2019

Using 3-dimensional FFTs to simulate CMB maps on a spherical cap

Eric Goetz
University of Richmond

Using 3-Dimensional FFTs to Simulate CMB Maps on a Spherical Cap

by Eric Goetz

Honors Thesis

Submitted to :

Physics Department
University of Richmond
Richmond, VA

March 27, 2019

Advisor: Dr. Ted Bunn

The current methods of simulating the Cosmic Microwave Background (CMB) involve either simulating the entire sky using spherical transforms or simulating a flat patch with fast Fourier transforms (FFTs). For patches that are too large to be considered flat but much less than the full sky, the former method is inefficient and the latter is inaccurate. One alternative method of CMB simulation is to simulate the random processes behind the CMB in a 3-dimensional box that contains the part of the sphere that we want to measure. Then, we can select the points we want from the box. This method should be more efficient than previous methods because it performs simulations over a box instead of a sphere, allowing for the use of FFTs in place of much slower spherical harmonic transforms.

For this method to work, there must be a 3-dimensional power spectrum defined on the box that has the same correlation function as the angular power spectrum. Since the angular power spectrum is known, this becomes a linear programming problem, where the constraints for the 3-D power spectrum are that it matches the angular power spectrum over the observed region and that it be non-negative. If a power spectrum satisfying these constraints exists, we can use it to create maps of the CMB with the same statistical properties as the observed CMB. These maps can then be used to test theories about the early universe.

We have solved this linear programming problem for a variety of realistic scenarios. We have performed extensive statistical tests comparing the statistics of maps produced via FFT simulation with maps produced via standard spherical harmonic methods, and found the two methods to be statistically indistinguishable.

I. INTRODUCTION

Studying the cosmic microwave background requires making many simulations of the CMB, which are then used to test theories. Making many simulations is necessary because we believe that our observed CMB is the result of a random process. Theories about the early universe will make predictions about the features of the cosmic microwave background that would have occurred if the theory is correct. These theories will not predict precise values of the CMB; rather, they predict the statistical properties. We therefore do not want to test our theories against the exact values we measure, but instead against the statistical properties of the CMB we observe. For this reason, we want to make large numbers of maps, each drawn from the predicted distribution.

The focus of this project is on simulating partial sky cosmic microwave background maps. Because the CMB is isotropic and homogeneous, we can interpret these partial sky maps as maps covering a spherical cap of the sky. There are two current methods of simulating a map over a spherical cap. The first is by using a flat sky approximation. This method ignores the curvature of the surface and treats the spherical cap as a flat plane. The issue with this method is that it only produces accurate results for extremely small patches of the sky. These patches are too small to be of practical use in many cases. The second method in use today is to simulate the map over the entire sphere, then select just the points on the desired spherical cap. This method produces accurate maps but is very inefficient. It requires using computationally expensive spherical harmonic transforms over the entire sphere [1]. With new experiments only taking partial-sky measurements of the CMB on the horizon, there is a need for a more efficient method of making these simulations.

The idea behind this project is to simulate the map first in a three-dimensional box, and then select the points in the box that correspond to the points on the spherical cap. We do this so that we can use Fourier transforms instead of the slower and more computationally expensive spherical harmonic transforms. We do this by first finding a three-dimensional power spectrum inside the box that has the same correlation function as the observed angular power spectrum over the sphere. We then use this power spectrum to make maps inside the box. Finally, we select the points inside that box that correspond to points on the sphere. The diagram in figure 1 demonstrates the idea behind this project.

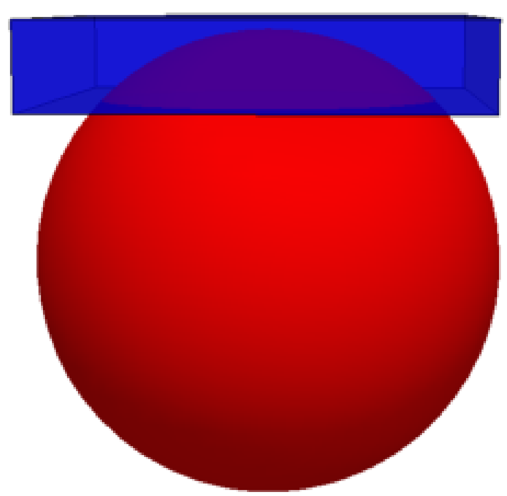

FIG. 1. Diagram of process. We first simulate the map inside the blue box. We then select the points in the box that lie on the red sphere. This becomes the spherical cap CMB map.

II. BACKGROUND INFORMATION

The cosmic microwave background is composed of photons that were emitted by the universe roughly 380,000 years after the Big Bang. These photons are the oldest light in the universe that we can observe. They form a surface of a sphere centered on the Earth called the last scattering surface. Because this light comes from right after the Big Bang, it contains information about the density and composition of the early universe [2]. This in turn gives us information about structure formation and the development of the universe.

After the Big Bang, the universe was very hot and dense [3]. All baryonic matter was ionized, so there was a high number of free electrons. Photons could not travel far without hitting these free electrons and scattering. As the universe cooled down, ions began collecting electrons. This point is called the epoch of recombination. This allowed the photons to travel unimpeded until the present day.

