, that will support the final BeyondPlanck analysis
- The LFI 44 GHz channel appears for the first time suitable for detailed scientific analysis, although further validation is warranted
- The statistical noise of the Planck LFI maps are 10-15% lower than in the official products, due to usage of more data and filtering of cold load reference data
- The optical depth of reionization appears more robust with respect to data cuts, both among LFI and HFI frequencies
- We have derived a new estimate of both synchrotron and thermal dust emission in polarization
- We have developed a new temperature sky model for use in BeyondPlanck processing that already represents a new state-of-the-art in the community
- We have established the first all-frequency estimate of the CMB solar dipole based on true component separation, and developed a new statistical framework to estimate robust uncertainties

As should be clear from this overview, we are very pleased with the overall progress and status of the BeyondPlanck project, and are excited about presenting these results to the community at large in the near future.

WP	PM spent	PM budgeted	PM difference	Partner breakdown
1	11.4	14.5	-3.1	Oslo = 11; Trieste = 0.4
2	5.6	5.5	0.1	Trieste = 5.6
3	8.7	10	-1.3	Trieste = 8.7
4	7.0	3.5	3.5	Helsinki = 7.0
5	0.0	3.5	-3.5	Helsinki = 7.0
6	17	18.5	-1.5	Oslo = 17
7	29	30	-1	Oslo = 5.0; Milano = 24
8	1.7	2.0	-0.3	Milano = 1.7
9	27.6	29	-1.4	Planetek = 27.6
10	2	3	-1	Oslo = 1; Milano = 0.5; Trieste = 0.5
Total	110.0	119.5	-9.5	

Table 1: Overview of PM efforts spent per WP (second column), compared with the original expectation (third column). The fourth columns shows the difference between the real and projected PM counts. The fifth column provides a breakdown of the spent PMs for the individual partners.

1.2 Explanation of the work carried out per WP

Table 1 provides an overview of the total number of person months (PM) per work package spent during the first half of the project period, including a break-down of PMs per partner. The number of budgetted PMs are reproduced from the Work Package overview in Part A, and each number corresponds to half the total of budgetted PMs for the full experiment, recognizing that we are currently halfway through the project. Specific tasks are defined in a Gantt chart shown in Figure 7 in Part B.

Overall, we observe good agreement between the planned and actual realized work effort, with an overall differential of 9.5 PM, or 8% of the total work load. The most notable outlier is WP5 (beam deconvolution), for which no work has been performed. However, this work must be seen in light of WP4 (mapmaking), for which the same people are responsible. The mapmaking component plays a project critical role in the pipeline, and WP4 has therefore had a higher priority than WP5 in the first half of the project. WP5 will, on the other hand, become more important in the second half.

It is also important to note that less than half the work resources are spent at this time. This was done on purpose, since we know by experience that the intensity of the overall effort always increases towards the end of a project. We have therefore made save some resources in the first part of the project period, in anticipation of a busy finishing period.

As described in Annex 1, joint consortium meetings have taken place every six months, each lasting for one full week. The three first meetings were held in Oslo (March 2018; kick-off meeting), Helsinki (September 2018), and Athens (February 2019). We have found these meetings to be very useful in order to organize the overall work, and ensure that tasks are completed in a timely fashion. The work time spent on participation in these meetings are distributed among the respective work packages in Table 1, but we do not list this item explicitly in the detailed descriptions below.

We now provide a summary of the work done within each WP by each beneficiary:

WP1 – Gibbs sampling integration

Oslo

- Implementing Python-based integration framework for iterative analysis, including automatic parameter file generation code; this serves as our baseline approach, and is algorithmically safe and well-tested, but computationally sub-optimal, since data products are stored on disk between iterations
- Implementing Fortran-based infrastructure for time-domain processing within Commander; this serves as an experimental platform for now, as more development work is required, but potentially, significant computational savings can be made by eliminating intermediate IO.
- Set up and maintain computational infrastructure; providing access to and supporting all BeyondPlanck users in the computing cluster
- Performing initial tests and benchmarking analyses

Trieste

- Organized LFI data set on the Oslo cluster.

