Mitigating correlated *1/f* noise in map-making

A comparison of approximate and exact noise covariance models

Challenges at the CMB Frontier

29 June - 1 July 2026
Suwa University of Science

Avinash Anand

University of Rome “Tor Vergata”
INFN – Tor Vergata
avinash.anand@roma2.infn.it

Map-making for CMB experiments

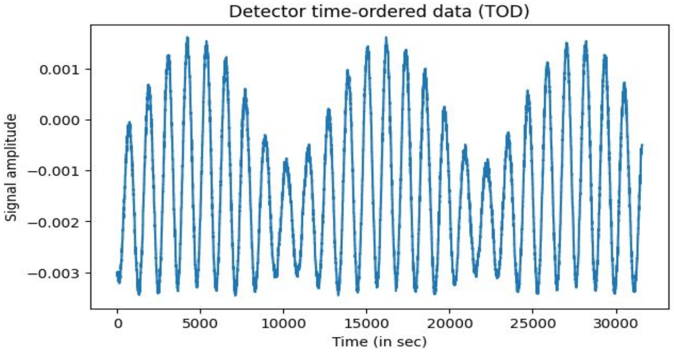

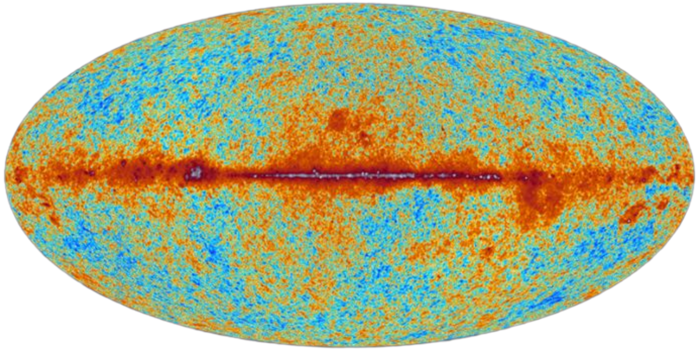

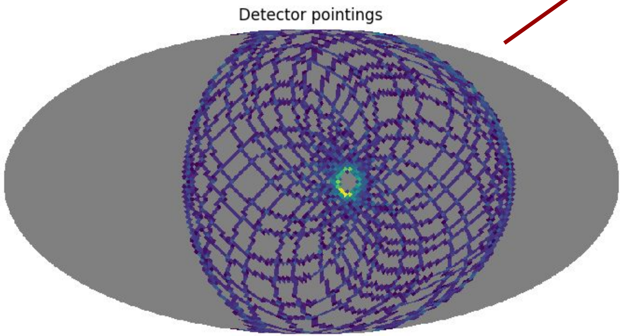

The Map-making step

- Mapping the time-series data onto the sky
- First scientific product of a CMB experiment