There have been several attempts to measure the cosmic microwave background. It was first measured by Arno Penzias and Robert Wilson in the 1960's. They could observe that the CMB was present but could not make a detailed measurement. The first full-sky map was made by the COBE satellite in the 1990's. Since then, more precise measurements have been made by satellites like WMAP and Planck. A map measured by the Planck satellite is shown in figure 2.

The cosmic microwave background fits a blackbody spectrum with a temperature of 2.7255 ± 0.0006 K. At the epoch of recombination, the temperature of the universe was 2970 K. The drop in temperature is due to the expansion of the universe. The blackbody spectrum of the cosmic microwave background is shown in figure 3.

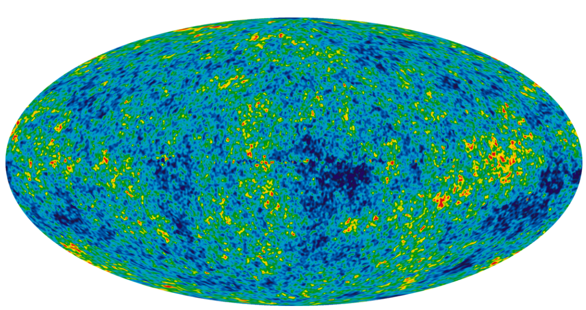

FIG. 2. Map of CMB from Planck satellite
https://en.wikipedia.org/wiki/Cosmic_microwave_background

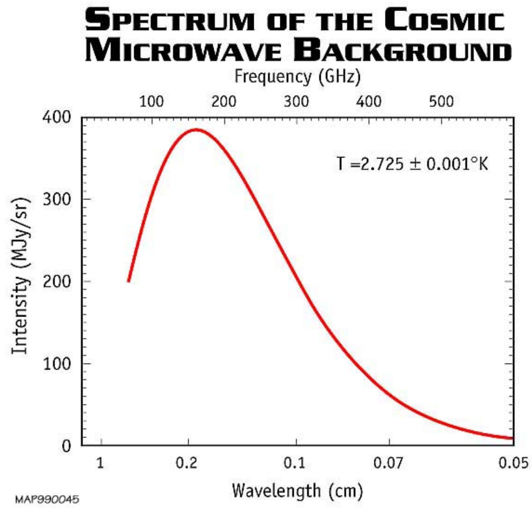

FIG. 3. Blackbody spectrum of the CMB.
https://wmap.gsfc.nasa.gov/universe/bb_tests_cmb.html

III. BASIC SETUP AND NOTATION

The data set that we used in this project consists of measurements of the angular power spectrum of the cosmic microwave background. The temperature of the CMB was measured by the Planck satellite. These temperatures are described by a function on the surface of last scattering

$$T(\theta,\phi) = \sum_{l=0}^{\infty}\sum_{m=0}^{\infty} \tilde{T}_{lm} Y_{lm}(\theta,\phi)$$

Here, Y_{lm} are the spherical harmonics and $\tilde{T}_{lm}$ are the coefficients of each spherical harmonic. In the standard theory of the CMB, the temperature measured is a realization of a Gaussian random process, so each $\tilde{T}_{lm}$ is an independent normally distributed random number with mean 0 and variance determined by $< |\tilde{T}_{lm}|^2 >$. We can define the angular power spectrum

$$C_l = \left\langle |\tilde{T}_{lm}|^2 \right\rangle$$

The standard theory of the Cosmic Microwave Background also states that the CMB is homogeneous and isotropic on large scales. In the spherical harmonics, l is related to the wave number k and m represents the direction of the wave. If the CMB is homogeneous and isotropic, we do not need to worry about the direction of the wave, so we can ignore m.

The data we used did not include the monopole and dipole, so it started at $l = 2$. We then set the monopole and dipole terms ($l = 0, 1$) to 0.

In this project, we also need to find a power spectrum in a 3-dimensional box. We can write the function for the temperature of the CMB in the box as

$$T(\vec{r}) = \sum_{\vec{k}}^{\infty} \tilde{P}_{\vec{k}} e^{i\vec{k}\cdot\vec{r}}$$

We can also define the power spectrum for this box

$$P_{\vec{k}} = \left\langle |\tilde{T}_{\vec{k}}|^2 \right\rangle$$

These two power spectra can be related in the following way:

$$C_l = \int_0^\infty P(k) j_l^2(k) k^2 dk$$

Here, the j_l are the spherical Bessel functions.

In this way we can make the two power spectra match over the entire sphere.

IV. CORRELATION FUNCTIONS

The two point correlation function describes how two different points on the sky are related. If two points have a high correlation, we expect that if the temperature measured at one point is higher than average, the temperature at the second point will also be high. Conversely, a low correlation means that the temperature at one point has little effect on the temperature at the second point. For the cosmic microwave background, we expect that two points very close to each other would have a high correlation and two points far away from each other would have a low correlation. We can calculate the correlation function for each of the power spectra that we use.

