

Finanziato dall'Unione europea NextGenerationEU







BrahMap: Map-making for the future CMB experiments

Avinash Anand¹, Giuseppe Puglisi²

¹University of Rome "Tor Vergata", ²University of Catania

Spoke 3 III Technical Workshop, Perugia 26-29 Maggio, 2025

ICSC Italian Research Center on High-Performance Computing, Big Data and Quantum Computing









Scientific Rationale

- Future CMB experiments: Targeting the B-mode polarization of CMB
- Detectors: $O(10^3) O(10^5)$ in number with a very high sampling rate
- Data acquisition: ~250 TB (from space) to ~10 PB (from ground)
- First step of analysis: Reduction of time-series data to sky maps *aka* Map-making
- Map-making goals:
 - Reduction of enormous amount of data in a reasonable timeframe
 - Mitigation of instrumental systematics
 - Removal of both un-correlated and correlated noise



- First step of analysis: Reduction of time-series data to sky maps aka Map-making
- Map-making goals:
 - Reduction of enormous amount of data in a reasonable timeframe
 - Mitigation of instrumental systematics
 - Removal of both un-correlated and correlated noise









Technical Objectives, Methodologies and Solutions



ICSC Italian Research Center on High-Performance Computing, Big Data and Quantum Computing









Technical Objectives, Methodologies and Solutions



ICSC Italian Research Center on High-Performance Computing, Big Data and Quantum Computing









Technical Objectives, Methodologies and Solutions

BrahMap : A scalable and modular map-making framework for future CMB experiments

Target

- Hardware-level optimization to squeeze the most of the computational resources
- Scalability across multiple CPUs and compute nodes
- Utilizing complex noise covariance models
- Offloading the computations to multiple GPUs

Features

- An intuitive Python 3 interface for linear operators used in map-making
- Interface to create new linear operators
- High performance code base with vectorization, and MPI+OpenMP based hybrid parallelization
- C++ based backend for compute intensive routines









Milestone	Target		
M6	 Conversion of code base from Python 2 to Python 3 Debugging and validation 		
M7	 Writing computationally extensive parts to C++ with pybind11 11x performance gain over previous version 		
M8	 Identification of bottlenecks Code refactoring and vectorization 2x performance gain over previous version 		
M9	- Code parallelization with MPI		
M10	 OpenMP + MPI based hybrid parallelization Implementing more realistic noise covariance operators 		









OpenMP based parallelization

- Some integration tests were failing randomly
- Output maps in case of some unobserved pixels were bad sometimes
- Output maps for N_{side} > 256 were always corrupted



ICSC Italian Research Center on High-Performance Computing, Big Data and Quantum Computing







repixelize pol IQU()



Main Results

OpenMP based parallelization

- The function repixelize_pol_IQU() maps the observed pixel indices to the new pixel indices
- Loop iterations are not independent it fails when there are two or more unobserved pixels in single vector lane
- Commenting the pragma directive fixes everything

40			
49	<pre>template <typename dfloat="" dint,="" typename=""></typename></pre>		
50	<pre>void repixelize_pol_IQU(</pre>	11	
51	const ssize_t new_npix,	11	
52	<pre>const dint *restrict observed_pixels,</pre>	11	
53	<pre>dfloat *restrict weighted_counts,</pre>	11	
54	<pre>dfloat *restrict weighted_sin_sq,</pre>	11	
55	<pre>dfloat *restrict weighted_cos_sq,</pre>	11	
56	<pre>dfloat *restrict weighted_sincos,</pre>	11	
57	dfloat *restrict weighted_sin,	11	
58	dfloat *restrict weighted_cos,	11	
59	dfloat *restrict one_over_determinant	11	
60) {		
61			
62	<pre>#pragma omp parallel for simd</pre>		
63	<pre>for (ssize_t idx = 0; idx < new_npix; ++id</pre>	4x) {	
64	<pre>dint pixel = observed_pixels[idx];</pre>		
65	<pre>weighted_counts[idx] = weighted_counts[pixel];</pre>		
66	<pre>weighted_sin_sq[idx] = weighted_sin_sq[pixel];</pre>		
67	<pre>weighted_cos_sq[idx] = weighted_cos_sq[pixel];</pre>		
68	<pre>weighted_sincos[idx] = weighted_sincos[pixel];</pre>		
69	<pre>weighted_sin[idx] = weighted_sin[pixel]</pre>	8. 9	
70	<pre>weighted_cos[idx] = weighted_cos[pixel]</pre>	*	
71	one_over_determinant[idx] = 1.0 / one_ov	ver_determinant[pixel];	
72	} // for		
73			
74	return;		
75			









