European Commission

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 01.03.19 to 28.11.19

Periodic report: 2nd

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. To make the software necessary for time-domain analysis available to the community under an Open Source license.

As reported during the first Periodic Review, the first 12 months of the BeyondPlanck project were largely spent on building infrastructure at the local computing cluster in Oslo. This included collecting and consolidating both software and data in one location, as well as disseminating knowledge and experience among the various team members, to the point that development work could be carried out efficiently. As of March 2019, we had established a single computational framework at UiO, in which a full LFI analysis loop could be undertaken in 1.5 days/iteration, as opposed to more than 2 weeks/iteration in the previous distributed form.

In addition to this main BeyondPlanck pipeline work, the BP consortium also played a central role in the development of the NPIPE analysis pipeline. The NPIPE results will constitute the final official data release from the Planck collaboration. These results significantly improve on the Planck 2018 (and 2015) products, which represented the state-of-the-art when the BeyondPlanck project was proposed. The NPIPE paper has very currently undergoing the final internal review within the Planck project, and is scheduled to be made public shortly, together with all necessary data products. To some extent, one could therefore argue that the objectives outlined above already were fulfilled with the NPIPE release.

However, if one reviews the detailed BeyondPlanck proposal, as opposed to just the primary objectives, it is clear that our original plan was more ambitious than what is implemented in NPIPE. While NPIPE clearly represents a significant improvement over earlier results, it still represents a traditional distributed approach to CMB analysis. Our goal, however, was to establish a single, tightly integrated, analysis framework that allows for fast turn-around, and, if possible, perform true Bayesian end-to-end analysis.

In this respect, it is worth recalling one of the main recommendations made by Prof. Ferreira during the March review meeting in Brussels. He strongly recommended that we focus on optimizing for computational speed and efficiency, in order to be able to explore much wider model spaces; as impressive as 1.5 days/iteration may sound, this is still quite expensive for performing true Bayesian analysis, and only a few runs would be possible with this computational cost.

Given this background, our main focus since the review meeting has been tight code integration and a strong focus on run-time optimization. This work has benefitted greatly from on-going Commander3 efforts, as supported by an ERC Consolidator grant led by Prof. Eriksen called “bits2cosmology” (see sect. 6) . Combining this low-level algorithm development work with the BeyondPlanck LFI-oriented production work has been extremely successful, and has led to further major improvements, even as measured against NPIPE. In fact, we believe that our current implementations and codes have the potential of transforming the field of CMB analysis once released to the public next summer, representing the first end-to-end MCMC sampler for CMB observations, from time-ordered data to cosmological parameters.

In parallel with code development and preliminary data analysis, we have also started the work of paper writing. As of now, a total of 21 papers are in development, not counting the already completed NPIPE and Thommesen et al. dipole papers that were reported during the first review period, or the S-PASS paper (Fuskeland et al. 2019) that already has appeared on the arxiv. The most advanced of the proper BeyondPlanck papers is the overview paper, which aims to provide an end-to-end overview of the method and main results. Rather than repeating the material from that paper in this report, we provide the latest draft as an attachment. However, we emphasize that all parts of that paper are preliminary and subject to change (and, indeed, most of the results shown in that paper are already obsolete). Rather than spending time on updating these preliminary results, we are currently focussed on finishing the main infrastructure. Thus the attached draft is intended only to provide a snapshot of the situation as it has existed at some point during the last few months, and to give an impression of where we expect to be next summer.

Some important highlights from the infrastructure-oriented work include the following:

- An entire analysis loop from raw TOD to the CMB likelihood has been established within the Commander code. The time-domain version of this code is now referred to as Commander3.
- The run-time for an entire time-domain LFI Gibbs iteration (plus component separation with HFI, WMAP and Haslam in the pixel domain) has been reduced from 1.5 days/iteration to 1 hour/iteration.
- The memory requirements for storing the LFI data (both on disk and in memory) have been reduced from 16 TB to 900 GB through pre-analysis discretization of pointing information combined with Huffman compression, and eliminating housekeeping data.
- We have implemented proper Gibbs sampling steps for gain estimation, absolute calibration, correlated noise, bandpass uncertainties, and main beam efficiencies.
- We have re-implemented the libConviqt formalism for 4pi beam convolution in terms of spin harmonics, significantly reducing run times and eliminating dependencies on external libraries
- We have coupled Commander TOD products with existing 4pi beam deconvolution codes through the use of 4D maps, and can now produce asymmetric beam deconvolved Commander-based frequency maps.
- We have interfaced the Planck LevelS simulation facilities with compressed Commander input files, and will soon be able to analyze Planck simulations with Commander.
- We are working on an extension in which Commander can output TOD simulations, exploiting the same infrastructure as is used for analysis; this may potentially in the future replace both the Planck Sky Model (PSM) and LevelS (TOD simulation) for the entire community.