For the angular power spectrum C_l, we can find the correlation function by taking the spherical harmonic transform of C_l.

$$\xi(\hat{r}_1, \hat{r}_2) = \sum_{l=0}^{\infty} \frac{2l+1}{4\pi} C_l P_l(\hat{r}_1 \cdot \hat{r}_2)$$

Here, $\hat{r}_1$ and $\hat{r}_2$ are the unit vectors corresponding to the two points. P_l is the lth Legendre polynomial. As mentioned in the previous section, we can ignore m because we are assuming that the CMB is isotropic and homogeneous.

We can also find the correlation function corresponding to the 3-dimensional power spectrum P_k by following a similar procedure. In this case, we can use Fourier transforms because we are working in Cartesian coordinates.

$$\xi(\Delta\vec{r}) = \int P(\vec{k}) e^{i\vec{k}\cdot\Delta\vec{r}} d^3k$$

We can then express k in spherical coordinates. Integrating over the angles yields

$$\xi(\Delta\vec{r}) = 4\pi \int_0^\infty P(k) \frac{\sin(k\Delta r)}{k\Delta r} k^2 dk$$

In order to create identical maps with the 3-dimensional power spectrum, we want to have the two correlation functions equal each other over the range of the sphere that we are trying to simulate. The specific setup of this problem is discussed in the next section.

V. LINEAR PROGRAMMING

The goal of this project is to simulate the CMB in a spherical cap of some angle θ_{max}. To convert this spherical cap to Cartesian space, we need to create a 3-dimensional box. The size of the box is dependent on the size of the spherical cap. We have found that the size of the box must be at least twice as big as the spherical cap in the x and y directions. The height of the box can be smaller than the x and y directions, so the box need not be a cube. The height must be large enough to encompass the entire spherical cap.

Once we have the size of the box, we can calculate the values of the correlation function at each point inside the box. We do this with the correlation function corresponding to the angular power spectrum C_l. These values will form the right hand side of the linear programming problem we will solve.

The linear programming problem we are trying to solve can be expressed as:

$$Ax \geq b$$

$$Gx \geq 0$$

Here, x is the 3-dimensional power spectrum we want to find. As mentioned above, b is the list of correlation function values that correspond to the points in the box that we want to simulate. A is the Fourier transform operator. This is because we want to match the correlation functions, and the Fourier transform of the power spectrum $P(k)$ is the correlation function $\xi(\Delta r)$.

G is the identity matrix. The purpose of this is to ensure that the power spectrum is non-negative at every point.

When we solve this linear programming problem, we want to minimize the difference between Ax and b.

This ensures that the two correlation functions match as closely as possible, subject to the constraint that $x \geq 0$.

The code for this project was done in Python and the linear programming problem was solved using the CVXOPT solver [4].

The CVXOPT solver returns the values of $P(k)$ for integer values of k. This is a one-dimensional function of a scalar number. Because the CMB is homogeneous and isotropic, we can assume that the value of $P(\vec{k})$ is the same for each direction of $\vec{k}$. This means that instead of being a function of the vector $\vec{k}$, the power spectrum is a function of the magnitude of $\vec{k}$. We can express this as $P(|\vec{k}|)$.

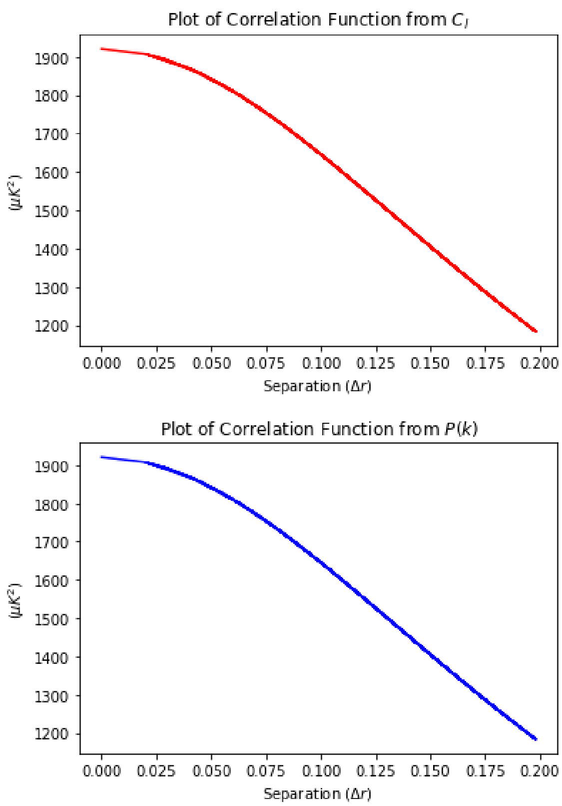

FIG. 4. Top: $\xi(\theta)$, correlation function for observed angular power spectrum. Bottom: $\xi(\Delta r)$ correlation function for generated three-dimensional power spectrum. The two correlation functions match over the region of the sky that we are trying to simulate, so the linear programming solver did find a valid solution.

Now that we have a one-dimensional power spectrum $P(|\vec{k}|)$, we can find the three-dimensional power spectrum by creating a box in k space and filling it with values from $P(|\vec{k}|)$. For each point in the box, we calculate the magnitude of the vector $\vec{k}$ corresponding to its position. We can then use interpolation to find the value of $P(|\vec{k}|)$ at every point in the box. We can then use this three-dimensional power spectrum to make spherical cap maps of the CMB.