WP2 – Data selection

Trieste

- Raw data extraction from Trieste internal DB to HDF5 format. During the operation all the flags associated to each sample (e.g. not usable data, planets transits, slewing period) have been checked and exported. The new HDF5 files has been transferred to the Oslo computing system.
- Satellite pointing information creation based on Attitude History File corrected for wobble angle, saved in fits format and transferred to the OSLO computing system.
- Detector pointings creation from the Attitude History File.
- Known systematics effect removal: ADC non linear response correction and 1Hz spikes removal.
- Data differentiation, step necessary to reduce the $1/f$ instrument noise. The operation consists in computing the ratio between the average of the data signal coming from the sky and the corresponding 4k reference signal over each pointing period (roughly 45 minutes) and then computing the difference between the sky signal and the reference signal multiplied by the Gain Modulation Factor.
- Differentiated data and detector pointings saved in HDF5 format. This is the input to the subsequent step *Gain Estimation*.

As noted in Part A, task 2.3 defined by the Gantt chart in Part B has not yet been concluded, despite the fact that it was originally scheduled to end at Month 12. This is because the overall BeyondPlanck results are still in development, and it is difficult to fix the timelines at this time. Indeed, in retrospect we believe it was a mistake to assume that WP2 would conclude already at Month 12, just when the heavy processing starts. Instead, we would like to extend Task 2.3 until Month 21 in the current schedule. At the same time, it is clear that the file transfer went more smoothly than originally anticipated, requiring only about 5PM in total instead of 10 PMs. Another 5 PMs are not required for the final preparation. Instead, we propose to reallocate 3 PMs from WP2 to WP3, to optimize the gain sampling step, which currently is the main driver of the Gibbs loop in terms of CPU time. WP2 and WP3 are organized by the same team, and such a reallocation is therefore seamless in terms of funding.

WP3 – Gain estimation

Trieste

- Sky Model to be used in calibration creation: the sum of the timeline obtained convolving the Solar and Orbital Dipole using convolution parameters and the timeline of the Galactic components extracted from the Galaxy components ringsets. Then the resulting sky model and the data coming from WP2 (*Data selection and anomalies flagging)* are subsampled to reduce their size.
- First conversion factor computation: perform a linear fit between the subsampled data and the subsampled sky model to compute a first guess of the conversion factor.
- DaCapo: uses destriping techniques to iteratively improve the conversion factors. It produces an estimation of the sky signal from the previous conversion factors, then it applies a destriping algorithm and use the results to apply a correction. The iteration continues until a convergence criterion is satisfied. The output are the raw conversion factors.
- Smoothing of the raw conversion factors: a de-noising procedure is applied to the conversion factors stream. The results are the final conversion factors to be applied.
- Conversion Factors application: multiplies the data coming from WP2 to convert them into thermodynamic units and removes the systematic signal coming from the optic of the Planck telescope and the Orbital Dipole.

WP4 – Mapmaking

Helsinki

- Integrating the Madam map-making code into the BeyondPlanck pipeline. The data input/output routines were adapted for the change in the input data format. The updated Madam has been tested and validated.
- A new framework for map-making in Gibbs sampling has been developed, and the idea has been tentatively coded and tested.

As seen in Table 1, a total of 7 PMs were spent on WP4 during the first 12 months of the project. This corresponds to the full resource allocation for this WP. However, tasks 4.3 and 4.4 (defined by the Gantt chart in Part B) are not yet concluded. To solve this issue, we note that WP4 and WP5 are scientifically very closely connected (corresponding to standard and deconvolution map making, respectively), and it would in fact have been better to define these two as a single work package in the original proposal, simply called "Mapmaking". During the first year, however, WP4 had to take priority, since the entire pipeline depends directly on that module, whereas WP5 is less mission critical. Given the fact that we have run out of time on WP4, we therefore propose to formally reallocate 2 PMs from WP5 to WP4, to support the final work in WP4 though the project period.