Map-making for CMB experiments

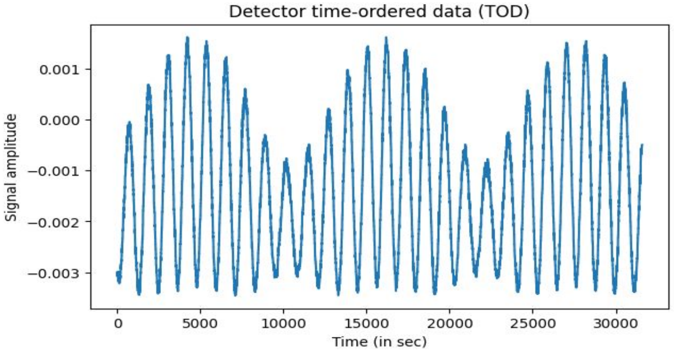

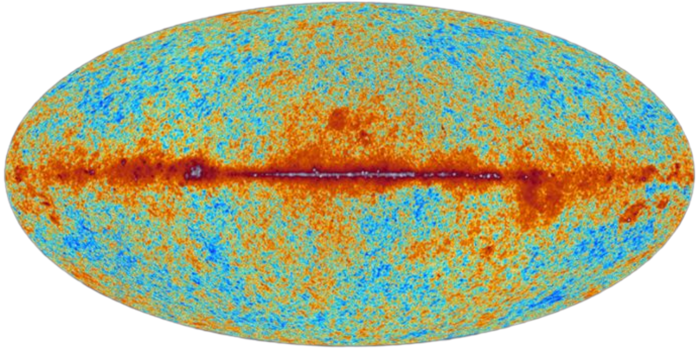

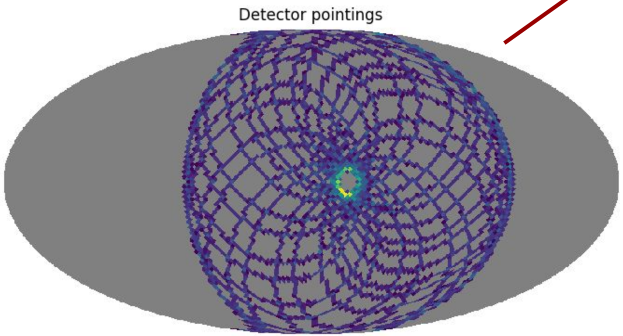

Goals

- Transfer all cosmological information from time-domain to pixel-domain
- Mitigation of instrumental systematics
- Removing both uncorrelated and correlated noise

Computing the sky map

Data model in matrix equation form

$$d = P \cdot s + n$$

- $d \rightarrow$ Signal vector
- $n \rightarrow$ noise vector
- $s \rightarrow$ Sky map vector
- $P \rightarrow$ Pointing matrix (sparse matrix with 3 non-zero elements per row)

General unbiased linear solution

$$\hat{s} = \left(P^T M^{-1} P\right)^{-1} P^T M^{-1} d$$

- $M \rightarrow$ A symmetric positive-definite matrix

$$\left(P^T M^{-1} P\right) \cdot \hat{s} = P^T M^{-1} d \longleftrightarrow A \cdot x = b$$

$$A = P^T M^{-1} P \qquad x = \hat{s} \qquad b = P^T M^{-1} d$$

Solve $A \cdot x = b$ using iterative linear solvers, like PCG (pre-conditioned conjugate gradient)

Correlated noise in CMB experiments

The 1/f noise model

$$P(f) = \sigma^2 \left[1 + \left(\frac{f_{knee}}{f}\right)^{\alpha}\right]$$

σ → white noise level

f_{knee} → knee frequency

α → spectral index

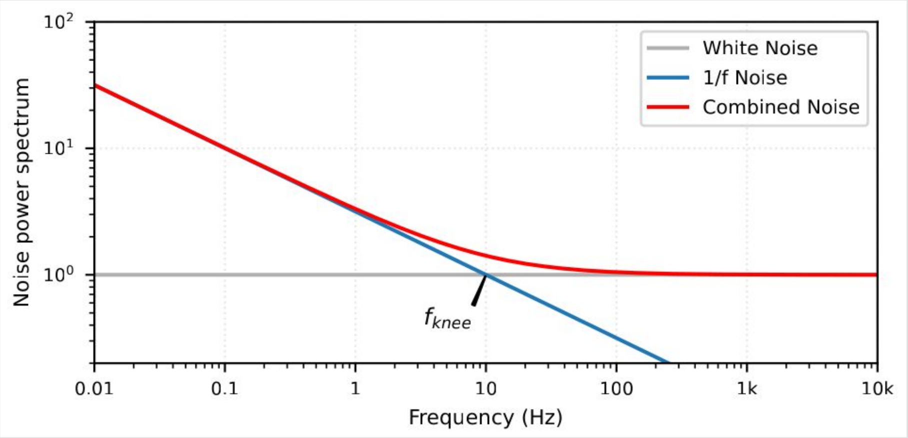

Noise parameters

Spectral Index (α)

- $\alpha \sim 1$: Thermal instabilities in electric circuitry
- $\alpha \sim 2$: Atmospheric fluctuations (ground-based)
- **Larger** α: Higher noise at lower frequencies

Knee Frequency (f_{knee})

- Determines the transition from *1/f* to white noise
- **Higher f_{knee}**: Larger noise correlation length

Correlated noise in CMB experiments

General unbiased linear solution $$\hat{s} = \left(P^T M^{-1} P\right)^{-1} P^T M^{-1} d$$

1. TOD to noise correlation in map

Any choice of matrix *M* leaves correlated noise in the pixel domain

$$N_{pp}(M) = (P^T M^{-1} P)^{-1} P^T M^{-1} N M^{-1} P (P^T M^{-1} P)^{-1}$$

2. Map to power spectrum residual

Correlated noise in map leaves non-zero noise residuals in the power spectra

$$N_{pp} = \frac{1}{4\pi} \sum_{\ell=0}^{\infty} (2\ell+1) P_\ell(\cos\theta) C_\ell$$

Choosing the weight matrix *M*

- Minimizes the noise residuals in power spectra → **The noise covariance matrix in time-domain**
- Computationally efficient for large TODs