Before we can use this three-dimensional power spectrum, we must make a correction. Even though one of the constraints in the linear programming solver is that $P(k)$ must be greater than or equal to 0, some values of $P(k)$ come out to be slightly negative due to numerical inaccuracies. This would cause issues later on when we create the maps. To fix this issue, after creating the three-dimensional box with $P(|\vec{k}|)$ values, we set all negative values to 0. Any negative values are on the order of 10^{-6} or smaller, so setting them to 0 would not have a noticeable effect on our results. It does, however, avoid issues later on in the process.

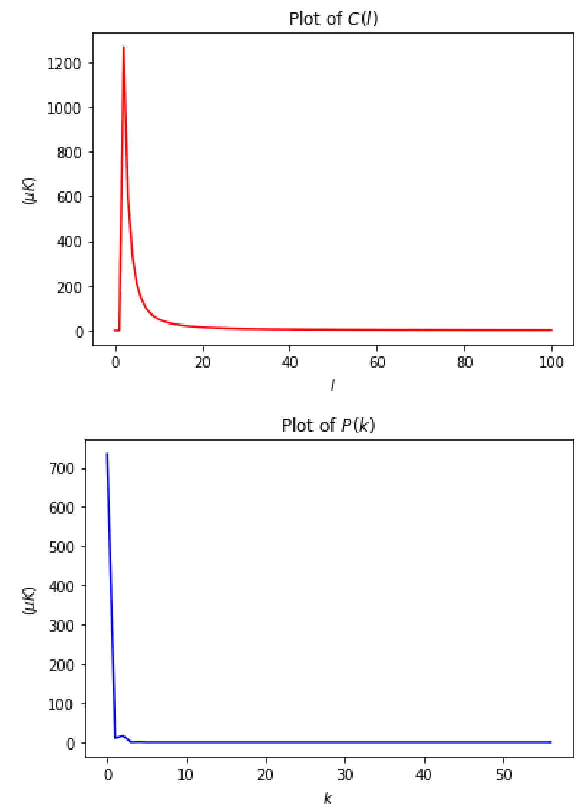

FIG. 5. Top: $C(l)$, observed angular power spectrum. Bottom: $P(k)$, generated three-dimensional power spectrum.

When we solve the linear programming problem, the generated three-dimensional power spectrum often does not match the angular power spectrum. This is expected because of the method we used to find the three-dimensional power spectra. The constraints on the linear programming problem concerned the matching of the two correlation functions, not the power spectra. This is not an issue because the goal of this project is to make maps of the CMB. We can use any three-dimensional power spectrum to do so as long as it has the correct correlation function.

VI. USING THE POWER SPECTRUM

The method outlined in the previous section allows us to find a 3-dimensional power spectrum that matches the observed angular power spectrum over a designated spherical cap. We now can use this power spectrum to generate CMB maps. To do this, we need to convert our three-dimensional box in Fourier space to a three-dimensional box of map values in real space. We start by taking the one-dimensional Fourier transform of the three-dimensional power spectrum in the z direction. This converts the three-dimensional box in Fourier space to a three-dimensional box in (k_x, k_y, z) space. This is done because we do not need to simulate the entire box in the z direction; we only need to simulate enough to encompass the spherical cap. We needed padding to ensure that the linear programming solver could find a valid solution. Once we have the solution, we no longer need the padding. This reduces the size of the problem, helping with the speed of the process.

Once we have reduced the size of the box in the z direction, we can begin the process of calculating the map values. The standard theory of the cosmic microwave background states that the numbers that we measure are normally distributed random numbers with mean $\mu = 0$ and variance given by the power spectrum. To make maps, we need to generate random numbers with mean $\mu = 0$ and variance given by our generated three-dimensional power spectrum.

We do this with a layer-by-layer procedure. Each of the random numbers in our map should be uncorrelated in Fourier space. In real space, they should be correlated with each other. Because we are generating them in (k_x, k_y, z) space, the random numbers we generate should be correlated in the z direction only. We can accomplish this by generating a covariance matrix Γ for each (k_x, k_y) pair in our three-dimensional box. These covariance matrices depend on z and are different for each (k_x, k_y) pair. We generate these covariance matrices in the following way using the definition of a covariance matrix:

$$\Gamma^{k_x,k_y}_{z_1,z_2} = \left\langle \tilde{T}_1(k_x, k_y, z_1)\tilde{T}_2{}^*(k_x, k_y, z_2)\right\rangle$$

$$\Gamma^{k_x,k_y}_{z_1,z_2} = \left\langle \left(\sum_{k_z} \tilde{T}_1(k_x, k_y, k_z)e^{ik_z z_1}\right) \left(\sum_{k'_z} (\tilde{T}_2(k_x, k_y, z_2)e^{ik_z z_2})^*\right)\right\rangle$$

$$\Gamma^{k_x,k_y}_{z_1,z_2} = \sum_{k_z,k'_z} \left\langle \tilde{T}(k_x, k_y, k_z)\tilde{T}^*(k_x, k_y, k_z)e^{i(k_z z_1 - k'_z z_2)}\right\rangle$$