WP5 – Beam deconvolution

No activity yet. First efforts started in March 2019, shortly after the conclusion of the first version of this review.

WP6 – Component separation

Oslo

- Development of both Planck 2018 (v1) and NPIPE-based (v2) baseline astrophysical sky model in both temperature and polarization, for use in final BeyondPlanck analysis
- Extending baseline model to additionally include WMAP and Haslam 408 MHz observations
- Implementing support for multiple CO isotopes (12CO and 13CO) in foreground model, as well as Doppler shift due to the Milky Way rotation; will be used to correct HFI observations in final BeyondPlanck analysis
- Development of new CMB dipole estimation methodology, and application to NPIPE observations. Drafting dipole paper.
- Preparing public NPIPE data delivery, including both data products and paper

WP7 – Science exploitation

Milano

- Implementation of a likelihood pipeline for the analysis of low-resolution Planck polarization maps. The BP pipeline is an evolution of the one used for the analysis of LFI low-resolution data in the Planck 2016 and 2018 releases, and it incorporates several improvements geared mainly toward increasing the flexibility and speed of the software.
- Analysis of low-resolution NPIPE polarization maps. The new maps show promising results on the optical depth to reionization, in particular for the first time 44GHz estimates are in line with those from 70GHz and 100GHz.

- Development of CosmoMC/CAMB plugins which explore non-minimal models of reionization.

Oslo

- Astrophysical analyses for NPIPE publications and results

WP8 – Systematic effects

Milano

- Analysis of the signal produced by moving objects (planets) to pinpoint inaccuracies in the beam and bandpass reconstruction, using DX12 timelines (Planck 2018 release). The aim is to re-apply this analysis to the new NPIPE timelines, once these will be released.
- Retrieval of all the information and technical documents stored in the now-dismissed Planck/LFI wiki, containing details about the 2013, 2015, and 2018 systematic analyses.

WP9 – Reproducible research

Planetek

- Set up the required project management infrastructure of the project. That included the creation of:
  - A git repository
  - Wiki pages and
  - Bug/Issue pages

  on a gitlab shared account. Appropriate access rights were given to all project members.
- Designed, created and maintained the BeyondPlanck website:
  - Hosted on https://beyondplanck.science
  - Website has been setup as a separate git repository with the contents of the website
  - Accessible and editable by all project members
  - Automated the deployment of the website, by automatically uploading edited code to the hosting server.
  - Maintenance of the website, by updating pages for events, talks and presentations
- Conducted survey on reproducibility in science in order to gather valuable information for the views of scientists on reproducibility. For this purpose we have created online questionnaires and distributed them to scientists including the consortium members, scientists we have previously collaborated with, connections in other scientific fields including research fellows, scientific personnel, professors, etc.
  At the same time, we have created separate questionnaires targeting only the consortium members and addressing specifically each work package.
- Evaluated existing Reproducibility services and tools.
- Reproducibility utility tool: Initiated development efforts on a command line utility tool that will allow scientists to make their work reproducible combining several existing services.
  - The goal of the tool is to provide a frictionless way for original authors to document and set up an automated reproducible workflow, and an even more easier way, for any other interested parties to locally recreate the same project and expect to get the same results as the original author.
  - We have defined the basic functionality, workflows, and features of the tool.
  - We currently have a working version of the reproducibility utility tool that demonstrates the basic functionality and implements sample use cases. HTTP, git, and Docker plugins have been implemented offering integration with these services
  - Code produced under BeyondPlanck Work Package 7 has been identified as the first candidate to be made reproducible using the reproducibility utility tool and efforts have been made to this direction.

GPU Enhancements

We focused on analysing the BeyondPlanck processing pipeline for determining potential “steps” suitable for optimization by using heterogeneous computing systems (i.e. those composed of both multi-core CPUs and many-core GPUs). We provided a preliminary analysis of state-of-the-art GPU technologies, documented in specific deliverables. Moreover, we also examined and benchmarked state-of-the-art basic linear algebra and generic computation libraries, as provided by main processors producers and scientific entities active in scientific software developments.