OpenMP based parallelization



ICSC Italian Research Center on High-Performance Computing, Big Data and Quantum Computing











































10⁹ rows









- 1. Noise covariance (and their inverse)
 - a. Diagonal noise covariance
 - b. Circulant noise covariance
 - c. Band-diagonal noise covariance
 - d. Toeplitz noise covariance (inverse computed through PCG)
- 2. Block diagonal noise covariance (incorporating the noise covariances listed above)













$$C = \begin{bmatrix} c_0 & c_{n-1} & c_{n-2} & \dots & c_2 & c_1 \\ c_1 & c_0 & c_{n-1} & \ddots & c_2 \\ c_2 & c_1 & c_0 & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & c_0 & c_{n-1} & c_{n-2} \\ c_{n-2} & & \ddots & c_1 & c_0 & c_{n-1} \\ c_{n-1} & c_{n-2} & \dots & c_2 & c_1 & c_0 \end{bmatrix}$$

- 1. Noise covariance (and their inverse)
 - a. Diagonal noise covariance
 - b. Circulant noise covariance
 - c. Band-diagonal noise covariance
 - d. Toeplitz noise covariance (inverse computed through PCG)
- 2. Block diagonal noise covariance (incorporating the noise covariances listed above)











$$B = \begin{bmatrix} b_0 & b_{-1} & \dots & b_{-s} & 0 & \dots & 0 \\ b_1 & b_0 & b_{-1} & \ddots & \ddots & \ddots & \vdots \\ \vdots & b_1 & b_0 & \ddots & \ddots & \ddots & 0 \\ b_r & & \ddots & \ddots & \ddots & b_{-s} \\ 0 & \ddots & \ddots & \ddots & b_0 & b_{-1} & \vdots \\ \vdots & \ddots & \ddots & b_1 & b_0 & b_{-1} \\ 0 & \dots & 0 & b_r & \dots & b_1 & b_0 \end{bmatrix}$$

- 1. Noise covariance (and their inverse)
 - a. Diagonal noise covariance
 - b. Circulant noise covariance
 - c. Band-diagonal noise covariance
 - d. Toeplitz noise covariance (inverse computed through PCG)
- 2. Block diagonal noise covariance (incorporating the noise covariances listed above)











$$T = \begin{bmatrix} t_0 & t_{-1} & t_{-2} & \dots & t_{-(n-2)} & t_{-(n-1)} \\ t_1 & t_0 & t_{-1} & \ddots & & t_{-(n-2)} \\ t_2 & t_1 & t_0 & \ddots & \ddots & & \vdots \\ \vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & t_0 & t_{-1} & t_{-2} \\ t_{n-2} & & \ddots & t_1 & t_0 & t_{-1} \\ t_{n-1} & t_{n-2} & \dots & t_2 & t_1 & t_0 \end{bmatrix}$$

- 1. Noise covariance (and their inverse)
 - a. Diagonal noise covariance
 - b. Circulant noise covariance
 - c. Band-diagonal noise covariance
 - d. Toeplitz noise covariance (inverse computed through PCG)
- 2. Block diagonal noise covariance (incorporating the noise covariances listed above)











Noise covariance operators

Validation against SANEPIC

- Output maps
- Power spectra



ICSC Italian Research Center on High-Performance Computing, Big Data and Quantum Computing









Final Steps

- Documentation update

June – J	luly	2025
----------	------	------

- Profiling OpenMP implementation
- KPI: Offloading to GPUs
 - cupy arrays on Python side
 - Exposing cupy arrays to C++ with Array API
 - OR use in cupy.RawKernel Python

August - Sept 2025

- Benchmarking with large simulations
 - Map-making with full focal plane simulation
- Preparing the final results