(Tegmark, 1997)

Noise covariances in map-making

Stationary observation period

- The noise covariance results in a **symmetric Toeplitz structure**

$$T = \begin{bmatrix} t_0 & t_1 & t_2 & \cdots & t_{n-3} & t_{n-2} & t_{n-1} \\ t_1 & t_0 & t_1 & \cdots & t_{n-4} & t_{n-3} & t_{n-2} \\ t_2 & t_1 & t_0 & \cdots & t_{n-5} & t_{n-4} & t_{n-3} \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ t_{n-3} & t_{n-4} & t_{n-5} & \cdots & t_0 & t_1 & t_2 \\ t_{n-2} & t_{n-3} & t_{n-4} & \cdots & t_1 & t_0 & t_1 \\ t_{n-1} & t_{n-2} & t_{n-3} & \cdots & t_2 & t_1 & t_0 \end{bmatrix}$$

- Computing $T^{-1}x$ is **computationally expensive**
- Feasible approximation of T → **circulant matrices**

Circulant approximations of noise covariance

Circulant matrix

A special class of Toeplitz matrix where the each row is a right cyclic shift of the previous row

$$C = \begin{bmatrix} c_0 & c_{n-1} & c_{n-2} & \cdots & c_3 & c_2 & c_1 \\ c_1 & c_0 & c_{n-1} & \cdots & c_4 & c_3 & c_2 \\ c_2 & c_1 & c_0 & \cdots & c_5 & c_4 & c_3 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ c_{n-3} & c_{n-4} & c_{n-5} & \cdots & c_0 & c_{n-1} & c_{n-2} \\ c_{n-2} & c_{n-3} & c_{n-4} & \cdots & c_1 & c_0 & c_{n-1} \\ c_{n-1} & c_{n-2} & c_{n-3} & \cdots & c_2 & c_1 & c_0 \end{bmatrix}$$

Key properties

- If C is circulant, then C^{-1} is also circulant
- C is diagonalized by the Fourier operators

$$C = \mathcal{F}^{-1}\Lambda\mathcal{F}$$

$$C^{-1} = \mathcal{F}^{-1}\Lambda^{-1}\mathcal{F}$$

Matrix vector multiplication via FFTs

Matrix-vector multiplication requires 2 DFTs

$$\begin{aligned} Cv &= \mathcal{F}^{-1}\Lambda\mathcal{F}v \\ &= (\mathcal{F}^{-1}(\Lambda(\mathcal{F}v))) = (\mathcal{F}^{-1}(\Lambda\tilde{v})) = \mathcal{F}^{-1}\bar{v} \end{aligned}$$

Computational complexity → *O(n log n)*

(Chan, R., et al., 1996)

Circulant approximations of noise covariance

$$T = \begin{bmatrix} t_0 & t_1 & t_2 & \cdots & t_{n-3} & t_{n-2} & t_{n-1} \\ t_1 & t_0 & t_1 & \cdots & t_{n-4} & t_{n-3} & t_{n-2} \\ t_2 & t_1 & t_0 & \cdots & t_{n-5} & t_{n-4} & t_{n-3} \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ t_{n-3} & t_{n-4} & t_{n-5} & \cdots & t_0 & t_1 & t_2 \\ t_{n-2} & t_{n-3} & t_{n-4} & \cdots & t_1 & t_0 & t_1 \\ t_{n-1} & t_{n-2} & t_{n-3} & \cdots & t_2 & t_1 & t_0 \end{bmatrix}$$

Standard circulant matrix

$$C_{std} = \begin{bmatrix} t_0 & t_1 & t_2 & \cdots & t_3 & t_2 & t_1 \\ t_1 & t_0 & t_1 & \cdots & t_4 & t_3 & t_2 \\ t_2 & t_1 & t_0 & \cdots & t_5 & t_4 & t_3 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ t_3 & t_4 & t_5 & \cdots & t_0 & t_1 & t_2 \\ t_2 & t_3 & t_4 & \cdots & t_1 & t_0 & t_1 \\ t_1 & t_2 & t_3 & \cdots & t_2 & t_1 & t_0 \end{bmatrix}$$

Strang matrix

Minimizer of $\|C - T\|_1$ and $\|C - T\|_\infty$ where $C \in$ set of all Hermitian circulant matrices