We identified the following scenarios for GPU usage in the BeyondPlanck project:

- Generic algebra operations;
- Spherical harmonics transformations;
- Components that execute per-pixel operations;
- The application of satellite attitude data;

- Sub cases of the map-making process.

During the past months, according to the Consortium’s suggestions and steering, we studied the “Generic algebra operations” scenario and in particular the pervasive Cholesky factorization method. After having analysed all variants of the aforementioned algorithm, we implemented the Cholesky-Crout one in OpenCL (chosen as a reference, portable, high performance computing platform for allowing testing on different vendor provided hardware) and executed it in a wide range of processing environments so to provide an initial benchmark. We also compared it with the main algebra libraries like LAPACK.

The first results did not show appreciable advantages of GPU execution over standard CPU, showing high end CPUs (the one with dozens of core and hardware threads) behaving better than any other architecture. Nevertheless, our benchmarks were limited to low-end GPU boards (provided by both AMD and NVIDIA).

Moreover, we analysed the implemented code and considered it improvable both in terms of adopted computing platform (CUDA and OpenCL) and in terms of memory usage (to be further analysed, since quite complex matter).
As a follow up to the described achievements, we identified the following next steps to improve the execution of the Cholesky algorithm on GPUs:

- Use NVIDIA devices with CUDA;
- Evaluate and integrate the MAGMA libraries;
- Explore serial block algorithm + GPU Cholesky decomposition in each block;
- Stress block mode by employing 3x3 computations and then proceed with parallelizable operations.
- Improve cache friendly-ness by implementing row packing.

The listed actions emerged from further experiments we made to improve implementation cache friendliness, computer cache usage in general, computing platform exploitation.

WP10 – Administration

All

- Preparation of documentation for deliverables and reports
- Participation in bi-weekly teleconferences

1.3 Impact

Although the primary BeyondPlanck products are not scheduled to be publicly released until 2020, the BeyondPlanck project and results have already had a massive impact in the internal Planck community through the NPIPE Planck data release. As a direct result of this work, every single Planck frequency channel has been significantly improved both in terms of sensitivity and systematic residuals. These maps are currently being subjected to detailed scientific analysis within the old Planck consortium, with very encouraging results: All component separation products appear cleaner from systematic errors; low-multipole power spectra looks for the first time visually consistent with LCDM predictions, without the need of high-level likelihood analysis; correspondingly, parameter constraints on the optical depth of reionization appear significantly more robust; the CMB dipole (and therefore absolute calibration of the entire Planck experiment) looks stable on sky fractions ranging between 20 and 95%, and derived uncertainties looks reasonable; and much, much more. These results will be released to the public in mid-2019, and will represent the first major data release for the BeyondPlanck project.

This NPIPE data release is likely to become the *de-facto* legacy release for Planck HFI data into the foreseeable future, and the corresponding maps will form the high-frequency basis for the remaining BeyondPlanck work. However, the NPIPE LFI results will eventually be superseded by the new BeyondPlanck LFI maps that are our main target for the remainder of the project period.

The first public presentation of the BeyondPlanck project to the general community took place at the European Space Astronomy Center in Madrid on December 3rd, 2018, at the last general Planck meeting. The project was well received, and great interest was shown regarding the scope of the program.

1.4 Background on the NPIPE project

As should be clear from the above descriptions, a significant amount of effort has been spent within the context of BeyondPlanck+NPIPE, and it is therefore useful to provide some background and specification of this work.

The NPIPE project was initiated near the end of the official Planck analysis period by Dr. Reijo Keskitalo, as a collaboration between the US and Oslo analysis teams, who have been working closely together for the last 15 years. The original main goal of this project was to perform a joint analysis of the Planck LFI and HFI observations within the high-performance computing context defined by the NERSC supercomputer facilities at Lawrence-Berkeley National Laboratories. The fundamental idea was to exploit the best ideas developed by the LFI and HFI data processing centers during the last 15 years, and apply them to both data sets. In particular, one of the most critical ideas in this respect was iterative application of calibration, mapmaking and component separation, as pioneered by the Planck LFI DPC, and that forms the scientific basis for the BeyondPlanck project. However, the NPIPE project did not aim at automating the analysis project, but simply applying the same procedure that was developed for LFI also to the HFI data set. At the same time, it would apply several important ideas developed by the HFI team to the LFI data set.