- Work on the reproducibility tool and documentation is on-going.

In terms of science results, some highlights include the following:
- We achieve 10-20% lower noise than NPIPE at intermediate angular scales. This is effectively due to the fact that noise estimation at a given frequency exploits information from other channels to estimate the signal, thereby allowing deeper noise analysis.
- A strong degeneracy between gains, bandpass corrections, foregrounds and large-scale CMB has been identified when analyzing a minimally constraining data set. We emphasize that this result is highly preliminary, but if confirmed, then previously quoted uncertainties on the CMB dipole, and therefore the overall CMB power spectrum calibration, have been significantly underestimated, perhaps by a factor of 5 or more. This will have important implications with respect to LCDM parameters, including H_0, sigma_8 and others, and could potential help relieve several important tensions in the current field.
- Overall consistency with WMAP is significantly improved.
- A new CMB likelihood is well underway.

WP	PM spent	PM budgeted	PM difference	Partner breakdown
1	10.78	10.9	-0.12	Oslo = 10, Trieste 0.78
2	0.63	4.1	-3.47	Trieste = 0,63 (note that 3PM has been transferred to WP3)
3	6,74	7.5	-0.76	Trieste = 6,74
4	0	2.6	-2.6	None
5	6	2.6	3.4	Helsinki = 6
6	12.5	14.25	-1.75	Oslo = 12.5
7	30.5	22.5	8	Oslo = 13.5; Milano = 17
8	0	0	-1.5	
9	15	21.75	-6.75	Planetek = 15
10	1.32	2.25	-0.93	Milano = 1, Trieste 0.32
Total	83.47	89.95	-6.48	

Table 1: Overview of EU-funded PM efforts spent per WP (second column), compared with the original expectation (third column). The fourth column 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 budgeted PMs are reproduced from the Work Package overview in Part A, and each number corresponds to a projected period of 9 out of 24 months. Thus, the reported budget numbers do not account for the approved six month project extension, but rather refer to the original 24-month project duration. 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, although with a slight overall under-usage, which is intended in order to save resources for the final period. The total differential for this period is -6.48 PM, while the corresponding under-usage in the first reporting period was -9.5 PM. The total of extra PMs available for the final project period is thus 15.98 PM.

As described in Annex 1, joint consortium meetings have taken place every six months, each lasting for one full week. The first four meetings were held in Oslo (March 2018; kick-off meeting), Helsinki (September 2018), Athens (February 2019), and Milano (September 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 is distributed among the respective work packages in Table 1, but we do not list this item explicitly in the detailed descriptions below.

> We emphasize that most of the reported tasks are fundamentally iterative and incremental in nature. Few tasks are therefore “final” in the sense that they are never revisited after being reported as complete. As the analysis proceed, and the data become cleaner, new effects are discovered that warrants re-assessment of earlier stages. One particularly example of this is gain estimation: With the current μK precision that we achieve on the final maps, relative gain variations at the level of 0.01% become important. As such, further revisions of the gain estimation procedure are required and on-going; see attached paper for further details.

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

WP1 – Gibbs sampling integration

Oslo

- Developed integrated Commander3 TOD Gibbs sampler, including sampling steps for gain, absolute calibration, correlated noise, bandpass and beam mismatch corrections etc.
- Implemented new sidelobe convolution code
- Implemented data compression code
- Implemented support for 4D map format, to interface with deconvolution map maker.

WP2 – Data selection

Trieste

Most of the work in this work package was completed already during the first reporting period, and only 0.63 PM has been spent during the second reporting period. This work was spent on correcting a bug in the original flags, as well as performing tests on the quality of repointing periods for data analysis. This quality appears good, and the data are currently part of the BeyondPlanck baseline, adding 8% new data as compared with the official Planck analysis.