We can rewrite $\left\langle \tilde{T}(k_x, k_y, k_z)\tilde{T}^*(k_x, k_y, k_z)f\right\rangle$ as $P(k_x, k_y, k_z)$, which is the calculated three-dimensional power spectrum. So,

$$\Gamma^{k_x,k_y}_{z_1,z_2} = \sum_{k_z} P(k_x, k_y, k_y)e^{ik_z(z_1 - z_2)}$$

This is the Fourier transform of $P(k)$ in the z direction. This means that the covariance matrix Γ is a Toeplitz matrix of the form

$$\Gamma^{(k_x,k_y)} = \begin{pmatrix} f(0) & f(1) & f(2) & \dots \\ f(1) & f(0) & f(1) & \dots \\ f(2) & f(1) & f(0) & \dots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix}$$

where f is the one-dimensional Fourier transform in the z direction of the three-dimensional power spectrum. Finally, we must add small multiples of the identity matrix to each covariance matrix to ensure that each matrix is positive definite. Due to numerical inaccuracies in the calculations, some of the eigenvalues of the Γ matrices are 0 or slightly negative. Covariance matrices must be positive definite, so we get around this issue by adding the small multiples of the identity. These multiples are on the order of 10^{-10}, so they do not have an effect on the covariances of the random numbers.

Once we have a Γ for each (k_x, k_y) pair, we can use them to generate the map values. We start by generating a vector $\vec{z}$ of n normally distributed random numbers with mean 0 and variance 1, where n is the number of points in the z direction that we have in the box. We then need to apply the correct covariances to these numbers. We do this by first taking the Cholesky decomposition of each Γ matrix, decomposing it into LL^T. Then, we calculate $x = Lz$ to get a vector $\vec{x}$ of n random normally distributed random numbers with mean 0 and covariances given by Γ.

We now can take a two-dimensional Fourier transform in x and y of these random numbers to convert from (k_x, k_y, z) space to real space. This gives a box of random numbers in real space that we can use as the values in the map.

The next step in this process is to get the coordinates of each pixel that we want to include in the spherical cap map. The number and locations of these pixels depends on the size of the spherical cap and the resolution of the map. We can use HealPy to find the vector coordinates of each pixel on the sphere that we need to simulate. We can then use interpolation to find the value of the map at each of those locations.

The final step in this process is to remove the monopole and dipole from the maps. This is done for simplicity. When comparing maps to data, we will need to remove the monopole and dipole terms from the map. We chose to remove them when making the maps as well.

We display the maps we create using the orthview option in HealPy. The center of the map represents the north pole of the sphere. Some examples of the maps we generated with a spherical cap of 0.3 rad are shown in figure 6.

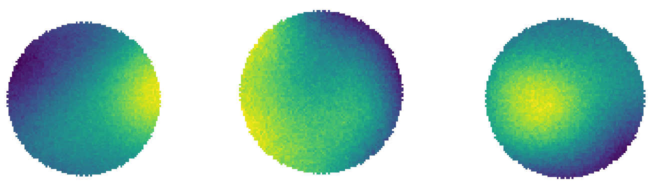

FIG. 6. Maps of spherical cap of CMB generated from calculated three-dimensional power spectrum.

Once we have found this power spectrum, we can use it to produce simulations of the CMB and compare those maps to those generated from the angular power spectrum. We generated the maps from the angular power spectrum by using the synfast method in the HealPy package. After simulating the entire map, we selected the spherical cap. Some examples of these maps are shown in figure 7.

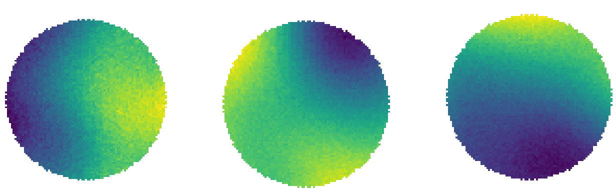

FIG. 7. Maps of spherical cap of CMB generated from observed angular power spectrum.

The two sets of maps seem to be similar. Each is random, and there do not seem to be any patterns in one set compared to the other. However, to prove that our method produces the correct maps, we must analyze and compare the statistical properties of each set of maps.

VII. STATISTICAL TESTS

Once we have created the maps, we need to test their statistical properties to ensure that they are identical to the maps produced by the standard HealPy method. There are four tests that we perform on our maps. If we pass each test, we are confident that the statistical properties of the two maps are the same. We performed each statistical test on both sets of maps, those created by the standard HealPy method and those created by our method. We compared the results both to each other and to the expected results.

Before we begin the statistical tests, we first remove the correlations in the maps. We did this by first calculating the eigenvalues of the covariance matrix. We then take the dot product of the map with the eigenvalues. This removes the correlations between the map values.

The four statistical tests that we performed are variance, covariance, χ^2, and Kolmogorov-Smirnov.

Each example shown below is for a spherical cap of 0.3 radians.