Two of the most critical computational components in this process was the Madam mapmaker, that has been developed at the University of Helsinki, and the Commander analysis framework, that has been developed at the University of Oslo. These form a common computational core for both NPIPE and BeyondPlanck, and defines the state-of-the-art of CMB processing today.

When the BeyondPlanck project was proposed in early 2016, it was not clear whether the NPIPE program would be successful. However, a new picture emerged quite shortly after the time BeyondPlanck started its operations in March 2018, and by early summer 2018 it was obvious that our NPIPE efforts had indeed turned out highly successful; not only for the HFI analysis, as we originally anticipated, but also for improving the overall quality of LFI data. By summer 2018, it was therefore clear that the only meaningful way forward for the BeyondPlanck collaboration itself was to join efforts with NPIPE, and take direct advantage of the improvements made during the last year.

Practically speaking, this collaboration took the form of a close collaboration between Dr. Keskitalo and, primarily, the BeyondPlanck component separation team, in particular Andersen, Eriksen, Svalheim, Thommesen, and Wehus . Specifically, while Keskitalo performed the actual time-line operations (calibration and map-making), the BeyondPlanck team performed the component separation work and quality assessment. Indeed, it was precisely in the latter step that many of the fundamental improvements were proposed and made, through inspection of various goodness-of-fit statistics. Based on the BeyondPlanck input, Keskitalo then implemented improvements in the calibration procedures that allowed further progress.

At the current stage, where the NPIPE processing is near its completion, the BeyondPlanck team is heavily involved in paper writing and preparations of final products. These are anticipated to be released in mid-2019, pending final approval of the Planck Science Team.

Regarding specific improvements developed for the NPIPE project, the following are among the most noteworthy:

1. Implementation of joint calibration, mapmaking and component separation.
2. Exploitation of 8% additional data taken during Planck repointing maneuvers, resulting in notably lower instrumental noise.
3. Smoothing of the LFI reference load measurements, reducing the instrumental noise in the final maps by about 10%.
4. Explicit bandpass correction prior to mapmaking, reducing polarization specific systematic effects.
5. Fitting polarization angles and efficiencies for individual detectors, improving the overall polarization reconstruction.
6. Producing new maps with the CMB dipole retained, allowing joint relative calibration and component separation.

These improvements, and several other of more minor nature, leads to significantly improved sky maps in both temperature and polarization, and the scientific exploitation of these maps is currently on-going. In particular, we note that our current estimates of both the CMB solar dipole and the optical depth of reionization defines a new state-of-the-art in the field today, both with important consequences for cosmological parameter estimation.

2. Update of the plan for exploitation and dissemination of result (if applicable)

While the main elements of the original plan remain unchanged at this time, we note that the original plan did not account for the NPIPE results. In short, we did not expect this line of work to be as successful as it has been, when writing the original application. However, given the internal dynamics of the Planck consortium, it has now become clear that these results will be published as a separate data release organized by the Planck project. In practice, this implies that the original BeyondPlanck dissemination plan will be augmented with an intermediate data release, including both basic data products (published through the Planck Legacy Archive) and a publication. The paper has already been drafted, and is currently undergoing internal review under the supervision of the Planck Editorial Board.

After this data release has been completed, we anticipate a number of scientific applications to be published as well, exploiting the improved Planck data quality. These include:

- Dipole methodology paper; draft already in place
- Analysis of multiple CO isotopes in Planck data; code and preliminary results already in place
- Analysis and mitigation of the effect of Doppler shift due to the Milky Way’s rotation on CO reconstruction; code and preliminary results already in place

In addition to being scientifically interesting in their own right, each of these topics will directly impact the upcoming BeyondPlanck analysis, in that they 1) improve our ability to calibrate the Planck data, and 2) derive more robust thermal dust extrapolations to low frequencies by breaking the degeneracies between thermal dust and CO emission.