T. Chan matrix

Minimizer of $\|C - T\|_F$ where $C \in$ set of all circulant matrices

R. Chan matrix

Embed the Toeplitz matrix into a larger circulant matrix

$$\begin{bmatrix} T & B \\ B & T \end{bmatrix}$$

$$C_{RChan} = T + B$$

(Chan, R., et al., 1996)

Circulant approximations of noise covariance

LiteBIRD PTEP configuration (M-119 channel)

- 6 detectors
- 3 years
- 19 Hz sampling
- 24 hrs stationary period

$$P(f) = \sigma^2 \left[1 + \left(\frac{f_{knee}}{f}\right)^{\alpha}\right]$$

Case:
α = 1

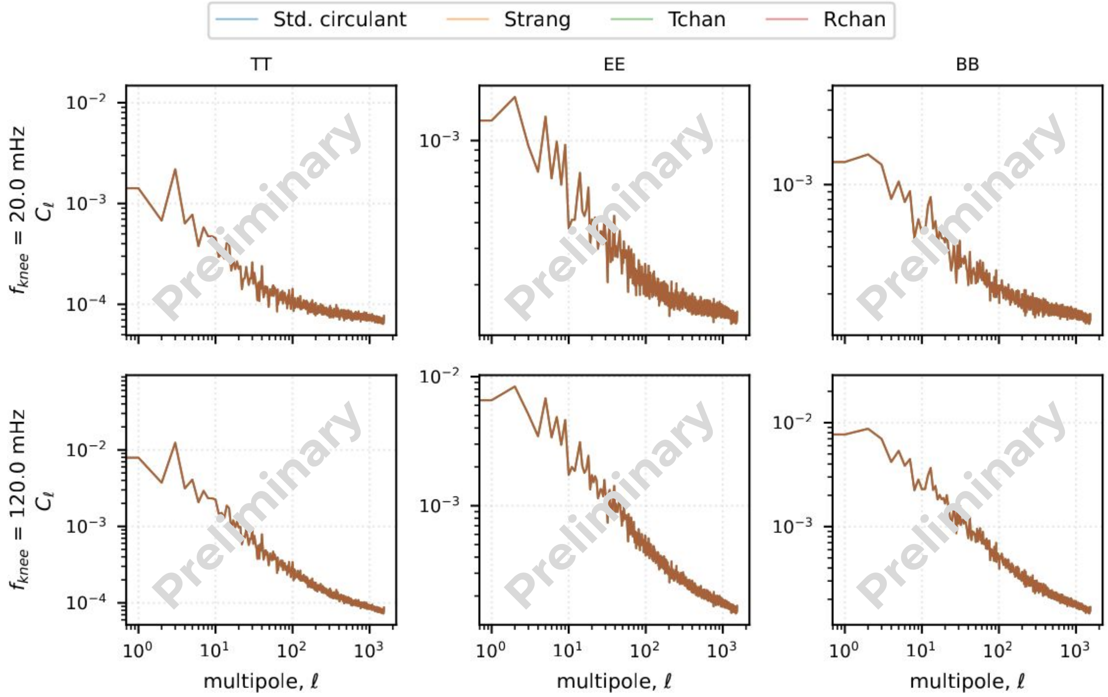

Circulant approximations of noise covariance

LiteBIRD PTEP configuration (M-119 channel)