A. Variance

Because we have removed the expected correlations from the maps before conducting these tests, we expect

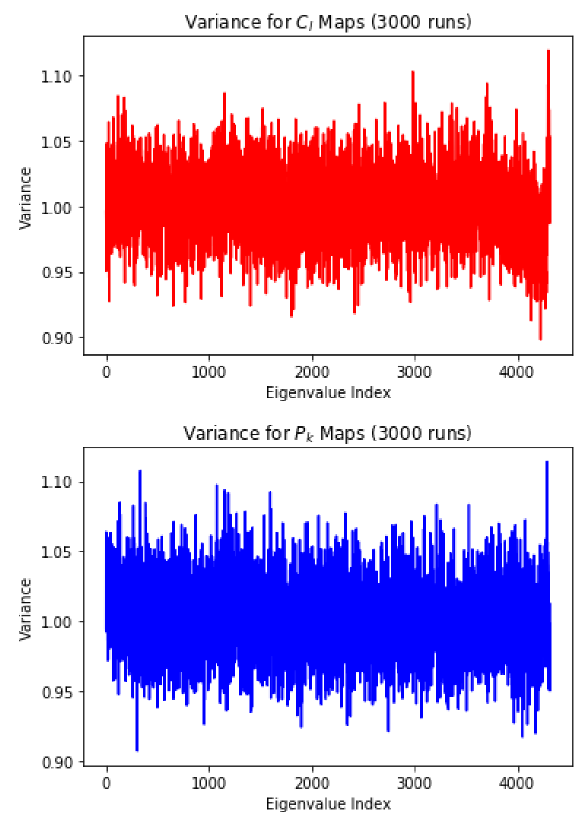

FIG. 8. Plot of the variances for 3000 maps. The top panel shows the variances for maps made with the angular power spectrum and the bottom panel shows the variances for maps made with the calculated 3-dimensional power spectrum. Both plots have a mean of 1.

the variance to be 1. This is the case for both sets of maps. As can be seen in the figure 8, the average variance of each map is 1. The variations are due to the limited number of maps. We expect the standard deviation of the variances $\sigma = \frac{1}{\sqrt{N}}$, where N is the number of maps. In this case, $N = 3000$, so we expect $\sigma = 0.018$. This matches the results of these variance tests. This means that we can say that the variances of the two sets of maps are both the same and match the expected results.

B. Covariance

The next test we performed involved the covariances of the maps. The maps have covariances described by the covariance matrices Γ that we calculated. However, as we removed this covariances before doing this test, we expect the covariances of each set of maps to be 0. To test this, we calculated the covariances for each set of maps and plotted a histogram of the covariance values. The histograms for each set of maps are shown in figure 9. Each is centered at 0 and falls off quickly. We expect that the standard deviations should be 0.032 because we have 1000 maps of each type. This matches the histograms, so we can say that the two sets of maps have the covariances we expect.

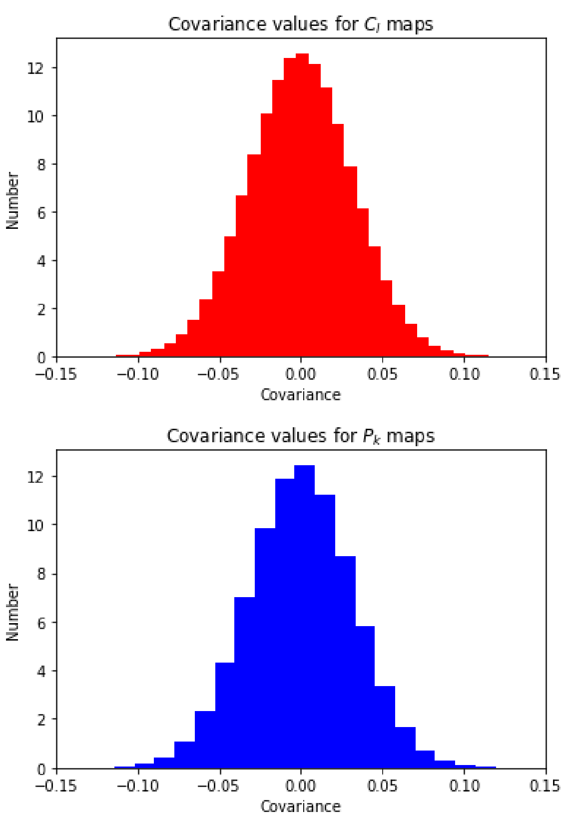

FIG. 9. Histogram of the covariances for 1000 maps. The left panel shows the covariances for maps made with the angular power spectrum and the right panel shows the covariances for maps made with the calculated 3-dimensional power spectrum. In both cases, the distribution of covariances is centered at 0.