3. Update of the data management plan (if applicable)

Same comments apply here as above; the plan remains intact, with the exception that we will have an additional intermediate data release in mid-2019.

4. Follow-up of recommendations and comments from previous review(s) (if applicable)

Not relevant.

5. Deviations from Annex 1 and Annex 2 (if applicable)

The realized BeyondPlanck work efforts are overall well aligned with the original proposal, and we do not consider any of the variations described above to represent a significant scientific deviation from the plan laid out in Annex 1. However, we would like to make a few comments regarding the scheduled end date for the project.

First, when the BeyondPlanck proposal was originally developed, the official Planck project was scheduled to end formally on December 31st 2017. However, we know by experience that these types of projects tend to run overtime, and we anticipated this by choosing a start-up date for BeyondPlanck on March 1st 2018, allowing two months delay for Planck. What we did not expect at the time, though, was that both ASI and CNES would provide additional funding for Planck analysis to continue until June 30th 2018, pushing the official Planck end date back by a full six months.

Second, as outlined above, we also did not anticipate that our LFI+HFI NPIPE analysis work to be as successful as it actually turned out. The fact that these results are publishable in their own right, and may turn into one of the long-lasting legacies of Planck, represents a massive success for the program, and it would be reckless not to exploit this success for science exploration.

On this background, we would take this opportunity to ask, if technically possible, for a so-called «no-cost extension» of the BeyondPlanck project for six months, delaying the end date from February 28th 2020 to August 31st 2020. This will provide sufficient time for paper writing and proper science exploitation of the avenues described above. Furthermore, funding for this additional time is already available through other local grants, and no additional H2020 COMPET-4 funding is required.

We therefore ask that the already allocated BeyondPlanck funding can be spent during the last six months as well, spreading the total effort out across the remaining 18 months in this scenario, with a final project end dateof August 31st 2020.

5.1 Tasks

All performed major tasks have been outlined in Annex 1. However, as already noted, the preparatory HFI work taking place with respect to NPIPE products has taken a more prominent role than originally anticipated. Fortunately, we have managed to perform this extra work within the limits of the already allocated resources, in parallel with the originally scheduled work. No delays have been introduced because of this work.

	Partners	WP1 person month		WP2 person mont		WP3 person mon		WP4 person mon		WP5 person mont		WP6 person months		WP7 person months		WP8 person months		WP9 person months		WP10 person month	
		EU funded	In-kind	EU funded	In-kind	EU funded	In-kind	EU funded	In-kind	EU funded	In-kind	EU funded	In-kind	EU funded	In-kind	EU funded	In-kind	EU funded	In-kind	EU funded	In-kind
1	UiO	10	1											3	2,5						1
2	UMIL													23,5			1,7			0,5	
3	INAF	0,3	0,1	5,1	0,5	7,6	1,1					15,5	1,5							0,5	
4	UH							7													
5	PLANETEK																	27,6			
	Total	10,3	1,1	5,1	0,5	7,6	1,1	7	0	0	0	15,5	1,5	26,5	2,5	0	1,7	27,6	0	1	1
	Grand Total	11,4		5,6		8,7		7		0		17		29		1,7		27,6		2	

Table 2: Breakdown over PMs spent per institution, divided into EU and in-kind funding.

5.2 Use of resources

The use of personell resources are summarized in Table 1. Note that the column marked by «PM budgetted» corresponds to half the total amount dedicated to the respective WP for the full project period. Thus, while the total number of PMs for the project as a whole is 239 according to the overview given in Part A, only 119.5 budgetted PMs are accounted for in Table 1.

Table 2 provides a similar overview, but additionally distinguishing between EU-funded and in-kind PMs.

5.2.1 Unforeseen subcontracting (if applicable)

Not relevant.

5.2.2 Unforeseen use of in kind contribution from third party against payment or free of charges (if applicable)

Not relevant.