- 6 detectors
- 3 years
- 19 Hz sampling
- 24 hrs stationary period

$$P(f) = \sigma^2 \left[1 + \left(\frac{f_{knee}}{f}\right)^{\alpha}\right]$$

Case:
α = 2

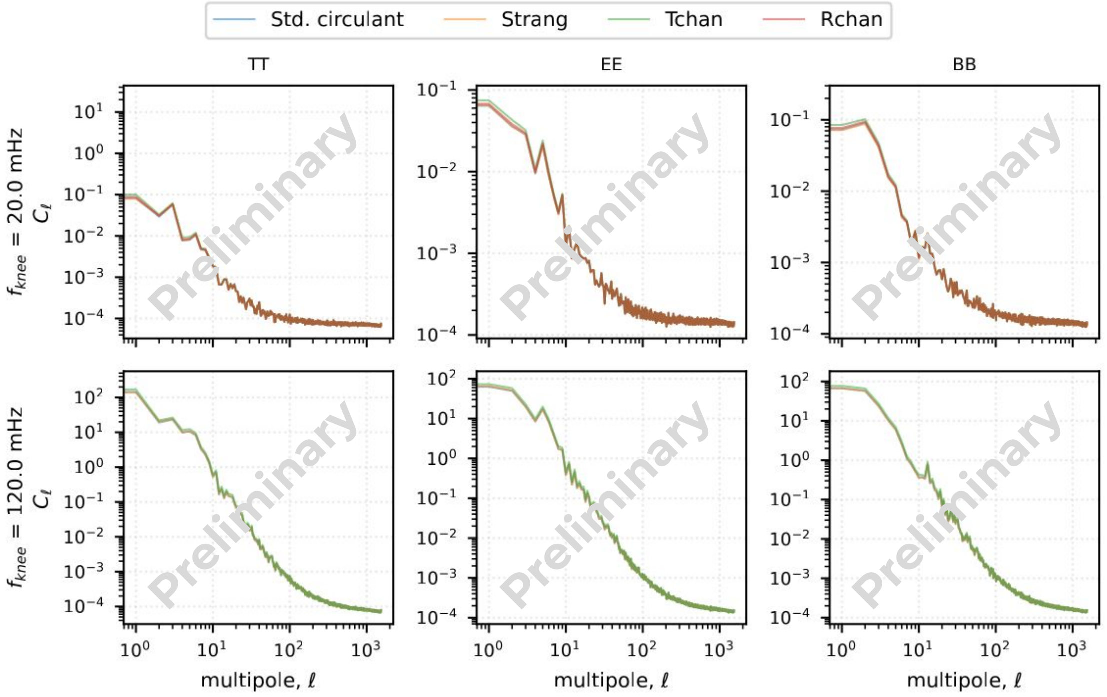

Full Toeplitz noise covariance implementation

PCG based Toeplitz solver with circulant preconditioner

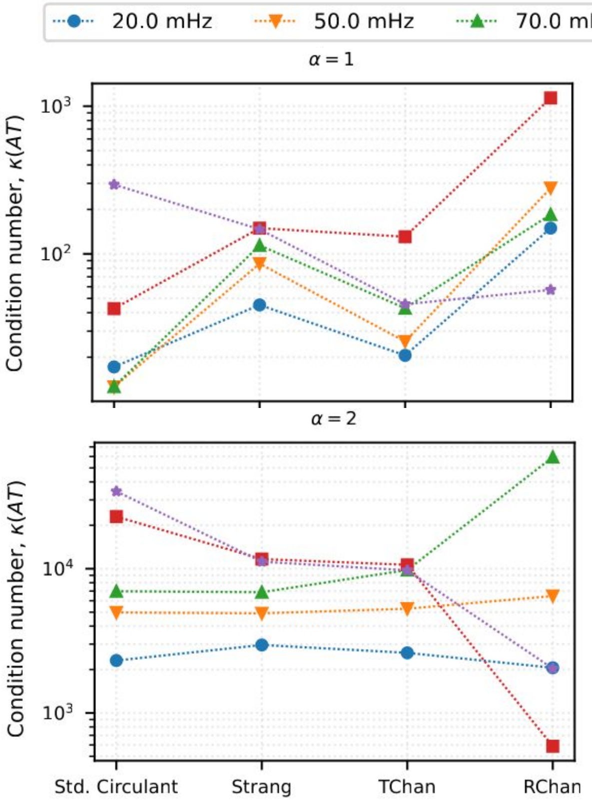

$$P(f) = \sigma^2 \left[1 + \left(\frac{f_{knee}}{f}\right)^{\alpha}\right]$$