WP3 – Gain estimation

Trieste

- Implementation and testing gain estimation module.
- Exploration of 4K load smoothing, with the goal of reducing statistical uncertainties further. This was first implemented in NPIPE, and demonstrated to be beneficial.
- Comparison between and validation of traditional and Bayesian gain estimates.

WP4 – Mapmaking

Helsinki

- The Madam map-making code has been interfaced with the BeyondPlanck data structure, and is now fully functional. The integration was completed already during the previous reporting period. No FTE has therefore been counted for the current period.

- The code is being routinely used for production of pixelized sky maps from calibrated timelines. These maps are used for cross-checking and validation. Since the BP project has moved over to a new data approach where removal of noise and construction of sky signal both are performed inside Commander3, Madam maps are not considered primary data products any more.

WP5 – Beam deconvolution

Helsinki

- ArtDeco beam deconvolution code has been installed and tested on the Oslo cluster, and sample deconvolution maps have been produced from calibrated inputs. Since at the point it is unclear what will be the role of noise covariance matrices (NCVM), or if they will be replaced by a sampling-based approach, we have postponed their production until the situation is clarified.
- Instead, we have been concentrating on two tasks: First is the production of simulated data for the use of the BeyondPlanck collaboration. The simulations are performed with the LevelS software. Secondly, we have started writing a publication on the noise sampling algorithm, which plays a central part in the new approach, and developing a related test code.

WP6 – Component separation

Oslo

- Completed NPIPE analysis effort. The publication is now in final review at the Planck Editorial Board.
- Continuously testing and debugging end-to-end Commander3 Gibbs sampler

- Drafting BeyondPlanck publications
- Performing WMAP versus S-PASS consistency analysis, in order to assess whether S-PASS is suitable for integration into BeyondPlanck; paper is available at arxiv/1909.05923.
- Started exploratory C-BASS temperature analysis.

WP7 – Science exploitation

Milano

- Finalized Likelihood pipeline for the analysis of Commander2 and NPIPE results. Due to the switch to Commander3 framework, the pipeline is being updated to take advantage of the new capabilities provided by Commander3. A set of ancillary tools to process Commander3 is also being developed.
- Finalized extended reionization modules for CAMB/CosmoMC
- Started preliminary analysis of Commander3 products.
- Final analysis of NPIPE maps. The corresponding paper is undergoing Planck internal review before submission.

Oslo

- Preliminary BeyondPlanck analysis

WP8 – Systematic effects

Milano

Finalized the systematics assessment scheme for Commander2 results. As the Commander2 pipeline was a direct extension of the pipeline used to produce the LFI legacy results, this scheme was built on the systematics analysis work done for the PR3 papers. Given the changes introduced by the new Commander3 pipeline, the scheme is being reworked to deal with the peculiarities of the new framework.

Produced a synthetic description of instrument-generated systematic effects that will be treated specifically by the BeyondPlanck analysis.

WP9 – Reproducible research

Planetek Hellas

Reproducibility Tool

- After feedback from the project team as well as Prof. Ferreira during the March review meeting in Brussels, we concentrated on a BeyondPlanck-specific reproducibility tool, rather than a generic reproducibility tool.

- The current state of the Commander 3 pipeline has been documented. The pipeline has been broken down into 4 execution stages, as well as 3 additional pre-calculation stages. For each stage, required input files have been identified, as well as the list of expected output files.
- The BeyondPlanck reproducibility tool is currently able to download all required input data sets.
- Separate docker images are been created for each pipeline execution stage (work in progress).
- An effort on documenting Commander3 has been initiated. Commander3 parameters have been grouped and an initial documentation draft has been produced. Required input files and execution stages are also being documented.