5.3 Risks and unforeseen events

Several general risks were identified in Table 3.2b in Annex 1. These are the following:

1. *Delayed C-BASS delivery*. As described in Annex 1, the C-BASS data taken at 5 GHz were anticipated to be used to improve the polarization reconstruction at low frequencies. However, there was a definitive risk that these would not be made public in time to be useful for BeyondPlanck. Indeed, this has turned out to be the case, as C-BASS data are still not publicly available, and rumours say that they are unlikely to be so during the next year. However, in March 2019 another similar data set called S-PASS data, observing at 2.3 GHz, were made public. These are currently being analyzed jointly with Planck LFI, and a first assessment of systematic effects in S-PASS is on-going. If this turns out favorably, S-PASS will replace C-BASS in the final BeyondPlanck analysis.
2. *Securing good PhD and postdoctoral fellows*. This turned out to be not an issue; we are extremely happy with the team that has been assembled.
3. *Computer system failures*. This risk has not materialized; the current system in Oslo performs very well, with very low down-times.
4. *Difficult in reproducing computational efforts*. This is too early to assess, as the general pipeline is still under construction, and reproduction efforts are still in progress.
5. *Loss of personell*. No personell has withdrawn from the project during the project period.
6. *Partner withdraws from project.* No partners have withdrawn from the project during the project period.

6 Synergies with other on-going projects

The BeyondPlanck project takes place at the cutting edge of CMB research, and is a central component within a larger community. In particular at the University of Oslo multiple synergistic efforts are currently on-going, and, indeed, the BeyondPlanck project has played an important part in securing funding for some of these. In this section, we summarize these on-going projects, and detail the synergies between the various projects. Of course, we emphasize that there is no overlap or expense double-counting among any of these projects, even though there are obviously scientific synergies between them that will improve each project individually.

Bits2cosmology is an ERC Consolidator project led by Prof. Hans Kristian Eriksen. This project is algorithmic oriented, with a main goal of developing an integrated time-domain Gibbs sampler that is able to analyse the combination of Planck, WMAP and SPIDER data, as well as future LiteBIRD simulations. The main output from the project is a Fortran code we refer to as «Commander3» that will read raw time-ordered data and output final cosmology products, in particular constraints on B-mode polarization and the tensor-to-scalar ratio. This is in contrast to BeyondPlanck, for which the main goal is to produce better Planck LFI products by pipelining individual exisiting codes. Of course, bits2cosmology will benefit from the knowledge and experience gained during the BeyondPlanck project, but its scope in terms of code development, experiment selection and science target are different.

Cosmoglobe is an ERC Consolidator project led by Prof. Ingunn Wehus. This project aims to establish a new state-of-the-art astrophysical model of the radio, microwave, and sub-mm sky, covering frequencies between 100 MHz and 10,000 GHz, by combining observations from many leading experiments, including AKARI, C-BASS, COMAP, DIRBE, FIRAS, IRAS, Planck, S-PASS, SPIDER, WISE, WMAP and many others. As such, this project will benefit from the improved Planck maps that will result from BeyondPlanck. At the same time, Cosmoglobe may in principle provide an improved sky model for BeyondPlanck calibration purposes, although the timing for this will be challenging, given that Cosmoglobe starts in June 2019, while BeyondPlanck ends in 2020.

Global Component Separation Network is a research and education network led by Prof. Ingunn Wehus, and funded by the Research Council of Norway. This project aims to optimally exploit educational and scientific synergies between COMAP, LiteBIRD, PASIPHAE and SPIDER, and build a long-lasting academic network between top international educational and research institutions in Canada, India, Japan, Norway, South Africa and USA, and currently includes Caltech, CITA, IUCAA, kwaZulu-Natal, Oslo, Princeton, SAAO and Tokyo. The network does not include funding for research per-se, but only for travel, student exchanges, conference organization etc. For BeyondPlanck, this network will provide unique opportunities to disseminate its results efficiently to world-leading experts, and ensure that the BeyondPlanck products are fully integrated in the research community.