C. χ^2 Test

The next statistical test we performed on the maps was the χ^2 test. The definition of the normal distribution is that if we have a list of random numbers $\vec{x_j} \sim N(0, \Gamma)$, then the likelihood of measuring $\vec{x_j}$ given Γ is

$$P(\vec{x_j}|\Gamma) = \frac{1}{(2\pi)^{\frac{N}{2}} det\Gamma^{\frac{1}{2}}} e^{-\frac{1}{2}\vec{x_j}^T \Gamma^{-1} \vec{x_j}}$$

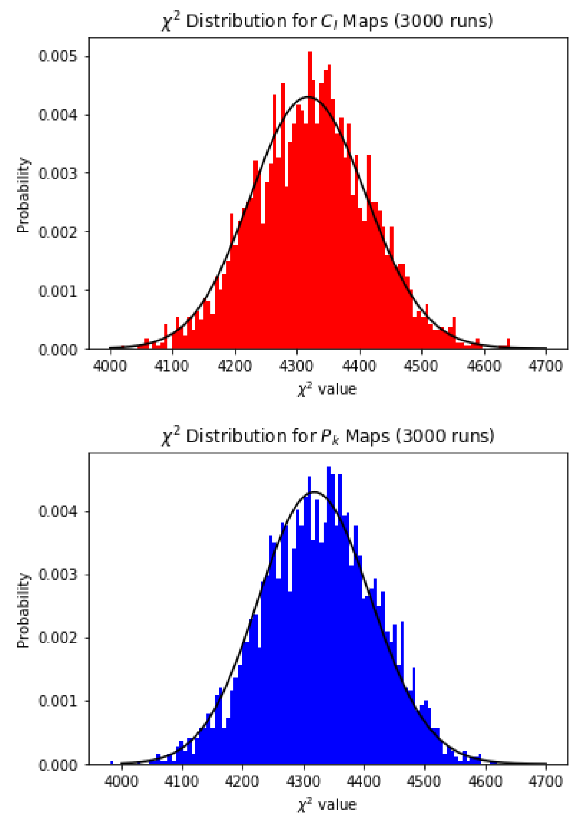

FIG. 10. χ^2 distributions for maps made with methods described in Section . The left panel shows the χ^2 distribution for maps made with the angular power spectrum and the right panel shows the χ^2 distribution for maps made with the calculated 3-dimensional power spectrum. The black curves in each are the theoretical distributions for these maps. These two distributions are very similar to each other, providing evidence that the two methods produce maps with identical statistical properties.

The $\vec{x_j}$ values are the values in the maps. From this, we can get an expression for the χ^2 value of the maps

$$\chi^2 = \vec{x_j}^T \Gamma^{-1} \vec{x_j}$$

If our map values were drawn from the correct distribution, then we expect the χ^2 values to be drawn from a χ^2 distribution with N degrees of freedom. N is the number of random numbers, which in this case is the number of pixels in the map. Figure 10 shows the χ^2 distributions for both sets of maps. The black curves represent the χ^2 distribution for 4320 degrees of freedom. Both histograms match the expected χ^2 distribution, which suggests that the map values are drawn from the correct distributions in both cases.

D. Kolmogorov-Smirnov Test

Our final statistical test of the maps is the Kolmogorov-Smirnov test. This test examines the cumulative distribution functions of each set of map values. Because we removed any correlations between the map values before doing this test, we expect that the cumulative distribution functions should match that of a set of random numbers $\vec{x} \sim N(0,1)$. The Kolmogorov-Smirnov test finds the maximum discrepancy between the two cumulative distribution functions [5]. From this, it determines a p value for the null hypothesis that the two sets of random numbers were drawn from the same distribution. A higher p value signifies that the two sets of random numbers were drawn from the same distribution.

Plots of the cumulative distribution functions are shown in figure 11. The green curves in each plot represent the cumulative distribution of a set of random numbers $\vec{x} \sim N(0,1)$. The cumulative distribution functions match for each set of maps, and the p values suggest that both maps were drawn from the correct distribution. Instead of comparing one distribution against the normal, we can compare the two sets of maps against each other as well. The results of this test further show that the two sets of maps are drawn from the same distribution.

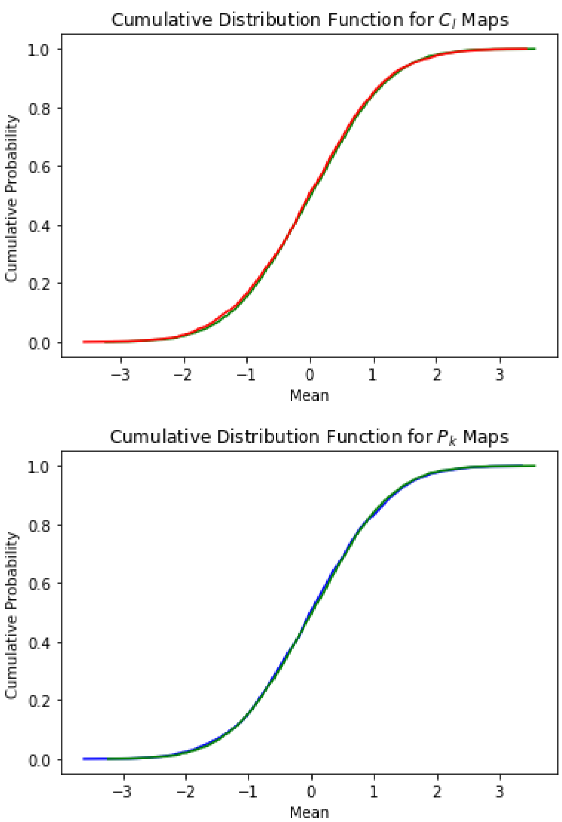

FIG. 11. Cumulative distribution functions for maps. The left panel shows the cumulative distribution function for maps made with the angular power spectrum and the right panel shows the cumulative distribution for maps made with the calculated 3-dimensional power spectrum. These two functions are very similar to each other, providing evidence that the two methods produce maps with identical statistical properties.