GPU Enhancements

The Planetek team performed the following activities regarding GPU usage in the scope of BeyondPlanck:

- Analysis of available implementations of SHT on GPU with performance assessment.
- Definition/Identification of a baseline reference value on OWL.
- Identification of the best performing SHT implementation on GPU (ARKCoS emerged as best for various reasons).
- Comparison of output maps and numbers versus reference implementation and its results.
- Analysis of value displacements, errors and unexpected emerging artifacts.
- ARKCoS output validation in the optimized and non-optimized versions.
- Investigation around techniques for adapting ARKCoS to Commander3.
- Preliminary identification of possible interface between Commander3 and ARKCoS
- Project configuration and fixing of compilation issues of Commander3 to gain experience on the platform.

Infrastructure

- Set up an automated paper compilation pipeline for continuously building the BeyondPlanck papers and publishing them online at https://papers.beyondplanck.science/.
- Updating and maintaining of the BeyondPlanck website as well as the GitLab infrastructure.

WP10 – Administration

All

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

1.3 Impact

The impact of the BeyondPlanck project continues to be strong in the cosmological community.

First, the NPIPE work is close to being completed, with the paper currently being in final review within the Planck Editorial Board. Products are currently being integrated in the Planck Legacy Archive, and both the paper and products are expected to be publicly released shortly.

Second, we have started exploring the possibility of integrating low-frequency surveys into the analysis. The first effort has revolved around the 2.3 GHz S-PASS survey, and the first publication on this has already appeared on the arXiv (1909.05923), and the results have gained the attention of important players in the field. Second, individual members of the Oslo group now have direct access to the C-BASS temperature observations, and are working on integrating this data set with Planck and WMAP. For now, this work is exploratory in nature, but if successful and agreed by the C-BASS collaboration, this data set may be integrated in the first publicly released version of the BeyondPlanck model.

Third, BeyondPlanck will be presented at the next LiteBIRD face-to-face meeting in Garching in December this year through a dedicated talk. The aim of this talk is to introduce end-to-end Gibbs sampling as a complete analysis solution for LiteBIRD.

Fourth, given the very positive results we have obtained in the current project, we are currently exploring possibilities for applying for extending the project into a second phase, in which Planck HFI and/or WMAP data will play a central role. The consortium works very well at the moment, and there is a strong interest in maintaining this momentum beyond the end of the first phase.

Fifth and finally, the success and visibility of the BeyondPlanck project has led to a significant growth of the collaboration. Specifically, new members include Ragnhild Aurlien (Oslo; LiteBIRD), Ranajoy Banerji (Oslo; LiteBIRD); Sara Bertocco (Trieste; infrastructure), Marie Foss (Oslo; COMAP and mapmaking), Unni Fuskeland (Oslo; foreground analysis), Brandon Hensley (Princeton, foreground modelling), Daniel Herman (Oslo; foreground analysis), Håvard Tveit Ihle (Oslo; COMAP and noise estimation), Ata Karakci (Oslo; data management), and Bruce Partridge (Haverford; Planck). We fully expect the collaboration to continue growing in the future, especially if the work continues into a second phase, in which new expertise will be required. All new members are self-funded, and no costs will be declared for the BeyondPlanck project, except for travel costs to internal work meeting when their presence is considered important for the overall project.

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

The original BeyondPlanck proposal anticipated a total of nine refereed publications. However, this number will most likely be significantly increased before the end of the project.

First, we note that at the present time, one paper has already appeared on the arXiv (the S-PASS paper), while two more (the NPIPE and dipole papers) are completed, but are awaiting final approval from the Planck Editorial Board before they are released.

Then, as outlined above and described in Section 2 in the attached overview paper, a total of 21 BeyondPlanck papers are currently being drafted. Among these, 10 are defined as core papers (marked as red in Table 1 in the overview paper), and must appear in order for the BeyondPlanck release to be considered content complete. (A fall-back solution is to describe the same material in other papers, for instance in the overview paper, but this is not in our current plan.) Five of the papers are defined as important but not critical, and these are marked as orange in the table; these do not strictly need to be completed in order to go ahead with the release, as the material quite easily may be integrated into the overview paper – but we certainly plan on having separate papers for each topic. Furthermore, these are largely technical papers, and the main storyline of each is already well defined. We therefore do not foresee any major problems with finalizing these papers before the release. Finally, six papers are defined as non-critical, and are marked by green in the table. These represent interesting extensions of the analysis that would be «nice to have», but we fully anticipate that some (or all) of these will not appear together with the main release.