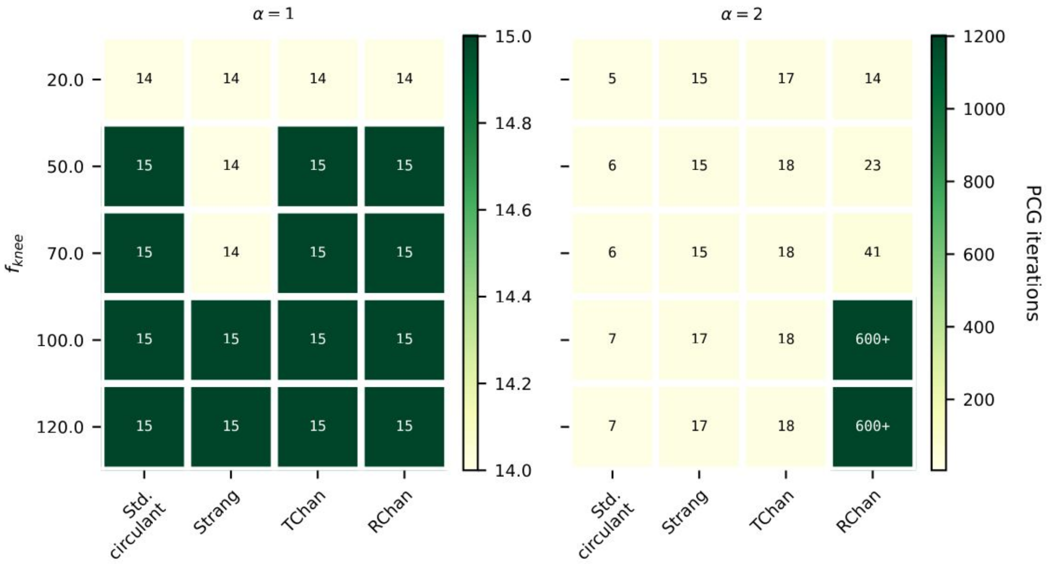

Noise residual power comparison

LiteBIRD PTEP configuration (M-119 channel)

- 6 detectors
- 3 years
- 19 Hz sampling
- 24 hrs stationary period

$$P(f) = \sigma^2 \left[1 + \left(\frac{f_{knee}}{f}\right)^{\alpha}\right]$$

Case:
α = 1

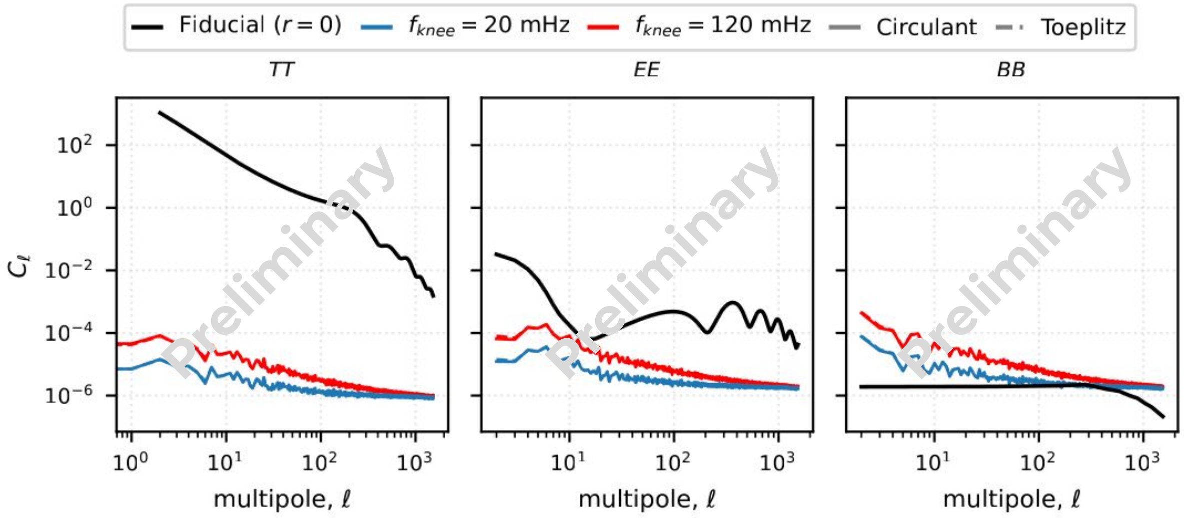

Noise residual power comparison

LiteBIRD PTEP configuration (M-119 channel)