VIII. RUNTIME

We have now shown that our method produces maps of the cosmic microwave background that have the same statistical properties as those generated from the standard HealPy method. We now need to show that our method is faster than the Healpy method. We can do that by analyzing the number of operations that each method must perform to make a map.

The HealPy runtime is known. It is proportional to the resolution of the map. The precise runtime is:

$$T = O\left(N_{pix}^{\frac{3}{2}}\right)$$

where N_{pix} is the number of pixels in the map. We can also express this as

$$T = O\left(\left(12N_{side}^2\right)^{\frac{3}{2}}\right)$$

because $N_{pix} = 12N_{side}^2$. N_{side} is the resolution of the map.

We can then compare this runtime to the runtime of our method. The two parts of our method that contribute most to the total runtime are the generation of random numbers and the two-dimensional Fourier transforms. The generation of N_z random numbers takes $O(N_z^2)$ time, and we have to do it N_{xy}^2 times for a total time of $O(Nxy^2 N_z^2)$. We then must do the two-dimensional fast Fourier transform of these random numbers in the x and y directions. Fast Fourier transforms take $O(nlog(n))$ time in one dimension. We do the two-dimensional FFT N_z times, so the total runtime for that part of the process is $O(N_z(N_{xy}^2 log(N_{xy})))$. This gives a total runtime of our method of

$$T = O(N_{xy}^2 N_z^2 + N_z(N_{xy}^2 log(N_{xy})))$$

We do not need to include the generation of the three-dimensional power spectrum or the creation of the N_{xy}^2 covariance matrices in the runtime calculation because these steps can be done once and stored. The same power spectrum and set of covariance matrices can be used to produce as many maps as needed for some given set of parameters.

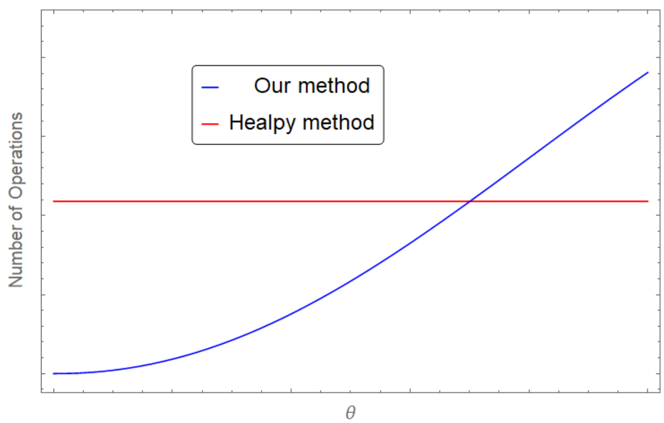

FIG. 12. Plot of runtimes of Healpy method (red) and our method (blue) for a given resolution. Our method is faster for small enough angles but becomes less efficient at larger spherical caps.

A plot of the runtimes of the two methods is shown above in figure 12 for a given resolution. For a constant map resolution, the runtime of the HealPy method is constant because it depends only on the resolution. On the other hand, for our method, an increase in the size of the spherical cap causes an increase in runtime because the size of the box and the number of points that must be simulated increases.

IX. CONCLUSIONS

In conclusion, we were able to successfully find a three-dimensional power spectrum $P(\vec{k})$ whose correlation function matched that of the observed angular power spectrum $C(l)$ over the spherical cap we wanted to simulate. We were then able to use this generated power spectrum to create maps of a spherical cap of the cosmic microwave background. These maps had the same statistical properties as the maps produced by the standard method of creating CMB maps. Our method was also faster than the existing method for reasonable sizes of the spherical cap.

X. FUTURE WORK

We have shown that our method works for creating temperature maps of the cosmic microwave background. We can extend this process to maps of the cosmic microwave background polarization. The difference here is that polarization is two-dimensional vector whereas temperature is a scalar quantity. The current method of simulating these polarization maps involves using spin-weighted spherical harmonics. However, we expect that we can apply our method of using Fourier transforms to this problem as well.

XI. REFERENCES

[1] K.M. Grski, Eric Hivon, A.J. Banday, B.D. Wandelt, F.K. Hansen, M. Reinecke, and M. Bartelmann, "HEALPIX - a Framework for High Resolution Discretization, and Fast Analysis of Data Distributed on the Sphere,Astrophysical Journal, 622, 759 (2005).

[2] E.F. Bunn, "Calculating Cosmic Background Anisotropies and Implications,' in The Cosmic Background Radiation, NATO ASI Series C, vol. 502, edited by C.H. Lineweaver et al., p. 135 (1997).

[3] B. Ryden, *Introduction to Cosmology*, 2nd ed., Cambridge University Press (2017)

[4] M. S. Andersen, J. Dahl, and L. Vandenberghe. CVXOPT: A Python package for convex optimization, version 1.1.6. Available at cvxopt.org, 2013.

[5] *Kolmogorov-Smirnov Test.* www.physics.csbsju.edu/stats/KS-test.html.