The goal is to complete the main BeyondPlanck papers before June 2020, and present them at the release conference taking place in the summer of 2020; exact dates and venue are still to be determined.

A copy of the PEDR has been posted on the project homepage.

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

No updates required.

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

During the review of the first reporting period, the external referee made one particularly strong recommendation, namely to focus on computational speed and optimization, in order to be able to perform broader model exploration. Implementing this recommendation has been a main focus during the second reporting period, and we have managed to reduce the runtime of a single iteration from about 1.5 days to 1 hour. This efficiency has been transformative for the project, and will increase the visibility of the project as a whole.

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.

5.1 Tasks

All performed major tasks have been outlined in Annex 1.

We note that Milestones 12 (Deconvolution mapmaker) and 13 (Beta release of the reproducibility tool) were reported later than originally scheduled. This was due to an administration error/misunderstanding during the Amendment process caused by the Coordinator. The actual work was completed on time, and the reports have been now been distributed as originally planned.

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

Year 2	01/03/2019-30/11/2019

	Partners	WP1 person month		WP2 person mon		WP3 person mon		WP4 person mon		WP5 person mont		WP6 person months		WP7 person months		WP8 person months		WP9 person months		WP10 person months	
		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	2									12,5	2	13,5							1
2	UMIL													17						1	
3	INAF	0,78		0,63	0,44	6,74	0,56													0,32	
4	UH									6											
5	PLANETEK																	15			
	Total	10,8	2,0	0,63	0,44	6,74	0,56	0	0	6	0	12,5	2	30,5	0	0	2,1	15	0	1,32	1
	Grand Total	12,78		1,07		7,3		0		6		14,5		30,5		2,1		15		2,32	

5.2 Use of resources

The use of personnel resources are summarized in Table 1. Note that the column marked by «PM budgeted» corresponds to 9 months out of a total of 24 months. Thus, while the total number of PMs for the project as a whole is 241 according to the overview given in Part A, only 89.6 budgeted PMs are accounted for in Table 1. The total number of available PMs for the final period is thus 3/24 times 241 PM, plus the 16 PMs saved from earlier periods, for a net total of 46 PMs. Nominally, considering that 83.47 PMs were spent during the previous 9-month period (see Table 1), this covers only about half of the required work effort required to complete the program. However, as noted in the Amendment application, once the BeyondPlanck resources have been spent, the remaining costs will be covered through external and in-kind funding.

Table 2 provides a similar overview, but additionally accounting for both 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*. We now have access to the C-BASS temperature data, and analysis is proceeding well. However, a careful assessment of systematic errors in the C-BASS data is essential, and preliminary results suggests that caution is warranted; we do see signatures of large-scale systematic effects. This will be assessed more fully in collaboration with the C-BASS team. We have also successfully analyzed S-PASS data, although it remains to be determined whether they will be integrated into BeyondPlanck; partial sky coverage and strong Faraday rotation represents some technical challenges that must be addressed before this can happen.
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*. During the last month, we have had (and still have) an issue with the Infiniband library, temporarily rendering intra-node analyses

impossible. As such, current work focus on half-mission data splits, which can be run on a single computing node. We expect this to be resolved shortly, as the problem has been identified by the IT support team, and a new Mellanox MPI library will be installed.

4. *Difficult in reproducing computational efforts*. This has not been an issue.
5. *Loss of personnel*. No one has withdrawn from the project during the project period. Rather, the team has grown significantly through voluntary contributions.
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 tool used for this is Commander3, which now also forms a backbone of BeyondPlanck processing. BeyondPlanck has as such greatly benefitted from the bits2cosmology code development efforts, while bits2cosmology has greatly benefitted from the BeyondPlanck project for access to raw Planck LFI, as well as important new ideas. Perhaps the most important of these is a new Gibbs sampling mapmaking idea proposed by Elina Keihänen, which now forms the backbone of the Commander3 TOD Gibbs sampler. Thus, ground-breaking synergies have resulted from collaboration between bits2cosmology and BeyondPlanck.

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. For now, the most direct benefit from Cosmoglobe for BeyondPlanck has been through the S-PASS and C-BASS analyses, which has been led by Dr. Unni Fuskeland and Mr. Daniel Herman, both of whom are now officially BeyondPlanck members.

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.