- 6 detectors
- 3 years
- 19 Hz sampling
- 24 hrs stationary period

$$P(f) = \sigma^2 \left[1 + \left(\frac{f_{knee}}{f}\right)^{\alpha}\right]$$

Case:
α = 2

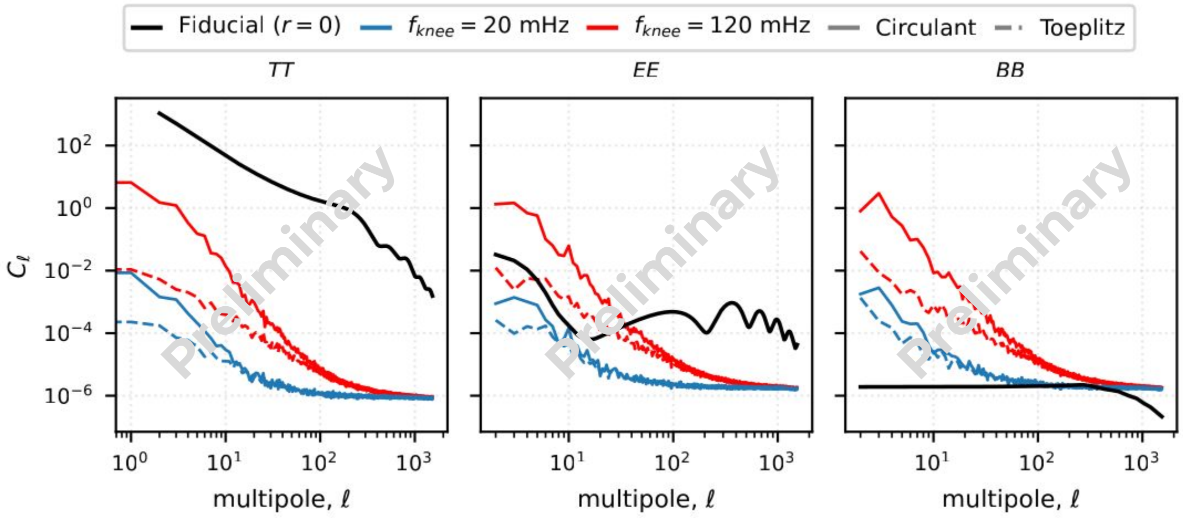

Noise residual maps comparison

Noise residual maps

f_{knee} = 120 mHz, $\alpha = 2$

Circulant approximation leaves prominent large scale residuals

→ large noise residual power

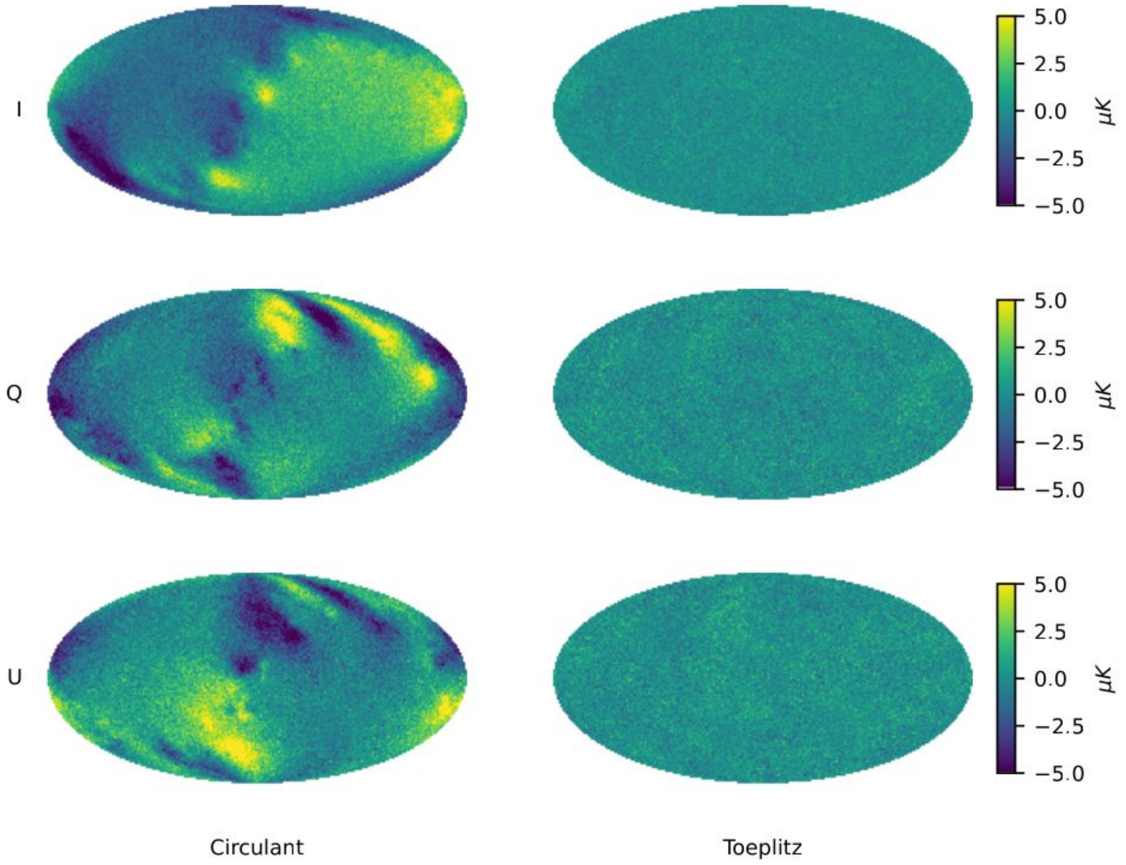

Bias on the estimation of tensor-to-scalar ratio, *r*

Defining the bias *Δr*

The bias Δr induced by 1/f noise is the value that maximizes the likelihood function:

$$\frac{dL(r)}{dr}\bigg|_{r=\Delta r} = 0$$

$$\log L(r) = \sum_{\ell=\ell_{min}}^{\ell_{max}} \log P_\ell(r)$$

$$\log P_\ell(r) = -f_{sky}\frac{2\ell+1}{2}\left[\frac{\hat{C}_\ell}{C_\ell} + \log C_\ell - \frac{2\ell-1}{2\ell+1}\log \hat{C}_\ell\right]$$

(LiteBIRD collaboration, 2023)

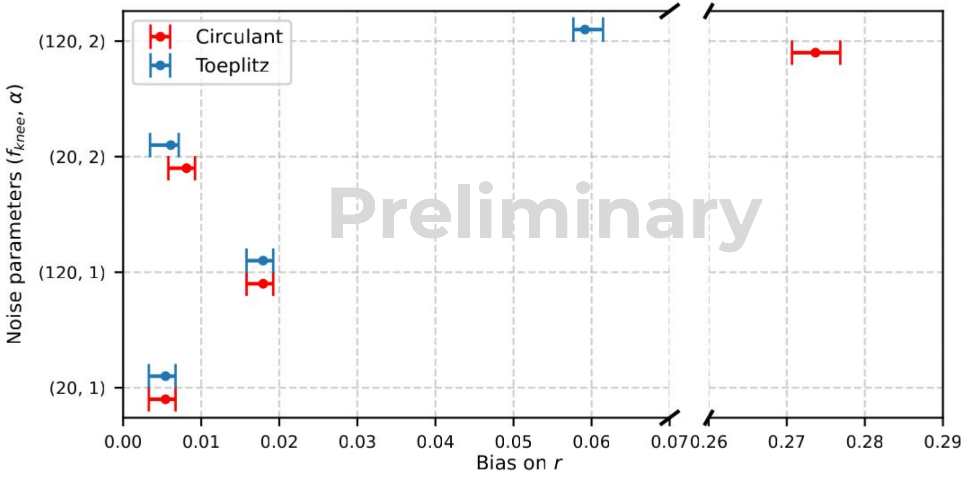

- Measured power spectrum

$$\hat{C}_\ell = C_\ell^{lens} + N_\ell$$

- Modeled power spectrum

$$C_\ell = rC_\ell^{tens} + C_\ell^{lens}$$

Performance consideration

Performance comparison

- Convergence properties for both cases are almost same

Map-making with full Toeplitz implementation is 6-12x slower than with circulant approximation

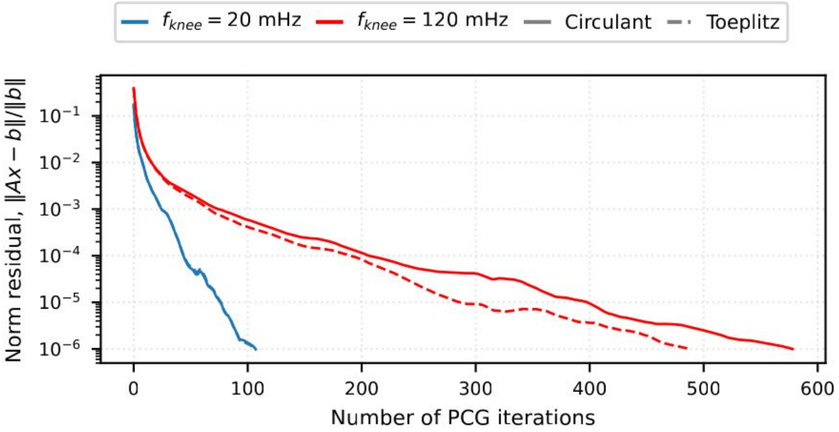

GPU offloading as solution

- GPU accelerated Toeplitz solver → offloading FFT computations to GPU
- Possible to compute the maximum-likelihood map with GPUs within the time-scale of circulant matrix based solver on CPU nodes

Summary

Noise treatment at map-making step

- Circulant approximations of noise covariance → **Similar performance across all approximations**
- Full Toeplitz solver implementation → **Reduces noise residual power by 2 orders of magnitude and bias on the tensor-to-scalar ratio significantly**

Implementation details

- All the circulant approximations and full Toeplitz noise covariance

  ➜ Implemented in **BrahMap** **(arxiv: 2501.16122) (Github: anand-avinash/BrahMap)**